The Depth of the Maze

How the Maintenance Ratio Fixes Route Depth, Step Distance and Tempo

J. Konstapel, Leiden, 18 August 2026

Context Header

Programme. MAZE works out what the vacuum.net model means for the human being. Vacuum.net is the model of the whole net. MAZE is its human correspondence.

This piece. It joins two lines that were developed apart: *Paths to the Knot* (the general maintenance theory) and the MAZE kernel (the ternary route address). The join is arithmetic, not analogy. Every claim below is either proved or labelled.

Settled. The Living Winding theory phase is closed at edition 8. The MAZE kernel implements the balanced-ternary bijection with tests against the paper's exact numbers. *Paths to the Knot* states the maintenance ratio and the depth recurrence.

Open before this piece. Route depth was a free parameter. The step-distance rule used a postulated constant, 1.618. The choice between the route ceilings 40 and 43 was undecided.

Conventions. All quantities are defined in the text. No reference needs to be fetched to follow a calculation. Status of each result is stated in the ledger at the end.

1. What the Maze does not yet say

The Amazing Maze computes three things from a birth date and a birth place: a profile, the next step at the right distance, and with whom you build.

Three numbers in that computation are still unearned.

The first is depth. A route in the Maze is a sequence of turns. How long may that sequence be? The kernel imposes no limit. But a person is not an unbounded register.

The second is step distance. The current rule sets the next step at 1.618 times the current one. The number comes from a fit to fifty years of transition data. A fit is not a derivation.

The third is the route ceiling. The kernel counts 1, 4, 13, 40, 121. The Bronze Mean line counts 1, 1, 4, 13, 43, 142. They agree up to 13 and split after it. Nothing decided which is right.

All three are fixed below, from a single quantity that a person can measure.

2. The two objects, stated in full

This section makes the article self-contained. A reader who knows neither line can follow from here.

2.1 The knot

A knot is a pattern of connections that persists because it is continually re-laid. Connections lapse when nothing runs across them. New connections are laid where flow runs. Persistence is therefore not storage. It is a balance between re-laying and lapse.

Two structural quantities describe a knot.

Form is which connections exist and how they are arranged. It answers: what makes this knot this one?

Depth is the number of times the pattern has been closed over constituent patterns. A knot of knots has depth one greater than its constituents.

Lay the contacts out along a line. Two contacts either *cross* (their endpoints alternate) or *nest* (one lies wholly inside the other). Crossing structure is form. Nesting structure is depth.

2.2 The four rates

Four quantities carry the whole theory.

- **C** — the number of contacts the maintained pattern requires.
- **λ** — the rate at which an idle contact lapses.
- **r** — the rate at which a lapsed contact is restored.
- **f** — the fraction of time in which flow is near enough to zero for repair to run.

The maintenance ratio is

$R \equiv r / \lambda$

R is dimensionless. It compares restoration against decay. Absolute speed does not appear.

2.3 The Living Winding form of the same ratio

The human projection of the model adds two rates. **μ** is the number of contacts laid per unit of flow. **ν** is the number of non-conserving re-lay events per unit of flow — drift. With throughflow J, edition 3 of the Living Winding derives

$C^*(J) = \mu rJ / [\lambda(\nu J + r)]$ and $f = \nu J / (\nu J + r)$

These will be used in Proposition 1.

2.4 The Maze address

A turn in the Maze is one of three: over (+1), straight (0), under (−1). A route is a string of turns. The kernel maps a route of length n to an integer by balanced ternary. The example in the source: $(+1, -1, -1, -1, +1) = 81 - 27 - 9 - 3 + 1 = 43$.

A route of length n therefore addresses 3^n distinct positions, symmetric about zero, with largest address

$(3^n - 1) / 2 = 1, 4, 13, 40, 121, 364, \ldots$

Rest is address zero. Mirroring a route negates its address. Leading zeros do not change the address, which is why the same identity appears at different scales.

3. Proposition 1 — R is the ceiling of the rest curve

Claim. The maintenance ratio R is the plateau of the load–capacity curve, and the stationary form satisfies $C^* = R \cdot f$ exactly.

