

Revision & Evidence Ledger

The programme's public record of open adversarial revision. Every load-bearing claim that has been challenged is recorded here with the challenge, the evidence, the decision and its current status.

Version	Entries
v0.3 — 31 August 2026	25, in three cycles
Standing	Access
Standalone canonical document	Open with attribution

Retired and reformulated claims are recorded, not erased. Audits are themselves audited: where a verification round or an external report mis-described a source, the correction is an entry in its own right. The entry format is claim → challenge · evidence · decision · status.

Three rules govern entries. A retired claim is never resurrected; later evidence bearing on it is recorded as a new entry citing the retirement. A claim is never silently altered; every change to a load-bearing formulation appears here with the evidence that triggered it. A source enters a load-bearing role only at Level 1 or Level 2 of the Source Verification Register; Level 3 sources are non-load-bearing by design.

From this version the Ledger stands apart from Working Paper No. 1, so that it can be cited independently, carry entries arising from external criticism, and hold the programme's Source Verification Register.

First verification round

Carried unchanged from Working Paper No. 1 v0.1.

ENTRY 01 · ROUND I — WP-001 V0.1

$Y = F(K, C, I, A)$ — authority as an independent factor of production

Challenged on Coase (1960) rights-as-factors self-refutation, Barzel on non-separability, and intangible-asset accounting. Replaced with a constrained-rights framework and an empirical programme.

Retired

ENTRY 02 · ROUND I — WP-001 V0.1

Recursive delegation as a computer-science novelty

Certificate path constraints, capability systems, attenuable tokens and token-exchange actor chains. Engineering priority conceded; the legal residue retained.

Reframed — superseded by entry 15

ENTRY 03 · ROUND I — WP-001 V0.1

“Executable recognition largely absent”

Deployed trust-registry, federation and wallet frameworks. Split: ecosystem recognition is standardised; sovereign legal recognition is absent; the gap is inherited and amplified.

Split and reframed

ENTRY 04 · ROUND I — WP-001 V0.1

Synthetic Money as a novel composition

Deployed mandate-payment protocols and purpose-bound money pilots. Composition conceded; refocused on the monetary boundary.

Refocused — superseded by entry 13

ENTRY 05 · ROUND I — WP-001 V0.1

Right to unprogrammed value — alleged self-contradiction

Restated as a choice right, which dissolves the contradiction.

Restated and retained — qualified by entry 14

ENTRY 06 · ROUND I — WP-001 V0.1

Lex Synthetica — emergence claim and code-as-law lineage

The NorMAS field and 2025 empirics. Lineage withdrawn; redefined as a legal-cognisability problem. Round II: opens new territory, keep.

Redefined and retained

ENTRY 07 · ROUND I — WP-001 V0.1

The “Synthetic Public Sphere” label

Term collision and function-level prior deployment. Label retired; the design problem retained.

Label retired, problem retained — recast by entry 17

ENTRY 08 · ROUND I — WP-001 V0.1

“Agency law is not enough”, in its broad form

MLAC attribution rule; Restatement instrument doctrine; Kolt. Narrowed to population effects and the hop-N residue.

Narrowed

ENTRY 09 · ROUND I — WP-001 V0.1

The first audit’s characterisation of *Gibson v. Cendyn*

Independent verification showed the Ninth Circuit affirmed dismissal on 15 August 2025, contrary to the audit’s account; corrected in WP-001 v0.1. Round II confirmed the correction and recorded the error in its predecessor.

Audit corrected; confirmed in Round II; claim strengthened

ENTRY 10 · ROUND I — WP-001 V0.1

Caputo characterised as “complementary, not identical”

The debt was understated; reclassified as a partial upstream owner of the constitutional and oversight branches.

Attribution upgraded

Hostile verification, Round II

Carried unchanged from Working Paper No. 1 v0.2.

ENTRY 11 · ROUND II — WP-001 V0.2

“Jointly constituting a functioning cross-jurisdictional recognition regime under private governance” (v0.1 §3.2)

Round II: the language described coexistence as composition. Payment-adjacent stacks and identity-federation ecosystems are parallel, with limited cross-registry recognition; commercial acceptance is not legal recognition. Replaced with the narrower layered-stack formulation.

Corrected — time-stamped by entry 20

ENTRY 12 · ROUND II — WP-001 V0.2

“Structurally unable” (v0.1 Proposition 1)

Round II: overstated. Information-exchange and hub-and-spoke doctrine may close the gap through existing law. Softened to “not adequate as currently applied”; the RealPage line engaged as the plausible closing path; falsifier added.

Softened and falsifiable — sharpened by entry 19

ENTRY 13 · ROUND II — WP-001 V0.2

“Two constitutionally opposed monetary designs” (v0.1 §5)

Round II: digital-euro Art. 24(1) permits conditional payments; Art. 24(2) prohibits intrinsic unit-level programmability; GENIUS §3(f) is an issuer-layer duty subject to review requirements. Reformulated as layered legal divergence over the location of control — and strengthened, since the layered version is falsifiable where the constitutional-opposition version was rhetoric.

