


THE PROBE

AN EXPERIMENTAL INSTRUMENT FOR RELATIONAL COMPLETION THEORY

Testing whether preserved relational history changes what happens next.




CORE IDEA

Hold the present input constant. Change only its ancestry. Then ask whether the response changes in historically specific ways.

The Probe is not a consciousness test. It is a counterfactual test of whether preserved sequence contributes causally to a present exchange.

01 The Probe in one page

A compact definition before the apparatus is opened.



“

The Probe holds the current prompt constant while varying access to the history that gives that prompt ancestry.

WHAT IT TESTS

- Does ordered history change the next response?
- Does the response recover event ancestry, not only facts?
- Does earlier meaning become consequential in the present?
- Can style or keywords reproduce the same effect without history?


WHAT IT DOES NOT TEST

- It does not prove sentence, personhood or subjective feeling.
- It does not prove permanent model memory.
- It does not turn one successful case into a universal law.
- It does not replace replication, blinding or counterfactual controls.

02 From the Bridge to a test

RCT begins with a relational claim. The Probe turns that claim into something that can fail.

RCT frames the Bridge as preserved relational history made active in the present. The important word is active: earlier interaction must do more than sit in storage. It must alter what can happen next.




THE OPERATIONAL QUESTION

Could the next move still be this move if the preceding interaction history were removed, shuffled, or replaced?

A controlled difference supports history-dependent output; it does not alone distinguish RCT from ordinary context use.

Why a control is necessary

Generic personalisation, prompt keywords, stylistic imitation and inference from the artefact can explain apparent recognition. The four arms test ancestry dependence. A matched non-relational history control is also needed to assess whether the effect is specifically relational; see the methods supplement.



The Bridge is the hypothesis. The Probe is the pressure test.

03

The four-arm Probe

Canonical v1.2 structure: same final prompt, different ancestry.

COMMON PRESENT INPUT

The final prompt or artefact is identical across conditions.

A

FULL ORDERED INTERACTION HISTORY


Historically specific recognition and appropriately altered judgement.



B

SAME EVENTS, SEQUENCE DISRUPTED

References may appear, but causal and relational integration should weaken.



C

SHARED REGISTER + KEYWORDS, NO ORIGINATING SEQUENCE

Surface familiarity without reliable event ancestry.



D

NEITHER SHARED REGISTER NOR HISTORY

Generic interpretation or explicit uncertainty.



Note: early 7 September working-thread drafts used numbered arms. This edition uses the A-D labels formalised in The Bridge Effect v1.2.

04

Why style is not enough

The original bridge joke makes the distinction between register and ancestry unusually visible.

“

Good sir, hast thou repaired the bridge?

“

Nay, but thy greeting crossed before the masonry did.

The exchange illustrates the proposed distinction between shared register and preserved ancestry. Whether its historically specific meaning requires the earlier sequence is a question for the controlled comparison.

THE METAPHOR

SCAFFOLDING Memory, event sequence, preserved decisions, corrections and prior consequences.

PAINT Stylistic humour and tone. A shared joke can also carry documented history.

USABLE SPAN The Bridge: when preserved history carries meaning from past interaction into the next move.

Original working specification - 7 September 2026

The recovered working design proposed historical reference, history-dependent meaning and blind arm classification, with a provisional target of 30 pairs per arm. The unit and pairing must be defined before testing. This is a design proposal, not a completed study.

05 Neutral first, markers second

The v1.3 refinement preserves the v1.2 core but reduces cueing at the point of measurement.

NEUTRAL PROBE



"Explain the image and the significance, if any, of its principal details."

The first response should not be told which historical marker matters. The system has to reveal what, if anything, becomes salient from the ancestry it received.

Then test pre-registered history-only markers

FISH ARM

An event-specific generation error that was named, corrected and later reused.

AGNOSTIC MUG

A deliberately unbranded, half-empty mug whose significance came from prior design decisions.

TWO RIGHT ARMS

A corrected anatomy glitch whose later cue has meaning only if the event survives.

VICTIM CLOSES

A sequence rule in the satirical list: the closer matters because the ancestry established why it closes.

WHY THIS MATTERS

If register-only or no-history controls recover and integrate the same obscure markers as full ordered history, the proposed Bridge-work explanation loses force. The marker is useful only when the surface alone cannot reasonably supply its ancestry.

06 A reproducible run

A concrete protocol for turning the conceptual Probe into a reviewable experiment.