Proof. Take the two Living Winding expressions of §2.3. Divide the first by the second:

$$C^*/f = [\mu r J / (\lambda(\nu J + r))] \cdot [(\nu J + r) / (\nu J)] = \mu r / (\lambda\nu)$$

The factor $(\nu J + r)$ cancels, and J cancels. The right-hand side contains no J. Define

$$\mathbf{R \equiv \mu r / (\lambda\nu)}$$

Then $C^* = R \cdot f$ for every J. As $J \to \infty$, $f \to 1$, so $C^* \to R$. R is therefore the asymptotic plateau of $C^*(J)$. ∎

Consequence. The knot paper writes $R = r/\lambda$. That is the same object with μ and ν absorbed into the units of r and λ. The convention should be fixed on the plateau, because the plateau is what can be measured.

Why this matters for MAZE. The Self-Calibration Guide already instructs a person to trace the load–rest curve and read off its ceiling, and already calls that ceiling the most valuable number in the exercise. Proposition 1 says what that number is. It is R. Nothing new needs to be measured.

4. Proposition 2 — the operative bound is C ≤ R·f, not C ≤ R

Claim. The maintainable form is bounded by $R \cdot f$, and the bound $C \leq R$ is its unreachable limit.

Proof. By Proposition 1, $C^* = R \cdot f$. A pattern persists only if the required C does not exceed the stationary maintained C^*. Hence $C \leq R \cdot f$.

Now $f < 1$ strictly. Repair runs only in low-flow phases, and flow is what lays contacts. A system at $f = 1$ lays nothing and has $C = 0$. Therefore $R \cdot f < R$, and $C \leq R$ is attained only in a limit that destroys the system it describes. ∎

Consequence. Two quantities set the ceiling, not one. Capacity and rest multiply. A person with a high plateau and no rest holds no more form than a person with a modest plateau and a protected rest fraction.

Paths to the Knot states $C \leq R$ in §4.3 and then uses $R \cdot f$ in the depth formula of §4.5. The second use is the correct one. §4.3 should be corrected to match it.

5. Proposition 3 — the Maze route count is the depth recurrence

Claim. The Maze route ceiling $(3^n - 1)/2$ is exactly the knot-theoretic form recurrence with q = 3 constituents and one new contact per closure.

Proof. The knot recurrence for form under closure is

$C(k+1) = q \cdot C(k) + \zeta q$

where q constituent forms are gathered and ζq new contacts are laid. Set q = 3 and ζq = 1:

$C(k+1) = 3 \cdot C(k) + 1, C(0) = 1$

By induction, C(k) = (3^(k+1) − 1)/2. Base: k = 0 gives (3 − 1)/2 = 1. ✓ Step: assume C(k) = (3^(k+1) − 1)/2. Then

C(k+1) = 3·(3^(k+1) − 1)/2 + 1 = (3^(k+2) − 3 + 2)/2 = (3^(k+2) − 1)/2 ✓ ∎

The sequence is 1, 4, 13, 40, 121, 364, 1093.

Consequence. The Maze address space is not merely *like* a nested knot. It *is* the depth recurrence of a ternary closure. The three turns are the three constituents. The single new contact per closure is the minimal closure that still makes a knot of knots.

This also names what a Maze ring is. A ring is one closure. Passing from ring k to ring k+1 is one act of closing three constituent routes into a higher route.

6. Proposition 4 — maximum route depth from one measured number

Claim. The greatest depth a knot can hold under ternary closure is

$$\mathbf{m = \lfloor \log_3(2C + 1) \rfloor - 1}$$

where C is the maintained form, itself bounded by R·f.

Proof. Depth m is available when the form at that depth is still maintainable: $C(m) \le C$. By Proposition 3, C(m) = (3^(m+1) − 1)/2. So

$(3^{(m+1)} - 1)/2 \le C \iff 3^{(m+1)} \le 2C + 1 \iff m + 1 \le \log_3(2C + 1)$

The greatest integer m satisfying this is $\lfloor \log_3(2C + 1) \rfloor - 1$. ∎

The bands. Because the ceilings are sparse, m is a step function with wide flat stretches.

maintained form C	route depth m	route ceiling at m
1 – 3	0	1
4 – 12	1	4
13 – 39	2	13
40 – 120	3	40

121 – 363	4	121
364 – 1092	5	364

Three consequences follow immediately.