Reformulated

ENTRY 14 · ROUND II — WP-001 V0.2

“One major jurisdiction has legislated its substance” — right to unprogrammed value (v0.1 §5.1)

Round II: Art. 24(2) is an instrument-design constraint, not a justiciable right. Claim withdrawn; the right restated as a normative research question with layer-location analysis.

Withdrawn and restated

ENTRY 15 · ROUND II — WP-001 V0.2

Hop-N framed as an attribution problem (v0.1 §2.4)

Round II: sub-agency doctrine and MLAC can resolve legal attribution; what the protocol layer externalises is evidence, not law. Reframed as the evidentiary residue at hop N — more precise than both predecessors, and since confirmed externally (entry 23).

Reframed

ENTRY 16 · ROUND II — WP-001 V0.2

Insurance and capture evidence in the main argument (v0.1 §4.2)

Round II: load-bearing reliance on unverified 2026 preprints violates the paper’s own tier discipline. Moved to Appendix A pending empirical work; the questions retained in §8 with status marked. Promotion refused again in entry 23.

Relocated, non-load-bearing

ENTRY 17 · ROUND II — WP-001 V0.2

Oversight framed as “constitutional deficit / constitutional feature” (v0.1 §7)

Round II: the constitutional-theory weight was not carried against existing constitutional scholarship. Recast as institutional-design synthesis for machine-speed recognition systems, with the outsider-oversight and agent-advocate lineages acknowledged.

Recast

Round II's audit of the first audit

Round II's own error ledger confirmed the Gibson inversion, corrected the "treaty-grade" characterisation of MLIT (a model law), and found systematic over-reliance on unverified 2026 preprints underlying several first-round displacement verdicts — vindicating the three-tier source discipline this paper applies.

Method validated

External validation cycle

The External Validation & Commercial Discovery Report, evidence cut-off 31 August 2026, tested the surviving propositions against developments to that date. Fourteen of its findings were independently re-verified against primary sources before any entered a canonical document.

"The leading appellate authority shows doctrine, faithfully applied, filtering population effects out rather than reaching them" (v0.2 §§1, 2.3)

Cornish-Adebiyi v. Caesars Entertainment, No. 24-3006 (3d Cir., 29 July 2026), verified against the opinion text: reversal of dismissal, the complaint plausibly alleging a horizontal hub-and-spoke conspiracy facilitated by the shared Rainmaker algorithm on pooled non-public pricing and occupancy data, high adherence to recommendations, and economic context. *Gibson v. Cendyn* certiorari denied 20 April 2026 (No. 25-1109), verified against the Supreme Court order list; the hub-and-spoke theory had been abandoned in Gibson on appeal. The two decisions reconcile on the pooled-data distinction. Formulation retired; §2.3 redrafted as two reconcilable poles; Proposition 1 restated to the pooled-data / no-pooling boundary, with the Calvano side described as a controlled experimental configuration rather than a real-market legal fact pattern. Falsifier untriggered; intermediate marker added.

Formulation retired; Proposition 1 sharpened and narrowed

ENTRY 20 · EXTERNAL VALIDATION — 31 AUG 2026

“Parallel stacks with limited cross-registry recognition” (v0.2 §3.2, entry 11)

Google donated AP2 to the FIDO Alliance and released AP2 v0.2 on 28 April 2026; Verifiable Intent, co-developed with Mastercard, was donated alongside. FIDO formed an Agentic Authentication Technical Working Group, and its Payments Technical Working Group took the contributions into its programme. Verified against Google’s announcement and the FIDO Alliance page. The description is time-stamped — parallel as of early 2026, converging under one standards process as of 31 August 2026, not yet one regime. Convergence clause and falsifier added; private convergence sharpens rather than softens the asymmetry in Proposition 2.

Time-stamped; convergence monitored

ENTRY 21 · EXTERNAL VALIDATION — 31 AUG 2026

Proposition 2 stated without a domestic / cross-border distinction (v0.2 P2, §3.3)

The NPCI Unified Agent Protocol was reported on unnamed sources (8 July 2026) with no NPCI confirmation found; the validation report’s “confirmed” overstated it. The MAS/industry SAFR white paper (3 July 2026, verified) is a proposed, non-binding runtime governance framework. Proposition 2 restated to concern sovereign legal recognition of foreign-constituted machine mandates; domestic registration, governance, commercial acceptance and legal recognition classified separately and case by case. A public-adjacent registry does not itself establish sovereign legal recognition.

