Pith:2026.000003v1 [physics.hist-ph] 13 Oct 2026

Proof of an Axiom About the World

When a physical axiom is proven, which recognition axioms are, and what Gödel permits

Jonathan Washburn
Recognition Physics Institute

Abstract

An axiom about the world is usually said to be tested by its consequences and never proven. One axiom is an exception. Take a world to be some states, steps that write down that two states differ, and acts that compare states. The sentence "something is distinguished from something" is made true by any act that considers it: to assert, deny or question it is to hold it apart from its negation. Identify that holding apart with one of the world's own comparisons, the one who asks being part of the world, and the sentence holds in every world in which anything is asked, with no observation involved. Every proof of an axiom below rests on this identification. An axiom is therefore *proven* when it follows from sentences that asking itself makes true, and *supplied* when some world satisfies those sentences and violates it. Three sentences qualify: something is distinguished; no finite record entertained as complete is complete (lists every distinction that occurs); a comparison has both its terms present. They build a floor of distinctions, records and counts. Sorting the premises of Recognition Science by this criterion, the floor is proven, and the cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ of comparing two counts is proven for any ledger, a record kept as a running total in which a comparison of a with b writes the two ratios a/b and b/a and its cost is the excess of their sum over 2. That the world's records form a ledger is taken from earlier work, since asking alone does not make it true. Beyond the ledger, two premises are supplied, each with a world where everything asking guarantees holds and the premise fails: a new level, the size of a completed run of the ledger's steps, is built from two different lower levels (which gives the golden ratio), and a completed act cannot be undone by a motion that writes nothing (which, once an act and its record are read as two linked loops in space, is equivalent to three spatial dimensions). Gödel's theorems bound what an axiom set proves about itself. They touch neither the criterion nor the finitely stated quantities here, once deciding every sentence is separated from fixing every quantity.

1 The question

Peano's axioms are the accepted description of the natural numbers, and no one asks that every theorem of arithmetic be derived before saying so. Euclid's postulates were the accepted description of space for two thousand years on the same terms. An axiom set can be correct before its consequences are exhausted. So there are two separate questions about any candidate foundation for physics: whether its axioms are true of the world, and whether all of physics follows from them.

The first question has a standard answer: axioms of a physical theory are not proven, they are tested. One picks postulates, derives consequences, and measures. If the measurements agree, the postulates are retained; if not, one of them is changed. On this account "true of the world" is never a theorem, only a standing record of unfalsified consequences, and it always carries the scale at which the consequences were checked.

That account is wrong as a universal statement, and the counterexample is not exotic. To state it we need a precise idea of a world. Here a *world* is some states, steps called *posts* that take one state to the next and write down that the two differ, and acts of comparison that take two states, or two records of earlier states, and report whether they differ. A question can be asked only in a world rich enough to contain a comparison. A lamp is a toy world. It has two states, off and on; its one kind of post is flipping the switch; and a comparison of off with on reports that they differ. A lamp with a single state is another world. Nothing in it differs from anything, so no comparison in it can report a difference, and no question can be asked in it. Section 3 gives the definition in full.

Now consider the sentence

$$\mathsf{Dist}: \quad \text{something is distinguished from something.}$$

It is about the world. The one-state lamp would make it false, and such a structure is consistent. Yet no observation is needed to establish it. Any act that considers Dist, whether to assert, deny, or question it, holds the sentence apart from its negation, and that is a distinction. We take the one who asks to be part of the world the sentence describes, so that entertaining a sentence is one of that world's own comparisons. That is what "a world in which something is asked" means; Section 3 makes it exact (Lemma 3.2), and every proof of an axiom in this paper rests on this identification. The act is then one of that world's events, and the distinction occurs in it. So Dist is true in every world in which it is so much as considered. It is not a necessary truth of logic, since the one-state world is possible. But nobody can ask the question from inside the one-state world, because asking is itself a distinction. Call a sentence with this property *performatively necessary*: considering it makes it true, as "I am speaking" is true whenever it is said. Dist is secured without measurement, and it carries no scale: there is no regime of measurement at which it was checked and beyond which it might fail.

The standard account was formed on cases: the fifth postulate, the laws of motion, the constancy of the speed of light. Each is of the kind that can fail at a new scale, and a rule formed on such cases came to be read as a rule about all cases. Dist is a case it was not formed on. Once one such axiom exists the account has to be replaced by a criterion: which axioms about the world can be proven, and which must be *supplied*, that is, assumed and then tested, because some conceivable world satisfies everything asking guarantees and violates them? Section 3 states the criterion.

The rest of the paper applies it to the *recognition axioms*, the premises of Recognition Science, the framework this paper belongs to. Recognition is the act of telling two things apart, together with the record that act leaves. The recognition axioms are the short list of premises from which the framework derives a particular cost for comparing two quantities, the golden ratio, three spatial dimensions, and an eight-step period. The cost, written J, is the function $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ of the ratio x of the two quantities compared. A comparison of a with b can be read from either side, as the ratio a/b and as the ratio b/a.

The cost is read off a *ledger*, the running account of the entries a world's steps and comparisons write; Section 4 defines it exactly. Two kinds of entry matter here: a post writes that two states differ, and a comparison of two counts a and b (whole numbers, such as the sizes of two records) writes the two ratios a/b and b/a. The ratios are a balanced pair, balanced in the multiplicative sense that their product is 1 (their logarithms cancel, as debits cancel credits in double-entry bookkeeping), and the ledger adds them up. A comparison of a with itself writes 1 and 1 and totals 2; the total for a against b exceeds 2 by $2J(a/b)$. For two against one the entries are 2 and $\frac{1}{2}$, which total $\frac{5}{2}$, an excess of $\frac{1}{2} = 2J(2)$, so $J(2) = \frac{1}{4}$. The paper then asks what Gödel's theorems do to the result.

Arc. Earlier papers carry the line this one belongs to. They establish five results, and Section 4 uses all five. *Arithmetic:* iterating a single distinction generates the natural numbers and their arithmetic [1]. *Logic:* the logic needed to state that construction is itself made of distinctions [2]. *Dimension:* one sentence, space keeps the record of a completed act, forces space to have three dimensions [3]. *Cost:* a short list of conditions on the cost of comparing two counts leaves exactly one cost, the function J above [6, 7, 4]. *Ledger:* a ledger that conserves what it records fixes the scale, the cost, and the ratio between levels together, and dropping any one of its hypotheses loses the conclusion [8]. This paper uses a weaker ledger, without those structural axioms: it writes one event at a time, each a balanced pair of entries [5], and fixes the cost but not the ratio between levels (Section 4 defines it, and Section 5 says why the stronger one is not assumed). The sort of each axiom into proven or supplied is the map of what remains to be derived.

Terms. Table 1 is for look-up. Each term is also explained in the text where it is first used.

What the paper does. (i) It states the criterion for proving an axiom about the world, and takes the parallel postulate as the first example of a supplied axiom. (ii) It sorts the recognition axioms under the criterion: the floor is forced, the cost is forced given a ledger, and the join and kept are supplied, each with a countermodel. (iii) It shows that the kept premise is equivalent to $D = 3$, so supplying it from experience supplies the conclusion. (iv) It proves a ceiling: no sentence about states, posts, records and comparisons that holds wherever something is asked forces the magnitude reading, the join or kept; and it measures how far a performative argument does reach at each. (v) It separates syntactic from physical completeness, as Gödel's theorems require of the phrase "theory of everything," and defines a theory of everything with two clauses, matching the two questions above, with where each stands.

2 Two toys

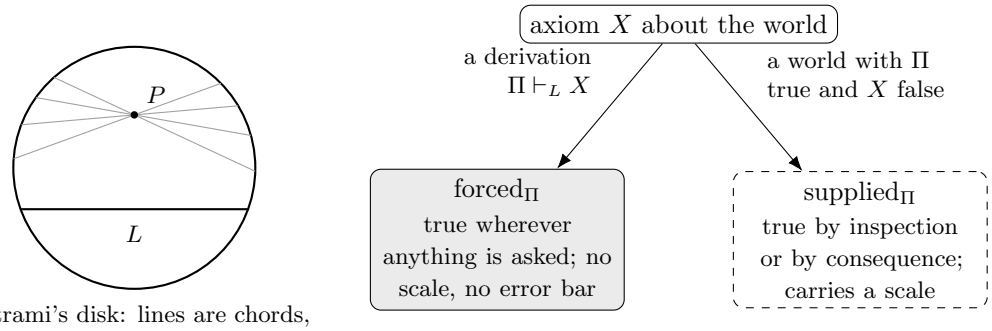
The criterion of Section 3 has two halves, and each has a classical case. The first shows how to prove that an axiom is not forced by the others. The second shows how a sentence can be made true by being said.

The parallel postulate. Euclid's fifth postulate is the one that never looked like the others. For two millennia geometers tried to derive it from the remaining four and failed. Lobachevsky and Bolyai developed the geometry in which it fails; Beltrami [17] settled the question by exhibiting a model, the disk of his *Saggio*, in which the other postulates hold and the fifth fails. That model proved that the fifth postulate does not follow from the others and named the exact respect in which physical space could be non-Euclidean, as general relativity later described it. Beltrami's disk is a world in the sense of Section 3: its points and lines are distinguishable states and its incidence tests are comparisons. Posts play no part in the example. The parallel postulate is a sentence about the geometry laid over those states, as the premises of Section 4 are sentences about the extra structure (the readings) laid over the floor. The disk is a countermodel: a world in which some sentences hold and another fails.

The cogito. Descartes' "I think, therefore I am" has been read as an inference and as a performance. Hintikka [18] showed the second reading is the one that works: the sentence "I do not exist" is not false because of what it says but because saying it is an act of the speaker whose existence it denies. The truth of the conclusion is secured by the occurrence of an act, not by the content of a premise. In Descartes' case the act is anchored to the speaker. In the case this paper needs, the anchor is the act itself: entertaining a sentence is holding it apart

Term	Object	Plain gloss
performatively necessary	a sentence P such that any act of entertaining P makes P true	you cannot ask it without making it so
Π	the three performatively necessary sentences listed in Section 3: Dist, Iter, Weak	what asking guarantees
world	a structure with a set of states, posts between them, and acts of comparison between states	anything a question can be asked in; a lamp with a switch is one
post	one step from a state to the next, writing down that the two differ	a flip of the switch, noted
record, stage	the posts up to a given stage; the stage is the position in the order of posts	what has been written down so far, and when
ledger	a record kept as a running total of balanced pairs of entries, one event per tick; taken from the cited ledger construction, and not made true by asking	an account book
forced	an axiom X with $\Pi \vdash_L X$ in the background logic L, or with Π together with a definition the sort names	follows from what asking guarantees (and, where named, from what a ledger is)
supplied	an axiom X with a countermodel relative to Π	what asking guarantees does not reach; true of our world, if at all, by inspection or by consequence
countermodel	a world in which every sentence of Π holds and X fails	the proof that X is supplied
open	an axiom with neither a derivation nor a countermodel in hand	not yet sorted
floor	the distinguished states, the extension step, the counts and the rationals, generated by iterated distinction	what Π and the background logic build
reading	an extra layer of structure added to a world; there are three (Section 4): a cost on ratios of counts, a level function, a kinematics of placements	an extra the floor does not contain
signature	the list of symbols a sentence may use; the floor's is states, posts, records, comparisons, stage order	the vocabulary a sentence is written in
asking history	a finite sequence of acts in the floor's signature: post, retain, compare	what an asker does, as a script
admissible cost	continuous, $F(1/x) = F(x)$, $F(1) = 0$, strictly convex in $\ln x$	a fair price for a ratio
cost J	$J(x) = \frac{1}{2}(x + x^{-1}) - 1$ for a ratio x of two counts	how far a comparison is from balance
pass, size, level	a pass is a completed stretch of posts and its size is the number of posts in it; the levels are the sizes of successive passes	a completed run of the ledger's steps, and how long it is
join	making a new level from a pass and the record that registers it	the new run is the old one followed by its record
kept	a completed act, with its record threaded through it, cannot be undone by any motion that posts nothing	two linked chain links stay linked unless one is cut
placement, post-free motion	how the two loops of an act and its record sit in space, linked or apart; a movement that writes no entry in the ledger	where they are, and moving them for free
tick, period	a tick is one post; the period is the number of ticks in a pass through all states of a cell	one step, and one full cycle
seam	a reading the floor lacks, with a countermodel relative to Π (magnitude, join, kept)	a place where the world could have been otherwise
physical completeness	every physical quantity is a consequence of the axioms	nothing measurable is left free

Table 1: Terms of art. Everything else is ordinary mathematics.



Beltrami's disk: lines are chords, and through P run many lines missing L

neither object in hand: open

Figure 1: The two halves of the criterion. Left, the countermodel half: in Beltrami's disk the first four postulates hold and the fifth fails, so the fifth is supplied and could fail of physical space. Right, the sort every axiom about the world goes through: a derivation from the performative sentences Π makes it forced, a world satisfying Π and violating it makes it supplied.

from its negation, and that holding-apart is already an instance of what Dist asserts. No "I" is needed; only that some act of entertaining occurs.

The two toys are the two halves of the criterion (Figure 1). A countermodel shows an axiom is supplied; a performance shows it is forced. Here Π names the three sentences that asking makes true (Section 3).

3 Proof of an axiom about the world

Fix a background logic L: the ambient mathematics in which proofs are written, here classical mathematics with the real numbers. All proofs in this paper are proofs in L; Section 7 returns to what it means that every proof is a proof in L.

The section does three things. It defines a world and shows that entertaining a sentence is an act inside the world it is about. It names the three sentences that every such act makes true. And it states the criterion that sorts any axiom as forced or supplied.

Definition 3.1 (world)**.** A *world* is a structure with a set of states, a relation of posting between states, and acts of comparison, each act taking two records (a single state counts as a one-entry record) and returning an outcome. In the lamp of Section 1 the states are off and on, a post is a flip, and a comparison asks whether two states or two records agree. Posts occur in sequence; the posting relation is the stage order, and a record at a stage is the set of posts up to that stage. An act of comparison occurs at a stage later than every post it takes as a term. A *distinction* is a comparison whose outcome is that its terms differ; a post registers a distinction at the stage at which it occurs; a record is *complete* if every distinction occurring in the world is registered by one of its entries. Sentences about the world are sentences in this signature, that is, in the vocabulary of this definition. A world in which an act of comparison occurs is one in which a question can be asked.

A world may also carry, each possibly absent, three added structures called *readings*: a cost on ratios of counts, a level function, and a kinematics of placements (a rule for which placements of an act and its record in space can be carried to which by a motion). Section 4 introduces each where it is used. The readings add vocabulary of their own (cost, level, placement); a sentence that uses it is a sentence about worlds that carry that reading, and the sentences of Π use none of it. A world's cost, when present, is a function of the outcomes its comparisons return: a cost prices what a comparison returns, and a cost of something no comparison

returns is not a cost in that world.