Depth is small. For any human-scale C between 13 and 363, depth is 2, 3 or 4. The Maze is short. Not because anyone chose it short, but because form grows geometrically and capacity does not.

Depth is sticky. Moving from depth 3 to depth 4 requires tripling C. Small gains in capacity buy nothing. This is why deep personal change resists incremental effort, and it is a statement about arithmetic, not about willpower.

Depth is countable without instruments. C can be counted. Define it operationally: the number of distinct commitments a person maintains simultaneously without lapse over the retention horizon. Roles, relationships, skills, practices, obligations — anything that decays if left alone. That count gives m directly.

7. Worked example

A person counts the commitments that would visibly lapse if abandoned for ninety days. The count is **C = 47**.

Depth. $2C + 1 = 95$. $\log_3 95 = \ln 95 / \ln 3 = 4.5539 / 1.0986 = 4.145$. Floor is 4. So $m = 3$. The person holds four nested levels, $k = 0$ through 3.

Rest fraction. From the twice-daily log over three weeks, the fraction of readings in the low-flow band is **f = 0.32**.

Required plateau. By Proposition 2, $C \leq R \cdot f$ requires $R \geq C/f = 47 / 0.32 =$ **147**.

Measured plateau. The load–rest curve levels off at **R = 120**.

Verdict. $R \cdot f = 120 \times 0.32 = 38.4$. The maintained form, 47, exceeds it. The person is above the bound by 22 per cent.

What the theory predicts next. Loss begins at the deepest maintained layer, not at the smallest item. The most integrated commitment goes first. This is the ordering prediction of the knot theory, and it is the opposite of the intuitive expectation that trivial things are dropped first.

Two exits, both computable.

Reduce form. Drop C below 38.4. But 38.4 falls in the band 13–39, so depth drops from 3 to 2. One level is lost. The edge at 40 is sharp: at $C = 40$ the person holds four levels, at $C = 39$ only three.

Raise rest. Keep $C = 47$ and raise f to $47/120 =$ **0.392**. That is roughly one extra low-flow hour in sixteen waking hours. It keeps the fourth level.

The second exit is cheaper and it is invisible to any account that measures only workload. That is the practical content of Proposition 2.

8. Proposition 5 — the fork between 40 and 43

Claim. The Bronze Mean sequence 1, 1, 4, 13, 43, 142 is the depth recurrence of a closure that also retains contacts to the level two below. The fork between 40 and 43 is exactly the presence or absence of that second-order term.

Proof. Let a closure gather $q = 3$ constituents at level k and additionally retain the contacts of level $k - 1$. The recurrence becomes second order:

$$C(k+1) = 3 \cdot C(k) + C(k-1), \quad C(0) = C(1) = 1$$

Compute: $C(2) = 3 \cdot 1 + 1 = 4$. $C(3) = 3 \cdot 4 + 1 = 13$. $C(4) = 3 \cdot 13 + 4 = \mathbf{43}$. $C(5) = 3 \cdot 43 + 13 = \mathbf{142}$. $C(6) = 3 \cdot 142 + 43 = 469$. ∎

Closed form. The characteristic equation is $x^2 = 3x + 1$, with roots

$$\beta = (3 + \sqrt{13})/2 = 3.302776 \text{ and } \beta' = (3 - \sqrt{13})/2 = -0.302776$$

Fitting $C(0) = C(1) = 1$ gives $C(k) = A \cdot \beta^k + (1-A) \cdot \beta'^k$ with $A = (\sqrt{13} - 1)/(2\sqrt{13}) = 0.361325$. This reproduces 1, 1, 4, 13, 43, 142, 469 exactly.

Since $|\beta'| < 1$, the second term vanishes and $C(k) \to 0.361325 \cdot \beta^k$. The growth constant is the Bronze Mean, 3.302776, in place of 3.

Consequence 1 — the fork is now one structural question. Does a closure keep contacts to its constituents' constituents? Yes gives 43. No gives 40. This is not a choice between two attractive sequences. It is a property of the closure operation, and it is implementable: in the kernel, does an address at depth n retain a reference to depth $n - 2$?

