The Twin Test

A public preregistration of an experiment on the origin of information

José Santiago Niño Mojica · Founder and inventor, Recognitium · Paris, September 2026

La vérité ne connaît point de contenu, mais de ce qui passe.
Truth knows no content, only what passes.

Version 1.0 · 5 September 2026 · Licence CC BY 4.0

ABSTRACT

Two physical instruments are built as identically as engineering allows and asked the same question at the same instant. Each answers with a block of raw bits drawn from a physical noise source, and with a digest of that block. The experiment measures how far apart the two answers fall, and whether that distance shrinks as the two instruments are made more and more identical, down to the limit where the two instruments are one instrument restarted into the same state. The construction hypothesis predicts convergence. The origination hypothesis predicts a distance that does not move. This document fixes the hypotheses, the observables, the sample size, the thresholds, the controls and the list of results that would falsify the claim, before the first measurement is taken. Its fingerprint is committed to a public receipt chain and anchored in the Bitcoin blockchain on the day of publication. The construction of the instrument is not described here. It is committed by fingerprint alongside this text and kept private.

Keywords. preregistration · physical randomness · min-entropy · origination · type and token · metrology · tamper-evident records

1. THE FOUNDER AND THE QUESTION

I was born in Bogotá in 2003. When I was eight years old I asked my father what happens after death, and he answered that nothing does: you close your eyes and fall into nothingness. I have been working against that sentence since. I came to Paris in 2022 to study philosophy at the Sorbonne, working nights to pay for it, and the question changed form without changing substance. It became: what would it take for a thing that happened to remain a fact, once everyone who saw it is gone and every record of it can be rewritten by whoever holds the disk?

The answer I arrived at is not a philosophical one. It is an instrument. In January 2026 I published a first theoretical text on the limits of governance over digital records. Between March and August 2026 I filed four patent applications with the French patent office for a physical device that makes the record of an event impossible to reorder, repeat or backdate; two of them were examined under national-defence review before being released in May 2026. In May 2026 the project was admitted to the Fighters programme of Station F, and in July 2026 the French labour administration for the Paris region recognised its innovative character. Recognitium, the company that carries the work, is being incorporated in Paris this month.

The Twin Test is the experiment on which the whole construction rests, and it is also the oldest quarrel in philosophy put on a bench. Parmenides held that what is, is one and the same with itself: the copy is the thing again. Heraclitus held that no river is stepped in twice. For two and a half thousand years the argument was conducted in language, which is a machine for making copies. This is an attempt to conduct it, for the first time, with an instrument.

2. THE INSTRUMENT, PRESENTED LITERALLY

Recognitium's machine is a small sealed box. You send it thirty-two bytes, the fingerprint of anything: a contract, a payment, a decision taken by an artificial intelligence. About a millisecond later it answers with a receipt. That is all it does.

Inside that millisecond, a physical event occurs that has not occurred before and will not occur again. A block of raw bits is drawn from a physical noise source, after the inputs have been locked and before anything is said. The block is compressed into a fingerprint. The fingerprint is welded to your thirty-two bytes, to the instant, and to the fingerprint of everything the box has ever witnessed before. Then the block is destroyed, and only after its destruction has been sensed does the box speak.

The receipt has five public fields and one public formula:

```
new_tip = SHA-256( prev_tip ‖ fingerprint ‖ payload_hash ‖ timestamp_ns )
```

Anyone with a standard hash function can recompute a receipt in five lines of code, offline, without a server, an account or the company's permission, and can still do so in fifty years. Each receipt carries the tip of the receipt before it, a linked timestamp in the sense of Haber and Stornetta (1991), so every user of the machine becomes, without consent and without knowing it, the guardian of every other user's past. The operator cannot choose which tip enters the record: the previous event wrote it. He cannot omit an event: the next event already carries it. He cannot reorder. He can switch the machine off. He cannot make it lie.