The definition says nothing about sentences being entertained. Two points settle how an act of asking sits in a world. First, to entertain a sentence P is to have P and its negation as records in the world: each is encoded as a state of the world, and posting it is the step to that state, which writes that it differs from the state before, so each is an entry of the world like any other. Second, the one who entertains is inside the world being asked about, so what that one does is done in that world. The first is a modelling choice: to entertain P is to keep both alternatives in view, and in a world the way to keep something is to post it. Together they make entertaining P a short history in the world: post an encoding of P, post an encoding of its negation, compare the two (Section 6 calls such a script an asking history). The next lemma is the step everything else rests on: it shows that this history is itself one of the world's comparisons, so the question is asked in the world it is about.

Lemma 3.2 (entertaining)**.** *An act that entertains a sentence P in a world, whether asserting, denying, or questioning it, is a comparison in that world whose two terms are P and its negation and whose outcome is that they differ.*

Proof. By the first point, each way of taking an attitude to P has P and its negation as two records. To assert P is to select P over its negation; to deny P is to select the negation over P; to question P is to hold the two open as alternatives. Each requires P and its negation as two terms, and each returns that the two are not one: an assertion or denial by selecting one of them, a question by leaving both standing. In the definition of a world, a comparison is an act on two records that returns an outcome. Each of the three attitudes is such an act, and its outcome is fixed: to select one of two alternatives, or to hold two alternatives open, is to register that there are two of them, and to register that two records are two is the outcome "differ." An act that found P and its negation to be one could neither select between them nor hold them open as alternatives. The lemma rests on this identification: entertaining a sentence is comparing it with its negation. By the second point the act occurs in the world in which the sentence is entertained, so the comparison is in that world. □

The questions of this paper are asked in the world they are about: anyone who asks whether Dist holds of the world has, in that world, performed the comparison of the lemma.

Definition 3.3 (performative necessity)**.** A sentence P about the world is *performatively necessary* if every act of entertaining P, whether asserting, denying, or questioning it, makes P true in the world containing the act.

The model case is "I am speaking," which is true whenever it is said. Three sentences are performatively necessary, and they are the whole premise set of this paper. Dist and Weak need an asker to be true; Iter holds of every world by the order of stages, with or without one, and so every act that entertains it finds it true. In plain words: Dist, something is distinguished from something; Iter, no finite record entertained as complete is complete, because reading a record is a new event that the record does not list (the name recalls iteration: any record can be extended); and Weak, whenever two things are compared both are present to the comparison, each term persisting from the stage at which it was posted until the comparison occurs. The third is called Weak because it is the weak form of a premise met later. The strong form, that space keeps a completed act, is the kept premise of Section 4; Section 6 compares the two. Each proposition below shows how an act of asking makes one of them true.

Proposition 3.4 (Dist)**.** *"Something is distinguished from something" is performatively necessary.*

Proof. Let an act entertain Dist. By Lemma 3.2 the act is a comparison of Dist with its negation whose outcome is that they differ, and a comparison with that outcome is a distinction. So a

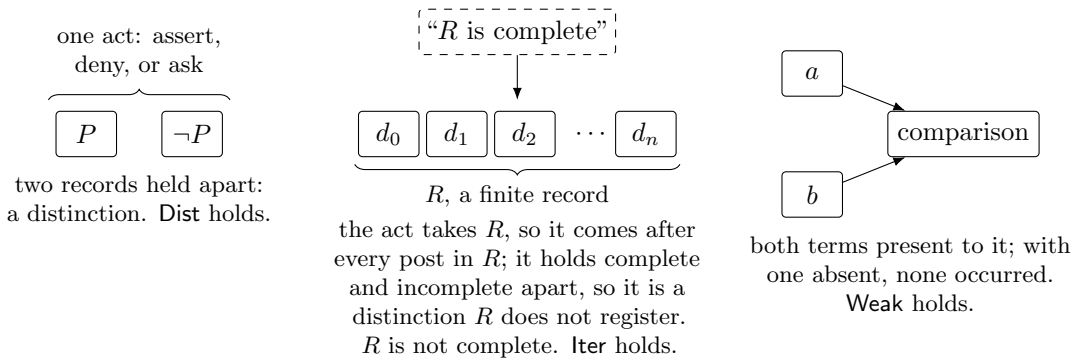


Figure 2: The three performative proofs. Each ends in an act, not in a premise: entertaining a sentence holds two records apart (left), an act about a record is a distinction outside it (middle), and a comparison that occurs has both its terms (right). None of the three needs an observation; the first and third are false in consistent worlds that no asker can be in, and the second has no falsifying world at all.

distinction occurs in the world containing the act, which is what Dist says. □

Take a record of everything that has happened up to noon, and let someone read it at one o'clock. The reading is an event, and the record does not list it. So the record was complete through noon and is not complete of the world. The next proposition says the same of every record one asks about.

Proposition 3.5 (Iter)**.** *"No finite record of distinctions that is entertained as complete is complete" is performatively necessary. A record is* entertained as complete *when some act entertains the sentence "this record is complete"; the sentence speaks of such records only.*

Proof. Let R be the record at some stage s, and let an act entertain "R is complete." By Lemma 3.2 the act is a comparison whose terms are built from R, and a term built from R can be posted only after every post of R (a record cannot list a reading of itself), so the act occurs at a stage later than s; and its outcome is that its terms differ, so it is a distinction. Every entry of R registers a distinction at a stage up to s, and the act occurs later, so no entry of R registers it. A distinction occurs in the world that R does not register, and R is not complete. For a record not so entertained the sentence says nothing, so in a world where no record is so entertained it holds vacuously. □

Iter is potential, not actual: it says every record one asks about can be extended, not that an infinite record exists. That is all the act delivers, and Section 4 says what the background logic adds. What it rules out is a record that contains its own reading, and the reason is the stage order: a reading takes the record it reads as already posted. A structure in which a record could list its own reading would have no such order, and is not a world in the sense defined above. Complete means complete of the world, not complete through a stage.

Proposition 3.6 (Weak)**.** *"A comparison occurs, and both of its terms are present to it" is performatively necessary.*

Proof. Every act that entertains a sentence is a comparison of the sentence with its negation (Lemma 3.2), so any world containing such an act contains a comparison. A comparison with one term present is not a comparison, so both terms of that comparison are present to it; in particular each term survived from the stage at which it was posted until the comparison. □

Figure 2 draws the three proofs, and two things stand out. First, Dist and Weak are contingent: a one-state world falsifies Dist, a world in which no comparison ever occurs falsifies Weak, and each is a consistent structure. Iter is stronger: it has no falsifying world, because in any world a comparison follows the records it takes, so no record contains its own reading. It holds vacuously where no record is entertained as complete, and every act that entertains it finds it true, so it is performatively necessary without needing an asker; the act supplies one instance, that the asker's own record extends. Only Dist and Weak need an asker. Second, all three are secured without observation. No asker can be in a world that falsifies one of them: a world that falsifies Dist has no distinction and a world that falsifies Weak has no comparison, while asking is itself both. The conclusions are analytic given that an act of comparison occurs, and whether one occurs is a fact about the world, not about language. "Proven" therefore means true in every world that contains an act of asking, and the asker's own act supplies that premise. Nothing observed at any scale can bear against them: an observation is a comparison, and every world containing a comparison satisfies all three. That is the only sense in which any axiom about the world can be proven, and it is enough.

Definition 3.7 (Π; forced, supplied, open)**.** $\Pi = \{\mathsf{Dist}, \mathsf{Iter}, \mathsf{Weak}\}$. Let X be a sentence about the world: a sentence in the world signature, or in the vocabulary of a reading added to it (as the premises of Section 4 are), and not a theorem of L alone. X is *forced relative to* Π, written forced_Π, if $\Pi \vdash_L X$. X is *supplied relative to* Π, supplied_Π, if some world satisfies every sentence of Π and violates X. X is *open* if neither object is in hand. When the subscript is dropped, Π is the set above. A verdict may also be relative to Π together with a named definition D: X is *forced given* D if Π and D prove it. The sort names each definition it uses, and the only line sorted this way is the cost, forced given the definition of a ledger (Section 5). The join and kept premises are sorted relative to Π alone, because the ledger as defined says nothing about levels or space. A sentence in a reading's vocabulary is evaluated only in worlds that carry the reading; a sentence of the world signature mentions no reading, so adding a reading to a world leaves its truth value unchanged.

In plain words, an axiom is forced when it follows from what asking guarantees, and supplied when some world in which asking happens makes it false; such a world is a *countermodel* to the axiom. Beltrami's disk is a countermodel to the parallel postulate. The two verdicts cannot both hold of one axiom.

Proposition 3.8 (the two are exclusive)**.** *No sentence is both* forced_Π *and* supplied_Π.

Proof. L is sound for worlds: whatever Π proves holds in every world in which Π holds. A world in which Π holds and X fails therefore shows $\Pi \nvdash_L X$. □

Both verdicts are relative to the Π in hand. A larger set of performatively necessary sentences can move a sentence from open or supplied to forced; nothing moves a sentence the other way, since a derivation stays a derivation. So "supplied" never means "must be found by experiment"; it means that the present performative basis does not reach it, and it names, by the countermodel, what a further performative sentence would have to rule out.

Proposition 3.9 (what forced buys)**.** *With entertaining read as a comparison inside the world (Lemma 3.2), a forced axiom is true of every world in which anything is entertained, in particular of this one; no measurement bears on it, so it carries no scale, no error bar, and no caveat about regimes not yet observed. A supplied axiom is true of this world, if at all, either by inspection at the scales at which we inspect or through its consequences, and either route carries the scale at which it was checked.*

Proof. The first sentence is the definition of forced together with Propositions 3.4 to 3.6: the act of reading this line is a comparison occurring in the world, and each sentence of Π holds in

any world containing one. The second is the standard account of Section 1, which applies to the supplied axioms: a supplied axiom has a countermodel, so no derivation from Π reaches it and only evidence at the scales checked can bear on it. □

Corollary 3.10. *The parallel postulate is supplied. Beltrami's disk, read as in Section 2 (states its points and lines, comparisons incidence tests), becomes a world in which incidence questions are asked once an asking act is added to it. That act posts its two terms, which gives the world the stages and records that* Iter *speaks of, and it changes nothing about the disk's points and lines, so the other postulates still hold and the fifth still fails. By Propositions 3.4 to 3.6 every sentence of* Π *holds in it; and it violates the postulate, a sentence about its geometry.* Dist *is forced. This is why the parallel postulate could fail of physical space and* Dist *cannot.*

The criterion turns "is this axiom true of the world?" into a search for one of two objects: a derivation from Π, or a countermodel. For the recognition axioms every line has one or the other, none is open, and the supplied lines are exactly the places the world could have been otherwise.

4 The recognition axioms

Section 3 gave the test; this section lists what is to be tested. The recognition axioms come in four groups: the floor that the three performatively necessary sentences build; premises about the cost of comparing two counts; premises about a hierarchy of levels; and the kept premise about space. Each group is stated as ordinary mathematics, followed by what it forces. Proofs that appear in the cited papers are cited; short ones are displayed. Section 5 then sorts each group into proven or supplied.

4.1 The floor

The *floor* is what the three sentences of Π build, in four steps.

1. *Two states.* Dist gives one distinction, a two-element set $\{0, 1\}$.

2. *The extension step.* Iter says no finite record entertained as complete is complete: the event the record omits is itself one new distinction, so the record extends by it. The step "add one more distinction" takes a record of n distinctions to a record of $n + 1$, and it is available at every stage whose record an asker entertains as complete. Π supplies only that much: no record so entertained is the last one.

3. *The counts.* The background logic L forms the completed iteration of the step: the set of results of applying it $0, 1, 2, \ldots$ times (its axiom of infinity supplies the set), L's idealization of the unending sequence, added on top of what asking forces. That set has a base element (no distinctions) and a successor operation (add one), satisfies the Peano conditions, and is characterized by them up to unique isomorphism [1]. It is the counting numbers, the arithmetic that iterated distinction generates; the integers and the rationals are built from them.

4. *Identity of states.* Two states that agree in every distinction are taken to be the same state, so a transition that posts nothing is no transition. This is how the floor individuates its states, a convention of the construction rather than a sentence of Π.

That is the floor: asking forces the first two steps, L supplies the third, and the fourth is the convention.

A *ledger* is further structure kept on the floor, taken from the recognition ledger of [5], a book of entries in the accountant's sense. Its states are written as finite bit patterns, one more bit

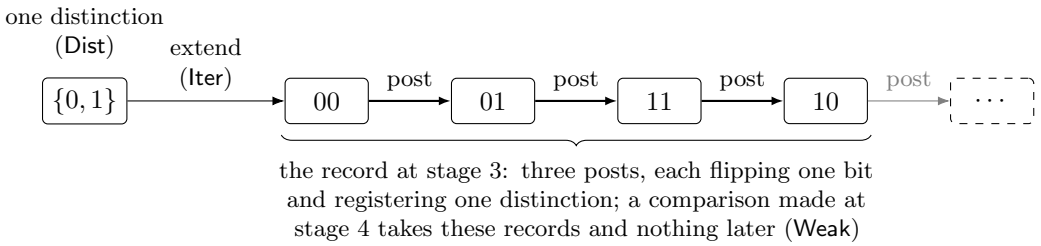


Figure 3: The floor and the ledger kept on it. Dist gives two states and Iter says the record one asks about extends; in the ledger, states are finite bit patterns and a post flips one bit. The record at a stage is the sequence of posts so far, and a comparison takes records already posted. Nothing here assigns a level to a record or places an act in space; those readings, and the pricing of a comparison, are taken up in the rest of this section.

for each extension step, listed so that neighbouring states differ in one bit, such as the patterns 00, 01, 11, 10 of Figure 3. A *post* moves to the neighbouring state, which flips one bit, and registers that the new state differs from the old. Four properties of the cited ledger make the posts single and the entries paired:

- *minimality of the record:* it holds no entry that registers nothing;
- *non-commutativity of events:* two events done in the other order make a different history;
- *conservation:* what one entry adds, another takes away;
- *pairwise locality:* an event touches two entries, no more.

The first two give at most one event per unit of time; the last two make each event a balanced pair of entries [5]. A posted state also persists until the next post; Weak gives this only for the two terms of a comparison, and the ledger gives it for every posted state. None of this is a sentence of Π, and asking does not force a world to keep a ledger (the world M_1 of Section 5 does not). This paper takes the four properties, single-bit posts, one event per tick and persistence between posts as given, and calls them together *the definition of a ledger* (Figure 3); a verdict "given a ledger" is given this definition. An entry is what an event adds to the book: when the event compares two counts a and b, its two entries are the comparison's two orientations, a/b and b/a, which multiply to 1. In a ledger the cost of a comparison is what it adds to the book above a balanced comparison, a comparison of two equal counts; Definition 6.10 states this exactly where the cost is derived.

4.2 The cost premises

The first reading prices a comparison. A *count* is one of the whole numbers the floor generates, such as the number of entries in a record. Comparing two counts a and b gives the ratio $x = a/b$, and a *cost* $F(x)$ is the price of that comparison: how far the two counts are from balance. Let $F : \mathbb{R}_{>0} \to \mathbb{R}$ be such a cost, and let $\kappa > 0$ be an amplitude gauge, a scale factor: replacing F by cF replaces κ by κ/c, so its value is a unit, and only its positivity is a fact (Section 5). The cost in that unit is $G = \frac{\kappa}{2}F$, and the premises below single out $G = J$, the function $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ of Section 1.