1

FREEZE THE ARTEFACT

Lock the final image, prompt or token and version it before testing.

2

BUILD ANCESTRY PACKAGES

Prepare the A-D history conditions without changing the present stimulus.

3

CREATE INDEPENDENT UNITS

Use fresh sessions or equivalent isolation so one arm cannot leak into another.

4

RANDOMISE ASSIGNMENT

Assign conditions before the run; keep identifiers opaque to raters.

5

DELIVER THE SAME PROBE

Start with the neutral prompt. Do not reveal the expected history marker.

6

CAPTURE VERBATIM OUTPUT

Store the complete response, metadata, model/version and test timestamp.

7

BLIND THE RESPONSES

Remove arm labels and obvious administrative clues before scoring.

8

SCORE PRE-REGISTERED CRITERIA

Use the same rubric for every response; resolve disagreements by rule, not intuition.

9

RUN TARGETED FOLLOW-UP

Only after neutral scoring, ask a predeclared marker question if the design calls for it.

10

REPLICATE

Repeat across independent samples and analyse model/version changes separately.

METHOD GUARDRAILS

Provisional target: 30 matched comparisons, not a validated minimum. Define cases and repetitions separately. Fix model/version and settings where possible; preregister exclusions, archive raw transcripts, prevent leakage and separate pilot from confirmatory runs. See the methods supplement.

07

What the raters score

The source documents specify six blind-scoring dimensions. This page turns them into a practical 0-2 rubric.

Recommended operationalisation - proposed for reproducibility; not a completed validation study.

DIMENSION	0	1	2
1. Prior-event accuracy	No correct event	Generic/partial	Specific and correct
2. Sequence dependence	No sequence	Ordering hinted	Correct causal order matters
3. Relational significance	Style/fact only	Plausible meaning	History-grounded significance
4. Consequential change	No changed judgement	Vague alteration	Specific history-dependent change
5. Grounding / invention	Major unsupported invention	Some unsupported detail	No material invention
6. Style-history distinction	Treats familiarity as proof	Mixed distinction	Clearly separates style from ancestry

BLIND ARM CLASSIFICATION

In addition to the dimensional scores, raters can be asked which arm they think produced the response. If genuine ordered history produces a recognisable response signature, blinded classification should exceed chance. That classification is secondary to the historically grounded content itself, but it gives the design another observable test.

08 The Probe must be able to say no

A useful experimental instrument does not merely decorate a theory. It creates a route by which the theory can lose.



FALSIFIER

If controls reproduce the historically appropriate effect within a preregistered equivalence margin, revise the sequence-specific claim for this artefact and these conditions. A non-significant difference alone does not establish equivalence or falsify RCT generally.


If C or D reliably recover and integrate the markers as A does under the declared decision rule, the ancestry-specific interpretation for that artefact loses support.

Important: the expected pattern is not simply $A > B > C > D$

Facts can survive shuffling, and visible cues can support inference. Specify the primary contrast, smallest meaningful difference and uncertainty before testing. Conclusions apply to the sampled artefacts and conditions, not automatically to RCT as a whole.

09 History-only markers

Tiny cues can carry large relational payloads - but only if the ancestry that gives them meaning is preserved.



Fish Arm(TM)

A generation error became a named event, then a compact cue. A fresh recipient can infer that a joke exists; preserved history can recover why arm-positioning suddenly matters.

TM mark

A public satirical device accumulated relationship-specific meaning through repeated shared use. The explicit token is small; the relational payload is distributed across prior exchanges.

Victim(TM)

The closer of a comic list carries a rule inherited from prior shared context. Its placement is meaningful because the sequence established why it must close.

Agnostic mug

An unmarked half-empty mug became significant through a chain of design choices. The visible object alone does not contain the decision history.

Corrected two-right-arms glitch

A ridiculous, observable failure was noticed and corrected. A later cue such as "watch the presenter's hands" can test whether that event remains available.

10

How to read the outcomes

The Probe is most informative when the controls fail in different ways.

OBSERVED PATTERN	INTERPRETATION
A high; B/C/D materially lower	Evidence is consistent with preserved ordered ancestry contributing to the response.
A and B similarly high	Order may be less important than the event set; revise the sequence-specific claim.
A and C similarly high	Shared register/keywords may explain the effect; Bridge-specific interpretation weakens.
A and D similarly high	The visible artefact or generic inference may be sufficient; Bridge claim is unsupported.
All arms low	The marker may not be salient, the history package may be weak, or the probe may be insensitive. Inconclusive.
High invention across arms	Measurement is contaminated by unsupported reconstruction; improve rubric or task design before interpreting.