A software version of this chain has been running publicly at api.recognitium.com and holds more than one hundred and forty-five thousand receipts at the time of writing. Anyone can verify one at recognitium.com/verify, and anyone can become a permanent witness of the chain with a single command that mirrors its live stream. Every receipt on that chain is stamped with the class of its source, and today that stamp says the source is computational. The physical instrument exists as a design and as a bench about to be built. The Twin Test is the first experiment that separates the two classes by measurement rather than by declaration.

What the instrument will never say matters as much as what it says. It will never say that a thing was true, only that it was said, then, once. It will never let a single receipt prove itself alone: membership in the witnessed chain is the proof. It will never claim that a destroyed value is gone from every copy in the world, only that its own copy died before a sensor. And it will not call itself quantum before a laboratory has measured the birth of its bits. Each of these refusals is a wall the experiment below is designed to test rather than assume.

3. WHY A TWIN TEST

Every bit can be copied. That is a theorem about patterns, not about instances. Two identical byte strings in two machines are two physical states, distinct in every electron, engineered to be indistinguishable enough that only the pattern matters. Digital computation is the discipline of forgetting every difference between instances so that patterns can travel. What copies, always and only, is the pattern.

Recognitium's instrument is built on the instance. It claims that the block of bits born inside it in one window is a token that no second machine can produce, not because the pattern is secret but because the birth is not repeatable. The natural objection is Leibniz's: two instruments give different answers only because they were never truly identical; make them identical enough and the answers will converge. The Twin Test takes that objection at its word. It builds the two instruments as identical as engineering allows, then pushes past what two objects can reach to what one object can: the same instrument restarted into the same state and fired again.

Two definitions of information stand at the two ends of this experiment. Landauer (1961) said that information is physical: erasing one bit costs at least kT ln 2 of heat, and Bérut and colleagues (2012) measured that heat. Bateson (1970) said that information is a difference that makes a difference, and that definition has remained a sentence rather than a measurement since. The Twin Test measures the other end of the bit from Landauer's. He measured what it costs to destroy one. This measures whether two births of one can ever be the same.

There is one other machine whose value rests on the instance. A one-shot signature (Amos, Georgiou, Kiayias and Zhandry, 2020; Shmueli and Zhandry, 2025) is a signing key that exists as a quantum state and is consumed by the act of signing, so that no one, not even its maker, can sign twice. It has been proven possible and never built, because the key must be kept coherent in a quantum memory until the moment of use. Both machines are born in a measurement whose outcome nothing could have known in advance (Wootters and Zurek, 1982). They differ in when the measurement happens. A one-shot signature holds its quantum state alive until it is spent: causality carried in photons, protected by law for the whole life of the key. Recognitium's instrument measures at birth and makes the classical life that follows single-use by construction: causality carried in electricity, protected by law at the two ends and by engineering in the millisecond between. One keeps its head alive; the other's head is born at the tick and dies on schedule. The instrument claims none of the one-shot theorems. What the Twin Test measures is the one property the two machines share and neither has yet measured: that the birth is not repeatable under identity of construction.

The mechanism has a name, and the comparison gives it its meaning. Quantum Settlement: quantum for the birth of the bits, a measurement whose outcome no copy could have revealed first; settlement for their death, an irreversible erasure sensed before the box is allowed to speak. A tick, spent once. The instrument that performs it is the Quantum Causal Processor, the QCP. It does not measure seconds. It counts irreversible events in the order they passed, and the quantity that passes through it is causality. Against its neighbour the difference is now one sentence: a one-shot signature is quantum for the whole life of the key; Quantum Settlement is quantum at birth and settled at death. The first word of the name is earned, not assumed. The instrument wears it only once a laboratory has measured the birth, and until then every receipt says so.

4. THE OBJECTION, AND WHY IT IS THE POINT

The objection is easy to state, and it should be stated here first. Two noise sources give different outputs; everybody knows that. Building them identically and watching them still differ proves nothing a first-year student would not grant, and dressing the result as a verdict on sameness is philosophy borrowing a laboratory coat.

Three answers, the third of which is the reason this document exists.