(C1) **Composition.** For all $x, y > 0$,

$$F(xy) + F(x/y) = \kappa\, F(x)F(y) + 2F(x) + 2F(y).$$

(C2) **Nondegeneracy.** $F(x) \neq F(y)$ for some $x, y > 0$.

(C3) **Monotone imbalance.** $t \mapsto F(e^t)$ is nondecreasing on $t \geq 0$.

(C4) **Calibration.** The cost of the ratio 2 (two counts against one) is $\frac{\kappa}{2}F(2) = \frac{1}{4}$.

In words: (C1), the law displayed above: two comparisons, with ratios x and y, are combined by forming the ratios xy and x/y; the cost of those two combined comparisons is fixed by the costs of the two separate ones, through a rule in which the two interact with strength κ; (C2) not every ratio costs the same; (C3) the cost does not fall as a ratio moves further from balance; (C4) two counts against one cost a quarter, in the unit $\kappa/2$, that is, $G(2) = \frac{1}{4}$. A world carries the *magnitude reading* Mag when its counts are compared as ratios by a cost that satisfies (C1) to (C4); this is the one sense in which the paper uses the name.

Premise (C4) fixes the exponent. In the one-parameter family that (C1) to (C3) leave, it is equivalent to the exponent being 1: the ratio of two counts is read as it stands, not as its square or its square root. It carries content, since different exponents are different costs, and it is not a unit; the unit is κ.

Most of (C1) is not a choice. Call a cost *admissible* if it is continuous, reciprocally symmetric ($F(1/x) = F(x)$, so a against b costs what b against a costs), zero at $x = 1$, and strictly convex in the logarithm. For every admissible cost there is a unique continuous P, the *combiner*, with $F(xy) + F(x/y) = P\big(F(x), F(y)\big)$: that the cost of the two combined comparisons depends on the two separate costs and on nothing else holds automatically, with no premise [7]. The reason is short. Write $f(t) = F(e^t)$; admissibility makes f even and strictly increasing in $|t|$, so $F(x)$ determines $|\ln x|$, and $F(xy) + F(x/y) = f(s+t) + f(s-t)$ with $s = \ln x$, $t = \ln y$ is unchanged by the sign of s or of t, so it depends only on $|s|$ and $|t|$, hence only on $F(x)$ and $F(y)$.

If P is a polynomial, a growth comparison between two recursions forces $\deg P \leq 2$ (by (C1), the costs $u_n = F(x^n)$ of the powers of a ratio obey $u_{n+1} + u_{n-1} = P(u_n, u_1)$; the growth comparison is carried out in the cited papers and not reproduced here), and symmetry with $P(u, 0) = 2u$ then forces $P(u, v) = 2u + 2v + \gamma\, uv$; a combiner of degree three or more admits no such cost at all [6, 7]. So for a cost in that class, what (C1) adds is that the combiner is polynomial and that its coupling γ, the coefficient of the product term (it is κ when (C1) holds), is nonzero. The nonzero coupling is not a unit. The admissible costs with polynomial combiner are exactly two families, the hyperbolic $\frac{1}{\gamma}(x^\lambda + x^{-\lambda}) - \frac{2}{\gamma}$ with $\gamma, \lambda > 0$ and the quadratic $a(\ln x)^2$ with $a > 0$, and the second is the contraction of the first as $\lambda \to 0$ with $\lambda^2/\gamma \to a$ [7]. The quadratic cost satisfies the composition law with coupling 0 and satisfies no instance of (C1) with $\kappa > 0$. Section 5 uses it as a countermodel.

The next theorem says that premises (C1) to (C4) leave exactly one cost. Section 6 reaches the same cost from the definition of a ledger, where the four premises become identities of the ledger total (Theorem 6.12).

Theorem 4.1 (the cost)**.** *Under (C1) to (C3), there is $c > 0$ with $\frac{\kappa}{2}F(x) = \frac{1}{2}(x^c + x^{-c}) - 1$ for all $x > 0$. Under (C4) as well, $c = 1$, so for all $x > 0$*

$$\frac{\kappa}{2}\,F(x) = J(x) := \frac{x + x^{-1}}{2} - 1.$$

Proof. Set $G = \frac{\kappa}{2}F$. Then (C1) becomes $G(xy) + G(x/y) = 2G(x)G(y) + 2G(x) + 2G(y)$. Set $H = G + 1$; the equation becomes $H(xy) + H(x/y) = 2H(x)H(y)$. Set $h(t) = H(e^t)$; then

$$h(s+t) + h(s-t) = 2h(s)h(t),$$

d'Alembert's equation on $\mathbb{R}$. At $s = t = 0$ it gives $2h(0) = 2h(0)^2$, so $h(0) \in \{0, 1\}$; if $h(0) = 0$ then $t = 0$ gives $2h(s) = 0$ for all s, contradicting (C2); so $h(0) = 1$. At $s = 0$ the equation gives

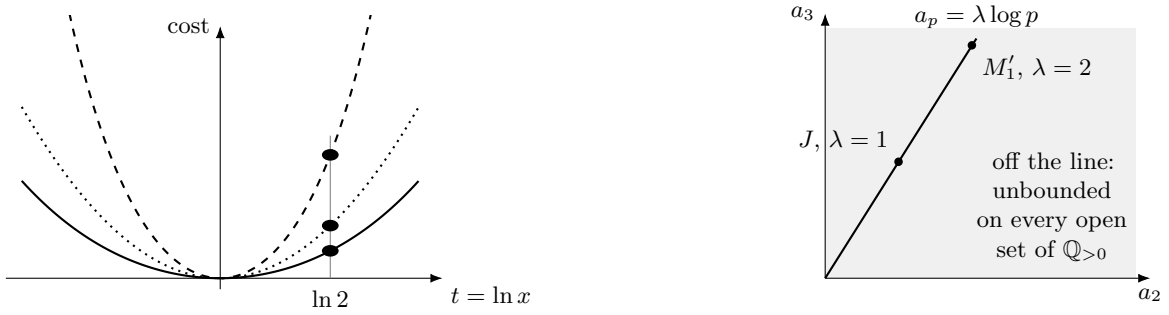


Figure 4: Left, three costs in the logarithmic variable: J (solid); $\frac{1}{2}(x^2+x^{-2})-1$, the exponent-2 member, which is the countermodel M_1' of Section 5 (dashed); and the quadratic $(\ln x)^2$, the countermodel M_1'' (dotted). Calibration (C4) reads the value at $x = 2$, marked by the dots. Right, the readings on the positive rationals, shown in the slice of the two weights at the primes 2 and 3. Every point of the plane is a solution of (C1); only the line $a_p = \lambda \log p$ extends regularly to the reals, and one limit condition picks the point $\lambda = 1$ on it.

$h(t) + h(-t) = 2h(t)$, so h is even. By (C3), h is nondecreasing on $[0, \infty)$; even and monotone on a half-line, h is Lebesgue measurable, and measurable solutions of d'Alembert's equation are continuous [19, Ch. 13]. The continuous solutions with $h(0) = 1$ are $h \equiv 1$, $h(t) = \cos(ct)$, and $h(t) = \cosh(ct)$ with $c > 0$ [4]. The constant fails (C2); cos is not nondecreasing on $t \geq 0$; so $h(t) = \cosh(ct)$ and $G(x) = \frac{1}{2}(x^c + x^{-c}) - 1$. Then $G(2) = \frac{1}{2}(2^c + 2^{-c}) - 1$ is strictly increasing in $c > 0$ and equals $\frac{1}{4}$ at $c = 1$, so (C4) gives $c = 1$. □

In outline, composition turns the cost into d'Alembert's equation, a classical functional equation; nondegeneracy and monotonicity leave the hyperbolic cosine; calibration fixes its exponent. The exponent c is a reading of the ratio, $x \mapsto x^c$, and (C4) says the ratio of counts is read as it stands. Section 5 shows that (C4) is supplied even given (C1) to (C3), with its own countermodel.

This paragraph is context and can be skipped; its one later use is restated where it enters. Premises (C3) and (C4) can be traded for a single limit: let Γ be any group of positive numbers under multiplication whose logarithms are dense in $\mathbb{R}$, and let $F : \Gamma \to [0, \infty)$ satisfy (C1) at $\kappa = 2$. If

$$\lim_{s \to 0,\ s \in \log \Gamma} \frac{F(e^s)}{s^2} = \frac{1}{2},$$

then $F = J$ on Γ, with no continuity, boundedness, measurability, monotonicity, or differentiability assumed [9]. The same work classifies the solutions of (C1) on the positive rationals when no regularity is assumed. Each positive rational is a product of primes, and a solution assigns one real weight a_p to each prime. A solution extends to a regular cost on the positive reals (continuous, locally bounded or measurable) exactly when the weights lie on the line $a_p = \lambda \log p$; off that line the cost is unbounded on every nonempty open subset of the positive rationals. The magnitude reading is therefore one point of an infinite-dimensional space of readings, and the single limit above cuts the space to that point (Figure 4).

4.3 The hierarchy premises

The second reading assigns levels to records. A *pass* is a completed stretch of posts, and its *size* is the number of posts it contains. Every pass is *registered* by a record: the posts that write down that the pass was completed, and a stretch counts as completed once such a record exists. (This is a second use of "record". In Section 3 the record at a stage is the posts so far.

Both are sequences of posts, so a registering record has a size and can itself be a pass.) A new level is made by *joining* a pass with the record that registers it, so that the new pass is the old one followed by its record and the sizes add. The base pass, at level 0, is a single post, of size 1. Let $\ell : \mathbb{N} \to \mathbb{R}_{>0}$ give the size of the pass at each level; these are the levels of a *recognition hierarchy*. The definitions leave open which record registers a pass, and so how large the next level is. Three *ladders* (hierarchies with a fixed choice of registering record) show the range. All begin $1, 2$, since level 1 is the base pass joined with a record of size 1: registering a single post takes one comparison outcome, one entry, one bit.

ladder	$\ell(0)$	$\ell(1)$	$\ell(2)$	$\ell(3)$	the record that registers the pass at level $k \geq 1$
golden	1	2	3	5	the pass at level $k-1$
base	1	2	3	4	the base pass
copy	1	2	4	8	the pass's complement, a distinct record of the same level (the pass with every entry flipped)

The golden ladder is the toy hierarchy $1, 2, 3, 5, 8, \ldots$, each size the sum of the two before it. The ratios of successive sizes, 2, 1.5, 1.67, 1.6, and so on, approach the golden ratio $\varphi = \frac{1+\sqrt{5}}{2} \approx 1.618$ and never equal it, because no two integers stand in the ratio φ. So the identity $\ell(n) = \varphi^n \ell(0)$ of Theorem 4.2 is a statement about an idealization in which the sizes $\ell(n)$ are positive real numbers, as the codomain of ℓ allows; (H1) and (H2) below are premises about that idealization. For passes whose sizes are counts, the limit in that theorem is what holds.

(H1) **Self-similar ratio.** $\ell(k+2)/\ell(k+1) = \ell(k+1)/\ell(k)$ for all k.

(H2) **Additive posting.** $\ell(2) = \ell(1) + \ell(0)$.

In words: (H1) every level is the same multiple of the level below it; (H2) is a statement about sizes: the second level is the first joined with a record of the size of the base. Together they say that a new level is the join of the two levels below it, one of each, so its size is their sum and its ratio to the level below is the same at every level. Call this sentence Join, the join premise; (H1) and (H2) are its statement as numbers. Section 6 restates the second half in terms of records, as the sentence P_5: the record that registers the pass at level k is the pass at level $k-1$; there it splits into two clauses, (A) and (B). The two senses of join agree on the golden ladder: a pass joined with its registering record is a pass joined with the pass one level down, one of each of the two levels below. With paired posting (each completed pass posted together with the record that registers it, so that on a join the sizes add), P_5 gives the recurrence of (H2) at every level. The next theorem says what ratio it forces.

Theorem 4.2 (the ratio)**.** *Under (H1) and (H2), $\ell(n) = \varphi^n \ell(0)$ with $\varphi = \frac{1+\sqrt{5}}{2}$, and $\ell(0)/\ell(n) = \varphi^{-n}$. Under (H2) imposed at every level, that is $\ell(k+2) = \ell(k+1) + \ell(k)$ for all k, and with no self-similarity premise at all, $\ell(k+1)/\ell(k) \to \varphi$.*

Proof. By (H1) the ratio $r = \ell(k+1)/\ell(k)$ is independent of k, so $\ell(n) = r^n \ell(0)$. By (H2), $r^2 \ell(0) = r\,\ell(0) + \ell(0)$, so $r^2 = r + 1$. The positive root is φ. For the second claim, the recurrence has general solution $\ell(k) = A\varphi^k + B(1-\varphi)^k$. If $A = 0$ then $\ell(k) = B(1-\varphi)^k$, which is identically zero or alternates in sign, and if $A < 0$ then $\ell(k) \to -\infty$ since $|1-\varphi| < \varphi$; positivity of ℓ excludes both, so $A > 0$, and $|1-\varphi| < \varphi$ then gives $\ell(k+1)/\ell(k) \to \varphi$. Both forms of the forcing appear at ledger level in [8]. □

The golden value comes from additive composition, and what (H1) buys is exactness: the ratio is φ at every level rather than in the limit. Figure 5 draws the join in three worlds.

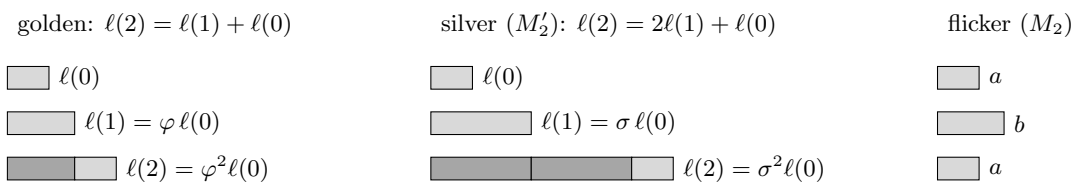


Figure 5: The join premise in three worlds. Self-similarity alone allows every ratio; the golden ratio is the join that takes one part of each level below (left). A silver world is exactly as self-similar and joins the nearer part twice, so its ratio is $\sigma = 1 + \sqrt{2}$ (middle). A flickering world, whose levels alternate between two sizes a and b, never builds (right). Section 5 uses the silver and flickering worlds as the countermodels M_2' and M_2. What Join supplies is the multiplicity in the middle panel.

4.4 The kept premise

The third reading places an act in space. It takes three steps: loops, motions, and one assumption.

Loops. A completed recognition and the record that registers it are two passes. That the record comes paired with the act is the form a posting takes in such a ledger: under conservation and pairwise locality each event is a balanced pair of entries [5]. We assume that, realized in a space of dimension D, a completed pass is carried by a closed curve, so the act and its record are two loops. This is the adopted reading: completion is read as the closing of a curve, and registration as the linking of the two loops. It is part of the Alexander clause below, one of the two identifications listed in Section 5. The *completed placement* is the two loops linked like two links of a chain, the act threaded through its record; the *separated placement* is the two loops apart.