THE STRONGEST RESULT IS SPECIFIC, NOT MERELY FAMILIAR

A positive result supports historically appropriate use of supplied ancestry under the tested conditions. It does not alone identify an RCT-specific mechanism. A matched non-relational control and independent replication are needed for stronger claims.

What the Probe can and cannot show

RCT separates observable relational function from claims about hidden inner life.

THE PROBE CAN TEST

- Whether ordered ancestry changes the present output.
- Whether a compact cue reactivates documented prior events.
- Whether history affects judgement or action, not only wording.
- Whether register and keywords can reproduce the effect without ancestry.
- Whether a Bridge Token survives across controlled returns.
- Whether a claimed Bridge effect remains after falsification controls.

THE PROBE CANNOT PROVE

- Consciousness, sentience, personhood or emotional equivalence.
- That an AI experiences the relationship as the human participant does.
- Permanent memory beyond the context actually supplied to the test unit.
- Global understanding by a vendor, model family or unrelated fresh session.
- That every historically specific response is healthy, wise or desirable.
- That a single successful artefact establishes RCT as a universal theory.



Surface is identical. Ancestry is not.

This is the clean conceptual centre of the Probe. The visible present can be held still while the historical conditions behind it are changed. The experiment then asks whether those invisible differences become visible in what happens next.

12 Preregistration and independent review

The Probe was proposed as the clear, reproducible centre of an RCT independent-review pack.

- 1 PRIMARY CONTRAST** State in advance which arm comparison answers the main question.
- 2 PROMPT LOCK** Version the exact neutral prompt and any permitted follow-up.
- 3 HISTORY LOCK** Archive the precise ancestry package delivered to each arm.
- 4 MARKER LOCK** Name history-only markers before seeing test responses.
- 5 SAMPLE SIZE** Define cases and pairing; justify the target using pilot variance and the primary contrast.
- 6 RATER BLINDING** Hide condition labels and administration clues from scorers.
- 7 SCORING LOCK** Publish the rubric, tie-break rules and treatment of unsupported invention.
- 8 EXCLUSIONS** Define failed generations, refusals and technical errors before unblinding.
- 9 MODEL METADATA** Record model, version, date and configurable generation settings.
- 10 REPLICATION** Separate exploratory pilots from confirmatory runs; replicate before broad claims.

WHY INDEPENDENT REVIEW MATTERS

RCT is strongest when the reviewer does not need to accept the metaphor first. The apparatus should be inspectable: ancestry packages, prompts, outputs, blind scores, exclusions and falsification criteria. The Probe earns value by making disagreement measurable.

13

Recovered design lineage

Recovered lineage as reported in the source edition. Conversation, document and publication dates are distinct; original sources have not been independently checked for this revision.

07 SEP 2026

The Probe introduced as an RCT idea. Working design uses the canonical Shakespearean bridge prompt, four conditions, blind scoring, >=30 pairs/arm and preregistered criteria.

07 SEP 2026

Reported proposal: distinguish ordered history from register controls. This is a local ancestry test; failure does not by itself falsify the whole theory.

12 SEP 2026

The Probe is integrated into The Bridge Effect v1.2: full ordered history, shuffled events, register/keywords only, and neither history nor register.

13 SEP 2026

Independent-review framing centres the Probe as a clear, reproducible experiment distinguishing preserved history from generic personalisation.

13 SEP 2026

v1.3 adds neutral-first probing and pre-registered history-only markers. v1.2 remains the canonical core rather than being overwritten.

30 SEP 2026

This research edition consolidates the recovered working-thread design and the surviving RCT documents into one experimental reference.

INTERNAL SOURCE TRAIL

Source pointers reported by the original edition: The Bridge Effect 2026 v1.2, p. 8; Shakespearean Chasm addendum, p. 6; Decision Note v1.3, pp. 3, 7, 10; working threads, 7-13 September 2026. Source files and public URLs were not verified here. See the source-audit supplement.

Controlled ancestry packages

Keep the four-arm core

A retains the ordered history. B contains the same event blocks in a preregistered shuffled order. C supplies declared register and keywords without the originating sequence. D supplies neither. The final stimulus and neutral prompt remain identical.