Consequence 2 — what 43 costs. The balanced-ternary bijection produces exactly $(3^n - 1)/2$. The number 43 is not of that form. Choosing memory-carrying closure therefore gives up the exact ternary address space. You cannot have both an exact ternary bijection and a closure with memory. That is the real price of the fork, and it should be paid consciously.

Consequence 3 — where the fork is decidable. The two sequences agree at 1, 4 and 13. They separate only at the fourth closure. Any system that never reaches depth 4 cannot distinguish them. By Proposition 4, human C must exceed 121 before depth 4 is held. The fork is therefore decidable only at the upper edge of the human range — which is precisely why it stayed open.

9. Proposition 6 — the tempo spectrum is rational, and ϕ is not in it

9.1 The tempo ratio

Suppose q constituent knots each hold b open contacts of comparable conductance. Closing them lays ζ new contacts per constituent and turns the corresponding open connections inward. The fraction of openness taken inward is $\phi\_in = 2\zeta/b$. Retention time is capacity divided by openness, so the ratio of adjacent retention times is

$\sigma = 1 / (1 - 2\zeta/b) = b / (b - 2\zeta)$

Any real closure has $\zeta \geq 1$, so $\sigma > 1$. A higher level always holds longer. And $2\zeta < b$ is required: a living knot must keep some boundary open.

9.2 The spectrum is discrete and its rungs are priced

Claim. σ is rational for every admissible closure, and each value carries a minimal cost b.

Proof. b and ζ are contact counts, hence positive integers. Write $d = b - 2\zeta \geq 1$. Then $\sigma = b/d$ is a ratio of integers. Note $d \equiv b \pmod 2$, since d differs from b by an even number. ∎

Cost. Given a target $\sigma = p/q$ in lowest terms, the cheapest closure realising it has

- $b = p$, $d = q$, if p and q are both odd;
- $b = 2p$, $d = 2q$, otherwise (parity must match).

The cost of a tempo ratio is the openness b it demands of each constituent. Cheap ratios need few open contacts. Expensive ratios need many.

9.3 The ladder

Enumerating all admissible (b, ζ) with $b \leq 16$ and listing the cheapest realisation of each value gives the following rungs in the band 1.2 to 2.5:

σ	cheapest (b, ζ)	cost b
1.200	(12, 1)	12
1.222	(11, 1)	11
1.250	(10, 1)	10
1.286	(9, 1)	9
1.333	(8, 1)	8
1.400	(7, 1)	7
1.500	**(6, 1)**	**6**
1.571	(11, 2)	11
1.600	(16, 3)	16
1.667	**(5, 1)**	**5**
1.800	(9, 2)	9
2.000	**(4, 1)**	**4**
2.200	(11, 3)	11
2.333	(7, 2)	7
2.500	(10, 3)	10

Three cheap rungs dominate the band: **2 at cost 4, 5/3 at cost 5, 3/2 at cost 6.**

9.4 φ is unattainable, and 5/3 is the cheapest thing near it

Claim. No admissible closure has $\sigma = \phi = (1+\sqrt{5})/2$, and the cost of approaching ϕ grows without bound.

Proof. ϕ is irrational. Every $\sigma = b/d$ is rational. Hence $\sigma \neq \phi$ for every admissible closure.

The convergents of ϕ are the Fibonacci ratios 2/1, 3/2, 5/3, 8/5, 13/8, 21/13, 34/21. Their costs by the rule of §9.2 are 4, 6, 5, 16, 26, 21, 68. The sequence is unbounded, since numerators grow geometrically and cost is at least the numerator. ∎

The gap. For every $b \leq 10$ there is no admissible σ strictly between 3/2 and 5/3. The first entry inside that interval is 11/7 = 1.571 at cost 11. The value 1.6 first appears at cost 16.

The postulated constant 1.618 lies inside this gap. It sits in the one stretch of the band that cheap closures cannot reach.

The reinterpretation. The step-distance rule was fitted to pooled transition data. Pooling people with different b smears a discrete ladder into a smooth peak. The peak lands between the cheapest rungs bracketing it. The two cheapest rungs bracketing 1.618 are 3/2 and 5/3. A pooled fit between them is exactly what a smoothed ladder produces.