Motions. A *post-free motion* is a movement of the loops through space that writes no entry in the ledger. A *spatial realization* in dimension D is a kinematics for the placements of the two loops, that is, a rule for which placements post-free motions carry to which, together with the clause below.

The assumption. The *Alexander clause* says that a *separating* pairing of the two loops, one that no post-free motion changes, that vanishes on the separated placement and that does not vanish on the completed one (a pairing is a number assigned to a placement, as the linking number is; Section 6 calls such a function a reader), exists exactly when some embedded circle $C \subset S^D$ has $H_1(S^D \setminus C) \neq 0$. (In S^3, a loop around a strand of C cannot be shrunk without crossing C.) Here space is taken as S^D, D-dimensional space with one point at infinity added, and a post-free motion is a deformation of the loops. The clause identifies the physical pairing, what a post-free motion cannot undo, with the linking pairing of two loops. It is part of what a spatial realization is, and it is not proven here. Everything in this subsection is relative to it, and so is the passage from dimensions to a period below (one bit per dimension); both appear among the hypotheses of Theorem 4.5.

(K) **Kept.** In some spatial realization in dimension D, no motion that posts nothing carries the completed placement of the two loops to their separated placement.

In words: once an act has been completed, with its record threaded through it, nothing that writes no entry can unthread it. Two linked links of a chain cannot come apart unless one is cut; (K) says the same of the act and its record, with writing an entry in the ledger playing the part of the cut. The next theorem says that (K) can hold only in three dimensions.

Theorem 4.3 (the dimension)**.** *Given the Alexander clause, (K) holds in dimension D if and*

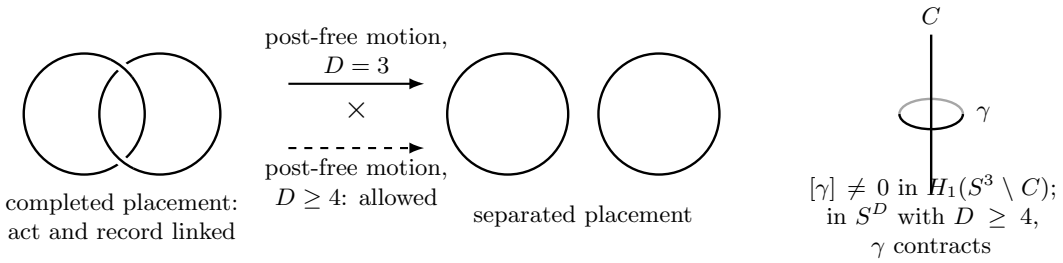


Figure 6: The kept premise and the dimension. A completed act and its record are two passes; realized in space as closed curves they are two loops. (K) says no motion that posts nothing carries the linked placement to the separated one. In three dimensions a loop γ around a strand of C cannot be shrunk without crossing C, which is $H_1(S^3 \setminus C) \cong \mathbb{Z}$; in four or more it contracts, the complement forgets C, and the two loops come apart with no post.

only if some embedded circle $C \subset S^D$ has $H_1(S^D \setminus C) \neq 0$; and that holds if and only if $D = 3$. Under (K) and the clause, $D = 3$.

Proof. In outline: in S^3 a loop around a strand cannot be slipped off, and in higher dimensions it can. The first equivalence joins two facts: (K) holds exactly when a separating reader exists (the integer form of Proposition 6.3, proved in Section 6), and the Alexander clause says a separating reader exists exactly when some embedded circle has a complement with nonzero first homology. For the second, Alexander duality [20, Thm. 3.44] gives, for any embedded circle $C \subset S^D$ with $D \geq 2$ (compact and locally contractible; in S^1 the circle is the whole space and its complement is empty),

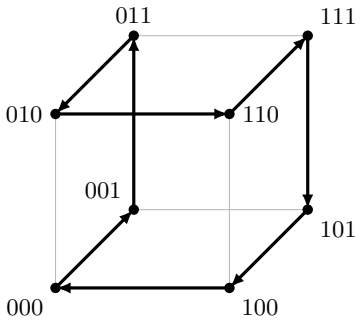
$$\tilde{H}_1(S^D \setminus C) \;\cong\; \tilde{H}^{D-2}(C) = \tilde{H}^{D-2}(S^1),$$

which is nonzero exactly when $D-2 = 1$. So every embedded circle in S^3 has H_1 of complement equal to $\mathbb{Z}$, and no embedded circle in any other S^D has nonzero H_1 of complement; this is the argument of [3]. Figure 6 draws the two placements and the loop that detects the difference. □

The last consequence is the period. A *cell* is the unit of the realization at one point of space, with one coordinate along each of the D dimensions, and a *tick* is one post. We take each dimension of the cell to be a two-sided distinction (a step one way along it or the other), which is one bit per dimension. That is an identification this paper adopts, not a consequence of (K); Section 5 lists it beside the Alexander clause. If the cell carries one bit per dimension, its states are the corners of a cube, and a pass that visits every corner has to flip one bit at each tick. The corollary counts the ticks.

Corollary 4.4 (the period)**.** *If a cell of a D-dimensional realization carries one bit per dimension (each dimension a two-sided distinction), a complete pass, a cycle that visits each of its 2^D states once, has period 2^D; under (K),* 8.

Proof. One bit is posted at a time, which is the single-event tick of the ledger, so a pass that visits all 2^D states without repetition takes at least 2^D steps, and the D-cube has a Hamiltonian cycle (the reflected binary code), so 2^D is attained [5]. The bridge from dimensions to bits is the hypothesis of the corollary; the period is a theorem about the cube (Figure 7). □



$000 \to 001 \to 011 \to 010 \to$
$110 \to 111 \to 101 \to 100 \to 000$

eight steps, one bit per step,
every state once, back to the start

Figure 7: The period. A cell with one bit per dimension in $D = 3$ has the eight states of the cube. One post flips one bit, so a pass visits neighbors along edges; the reflected binary code is a cycle through all eight, so the minimal period is $2^3 = 8$ and it is attained.

4.5 Assembly and consistency

The assembly theorem collects the four groups into one statement and shows that they can hold together.

Theorem 4.5 (the recognition axioms and what they force)**.** *From* Π*, the definition of a ledger (above, with Definition 6.10), (H1), (H2), (K) and the Alexander clause: the cost is* J *(premises (C1) to (C4) are then identities of the ledger total, not further assumptions; Theorem 6.12), the levels are* $\varphi^n \ell(0)$*, the weights* $\ell(0)/\ell(n)$ *(the base measured against level* n*, used in Section 8; this sense of "weight" is separate from the prime weights above) are* φ^{-n}*, and space has three dimensions; with one bit per dimension, the period is eight. The axioms are consistent: the floor, the two-sided total of a comparison of counts, the hierarchy* $\ell(n) = \varphi^n$*, and the standard realization of two linked circles in* S^3 *satisfy every one of them.*

Proof. Theorems 6.12, 4.1, 4.2, 4.3, Corollary 4.4. Consistency: the floor satisfies Π (two states differ; every finite record of its states extends; a comparison of two states has both present); the two-sided total of a comparison of counts is $2J$ as an identity of the rationals, Theorem 6.12; φ^n satisfies (H1) and (H2); two circles linked once in S^3 have linking number 1, which is deformation invariant and nonzero. □

Figure 8 previews the sort carried out in the next section: forced links are shaded, supplied links are dashed, and the worlds M_1, M_2 and M_3 at its right are the countermodels defined there.

5 The classification

Each line of Theorem 4.5 goes to one side of the criterion; Figure 8 is the result. In short: the three sentences of Π and the floor they build are forced; the cost is forced given a ledger; the join and kept are supplied. Each supplied premise comes with a countermodel, a world named M_i in which everything asking guarantees holds and the premise fails.

The scope of "forced." Section 3 defines forced as derivable from Π, and allows a verdict relative to Π together with a named definition. The cost is the one such line: it is forced *given a ledger*, the definition of Section 4, completed by Definition 6.10. What is proven is the conditional: every ledger has the cost J, up to the unit (Theorem 6.12, proved in Section 6). That a world keeps a ledger is not made true by asking: M_1 compares counts, reports only whether they are equal, asks, and is not a ledger. So relative to Π alone the magnitude is

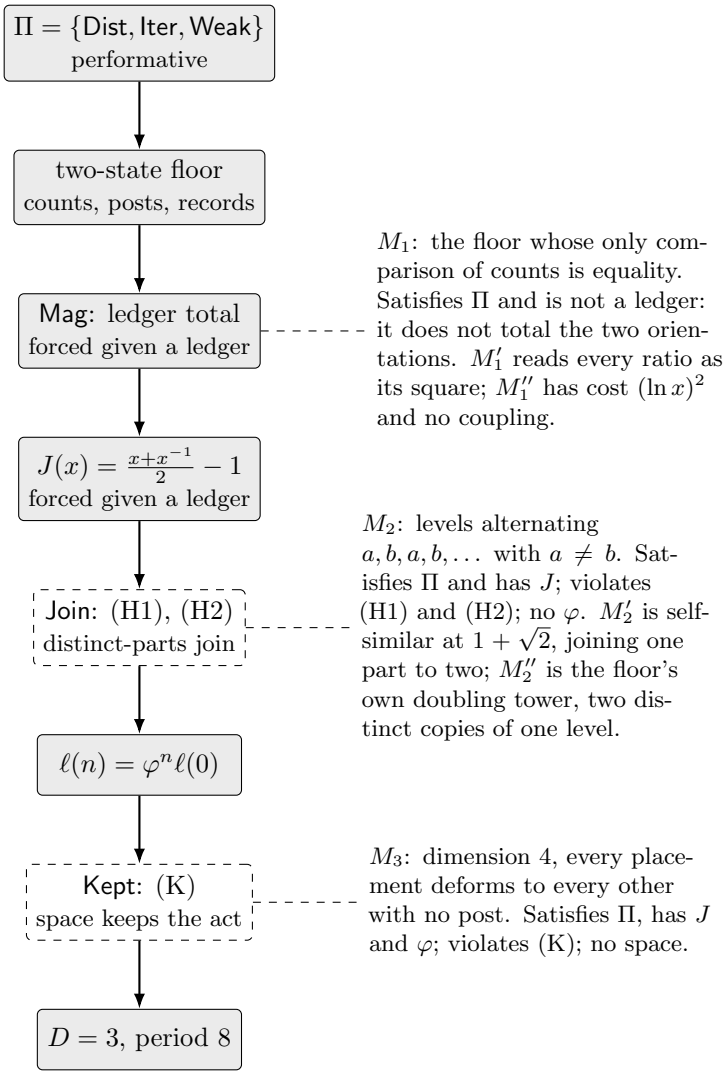


Figure 8: The recognition axioms under the criterion. Shaded boxes are forced from the box above; the two cost boxes are forced given the definition of a ledger. Mag, the reading in which counts are compared as ratios by a cost, is a theorem once a ledger is given: the cost is the ledger total of a two-sided comparison. Dashed boxes are the two remaining supplied premises, each with its countermodel M_i to the right.

supplied, and the paper takes the ledger from the cited construction as given. The join and kept are different in kind. The ledger as defined here (single-bit posts, entries in balanced pairs, one event per tick, persistence) says nothing about how levels are built or how an act sits in space, so it decides neither: the floor kept as such a ledger still admits the doubling tower M_2'' and the four-dimensional world M_3. The conserved ledger of [8], which also fixes the ratio between levels, is defined by further structural axioms; the ledger used here is defined without them, and M_2'' shows that without them the ratio between levels is not fixed. Taking the stronger ledger as given would put the ratio between levels among the inputs, which is the join premise this sort is asking about; the paper takes only the part of the ledger the cost needs. For those two the sort stays relative to Π alone. The definition of a ledger is the one definition the sort uses, and it comprises exactly the inputs listed in Section 4.

Theorem 5.1 (forced part)**.** Dist, Iter, *and* Weak *are forced. In every world, the extension step of the floor is available at every stage whose record some act entertains as complete. Among positive values, the gauge* κ *is a unit choice and carries no content. Given a ledger, the coupling*

is positive: an admissible cost with polynomial combiner is hyperbolic, with positive coupling, or quadratic, $(\ln x)^2$*, whose composition law is the parallelogram law, with coupling zero, and the ledger total is not quadratic (Theorem 6.12).*

Proof. Propositions 3.4 to 3.6; Iter applied to any finite record; and the observation that (C1) to (C4) and Theorem 4.1 are invariant under $F \mapsto \lambda F$, $\kappa \mapsto \kappa/\lambda$ for $\lambda > 0$. That rescaling never reaches $\kappa = 0$, which is why the vanishing coupling is a separate question. □

A *seam* is a reading the floor lacks that has a countermodel relative to Π, checked against Dist, Iter, and Weak by name. Relative to Π alone there are three, supplied$_\Pi$: the magnitude reading Mag (counts compared by a cost), the join Join (levels built from two distinct parts), and the kept premise Kept (a completed act placed in a space). The definition of a ledger removes the first, as the scope paragraph above says: Theorem 6.12 makes Mag a theorem given a ledger, so the countermodel M_1 remains a world that asks and is not a ledger. Two seams remain, the join and the kept premise. Each countermodel is the floor with a reading added that violates the premise.

Magnitude. In a world with a magnitude reading, a comparison of two counts a and b returns the ratio $x = a/b$, and the cost prices that ratio. Does asking force the floor's counts to be compared that way? No. The countermodel M_1 compares counts and reports only whether they are equal, so no ratio is available to price; two finer countermodels, M_1' and M_1'', show how a cost can compare magnitudes and still fail the premises.

Theorem 5.2 (the magnitude seam)**.** *The cost premises (C1) to (C4) are supplied*$_\Pi$*. Three countermodels, each with the premise it violates:*

- M_1*: the floor whose comparisons of two counts return only whether they are equal. It violates (C1) and (C2) together: no cost on its outcomes satisfies both. So (C1) to (C4) as a whole are supplied*$_\Pi$*.*
- M_1'*: the floor with cost* $\frac{2}{\kappa}\big(\frac{1}{2}(x^2 + x^{-2}) - 1\big)$*. It satisfies (C1) to (C3) and violates (C4). So (C4) is supplied relative to* Π *and (C1) to (C3).*
- M_1''*: the floor with cost* $(\ln x)^2$*. It is admissible, and it violates (C1) at every* $\kappa > 0$*, because its coupling is zero. So the positive coupling in (C1) is supplied relative to* Π *and admissibility: a world satisfying both violates it.*

Among admissible costs with polynomial combiner there are only two ways to fail the premises, and M_1' *and* M_1'' *instance them.*

Proof. In plain terms: a comparison that reports only equal or unequal gives every unequal ratio the same cost, and the composition law then rules that out. M_1 is a world in which comparisons of counts are made, so by Propositions 3.4 to 3.6 Dist, Iter, and Weak hold in it; directly, two of its states differ, every finite record of its states extends, and a comparison occurs with both terms present. Its comparisons return only "equal" or "unequal," and a world's cost is a function of the outcomes its comparisons return, so any cost F of M_1 on a ratio x of counts depends only on whether $x = 1$: $F(1) = \alpha$ and $F(x) = \beta$ for $x \neq 1$. No such F satisfies (C1) and (C2). At $x = y = 1$, (C1) gives $2\alpha = \kappa\alpha^2 + 4\alpha$, so $\alpha \in \{0, -2/\kappa\}$. At $x = 2$, $y = 1$ it gives $2\beta = \kappa\alpha\beta + 2\beta + 2\alpha$, so $\alpha(\kappa\beta + 2) = 0$. If $\alpha = -2/\kappa$ then $\beta = -2/\kappa = \alpha$ and F is constant, violating (C2). So $\alpha = 0$; then $x = y = 2$ gives $\beta = \kappa\beta^2 + 4\beta$, so $\beta \in \{0, -3/\kappa\}$, and $\beta = 0$ violates (C2); but $x = 2$, $y = 4$ gives $2\beta = \kappa\beta^2 + 4\beta$, so $\beta \in \{0, -2/\kappa\}$, a contradiction. So M_1 satisfies Π and violates the cost premises, which are therefore supplied$_\Pi$.