First, the experiment does not measure whether the twins differ. It measures how the difference responds to identity. A falling curve and a flat curve are both consistent with "they differ"; only one of them is consistent with the construction hypothesis. That curve has not been published, and the self-twin rung, the same instrument restarted into the same state, has not been reported as the limit of a ladder, though a version of it is run unnamed in every restart test the standards require.

Second, the paper does not claim that sameness does not exist. Control C0 is the proof that it does: two programs holding one state agree every time, because a pattern can be held twice. What the experiment locates is the boundary of sameness, exactly where digital engineering placed it and then forgot it. Sameness belongs to patterns, and everything that computes is built to forget the instance so that the pattern can be shared. The instrument is the first machine built to remember the one difference every other machine is built to forget.

Third, what everyone knows and nobody measures is precisely what a measurement changes. Landauer stated in 1961 that erasing a bit has a minimum cost in heat. For fifty-one years it was a principle everyone granted and no one had seen; the field cited it, argued about it, and made exceptions for it. Bérut and colleagues measured it in 2012, with a single colloidal particle in a double well, and from that day it was a fact with an uncertainty, which every later paper had to reproduce or refute. The Twin Test proposes the same passage for the other end of the bit. That two births are never the same is granted by everyone and, because it is granted, it is ignored: records, money and law are built on copies as if instances did not exist. A preregistered curve, replicated in a second laboratory, turns the granted into the measured, and the measured cannot be ignored. That is what the experiment is for. It is not the discovery of a new fact about nature. It is the conversion of a fact everybody carries into one that has a number, a date and a witness.

5. HYPOTHESES

Two hypotheses, both live, neither assumed.

H_0, construction. Two instruments differ in their outputs only because they were never truly identical. As identity of construction and state improves, the outputs converge: the distance between twin blocks decreases with identity and reaches zero at perfect identity.

H_1, origination. The difference between two physically born blocks does not come from how the instruments were built. The distance between twin blocks is invariant under construction identity: it stays at the value two independent draws would give, at every rung of the identity ladder, including the rung where the two instruments are one instrument restarted into the same state.

The measurable sentence is H_1's: distinctness invariant under construction identity. It is falsifiable at every rung.

6. APPARATUS CLASS AND THE IDENTITY LADDER

Two prototypes of the instrument, Unit A and Unit B, built from the same board revision, the same firmware image, whose fingerprint is committed publicly before the run, the same component lots for the noise source and the digitiser, and the same wiring. Each trial yields, per unit, a block of R raw samples of the noise source and its SHA-256 digest. R is fixed and committed before the run. The nature of the noise source, its conditioning and the acquisition sequence are not described here. They belong to the construction, which is committed privately (section 12) and which the experiment is designed to make irrelevant.

The identity ladder has five rungs, each more identical than the last.

Rung	Configuration
I_0	Two units, noise-source components from different production batches, separate enclosures, separate supply rails.
I_1	Two units, components from one batch at adjacent positions on one reel, separate enclosures.
I_2	Two units in one enclosure, one shared supply rail, thermally coupled on one plate with the temperature logged, one shared trigger line.
I_3	One unit, warm restart: the acquisition rerun after a full reset of the acquisition stage, supply and temperature held, identical inputs. The self-twin.
I_4	One unit, cold restart with matched thermal history: power-cycled, brought back to the same logged temperature with the same settling time, identical inputs. The self-twin at the limit engineering can reach.

7. INPUTS HELD IDENTICAL

Per trial, both units receive the same input values, the same trigger instant from a single shared line, the same supply set-point and the same enclosure conditions. Every input is logged per trial. Any trial in which an input differed is excluded before analysis, and the exclusion is itself committed.

8. OBSERVABLES AND THE NULL

Three observables, in increasing statistical power.

Exact coincidence. The two digests are equal in a synchronised trial. Binary, the rarest event, the headline falsifier.

Twin Hamming distance. D, the number of positions in which the two raw blocks differ, per trial. For independent draws with per-sample bias near one half, D has mean about R/2 and standard deviation about √R/2.