So the claim is not that 1.618 is wrong as a description of pooled data. The claim is that it is not a mechanism. **5/3 is a mechanism**: it is the cheapest closure whose tempo ratio approximates ϕ, and it costs five open contacts per constituent.

9.5 The test

This is decidable on data already held.

Take the fifty years of transition data. Do not pool. Split by any variable that plausibly tracks b — the number of simultaneously open commitments, for instance, which §6 already makes countable.

- **Ladder prediction.** Within homogeneous groups, transition ratios cluster at 3/2, 5/3 and 2, with a depleted stretch between 1.5 and 1.667.
- **Continuous prediction.** Within homogeneous groups, transition ratios remain a single smooth mode at 1.618, as they are in the pooled data.

The rival account is not weak. Continuous optimisation genuinely produces ϕ in phyllotaxis, through a dynamical mechanism of successive placement at maximal avoidance. If human transitions work that way, the smooth peak survives disaggregation. If they work by closure, it does not.

The test is a re-analysis. No new data collection.

10. Proposition 7 — ternary closure predicts a tempo ratio of exactly 3

Claim. If the Maze turn is genuinely ternary, adjacent retention times stand in the ratio 3.

Proof. The Maze constituent offers three directions: over, straight, under. So $b = 3$. The minimal closure lays $\zeta = 1$. Then

$\sigma = b/(b - 2\zeta) = 3/(3 - 2) = 3$ ∎

This is the cheapest closure that exists. No admissible closure has $b < 3$, since $2\zeta < b$ with $\zeta \geq 1$ forces $b \geq 3$. Ternary is the minimum, and it produces the maximum tempo ratio available at minimum cost.

The measurement. The Self-Calibration Guide already obtains retention time τ from a step response: $\tau = t\_half \times 1.44$. Run it at two scales — a day-scale disturbance and a season-scale disturbance — and divide.

Prediction: the ratio is 3, not 2.7 and not 4.

Worked forward. A person with a fastest level $\tau_0 = 0.5$ day and depth $m = 3$ should show retention times of 0.5, 1.5, 4.5 and 13.5 days. Those are four distinguishable timescales, and 13.5 days is a fortnight — the interval at which the deepest layer of an ordinary life turns over. That number is a prediction, and it is checkable with a notebook.

Why this test is decisive. No existing maintenance vocabulary predicts an integer here. Resilience, proteostasis, allostatic load and homeostasis are all compatible with any ratio whatsoever. An integer tempo ratio is a signature of closure with countable contacts. It is the one prediction in the framework that cannot be produced by relabelling.

11. Proposition 8 — real knots must show a nesting deficit

Claim. Living knots depart strongly from a random chord diagram, in a specific and measurable direction.

Proof sketch. In the combinatorics of chord diagrams, crossings and nestings are equidistributed: the joint distribution is symmetric under exchange. In a random matching, therefore, the expected number of nestings equals the expected number of crossings.

The theory of this article makes crossings and nestings grow differently. Form — crossing structure — is bounded by $R \cdot f$, a capacity that can be large. Depth — nesting structure — is bounded by $\lfloor \log_3(2C+1) \rfloor - 1$, which for C in the hundreds is 3 or 4.

So form scales with C while depth scales with log C. Their ratio diverges. ∎

The test. For any system whose contacts can be laid out along a line — a route, a schedule, a genome, a text, a dependency graph — compute the crossing number and the nesting number. Compare against the symmetric null.

Prediction: a systematic and large nesting deficit. Real maintained patterns are wide and shallow, not deep. A system that matched the symmetric null would be evidence against the closure account.

This is the eighth prediction of the framework and it was not previously stated.

12. What can now be done

Five items, in order of cost.

One. Decide the fork in the kernel. Implement both closure variants in `maze.ts`. First order, keeping only contacts to level k−1, ceiling 40. Second order, also retaining level k−2, ceiling 43.

Run the existing bijection tests against both. The variant that leaves the balanced-ternary address structure intact is the first. If the second is chosen, the address scheme must be rebuilt. One afternoon of work decides a question that has been open since the Bronze Mean line opened.