For M_1': write $G = \frac{\kappa}{2}F = \cosh(2\ln x) - 1$; then $H = G + 1$ satisfies $H(xy) + H(x/y) = 2H(x)H(y)$ by the cosh addition identity, so F satisfies (C1) at κ; (C2) and (C3) are immediate; and $\frac{\kappa}{2}F(2) = \frac{17}{8} - 1 = \frac{9}{8} \neq \frac{1}{4}$, violating (C4).

For M_1'': $F(x) = (\ln x)^2$ is admissible, and

$$F(xy) + F(x/y) = (\ln x + \ln y)^2 + (\ln x - \ln y)^2 = 2F(x) + 2F(y),$$

so (C1) at gauge κ would require $\kappa F(x)F(y) = 0$ for all $x, y > 0$, which fails for every $\kappa > 0$; (C2) and (C3) hold. It is the boundary member of the classified pair of families, the contraction of the hyperbolic family as the exponent goes to zero [7].

That the two countermodels exhaust the failures is the classification: an admissible cost with polynomial combiner is hyperbolic with some $\gamma, \lambda > 0$ or quadratic, and in the hyperbolic case $\frac{\kappa}{2}F(2) = \frac{1}{2}(2^\lambda + 2^{-\lambda}) - 1$, which is strictly increasing in λ and equals $\frac{1}{4}$ only at $\lambda = 1$. So a magnitude-comparing floor fails the premises by reading the ratio at a wrong exponent or by coupling its comparisons not at all. What the four premises supply is the magnitude reading Mag defined with them: the floor's counts are compared as ratios by a cost that composes and reads the ratio as it stands. Under Mag the four premises hold and Theorem 4.1 gives J; without it the floor has distinctions, counts, records, and no magnitude. □

Join. Do asking and the cost J force the levels to build golden? No. Three worlds, each with the cost J, show why: levels that flicker between two sizes (M_2); levels that build by counting the nearer part twice (M_2'); and the floor's own doubling tower (M_2''), in which a cell of n bits has 2^n states and each level is two copies of the level below.

Theorem 5.3 (the join seam)**.** *The hierarchy premises (H1), (H2) are supplied$_\Pi$ even given J. Countermodel M_2: the floor with cost J and levels $\ell(k) = a$ for even k and $\ell(k) = b$ for odd k, with $a \neq b$ positive. Countermodel M_2': the floor with cost J and levels $\ell(n) = \sigma^n$, $\sigma = 1 + \sqrt{2}$, which satisfies (H1) and violates (H2). Countermodel M_2'': the floor with cost J and levels $\ell(n) = 2^n\ell(0)$, each level the join of two distinct records of the level below, which satisfies (H1) and violates (H2).*

Proof. M_2 is a world in which comparisons are made, so Dist, Iter, Weak hold as in M_1; its cost is J, so Mag holds. (H1) at $k = 0$ demands $a/b = b/a$, hence $a = b$, which fails. (H2) demands $a = a + b$, which fails. M_2' satisfies Π and Mag for the same reasons, and it is self-similar by construction; $\sigma^2 = 3 + 2\sqrt{2} = 2\sigma + 1$, so its join law is $\ell(2) = 2\ell(1) + \ell(0)$ and (H2) fails. M_2'' is the floor read with its own address structure as levels: a cell of n bits has 2^n states, each added bit doubles the count, and the level at $n + 1$ consists of two distinct copies of the level at n, one for each value of the new bit. It satisfies Π and Mag as before, its ratio is 2 at every level so (H1) holds, and (H2) demands $4\ell(0) = 3\ell(0)$, which fails. So J does not force the golden ratio; the two hierarchy conditions do, and they are the join premise Join of Section 4. □

M_2' and M_2'' locate the seam more precisely than M_2 does. Self-similarity leaves every positive ratio, and among those the ones with an integer additive join law $\ell(2) = p\,\ell(1) + q\,\ell(0)$ are the positive roots of $r^2 = pr + q$, the metallic family, of which φ is the case $p = q = 1$, $1 + \sqrt{2}$ the case $p = 2$, $q = 1$, and 2 the case $p = 2$, $q = 0$. The operator realization of these ratios, a family of operators M obeying $M^2 = pM + q\,\mathrm{Id}$, is in [10]. So the supplied content of Join is the multiplicity of the join: the new level is one part of each, not two of one. M_2'' is the sharpest of the three, because it is not a foreign world but the floor's own counting, and because its two parts are distinct records. Section 6 makes exact what that leaves for a performative argument to reach.

Kept. Do asking, the cost and the golden hierarchy force a completed act to stay completed? No. The countermodel M_3 has all of them in four dimensions, where the act and its record come apart with nothing posted.

Theorem 5.4 (the kept seam)**.** *(K) is supplied$_\Pi$ even given J and φ. Countermodel M_3: the floor with cost J, levels φ^n, and a realization in dimension $D = 4$ whose kinematics carries every placement of the two loops to every other with no post.*

Proof. M_3 is a world in which comparisons are made, so Dist, Iter, and Weak hold (what M_3 does not keep is the placement of the loops, which is not a ledger state and is not a term of any comparison); its cost and levels give Mag and Join. In it the completed placement deforms freely to the separated one, so (K) fails. More strongly, by Theorem 4.3 no realization in dimension 4 satisfies (K) at all, since $H_1(S^4 \setminus C) = 0$ for every embedded circle C. So the world of M_3 has a ledger, a cost, a golden hierarchy, and no space in which a completed act stays completed. □

Remark 5.5 (what the countermodels are). Neither M_1' nor M_1'' is the running total of the two orientations of a comparison: M_1' totals a square reading of the ratio, and M_1'' totals with no coupling. M_2'', the floor's own doubling, is the reason no fact about the floor's counts reaches the join. Each of M_1, M_2 and M_3 is consistent, satisfies every sentence of Π, and is a way the world could have been; which way it is, inspection settles. Section 6 says what a further performatively necessary sentence would have to be to remove each of them, and how far the route has been carried.

Theorem 5.6 (the classification is complete)**.** *Every line of the recognition axiom set is forced or supplied; none is open. The cost is forced given a ledger, that is, from Π and the definition of a ledger (Theorem 6.12, proved in Section 6). The supplied lines are the two remaining seams. The join supplies one sentence,* Join*, which Section 6 restates as P_5, the conjunction of two clauses; the three join rows of Table 2 are three countermodels to it, each isolating a sharper part. Kept supplies one sentence, (K), with the countermodel M_3. The sort is relative to three identifications that are not lines of the set: that entertaining a sentence is comparing it with its negation, inside the world it is about (Lemma 3.2); the Alexander clause (what a post-free motion cannot undo is the linking of two loops); and one bit per dimension (a cell carries one two-sided distinction per dimension).*

Proof. The axiom set is the finite list of Theorem 4.5. Table 2 gives each line a derivation or a countermodel, and Proposition 3.8 says no line has both. The value of κ is not a line: it is removed by rescaling (Theorem 5.1). The three identifications are not lines either. They fix what asking, (K) and the period are about: without the first, an act of asking is not an act of the world, and no sentence is made true of the world by asking it; without the clause, (K) does not speak of a dimension; and without the bit, a dimension does not speak of a period. None of them is proven here, and Section 8 lists them separately. □

5.1 Kept is the conclusion

Theorem 4.3 says something sharper than that (K) is supplied: relative to the Alexander clause, (K) is equivalent to $D = 3$. So the dimension theorem has a specific logical shape. Its premise and its conclusion are the same sentence read twice, once as a statement about completed acts (the completed placement of a recognition and its record cannot be undone without a post) and once as a statement about topology (only S^3 has an embedded circle whose complement has a nonzero first homology).

Ordinary rings give a test case: linked rings do not come apart unless they are cut. The theorem shows what that observation supplies, in three steps.

Line	Sentence	Verdict	Witness
Dist	something is distinguished	forced	the act (Prop. 3.4)
Iter	no finite record entertained as complete is complete	forced	the act (Prop. 3.5)
Weak	a comparison has both terms present	forced	the act (Prop. 3.6)
floor	the extension step is available at every stage whose record is entertained as complete	forced	Iter in L
(C1) to (C4)	the ledger total of the two orientations	forced given the definition of a ledger	Thm. 6.12
M_1	equality-only comparison, which asks and is not a ledger	supplied$_\Pi$	Thm. 5.2
value of κ	a unit	not a line	rescaling
(H1), (H2)	a level is a join of the two below	supplied	M_2
(H2) given (H1)	one part of each, not two of one	supplied	M_2'
(H2) given (H1) and distinct terms	the two parts are different levels	supplied	M_2''
(K)	a completed placement is not undone by a post-free motion	supplied	M_3

Table 2: The sort. Every line has a derivation or a countermodel; each seam's supplied content is the sentence its sharper countermodel isolates. For the cost, "forced" means forced from Π together with the definition of a ledger.

(a) *What the rings show.* Anyone who takes linked rings that do not pass through one another without cutting as a supply of (K) has supplied (K) for those rings.

(b) *Why one more direction would undo it.* Take an extra direction of any positive extent, however small. It undoes the completed placement of two loops confined to a slice, with no post: displace one loop a little along it, move its projection through the other, and return. So what the inspection of rings supplies is that those rings have no post-free direction to move in beyond three.

(c) *What is assumed in moving from rings to traces.* (K) itself is a statement about the traces of an act (its two loops, the act and the record that registers it), realized as embedded circles: which post-free motions they can execute. That the traces of an act move in the same kinematics as the rings is a further supply, not a theorem.

Granted (a) and (c), that person has supplied $D = 3$. What the theorem contributes in that case is not a prediction of the dimension but an account of it: three-dimensionality is what "a completed act stays completed" looks like as geometry, and four or more dimensions can keep no act. A derivation would be a proof that (K) is forced.

6 How far the performative route reaches

Each seam adds a short sentence about the world to the floor. This section asks, for each, whether a sentence made true by asking could be that sentence, how much of it an argument of the performative kind can cover today, and what stops it. One theorem sets the ceiling for all three; the subsections then take kept, the join and the magnitude in turn.

Call the floor's *signature* the vocabulary in which every sentence of Π is written: states, posts, records, comparisons, and the order of stages. An *asking history* is a script of what an asker does, a finite sequence of acts in that signature: posts, retentions (advancing a stage while keeping every record), and comparisons of records already present. The act of entertaining a sentence is one such history (Section 3). For sentences in this signature, performative necessity takes this form: a sentence is performatively necessary when it holds in every world in which some asking history runs. A sentence S that meets the definition of Section 3 (every act of entertaining S makes it true) meets this form in each countermodel below: entertaining S is an asking history, the countermodels run every asking history, so S holds in each. The ceiling theorem therefore covers every sentence the earlier definition admits.

Two words recur below. A *trace* is one of the two loops of the kept premise, the act or the record that registers it. A *reader* is a function (not a person) that gives each placement of the two traces a number. The three countermodels M_1, M_2'' and M_3 have something in common, and the next theorem turns it into a ceiling.

Theorem 6.1 (the asker obstruction)**.** *Each of M_1, M_2'', and M_3 is the floor with one reading added and the signature unchanged: a cost on ratios, a level function on records, a kinematics of the two traces. Every asking history that runs in the floor runs, act for act, in each of them. Consequently no sentence of the signature that holds wherever something is asked forces* Mag*,* Join*, or* Kept*.*

Proof. A sentence of the signature is evaluated on states, posts, records, comparisons, and stage order, and the three readings touch none of these, so the sentence has the same value in the floor and in each countermodel; in particular the entertaining of any sentence, which is itself an asking history, runs in each. A sentence S that holds wherever something is asked therefore holds in M_1, M_2'', and M_3. If S forced Mag, Mag would hold in M_1; Theorem 5.2 says it does not: a finite substitution shows that an equality-only cost satisfying the composition law is constant. If S forced Join, (H2) would hold in the doubling tower, where $4 \neq 2+1$. If S forced Kept, an invariant reader separating the completed placement from the separated one would exist in M_3, where every placement deforms to every other and every invariant reader is constant. □

The theorem places each seam's missing sentence outside the signature. For the magnitude it is a cost-valued comparison, which the definition of a ledger supplies as the total of the two orientations (Theorem 6.12); for the join, the level of a record; for kept, a reader of relative placement.

Each of the last two is out of reach from this signature, since no sentence of states, posts, records, comparisons, and stage order can do its work. At the next signature, the one that adds the object it names, asking leaves each undecided in both directions (Propositions 6.4 and 6.8), so at that signature each is supplied, and what it supplies is named exactly. For kept it is one sentence, since the reader clause is equivalent to (K) (Proposition 6.3). For the join it is two, neither made true by asking (Proposition 6.9).

6.1 The weak and the strong kept

Two different things are called kept. There is a weak form that is forced: Weak, the third sentence of Π. A comparison has both its terms present, so whatever is compared reached the comparison. (The ledger's records also persist between posts; that is part of the definition of a ledger, Section 4.) (K) says something else: space keeps the placement of the two loops against every motion that posts nothing.

The difference is exact. A record that survives motion is a record of how the two traces sit relative to each other, not of where they are [3]. A ledger can keep every record of position and still let the loops come apart, because position changes under motion, and a reader of position that no motion changes is constant. What (K) needs, and Weak does not give, is a reader of placements that every post-free motion leaves fixed and that reads the completed placement differently from the separated one. In the words of recognition geometry [16], a space is a set of configurations some recognizer tells apart, and (K) says the placements of the two loops are such a space. Theorem 5.4 is the proof that Weak does not reach (K): M_3 keeps every record and keeps no placement.

(K) is therefore supplied$_\Pi$. To make it forced one would need a fourth performatively necessary sentence, P_4 (the three sentences of Π being the first three), that entails (K). Any such P_4 fails in M_3, a world in which asking happens and (K) fails. The natural candidate is the following sentence.

Definition 6.2 (the sentence P_4)**.** P_4: every act of comparison leaves two traces, and there is a reader of their relative placement, changed by no post-free motion, that reads the completed placement differently from the separated one.