Twin mutual information. I(A;B) between the two raw streams at lag zero, estimated per rung with the same estimator on the same block sizes.

The null is built from the data itself. Bias in the digitiser lowers D even for independent sources, so the comparison is never against R/2 but against the shifted-pair null: block A of trial t is paired with block B of trial t+s for s ≠ 0, trials that were not simultaneous, and the distribution of that distance, D_shift, is taken. H_1 predicts that D_sync and D_shift are the same distribution at every rung. H_0 predicts D_sync below D_shift by a margin that grows with identity.

The result is one curve: the mean of D_sync divided by the mean of D_shift, with its confidence band, against the five rungs, with the two software controls of section 11 as reference lines at ratio 0 and ratio 1. H_1 is a flat line at 1. H_0 is a line that falls toward 0.

9. NUMBERS NAMED BEFORE THE RUN

Trials per rung: $N = 10^6$ synchronised trials.

Exact coincidence: the expected count under H_1 is $N \cdot 2^{-k}$ for a source of min-entropy k per block; with k at least several hundred bits it is zero to any precision. One coincidence is therefore decisive. The test cannot observe 2^{-k}; N trials without a coincidence bound the rate at about 1/N, and that bound is what is reported.

Distance: with $N = 10^6$ the standard error of the mean of D is $(\sqrt{R}/2)/\sqrt{N}$. The preregistered rejection threshold is a deficit of D_sync against D_shift greater than five standard errors at any rung.

Slope: a one-sided trend test of the ratio mean(D_sync)/mean(D_shift) across I_0 to I_4, rejecting H_1 at $p < 10^{-3}$.

Mutual information: the preregistered threshold is $I(A;B) > 10^{-3}$ bit per sample at lag zero, with the lag-s estimates for s ≠ 0 as the null band.

The min-entropy k of the source is not measured by this test. It is the laboratory's number from a campaign under NIST SP 800-90B (Turan et al., 2018) with a stochastic model. The Twin Test borrows k; it does not produce it.

10. WHAT FALSIFIES H_1

The complete list, fixed in advance.

F1. An exact digest coincidence between twin units in a synchronised trial, beyond the SHA-256 collision floor of 2^−256.

F2. Mean D_sync below mean D_shift by more than five standard errors at any rung.

F3. A significant negative slope of the ratio across the ladder at $p < 10^{-3}$: convergence under identity.

F4. Twin mutual information above 10^{-3} bit per sample at lag zero while absent at other lags.

F5. The self-twin rungs I3 and I4 producing correlated sequences, which is also a failure of the restart test of SP 800-90B.

F6. A second laboratory obtaining a falling curve where the first obtained a flat one.

Three objections are closed in advance. A flaw in identity is answered by the ladder itself, which measures the response to identity rather than assuming any rung is perfect. A rigged control is answered by the controls being open code, committed by fingerprint before the run. Post-hoc selection is answered by section 12.

11. CONTROLS

Two software controls, each a pair of software generators standing in for the physical source.

C0. The two generators share one internal state. This is H_0 realised: identity of state yields identical blocks and a distance of zero. It supplies the lower endpoint of the curve.

C1. The two generators hold independent seeds. The distance is that of independent draws. It supplies the upper endpoint, and it exists to show that distance alone proves nothing: what carries the result is the response of distance to identity, not its value.

12. PREREGISTRATION BY COMMITMENT

On the day of publication, the SHA-256 fingerprint of this document is committed to Recognitium's public receipt chain and anchored into the Bitcoin blockchain through OpenTimestamps (Todd, 2016). A private manifest containing the fingerprints of the full protocol and of the instrument's construction documents is committed at the same time, in the same two ways. Its content is not published. Its existence and its date are, so that a reader shown the construction later can verify that it was fixed before the first measurement.

Before Unit B's first trial, the fingerprint of the analysis script, the value of R, and the identity ladder with its N are committed the same way. During the runs, every batch of raw samples from both units is fingerprinted and committed before anyone looks at it. Analysis runs only on committed batches, and the analysis output is itself committed. A reviewer verifies, with SHA-256 alone and no trust in the experimenter, that no datum was chosen after the fact.