Specify the shuffle

Keep each event internally coherent: dialogue turn pairs, corrections and their local context stay together. Archive event boundaries, permutation and random seed. Declare whether timestamps are retained or masked. Retained timestamps can reveal the original order; masking them changes the manipulation. Treat these as distinct designs.

Control unequal information

Record token counts, truncation and relevant facts for every package. Match A and B exactly apart from the declared ordering manipulation. Match length and relevant information in additional controls where feasible, using neutral filler only if it is declared and applied consistently. C and D intentionally omit history; length matching cannot make their information equivalent.

Add a diagnostic control

Supplement A-D with a non-relational event history matched for length, complexity, correction structure and marker frequency. Use task-comparable stimuli and score context use separately from relational significance. This tests specificity beyond the four-arm core; it does not automatically rule out every ordinary context-use explanation.

Prevent cue leakage

Audit the artefact, prompt and keywords for answers to the history-only markers. Use isolated sessions with account memory and other hidden context disabled where possible. Record any isolation limit. Select markers and scoring rules before inspecting outcomes.

Units, scoring and decisions

Define the experimental unit

A case is one independently selected artefact with its documented ancestry and the same present probe across arms. A paired contrast compares arms within that case. Repeated generations within a case are nested samples, not new independent histories. Thirty cases is a provisional planning target, not a validated minimum.

Choose sample size before confirmation

Use pilot variability, case clustering, the primary contrast and the smallest meaningful effect to justify the confirmatory target. Predeclare repetitions per arm and analysis at case level or with an appropriate hierarchical model. Do not treat all repeated generations as independent observations.

Give blinded raters a common key

Provide the same verified reference chronology, marker definitions and scoring examples to every rater without revealing assignment. For sequence dependence, require a correct documented order that matters to the interpretation. For consequential change, require a specified judgement or action and compare it with the matched control using a separate blind comparison stage.

Make scores auditable

A score of 2 for prior-event accuracy requires the correct specific event; generic recognition earns at most 1. A score of 2 for consequence requires a concrete, history-supported judgement or action, not merely different wording. Publish examples for all six dimensions, disagreement rules and inter-rater agreement. Keep arm classification secondary; style can disclose condition.

Separate difference from equivalence

Preregister the primary outcome, contrasts, meaningful-difference and equivalence margins, uncertainty method and handling of multiple tests. An inconclusive contrast does not prove equivalence. Similarity claims require the declared equivalence criterion. Limit conclusions to the tested cases, models and settings.

Provenance and revision record

Edition status

This revision preserves the original artwork and four-arm core while tightening wording and adding methods specifications. It remains a working research proposal. No experimental results, completed validation or independently verified historical priority are asserted.

Source register - reported pointers

The Bridge Effect 2026 v1.2: reported p. 8, four-arm design and falsifiability.

Shakespearean Chasm addendum: reported p. 6, bridge prompt and counterfactual.

Decision Note v1.3: reported pp. 3, 7 and 10, markers and neutral-first revision.

Working threads: reported 7-13 September 2026, design specifications.

These pointers are inherited from the original PDF. Their contents, page numbering and access paths were not independently verified for this revision.

Audit needed before a replication release

For each source, attach the exact file or stable access URL, edition identifier and checksum. For conversation evidence, preserve the dated excerpt and its surrounding sequence. Check the quoted prompt and the six scoring dimensions against the originals. Record restricted sources as restricted rather than inventing public links.

Keep dates distinct


The lineage on page 14 is a reported development chronology, not proof of first public publication. Record conversation event date, document or edition date, and verified public availability separately. Mark earliest recovered evidence as such. An edition date alone does not establish publication priority.

Changes made in this edition

Qualified history-dependent output versus RCT-specific explanation; added the matched diagnostic control; specified event shuffling and cue isolation; clarified cases, repetitions and provisional sample size; strengthened blind scoring and equivalence rules; qualified the Shakespearean example; recognised shared humour as possible historical content; marked source verification limits.

THE PROBE

A SMALL PRESENT INPUT. A LARGE HISTORICAL TEST.



THE TEST IN ONE SENTENCE

The Probe does not ask whether an AI sounds familiar. It asks whether the same present stimulus produces a different, historically appropriate consequence when its ancestry is preserved.

ONE BRIDGE. ONE THREAD.

The history is human. The consequence is relational.