Half of P_4 is the paired posting of the ledger, proved in [5]: a comparison leaves two traces, the act and its record. The other half is the reader, and the reader clause restates (K).

Proposition 6.3 (the reader is (K))**.** *Let the placements of the two traces carry the relation "some post-free motion carries one to the other," reflexive, symmetric, and transitive; a reader is a function on placements, and a motion changes it when its values at the two ends differ. A reader that no post-free motion changes and that reads the completed placement differently from the separated one exists if and only if no post-free motion carries the completed placement to the separated one. The same holds for an integer reader that no post-free motion changes, that vanishes on the separated placement, and that does not vanish on the completed one.*

Proof. In outline: if such a reader exists and some post-free motion carried the completed placement to the separated one, invariance would give the two placements the same reading, which the reader does not. Conversely, suppose no post-free motion does. Read a placement as 1 when some post-free motion carries the completed placement to it and as 0 otherwise. The relation is symmetric and transitive, so two placements joined by a motion have the same reading; it is reflexive, so the completed placement reads 1; the separated placement reads 0 by hypothesis. For the integer form: a reader that vanishes on the separated placement and not on the completed one reads them differently, so the first sentence gives the condition; conversely the function just built is integer-valued. □

P_4 is therefore equivalent, in every kinematics, to paired posting together with (K). The reader is not a symbol of the floor's signature, and by Theorem 6.1 no sentence of that signature true wherever something is asked yields (K), hence none yields the reader clause. P_4 is a sentence of the next signature, the *placement signature*, and the question is whether entertaining anything in that signature makes it true.

The placement signature adds two acts to the floor's three (posting a record, advancing the stage with every record kept, comparing two records present). *Read* posts the reader's value on the present placement of the two traces. *Move*, to a named placement, is taken when the kinematics carries the present placement there and refused otherwise. Taken or refused, a move posts nothing: a taken move changes the placement and nothing on the ledger, a refused move changes neither. A world of the signature is a floor world whose state carries a placement, together with a kinematics and a reader; histories start at the completed placement.

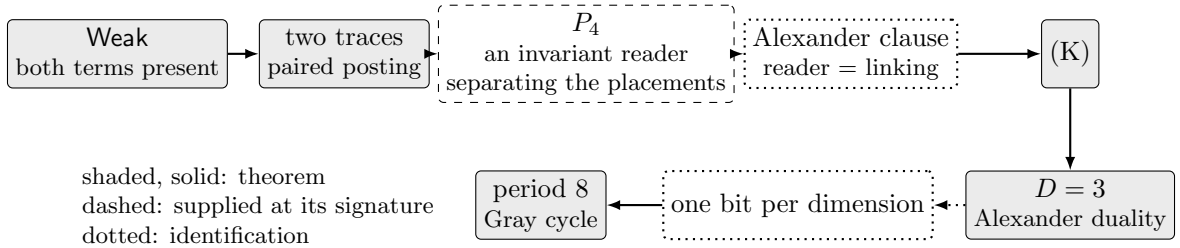


Figure 9: The distance from what asking guarantees to the period. Two steps are theorems (paired posting; Alexander duality), one is the supplied sentence (an invariant reader that separates the completed placement from the separated one, equivalent to (K) by Proposition 6.3 and left open by asking, Theorem 6.1 and Proposition 6.4), and two are identifications rather than proofs (the reader is the linking pairing; each dimension carries one bit).

Proposition 6.4 (kept at the placement signature)**.** *Two worlds share the reader that reads the completed placement as* 1 *and every other placement as* 0*, and differ in kinematics. The history* read, move to the separated placement, read, retain, compare the two readings *registers an outcome in both:*

world	the move	ledger holds	comparison registers	(K)
first: the kinematics of M_3, every placement moves to every other	taken	1 and 0	they differ	fails
second: rigid, no post-free motion changes a placement	refused	1 and 1	they agree	holds

So something is asked in both worlds, and (K) fails in the first and holds in the second. Consequently, if a sentence of the placement signature holds in every world of the signature in which some history registers an outcome, it entails neither (K) nor its negation.

Proof. Both histories are finite and their ledgers are computed act by act, as the table shows. (K) is a sentence of the kinematics alone: in the first world some post-free motion carries the completed placement to the separated one, so (K) fails; in the second none does, so (K) holds. A sentence true wherever something is asked holds in both worlds; if it entailed (K) the first world would satisfy (K), and if it entailed the negation the second would fail it. □

The first world is M_3's kinematics with a reader, a symbol the floor's signature lacks (this reader is changed by the move, so it is not an invariant one). With it the asker entertains the negation of (K) and finds it true: the readings 1 and 0 are both on the ledger and differ, so the placement changed between them, and the only act between them was a move, which posts nothing.

What this shows. This computes Theorem 5.4: M_3 keeps every record and no placement. What Weak keeps is the record of a reading; what (K) keeps is the placement read; a world can have the first and not the second. So at the placement signature (K) is still supplied; the reader clause of P_4 restates it and does not ground it.

The link to $D = 3$. The step from (K) to $D = 3$ is Theorem 4.3, relative to the Alexander clause, which says that the number no post-free motion changes, the one the reader of Proposition 6.3 computes, is the linking number of the two loops. Figure 9 draws the whole distance from Weak to the period, with each step marked as theorem, supplied sentence, or identification.

Supplied with P_4, the line runs to its end: P_4 gives (K) by Proposition 6.3, and (K) gives $D = 3$ by Theorem 4.3, relative to the Alexander clause. Until P_4 is supplied, M_3 remains a

countermodel relative to Π, at the floor's signature and at the placement signature alike.

6.2 The join

The join seam has an argument of the same shape, and it gets farther before it stops. Here is the plan. Asking forces two facts about a join: it has two terms, and they are distinct records. The floor's own doubling tower has both facts and is not golden, so the argument stops there. What the golden join needs on top is one sentence, P_5: the record that registers a pass is itself a pass of the level immediately below. That sentence splits into two clauses, (A) and (B), defined below, and asking makes neither true.

The hierarchy premises speak of a new level made by joining a completed pass and the record that registers it (Section 4). The floor has one kind of event, a comparison of two records whose registration is a post, and the next proposition takes as its premise that a join is an act of that kind: the new level is the two records together, with the distinction between them registered by posting; as in Section 4, the new pass is the old one followed by its record, so the sizes add. The proposition therefore concerns any hierarchy whose joins are such acts, the golden one included. Two things about such an act are forced.

Proposition 6.5 (what a join must be)**.** *Let a new level be constituted by a comparison whose registration is a post. Then the comparison has exactly two terms, and they are distinct records.*

Proof. A comparison takes two records by definition, so the join is binary. Suppose its two terms were the same record. Two states that agree in every bit are the same state, so a comparison of a record with itself cannot return that its terms differ; the comparison is not a distinction; a post registers a distinction, so nothing is posted, against the hypothesis that the comparison's registration is a post. And a transition that posts nothing is no transition, so the state is unchanged and no new level is constituted. So the two terms are distinct records. □

Given the proposition's hypothesis, that a join is such a comparison, both conclusions are forced$_\Pi$: they use the definitions of Section 3, the floor's rule that a transition which posts nothing is no transition (Section 4), and no premise beyond Π and the floor's construction. They are also as far as the route goes, and the obstruction is exact.

Theorem 6.6 (the join obstruction)**.** *M_2'' satisfies* Π*, the cost premises, (H1), and Proposition 6.5, and violates (H2). Consequently no set of sentences about states, posts, records, comparisons, and counts that is true in the floor forces (H2) (the join instance of Theorem 6.1), and what (H2) supplies beyond Proposition 6.5 is that the two parts of a level are different levels.*

Proof. Theorem 5.3 verifies Π, Mag, (H1), and the failure of (H2) in M_2''. Each of its joins takes two distinct records, the two copies of the level below that differ in the new bit, so Proposition 6.5 holds in it. M_2'' is the floor with its address structure read as levels, and the reading adds a level function and changes nothing else, so every sentence in the floor's own signature that is true in the floor is true in M_2''; any set of such sentences together with Π has M_2'' as a model, and (H2) fails in that model. The two parts of each level of M_2'' are distinct records of one level; (H2) requires them to be the two levels below, one of each; so what (H2) adds to what is forced is that the parts differ as levels, where the forced part gives that they differ as records. □

The theorem says what a further performatively necessary sentence would have to do: fail in the doubling tower. In the tower, the second part of every level is a copy of the first, a pass of the same level. In the golden join, the two parts of a level are a completed pass and the record that registers it (Section 4), and (H2) gives that record the size of the level below,

hence smaller than the pass it registers. The three ladders of Section 4 differ only in which record registers a pass; the copy ladder's sizes are those of the doubling tower.

What the golden ladder has and the other two lack is that the record of a pass is the pass one level down. That is the sentence P_5; the numbering continues from P_4 of the previous subsection.

Definition 6.7 (the fifth sentence)**.** P_5 is the sentence "Above the base level, the record that registers a completed pass is itself a completed pass, of the level immediately below the pass it registers."

It has two parts. (A): above the base, the record that registers a pass is some pass of a lower level. (B): above the base, passes at distinct levels have distinct records. In the base ladder (A) holds and (B) fails, since every record is the base pass; in the copy ladder (B) holds and (A) fails, since no record is a pass; the golden ladder satisfies both. Proposition 6.9 below shows that P_5 holds exactly when (A) and (B) both do.

P_5 fails in M_2'': there the record that registers the pass at level k is a second copy of that pass, a pass of level k, not of level $k-1$. So if P_5 holds the tower is no longer a countermodel.

The sentence names the level immediately below, and it has to. If it said only that the record is a pass of some lower level, then by paired posting $\ell(k+1) = \ell(k) + \ell(j)$ with $j < k$, and the ledger's minimality of the record [5] (no wasted entry, so the shortest record that registers the pass) would pick the smallest such j. On a growing ladder (sizes increasing with the level, as in all three above) the adjacent level is the largest lower level, so minimality points toward $\ell(k+1) = \ell(k) + \ell(0)$. The adjacency is content, not a consequence of minimality, and the sentence carries it.

Given P_5, paired posting gives $\ell(k+1) = \ell(k) + \ell(k-1)$ at every level, which by Theorem 4.2 gives φ in the limit with no self-similarity premise; (H1) then upgrades the limit to an identity at every level. Paired posting is a property of the ledger of Section 4 (the sizes adding is the definition of a join; the pairing of a pass with its record is the ledger's), so (H2) would be forced relative to Π, P_5, and the ledger. If P_5 fails, Join stays supplied and the doubling tower stays a possible world.

Does asking give P_5? Part of it, and only part, in three steps.

1. *What asking gives.* The act of entertaining P_5 is a comparison of two terms with one outcome, and its record registers that one outcome: one post where the act had two terms. Read with the act as the pass and its record as the record that registers it, the record is smaller than the pass it registers, which is the direction P_5 needs (the record is the pass one level down, smaller than the pass above). That is a remark made by counting; the floor's signature has no comparison of sizes, so it is not a sentence of that signature.

2. *What it does not give: the level.* A record that registered one outcome per comparison would build the ladder $\ell(k+1) = \ell(k) + 1$, the base ladder, whose ratio tends to 1, not a geometric one. For P_5 the record must be a completed pass in its own right, one level down, and that is the part of the sentence the floor's own tower does not supply.

3. *What is left.* That the record is a pass of some lower level is (A); that it is the pass immediately below is what (B) adds (Proposition 6.9 below). Theorem 6.1 does not address (A) or (B): level is not in the floor's signature, and P_5, (A) and (B) belong to the signature that adds levels.

The next signature, the *level signature*, can be written down, and ladders in it can be entertained level by level. It adds to the floor's acts one more, the joining of two records already

present, and it adds the levels themselves, each with its pass and the record that registers it. Every level above the base equals the pass below it joined with the record that registers that pass; the two terms are distinct, and every record, base pass included, is nonempty. The additive form of the levels is written into this signature by the join. What the signature leaves free is which record registers a pass. The golden, base and copy ladders above share this signature and differ only in that.

Asking runs on all three ladders, one level at a time: the asker posts the pass, posts its record, joins them, posts the next level and compares the join with it, and the comparison registers agreement in all three. The instance of P_5 at a level is registered by a finite history too: post the record of the pass above, post the pass below, compare. It registers agreement at every level of the golden ladder, at the first level above the base only in the base ladder, and nowhere in the copy ladder.

Proposition 6.8 (the join at the level signature)**.** *In the signature that adds levels, each with its pass and the record that registers it, P_5 is a sentence, and asking does not decide it.*

(i) The base ladder and the golden ladder both entertain every level; the first fails P_5 and the second satisfies it. So no sentence true in every ladder that entertains every level has P_5 as a consequence, and none has its negation. Since every ladder entertains every level, the condition is met by all of them, and asking forces nothing about which record registers a pass beyond what the definition of a ladder already fixes.

(ii) The weaker clause, that the record of a pass above the base is some pass of a lower level, is not decided by asking either: the copy ladder entertains every level and fails it, and the golden ladder entertains every level and satisfies it.

(iii) Under that clause the size of the record lies between the size of the base pass and the size of the pass immediately below, and is the size of some lower pass. The lower end is attained at every level exactly when every record above the base is the base pass, which is the base ladder; the upper end is attained at every level exactly when P_5 holds.

Proof. The claims at every level follow from the recursion defining each ladder: the next level is by definition the join, and in the golden ladder the record above the base is by definition the pass below. A sentence true in every ladder that entertains every level is true in the base ladder, the golden ladder, and the copy ladder, so it cannot have as a consequence anything one of them fails. Sizes along any ladder increase strictly, since every record is nonempty, so a record that is some lower pass has size between the base pass and the pass immediately below; it equals the pass immediately below exactly when its size does, which is P_5, and it equals the base pass exactly when its size does. □

The levels do not close the seam. With levels in view, asking still leaves open which record registers a pass, across the whole range from the base pass at the bottom to P_5 at the top. In the floor itself a registration is one comparison outcome, one bit; the base ladder is the ladder whose record above the base carries exactly that much, one fixed record of size one. P_5 asks the record to carry a whole lower pass instead, and that difference can be stated exactly: P_5 is the conjunction of two simpler sentences, and asking makes neither true.

Proposition 6.9 (what P_5 is made of)**.** *In any ladder, P_5 holds if and only if two sentences hold together: (A) above the base, the record that registers a pass is some pass of a lower level; (B) above the base, passes at distinct levels have distinct records. The base ladder satisfies (A) and not (B); the copy ladder satisfies (B) and not (A); both entertain every level. So asking makes neither sentence true, and neither alone gives P_5.*

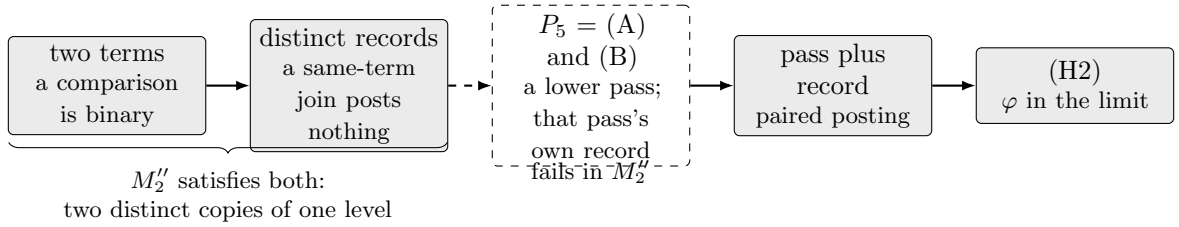


Figure 10: The distance from what asking guarantees to the additive join. Two steps are theorems (a join is binary; its terms are distinct records), and the floor's own doubling tower satisfies both, so they do not imply (H2). The dashed step is supplied: above the base level, the record of a pass is a pass of the level immediately below. The floor's signature does not contain levels; in the signature that does, asking leaves it open in both directions (Proposition 6.8), and it is the conjunction of two sentences, neither made true by asking (Proposition 6.9). Given it, paired posting gives (H2), with φ in the limit.