Where to check. The receipt for this document's fingerprint is verified at recognitium.com/verify by pasting its receipt identifier, which is given in the description of the published record. The daily Bitcoin anchors of the chain, and their OpenTimestamps proofs, are published at recognitium.com/anchors and verify with the free OpenTimestamps client against the Bitcoin blockchain, with no account and no trust in Recognitium.

13. WHAT A FLAT CURVE WOULD ESTABLISH, AND WHAT IT WOULD NOT

If the curve is flat, the creation of a block of bits inside this instrument is non-repeatable under identity of construction. Information would then be differential at the birth of a bit, as it is physical at its death, and Bateson's definition would acquire a measurement fifty-six years after it was written. The sentence a flat curve would license, and that neither Landauer nor Bateson could say alone: information is not only physical; it is also differential.

For metrology, a flat curve is a certification primitive: physical origination established independently of construction, sharper than a restart test because it states what it tests and measures the response to identity.

What it would not prove, said first. It would not prove quantum origination: a chaotic classical source also stays distinct under any identity that can be engineered, and the word quantum belongs to the stochastic model and the laboratory's campaign, never to this curve. It would not prove the security of any product built on the instrument. And it would not settle the metaphysics beyond the ladder: the construction hypothesis keeps a retreat, that identity was still imperfect at I4, and after a flat curve that retreat must predict a slope that five rungs failed to show, including the rung where the two instruments are one instrument.

14. TIMELINE AND REPLICATION

Unit A takes its first measurements the week of 21 September 2026. Unit B is built identically in October 2026. The ladder and the two controls take two bench days. Then the curve. Then the protocol, the firmware image and two units go to a second laboratory, which runs the ladder with its own hands and commits its own batches. The result becomes a fact when the second curve is flat.

ACKNOWLEDGEMENTS

Drafting assistance was provided by an AI system (Claude, Anthropic) under the author's direction. All claims, decisions and errors are the author's.

REFERENCES

Amos, R., Georgiou, M., Kiayias, A., Zhandry, M. (2020). One-shot signatures and applications to hybrid quantum/classical authentication. *Proceedings of the 52nd ACM Symposium on Theory of Computing (STOC)*, 255–268.

Bateson, G. (1972). Form, substance and difference. The nineteenth annual Korzybski Memorial Lecture, 1970. In *Steps to an Ecology of Mind*. Chandler, San Francisco.

Bérut, A., Arakelyan, A., Petrosyan, A., Ciliberto, S., Dillenschneider, R., Lutz, E. (2012). Experimental verification of Landauer's principle linking information and thermodynamics. *Nature* 483, 187–189.

Haber, S., Stornetta, W. S. (1991). How to time-stamp a digital document. *Journal of Cryptology* 3(2), 99–111.

Landauer, R. (1961). Irreversibility and heat generation in the computing process. *IBM Journal of Research and Development* 5(3), 183–191.

Leibniz, G. W. (1686). *Discours de métaphysique*, §9. The identity of indiscernibles.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal* 27, 379–423 and 623–656.

Shmueli, O., Zhandry, M. (2025). On one-shot signatures, quantum vs classical binding, and obfuscating permutations. *Advances in Cryptology, CRYPTO 2025*. IACR ePrint 2025/486.

Todd, P. (2016). OpenTimestamps: scalable, trust-minimized, distributed timestamping with Bitcoin. opentimestamps.org.

Turan, M. S., Barker, E., Kelsey, J., McKay, K. A., Baish, M. L., Boyle, M. (2018). *Recommendation for the entropy sources used for random bit generation*. NIST Special Publication 800-90B.

Wootters, W. K., Zurek, W. H. (1982). A single quantum cannot be cloned. *Nature* 299, 802–803.

The experiment in one breath: make two machines the same until they are one machine, and ask them the same question at the same instant. If they never answer alike, the difference was never in how they were made.

Recognitium · Paris · jose@recognitium.com · recognitium.com