Two. Add depth to the Self-Calibration Guide. One new instruction: count the commitments that would lapse if abandoned for the retention horizon. One new formula: $m = \lfloor \log_3(2C + 1) \rfloor - 1$. One new table: the bands of §6. The reader gets a number nobody has ever handed them — how many nested levels they are currently carrying, and how far they are from losing one.

Three. Run the overload check. With C counted and f read from the existing log, compute C/f and compare against the measured plateau R. Where C exceeds R·f, the theory predicts loss beginning at the deepest layer. That prediction is specific enough to be wrong, and a person can watch for it.

Four. Disaggregate the transition data. Split by open-commitment count. Look for spikes at 3/2, 5/3 and 2, and for a depleted stretch between 1.5 and 1.667. Either the ladder appears or the smooth peak survives. No new data is required.

Five. Measure the tempo ratio. Two step responses, two scales, one division. Expect 3.

13. What would falsify each result

The algebra of Propositions 1 through 6 is proved and cannot be falsified by observation. What can fail is the identification of these structures with the human case. Each identification has a failure condition.

Result	Fails if
R is the rest-curve	the load–capacity curve does not saturate, or saturates at a value that does
$C \leq R \cdot f$ binds (P2)	people are routinely observed maintaining form well above their measured
Maze count = depth recurrence (P3)	nothing; this is an identity of two sequences
Depth bands (P4)	measured level counts do not step at 13, 40, 121, or step at other values
Second-order closure	nothing; this is arithmetic. The *choice* fails if the kernel cannot carry level
Tempo spectrum is rational (P6)	nothing; b and ζ are counts. The *application* fails if disaggregated transition ratios stay smoothly unimodal at 1.618
Ternary gives $\sigma = 3$	measured adjacent retention ratios cluster away from 3 in homogeneous
Nesting deficit (P8)	measured crossing/nesting ratios match the symmetric null

14. Status ledger

Proved. Propositions 1, 2, 3, 4, 5, 6 and the arithmetic of 7. These are algebraic and combinatorial results. They follow from the stated definitions and nothing else.

Conditionally derived. Proposition 7 as a statement about human retention times. It follows if the Maze turn is genuinely ternary at the level being measured.

Identification. That R measured as the rest-curve plateau is the same R that bounds form. Proposition 1 proves this within the Living Winding equations. It carries to a real person only if those equations describe that person.

Hypothesis. Proposition 8, the nesting deficit, and the ladder reading of the transition data in §9.5.

Open. The absolute level count. Proposition 4 gives m from C, but confirming m independently requires counting nested levels in a person directly, and no procedure for that is yet written. This is the next thing to build.

Superseded. The step-distance constant 1.618 as a mechanism. It survives as a description of pooled data, which is a different claim.

Annotated references

Konstapel, J. (2026). ***Paths to the Knot: Complete Edition.*** **MAZE-project, Leiden.** The source of the maintenance object, the ratio R, the closure recurrence and the depth bound. *Why read it:* it supplies the recurrence $C(k+1) = qC(k) + \zeta q$ used in Propositions 3 and 5, and the tempo relation $\sigma = b/(b-2\zeta)$ used in Propositions 6 and 7. Read §4 first; the rest is application.

Konstapel, J. (2026). ***The Living Winding: One Model, Seven Variables, and a Mapping onto the Sciences of the Human.*** **Edition 8.** The human projection of the vacuum.net model, closed after seven adversarial test rounds. *Why read it:* Proposition 1 is proved from its equations for C*(J) and f. Read §5 for the window and §6 for the failure map.

Konstapel, J. (2026). ***The Self-Calibration Guide: Measuring Your Own Living Winding.*** The measurement procedures assumed throughout: the tension score, the step response for τ, the rest curve for f and its plateau. *Why read it:* it is the instrument this article adds a fifth reading to. Read the τ procedure before attempting Proposition 7.

Konstapel, J. (2026). "Everything Has Its Number in the Maze." MAZE-project, Leiden. The balanced-ternary address, the bijection route ↔ integer, the ceiling $(3^n-1)/2$. *Why read it:* it contains the worked example $81 - 27 - 9 - 3 + 1 = 43$ and the mirroring rule used in §2.4.