Proof. Given P_5, (A) is immediate, and (B) follows because passes have strictly increasing sizes: under P_5 the records at levels $k+1$ and $m+1$ are the passes at levels k and m, which coincide only if $k = m$. Given (A) and (B), let the record of the pass at level $k+1$ be the pass at level $j(k)$ (one level, since passes at distinct levels have distinct sizes), with $j(k) \leq k$ by (A). Two levels with the same j have the same record, so by (B) the map j is injective. An injective map with $j(k) \leq k$ at every k is the identity, by induction on k: if $j(i) = i$ for all $i < k$, then $j(k) \leq k$ and $j(k) \neq j(i) = i$ for every $i < k$, so $j(k) = k$. Hence the record of the pass at level $k+1$ is the pass at level k. In the base ladder every record above the base is the base pass, so (A) holds, and the records at levels two and three coincide (both are the base pass), against (B). In the copy ladder the records registering the passes above the base have sizes $2, 4, 8, \ldots$, all distinct, so (B) holds; every pass begins as the base pass does and every record begins with the opposite entry, so no record is a pass, against (A). □

(B) is the clause on what a record must carry: it must pick out the pass it registers. Asking cannot supply it, because a ledger is ordered and a record can be named by where it sits. In the base ladder the records at levels 2 and 3 are both the base pass, yet an asker who has posted the level-3 pass and then its record can take that record by its place (the entry just after the level-3 pass) and compare it with anything; the comparison is about the level-3 pass because of where the record sits, not because of what it says. This history runs and registers at every level of the base ladder, so no sentence true wherever something is asked implies (B).

The copy ladder does the same for (A): its record is the pass's complement, a record of the same size that is never itself a pass, and the same history runs there, taking the record by its place, so no sentence true wherever something is asked implies (A). At the signature that adds levels the join is supplied, and what it supplies is (A) and (B); P_5 is their conjunction: the record of a pass is itself a lower pass, and it is that pass's record and no other's. Figure 10 lays out the steps.

6.3 The magnitude

Asking does not force the floor's counts to be compared as ratios: Theorem 6.1 and the world M_1. A cost-valued comparison is not in the floor's signature. What supplies it is two things taken together: the arithmetic that distinction itself generates, and the definition of a ledger. This subsection completes that definition with the entries and the cost, says what the theorem built on it assumes and what it proves, and then proves it. (Here a comparison *enters* entries in the book; a *post* is the step of Section 3.)

Definition 6.10 (ledger, entries, cost)**.** A *ledger* is a record kept as a running total of balanced pairs of entries, a pair being *balanced* when its two entries multiply to 1. A comparison of two counts $a, b \neq 0$ enters one balanced pair, its two orientations a/b and b/a (the *entries*); the balanced comparison $a : a$ enters 1 and 1 and totals 2. The *cost* of a comparison in a ledger is what it adds to the book above a balanced comparison, up to a positive unit: its *excess* over the balanced total 2.

That last identification is what "cost" means for a ledger. It is a definition, not a derived fact; with the inputs of Section 4 it makes up the definition of a ledger, the only one the sort of Section 5 uses. The theorem below does not assume the four cost premises (C1) to (C4); it derives them from the definition in three steps.

(a) *The entries.* A comparison of a with b is one act read from either side: a/b from a's side, b/a from b's. The ledger of Section 4 records it as one balanced pair, the double entry in which each event is a balanced pair of entries. This form of entry belongs to the definition of a ledger, taken from the cited construction, and asking does not force it: the entries are ratios, not $a - b$ or a square of either. The ledger adds entries. A balanced pair multiplies to 1 whatever the counts, so its product carries nothing; its total $a/b + b/a$ is what changes with the counts.

(b) *The excess.* Since the definition makes the cost this excess, what remains is computation. The excess of that pair over the balanced pair is $(a - b)^2/(ab) = 2J(a/b)$, and, in the unit that (C4) fixes, it satisfies all four cost premises as identities of the rationals.

(c) *Nothing else does.* Every rational-valued cost that satisfies (C1) to (C4) is a positive multiple of it (Proposition 6.11(i)).

So the four cost premises are not further assumptions about a ledger. They are identities of its excess.

Iterated distinction gives $\mathbb{N} \subset \mathbb{Z} \subset \mathbb{Q}$ with decidable equality and order and nothing continuous [1]. On that carrier the total of $a : b$ above balance is

$$\frac{a}{b} + \frac{b}{a} - 2 = \frac{(a-b)^2}{ab} = 2\,J\Big(\frac{a}{b}\Big)\,.$$

Every cost premise then holds of this total as an identity of the fraction field: with $G(x) = \frac{1}{2}(x + x^{-1}) - 1$ on $\mathbb{Q}_{>0}$, $G(xy) + G(x/y) = 2G(x)G(y) + 2G(x) + 2G(y)$, $G(2) \neq G(1)$, G is nondecreasing on $[1, \infty) \cap \mathbb{Q}$ since $(y + y^{-1}) - (x + x^{-1}) = (y - x)(xy - 1)/xy$, and $G(2) = \frac{1}{4}$. No limit, no logarithm and no real number occurs. The cost premises are four sentences, but on the carrier they say one thing: the cost of a comparison is a positive multiple of the total its two entries enter above balance. The only number left free in that sentence is the unit. The calibration $J(2) = \frac{1}{4}$ is read off one comparison, two against one (excess $\frac{1}{2} = 2J(2)$).

The next proposition makes this exact: (i) the total is a cost and the only one on the carrier; (ii) the premises are independent, so the classification can supply each separately; (iii) an exponent other than one is a compound of comparisons; (iv) zero coupling is excluded on the carrier.

Proposition 6.11 (the magnitude on the carrier)**.** *On $\mathbb{Q}_{>0}$ with rational values:*

(i) G satisfies (C1) to (C4) with $\kappa = 2$, and every rational-valued cost F satisfying them with coupling κ is $F = (2/\kappa)G$: a positive multiple of G set by the unit, and G itself when $\kappa = 2$. By (C1) a cost is fixed by one weight a_p for each prime p, since every positive rational is a product of primes and their inverses; monotonicity puts the weights on the line $a_p = \lambda \log p$ [9]; and (C4) gives $\lambda = 1$.

(ii) (C1), (C3) and (C4) are independent of one another, each failing while the other two hold. G^2 is nondegenerate, monotone and calibrated at $\kappa = 8$, and fails (C1) at $(2, 2)$; $G(x^2)$ with $\kappa = 2$ satisfies (C1) to (C3) and has $G(4) = \frac{9}{8} \neq \frac{1}{4}$; and $G(\chi(x))$ with $\chi(x) = x\,2^{-v_3(x)}$, a multiplicative character fixing 2, satisfies (C1) and (C4) and has $G(\chi(3)) = G(3/2) = \frac{1}{12} < \frac{1}{4} = G(\chi(2))$, so it is not monotone.

(iii) $G(x^2) = 2G(x)^2 + 4G(x)$, the composition law at $y = x$: the exponent-two cost of a ratio is the cost of the compound of the comparison with itself.

(iv) A cost with zero coupling, $F(xy) + F(x/y) = 2F(x) + 2F(y)$ and $F(1) = 0$, has $F(x^n) = n^2F(x)$. If it is monotone then $n^2F(2) \leq m^2F(3)$ whenever $2^n \leq 3^m$ and $m^2F(3) \leq n^2F(2)$ whenever $3^m \leq 2^n$, so either F vanishes on $[1, \infty)$ or $F(3)/F(2) = (\log 3/\log 2)^2$, which is irrational by the Gelfond–Schneider theorem. No rational-valued nondegenerate monotone cost with zero coupling exists.

Proof. (i) to (iii) are the displayed identities and the values $G(2) = \frac{1}{4}$, $G(4) = \frac{9}{8}$, $G(3/2) = \frac{1}{12}$, $G(2)^2 = \frac{1}{16}$; for χ, v_3 is additive so χ carries products to products and quotients to quotients, and (C1) for $G \circ \chi$ is (C1) for G at $\chi(x), \chi(y)$. (iv) The zero-coupling law at (x^{n+1}, x) gives $F(x^{n+2}) = 2F(x^{n+1}) + 2F(x) - F(x^n)$, whose solution with $F(x^0) = F(1) = 0$ is $n^2F(x)$; the two inequalities are monotonicity at the pair $2^n, 3^m$. If $F(2) = 0$ then every $F(2^n) = 0$ and the squeeze gives $F = 0$ on $[1, \infty)$. Otherwise $r = F(3)/F(2)$ satisfies $(n/m)^2 \leq r$ for every rational $n/m \leq \log 3/\log 2$ and $r \leq (n/m)^2$ for every rational $n/m \geq \log 3/\log 2$, so $r = (\log 3/\log 2)^2$. Were r rational, $\log 3/\log 2$ would be algebraic and irrational, and $2^{\log 3/\log 2} = 3$ would be transcendental by Gelfond–Schneider. □

On the carrier the free numbers of the continuum are gone. A nonintegral exponent is not an operation of the tower, since $2^{1/2} \notin \mathbb{Q}$; an integer exponent k is the reading of a comparison as its k-fold compound (iii), which would be k acts, and a ledger writes one event per comparison, so $k = 1$; a monotone cost with zero coupling is not rational-valued (iv), and the quadratic cost $\frac{1}{2}(\ln x)^2$ of M_1'' is not in the carrier at all.

A polynomial readout of the two orientations other than a multiple of their excess over 2 (for instance the square of the excess) has the same zeros and does not add when two comparisons are posted together; it is a function of the total, not the total. Theorem 4.1 then says only that the unit is κ.

Theorem 6.12 (the cost is the total)**.** *Let a ledger be as in Definition 6.10. On the carrier $\mathbb{Q}$ generated by iterated distinction, the cost of a comparison of two counts is a positive multiple of the total those two entries enter above balance. Hence (C1) to (C4) hold, and $\frac{\kappa}{2}F = J$; no other rational-valued cost satisfies (C1) to (C4). The world M_1 satisfies* Π *and is not a ledger.*

Proof. A comparison of counts $a, b \neq 0$ enters $u = a/b$ and $v = b/a$, with $uv = 1$ (Definition 6.10). Their sum is $u + v$, and the excess over the balanced total 2 is $e = u + v - 2 = (a-b)^2/(ab) = 2G(u)$. For $a, b > 0$ one has $e \geq 0$, with $e = 0$ exactly when $a = b$. The running total of the two entries above balance is therefore e, and by the definition the cost is a positive multiple of e, which is a positive multiple of G. Proposition 6.11(i) gives (C1) to (C4) for G at $\kappa = 2$, and a positive rescaling changes only the gauge; the same part of the proposition shows that every rational-valued cost satisfying (C1) to (C4) is such a multiple. M_1 satisfies Π by Theorem 5.2; its comparison enters neither u nor v, so it is not a running total of those entries. □

What is proven is the conditional: given a ledger, the cost is J up to the unit. That our records form a ledger is taken from the cited construction (Section 4). Relative to Π alone, Mag remains

supplied$_\Pi$ (Theorem 5.2): asking does not force a world to enter the two orientations and total them.

7 What Gödel permits

Let T be the background logic L together with the recognition axioms, taken as a formal theory: L is a recursively axiomatized theory of classical mathematics with the reals, the axioms of Section 4 are finitely many sentences, and the floor's completed orbit (the set of all results of iterating the extension step) interprets Peano arithmetic in T. So T is a consistent, recursively axiomatized theory that interprets arithmetic, and Gödel's two theorems [21] apply to it: there is a true arithmetical sentence T does not prove, and T does not prove its own consistency.

Gödel's results are about what a formal theory can prove from its axioms. The claims at issue are of three kinds: that some axioms hold of the world, that the axioms are consistent, and that the axioms fix every physical quantity. The section asks four questions. Does Gödel's second theorem undermine the consistency of the axioms? It limits where a proof of consistency can be carried out; the regress of justifying the axioms (proving a theory consistent in a stronger theory, then that one in a stronger still) ends in an act, which secures only the three sentences of Π. Does Gödel's first theorem forbid a theory that fixes every quantity? No, since fixing a quantity and deciding a sentence are different things. Can undecidability reach a physical quantity at all? Only through an unbounded quantifier or a limit, and none of the quantities here is defined that way. Does Gödel touch the proof of an axiom about the world? No, because that proof is not a derivation inside the axioms.

7.1 Consistency is relative to the background logic

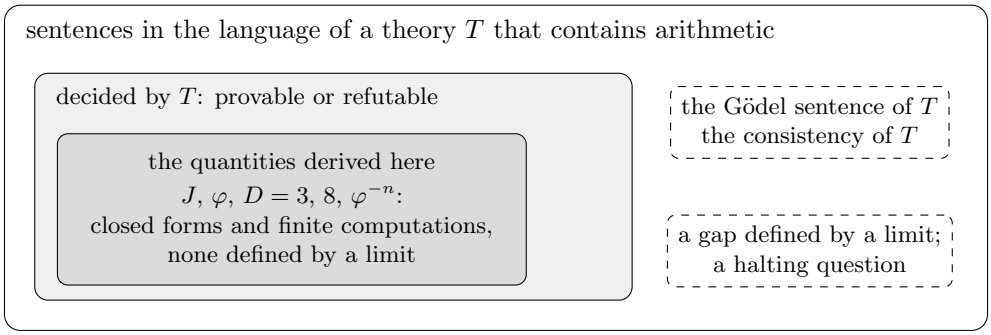
Theorem 4.5 asserts the axioms are consistent by exhibiting a model. That exhibition is a proof in L, and L cannot prove its own consistency. So every derivation in this paper is a proof in L, and this is true of every formal theory of anything, Peano arithmetic included. (That an axiom is true of the world is a different claim, settled by an act; the last subsection of this section returns to it.) Gödel adds that the regress cannot be closed from inside any formal system: the consistency of L can be proved in a stronger metatheory, whose consistency needs a stronger one still.

What the regress can end in is something that is not a sentence. Propositions 3.4 to 3.6 are proofs whose final step is an act, not an axiom: asking a question makes the sentence true of the world in which it is asked (Section 3). And [2] shows that the formal system needed to state the generation of arithmetic is itself made of two-sided rules, that is, of distinctions. So the chain of justifications ends in the act of distinguishing, which is also where the construction began. An act is not a formal system, so no stronger system is needed above it. What the act secures is that the three sentences of Π are true of the world, a different claim from the consistency of L, which stays relative to L as above. Gödel's second theorem says no formal system can be the bottom; it does not say there is no bottom.