Chen, W. Y. C., Deng, E. Y. P., Du, R. R. X., Stanley, R. P., & Yan, C. H. (2007). Crossings and nestings of matchings and partitions. ***Transactions of the American Mathematical Society*****, 359(4), 1555–1575.** Establishes that crossings and nestings are equidistributed in matchings and set partitions. *Why read it:* it is what makes form and depth countable on equal footing, and it supplies the symmetric null against which Proposition 8 predicts a deficit. Reading advice: the main symmetry theorem is enough; the bijective machinery can be skipped.

Eigen, M. (1971). Selforganization of matter and the evolution of biological macromolecules. ***Naturwissenschaften*****, 58(10), 465–523.** The error threshold: maintainable information is bounded by copying fidelity. *Why read it:* it is the one place where a bound of exactly this shape was established independently and confirmed empirically. It is the strongest existing support for Proposition 2, obtained without any knowledge of this framework.

Shannon, C. E. (1948). A mathematical theory of communication. ***Bell System Technical Journal*****, 27, 379–423, 623–656.** Reliable retention over a noisy channel, with capacity as the binding quantity. *Why read it:* the structural parallel to $C \leq R \cdot f$ is exact, and Shannon's treatment of the rate–reliability trade-off is the cleanest available model for how to state such a bound.

Dijkstra, E. W. (1974). Self-stabilizing systems in spite of distributed control. *Communications of the ACM*, 17(11), 643–644. Systems that return to a correct state from an arbitrary one without external intervention. *Why read it:* two pages, and it is the engineering proof that maintenance without a stored master copy is achievable rather than merely desirable.

Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society*, 106(6), 467–482. Near-decomposability: hierarchies with weak coupling between levels and strong coupling within. *Why read it:* it is the strongest existing motivation for the closure assumption behind Propositions 3 and 7, and Simon's parable of the two watchmakers is the shortest argument for why depth exists at all.

Balch, W. E., Morimoto, R. I., Dillin, A., & Kelly, J. W. (2008). Adapting proteostasis for disease intervention. *Science*, 319(5865), 916–919. The cell as continuously re-laid proteome, with disease as failure of maintenance capacity. *Why read it:* it is the clearest case of a field independently arriving at R without naming it, and the strongest evidence that the ratio is not a metaphor.

Schegloff, E., Jefferson, G., & Sacks, H. (1977). The preference for self-correction in the organization of repair in conversation. *Language*, 53(2), 361–382. Repair as the organising mechanism of conversation, with a structural preference for self-repair. *Why read it:* it shows repair rates being measured decades before anyone thought to divide them by decay rates.

Borbély, A. A. (1982). A two process model of sleep regulation. *Human Neurobiology*, 1(3), 195–204. Process S accumulates during waking and dissipates during sleep. *Why read it:* Process S is the rest fraction f under another name, and the distinction between time-based and load-based accumulation is the sharpest available test of Proposition 2 in an existing literature.

Douady, S., & Couder, Y. (1992). Phyllotaxis as a physical self-organized growth process. *Physical Review Letters*, 68(13), 2098–2101. The golden angle emerging from a continuous dynamical process of successive placement. *Why read it:* it is the serious rival to Proposition 6. It shows that ϕ can arise as a genuine continuous optimum rather than as a smeared ladder. Read it before running the disaggregation test of §9.5, so that the alternative outcome is understood in advance.

Hardy, G. H., & Wright, E. M. (2008). *An Introduction to the Theory of Numbers* (6th ed.). Oxford University Press. Continued fractions and the convergents of ϕ. *Why read it:* chapters X and XI supply the cost sequence 4, 6, 5, 16, 26, 21, 68 used in §9.4, and the standard proof that ϕ is irrational, which is what makes Proposition 6 a proof rather than an observation.

Hurwitz, A. (1898). Über die Composition der quadratischen Formen von beliebig vielen Variabeln. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, 309–316. The classification of normed division algebras at dimensions 1, 2, 4, 8. *Why read it:* it is the other branch of the fork noted in the blueprint line, splitting from the Bronze Mean sequence at 8 versus 13. Proposition 5 does not settle that fork, but it changes its shape: 13 is now derived from a closure rule rather than chosen from a sequence.

Prepared for the MAZE-project, Leiden. The Dutch companion blog follows.