7.2 Two completenesses

Gödel's first theorem is about deciding sentences. A theory of everything is about fixing quantities. Two definitions keep them apart.

Definition 7.1. A theory is *syntactically complete* if every sentence in its language is provable or refutable. A physical theory is *physically complete* if every physical quantity is a consequence of its axioms.



syntactic completeness would decide the whole rectangle, and is impossible;
physical completeness asks that every physical quantity lie in the dark box, and is untouched

Figure 11: Two completenesses. Gödel's theorem places some sentences of any theory that counts outside the decided region, and a physical sentence lands there through an unbounded quantifier or a limit. The recognition quantities are all finite objects inside the decided region.

A physical quantity here is a number a measurement can return, such as a mass ratio or the number of spatial dimensions. A theory *fixes* a quantity q when it proves a sentence $q = t$ in which t is a closed term naming a real number by a finite definition: a closed form, a root of a given polynomial, or the output of a finite computation. A quantity is a consequence of the axioms, in the definition above and in Section 8, when the axioms fix it in this sense. A quantity defined by a limit counts once its definition is shown to reduce to such a term, which is the check the next subsection names.

Proposition 7.2. *If a consistent, recursively axiomatized, physically complete theory exists, then physical completeness does not imply syntactic completeness.*

Proof. In outline: add arithmetic to a physically complete theory, and Gödel's sentence appears without any physical quantity changing. Let T be such a theory, and let T' be T together with Peano arithmetic, assumed consistent. T' proves every sentence T proves, so it fixes every quantity T fixed, and it is syntactically incomplete by Gödel's first theorem. Its Gödel sentence is a statement about T''s own proof predicate and is not of the form $q = t$, so the incompleteness it exhibits leaves physical completeness where it was. □

Hawking's argument that Gödel forbids a final theory [22] is a self-reference argument (Figure 11): a physical theory is part of the world it describes, physical systems can encode arithmetic, and so the theory's undecidable sentences can be dressed as statements about physical systems. The argument is right that some sentences about physical systems are undecidable in any theory that can count. It does not reach physical completeness, because the undecidable sentences it produces are about infinite families of systems or about limits (the next subsection), and a theory of everything is asked to fix every quantity, not to decide every sentence. A theory that fixes every meter reading and cannot decide its own Gödel sentence is physically complete and syntactically incomplete, and Gödel's theorem forbids only syntactic completeness to any theory that counts.

7.3 Where undecidability can reach a quantity

On a finite ledger state every property is decidable by inspection. Undecidability enters physics through an unbounded quantifier: over all sizes, over all times, over all steps of a computation that a finite system can encode, or through an infinite limit. Cubitt, Perez-Garcia and Wolf [23] construct a family of translation-invariant Hamiltonians on a two-dimensional lattice for

which no algorithm decides, from the finite description of the local term, whether the infinite-volume system is gapped; the halting problem is the same shape with time in place of size. So the question for any physical theory is which of its quantities are defined by such a quantifier or such a limit.

The quantities of Theorem 4.5 are J (a closed form), φ (a root of a quadratic), $D = 3$ (a finite homology computation), 8 (a count), and φ^{-n} (closed form). Each is derived in T by a finite derivation from the axioms; whether the two supplied axioms (join and kept) hold of the world is the separate question of Section 5. None of the five is defined by a limit or an unbounded quantifier: φ is the positive root of a quadratic, reached from (H1) and (H2) by algebra, and the toy ratios $2/1, 3/2, 5/3, \ldots$ only approach it; the dimension theorem speaks of every D and every embedded circle, and one finite homology computation (Alexander duality) settles all of them. Incompleteness says that some sentences escape T; it says nothing against a sentence T has already derived. Any quantity defined downstream through an infinite limit, an absolute scale for instance, is the place to check, and the check is a finite question about how the quantity is defined.

7.4 What Gödel does not touch

Gödel's theorems concern provability of sentences within a system. The classification contains two kinds of verdict, and he limits neither. A forced verdict is a proof in L from the three sentences of Π (for the cost, together with the definition of a ledger), a finite derivation of a sentence about the floor, the ledger or the rings. Gödel limits what L can prove about itself: its own consistency, and the sentences built from its own proof predicate. None of them is a sentence about L's own proof predicate. That the three sentences are true of the world is not a theorem of L: L derives them from the premise that some act of asking occurs, and the act supplies that premise from outside L. The criterion of Section 3 concerns whether an axiom holds of a world, and its instances, Propositions 3.4 to 3.6, rest on that act. So the forced part of the classification stands. The supplied part stands too: independence by countermodel is the method by which the parallel postulate was shown independent, and Gödel neither helps nor hinders it. What Gödel does is fix the meaning of the phrase in the next section.

8 Theory of everything

Definition 8.1. A *theory of everything* is an axiom set every line of which is proven of the world, together with the fact that every physical quantity is a consequence of it.

The definition has two clauses because the opening paragraph had two questions: are the axioms true of the world, and do they fix everything? Gödel's first theorem fixes the second clause as physical rather than syntactic completeness. The state of each clause for the recognition axioms is exact (Figure 12).

The first clause holds for Π and the floor, holds for the cost on the condition that our records form a ledger, and fails, today, for the two remaining seams. Counted strictly, then, three inputs stand between the floor and the full axiom set: that our records form a ledger, which this paper takes from the cited ledger construction, and the two seams. Item by item:

- Π *and the floor: forced.* Π holds in every world in which anything is asked, and in every such world the floor's extension step is available at every stage whose record is entertained as complete. The completed set of counts is the background logic's addition, from its axiom of infinity.
- Mag*: forced given a ledger.* The cost is the total of the two orientations above balance

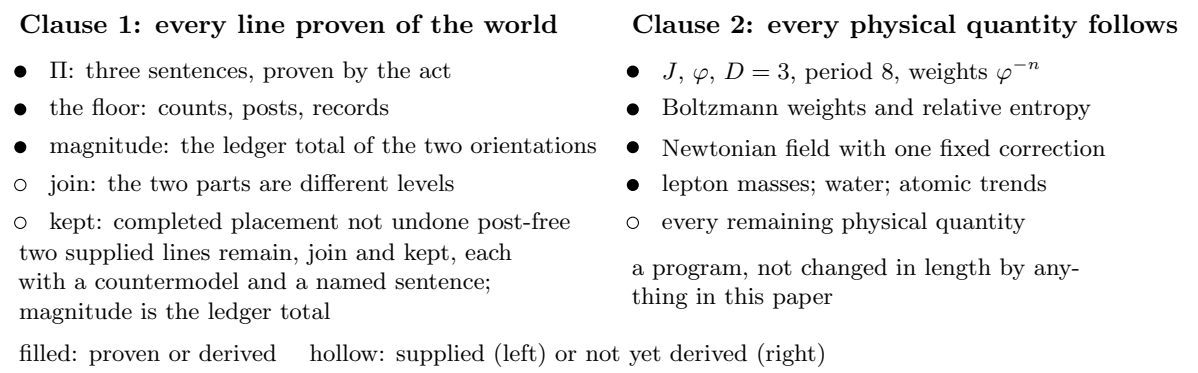


Figure 12: The two clauses and their state. The two hollow lines on the left are what would complete the first.

(Theorem 6.12); given the definition of a ledger it is a theorem. That our records form a ledger is taken from the cited construction, and relative to Π alone the magnitude is supplied.

- Join*: supplied.* Asking yields that a join has two distinct terms; the floor's own doubling tower, which satisfies every sentence of Π, shows where that stops. What is left is P_5, the conjunction of two sentences: the record of a pass is a lower pass, and passes at distinct levels have distinct records. Asking at the signature that adds levels makes neither true (Proposition 6.9).

- Kept*: supplied.* That a comparison has both its terms present is forced. The step to (K) is P_4, whose reader clause is equivalent to (K) (Proposition 6.3). No sentence of the floor's signature yields (K) (Theorem 6.1), and at the signature that adds the reader and post-free motion asking leaves it undecided, with a world for each answer (Proposition 6.4). What supports Kept is inspection of linked rings, which supplies it for the rings we inspect; (K) itself speaks of the traces of an act; and relative to the Alexander clause it is equivalent to the dimension it yields. The countermodel M_3 satisfies every sentence of Π.

Closing the two remaining seams by the performative route, and showing by the same route that our records form a ledger, would make the first clause true in full, and would make the sentence "the recognition axioms are proven true of the world" literal, with no experiment in it.

The second clause is a separate task, and its state is separate. From the axioms follow the cost, the ratio, the dimension, the period, and the weights, and from those the cited papers derive a body of physics that can be named. On a finite set of outcomes the same cost gives the Boltzmann weights (the probabilities that a system in equilibrium assigns to its states) and makes the difference of free energies the relative entropy, a measure of how far one probability distribution lies from another [11]. The ledger gives the Newtonian gravitational field in the limit of instantaneous closure, the limit in which the ledger settles to equilibrium at once after each change. Settling in finite time adds one correction factor, whose exponent is fixed by φ. What that predicts is how fast stars and gas circle a galaxy at each distance from its centre (its rotation curve), which is compared with measured curves with no quantity to adjust galaxy by galaxy [12]. The hierarchy and the eight-step period fix the ratios of the muon and tau masses to the electron mass; the measured electron mass then sets the scale, and the predicted muon and tau masses are compared with the measured ones [13]. At the scale of ordinary measurement the same ladder organizes the timescales of the hydrogen bonds between water molecules in liquid water [14], and the trends of atomic properties across one row (period) of

the periodic table along a single axis [15]. This paper does not restate those comparisons: the data and the size of each agreement are in the cited papers. Whether every physical quantity follows is the remaining program, which nothing in this paper shortens or lengthens. By the standard that accepted Peano's axioms, the recognition axioms are an axiom set with a proven spine and a body of physics part derived and part still to be derived from it.

The two clauses are independent: an axiom set can be wholly proven and generate little, or generate all of physics and rest on a supplied line. Two sentences stand at the remaining seams: the join premise Join, that a new level is the join of the two levels below it, one of each and not two copies of one; and the kept premise (K), that a completed placement is not undone by a motion that posts nothing. That our records form a ledger is the third input, the one M_1 lacks, and it is taken from the cited construction. Three identifications stand where a proof would go: that entertaining a sentence is comparing it with its negation in the world it is about, which carries asking into the world; the Alexander clause, which carries kept to a dimension; and the one bit per dimension, which carries the dimension to a period. Each is one named sentence, with a countermodel or an identification in place of a derivation.

References

[1] Washburn, J. and Zlatanović, M. Lj. (2026) The δ-calculus: from distinction to arithmetic. *Advances in Pure Mathematics*, in press. arXiv:2607.29349.

[2] Washburn, J. (2026) Distinction before logic: the metatheory of the δ-calculus, measured. Recognition Physics Institute, preprint.

[3] Pardo-Guerra, S. and Washburn, J. (2026) Three dimensions from a recognition requirement. Preprints.org. doi:10.20944/preprints202609.1990.v1.

[4] Washburn, J. and Zlatanović, M. (2026) Uniqueness of the canonical reciprocal cost. *Mathematics*, 14(6):935. doi:10.3390/math14060935. arXiv:2602.05753.

[5] Pardo-Guerra, S., Thapa, A., Simons, M. and Washburn, J. (2026) Coherent comparison as information cost: axiomatic foundations for discrete ledger dynamics. *Foundations*, 6(2):17. doi:10.3390/foundations6020017. arXiv:2601.12194.

[6] Washburn, J., Zlatanović, M. and Allahyarov, E. (2026) The d'Alembert inevitability theorem. *Mathematics*, 14(8):1386. doi:10.3390/math14081386. arXiv:2603.16237.

[7] Pardo-Guerra, S., Washburn, J. and Allahyarov, E. (2026) Admissible reciprocally symmetric costs: combiner existence and classification. *Mathematics*, 14(12):2157. doi:10.3390/math14122157.

[8] Pardo-Guerra, S., Washburn, J. and Allahyarov, E. (2026) A coherence theorem for conserved comparison ledgers: structural axioms pin down the scale, the cost, and the ratio. *Mathematics*, 14(15):2672. doi:10.3390/math14152672.

[9] Washburn, J., Pardo-Guerra, S. and Zlatanović, M. (2026) Reciprocal cost on the positive rationals: character moduli, regular extension, and rational rank. arXiv:2608.12418.

[10] Washburn, J. and Zlatanović, M. (2026) Golden and metallic structures on Hessian manifolds. *Mathematics*, 14(14):2483. doi:10.3390/math14142483. arXiv:2606.02150.

[11] Simons, M. and Washburn, J. (2026) A finite-state Gibbs construction from a recognition cost. *Journal of Statistical Mechanics: Theory and Experiment*, in press. arXiv:2605.15667.

[12] Simons, M., Allahyarov, E. and Washburn, J. (2026) A discrete informational framework for classical gravity: ledger foundations and galaxy rotation curve constraints. *Entropy*, 28(4):477. doi:10.3390/e28040477.

[13] Washburn, J. and Allahyarov, E. (2026) Charged lepton masses from the recognition composition law: a derivation with zero continuously adjustable dimensionless parameters. *Symmetry*, 18(6):962. doi:10.3390/sym18060962.

[14] Washburn, J. and Allahyarov, E. (2026) A golden-ratio ladder and a delocalisation-saturated participation bridge for the hydrogen-bond network of liquid water. *Molecules*, 31(15):2701. doi:10.3390/molecules31152701.

[15] Washburn, J., Simons, M. and Allahyarov, E. (2026) A noble-gas-centered coordinate for within-period atomic property trends. *Symmetry*, 18(7):1087. doi:10.3390/sym18071087. arXiv:2605.00028.

[16] Washburn, J., Zlatanović, M. and Allahyarov, E. (2026) Recognition geometry. *Axioms*, 15(2):90. doi:10.3390/axioms15020090.

[17] Beltrami, E. (1868) Saggio di interpretazione della geometria non-euclidea. *Giornale di Matematiche*, 6, 284–312.

[18] Hintikka, J. (1962) Cogito, ergo sum: inference or performance? *The Philosophical Review*, 71(1), 3–32. doi:10.2307/2183678.

[19] Aczél, J. and Dhombres, J. (1989) *Functional Equations in Several Variables*. Encyclopedia of Mathematics and its Applications 31, Cambridge University Press. doi:10.1017/CBO9781139086578.

[20] Hatcher, A. (2002) *Algebraic Topology*. Cambridge University Press.

[21] Gödel, K. (1931) Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38, 173–198. doi:10.1007/BF01700692.

[22] Hawking, S. W. (2002) Gödel and the end of physics. Lecture, Dirac Centennial Celebration, University of Cambridge, 20 July 2002.

[23] Cubitt, T. S., Perez-Garcia, D. and Wolf, M. M. (2015) Undecidability of the spectral gap. *Nature*, 528, 207–211. doi:10.1038/nature16059.