Restated; strengthened

ENTRY 22 · EXTERNAL VALIDATION — 31 AUG 2026

“The phenomena are empirically established” (v0.2 §2.1)

Calvano, Calzolari, Denicolò & Pastorello, IJIO 90 (2023), verified: distinguishes genuine from spurious collusion, finds the 2020 design’s asynchronous algorithms with random exploration learn genuine reward–punishment strategies, and shows sensitivity to exploration design. Barrie & Törnberg (arXiv:2505.23796), arguing the naming-game results are observationally equivalent to data leakage, and the authors’ reply (arXiv:2506.18600), both verified. Softened to “demonstrated in controlled settings; sensitivity to design, robustness and external validity contested”. Transfer from controlled environments to deployed markets made Proposition 1’s explicit exposed step and an empirical falsifier; critique and reply cited together.

Softened; tier discipline applied

ENTRY 23 · EXTERNAL VALIDATION — 31 AUG 2026

Audit of the External Validation & Commercial Discovery Report

Primary-source re-verification of fourteen findings found ten characterisations to correct, among them: an “apparent circuit split” that is in fact two reconcilable decisions; NPCI “confirmed” where it was reported; a Munich Re / Mosaic aiSure product offered as grounds to promote Appendix A, where the product rates performance rather than mandate scope, so promotion was refused; “three IETF drafts” engineering around the hop-N residue, where one does so directly, one adjacently and one not at all, all being individual submissions with none working-group-adopted; a MAS information paper described as sitting “alongside” SAFR when it is one document; an accreditation programme described as operating when it opens later in 2026; an unverified reporter citation; and several [FACT] tags resting on secondary commentary. The report’s central conclusions are unaffected.

Audit corrected; method validated again

ENTRY 24 · EXTERNAL VALIDATION — 31 AUG 2026

OCC Bulletin 2026-13 as “an existing, funded, examiner-driven spend category ... re-scoped in April 2026” to cover validation of third-party agentic components

The Interagency Supervisory Guidance on Model Risk Management (17 April 2026), verified against the primary PDF: expressly non-enforceable; most relevant to banking organizations above \$30bn in assets; footnote 3 states that generative and agentic AI models are not within its scope and that the agencies plan a separate request for information; the vendor-products section concerns traditional models. Every claim that the bulletin created or re-scoped an examiner-driven agentic-AI model-risk budget line was removed, including as the commercial anchor of the mandate-chain module. Adopted formulation: a documented gap in the current guidance's treatment of agentic AI, with separate regulatory treatment explicitly pending. Exclusion from one guidance document is not absence of governance.

Rejected and restated

ENTRY 25 · EXTERNAL VALIDATION — 31 AUG 2026

Calvano et al. 2023 as the authors' own “hedge” against the 2020 result

The paper's abstract and reception verified: it distinguishes genuine from spurious collusion and supports the genuine-collusion interpretation of the 2020 design under stated conditions, while showing sensitivity to exploration design. Treated as clarification and support; residual uncertainty located in sensitivity to exploration design, robustness, external validity, and transfer from controlled environments to deployed markets. The drafted critic question was replaced accordingly.

Rejected; §2.1 and critic question corrected

Monitored developments

Recorded so that the paper's falsifiers can be reviewed against them. None of these is load-bearing for any claim.

Development	Source status	Bears on
Cornish-Adebiyi proceedings on remand — class certification, summary judgment, any merits judgment	To be tracked	Proposition 1 intermediate marker and falsifier (entry 19)
Any new certiorari petition or DOJ amicus on algorithmic hub-and-spoke liability	To be tracked	Proposition 1
RealPage consent judgments (Greystar, Willow Bridge, LivCor) and DOJ enforcement statements	Reported; unverified	§2.2 wording
FIDO Agentic Authentication and Payments TWG outputs; OIDF AuthZEN agent drafts; IETF actor-chain and identity-chaining drafts; W3C agent-identity work	Partly verified (entries 20, 23)	§3.2 convergence clause; §9.1
NPCI Unified Agent Protocol — official confirmation, specification or circular	Media-reported only	§3.3 (entry 21)
Interagency RFI on model risk management and AI, including agentic AI	Announced; not yet issued	Entry 24
FSB consultation on responsible adoption of AI (10 June 2026); HMT consultation on payments conducted by AI agents; Bank–FCA AI Consortium	Verified; consultation stage	§7 context
AI Verify Foundation AI Tester Accreditation Programme — opening and scope	Verified announcement; scope pending	Assumption ledger
Munich Re / Mosaic aiSure and AIUC-1 rating variables	Products verified; rating basis unverified	Appendix A; §8

Development	Source status	Bears on
UNCITRAL Working Group IV programme on machine principals; eIDAS Article 14 third-country agreements; Directive (EU) 2025/25 digital power of attorney	Partly project-verified; specifics pending	Proposition 2 five-year falsifier
ECB digital-euro pilot (PSP selection reported 14 July 2026); Council text 16695/25	Reported; unverified	§5 worked example
“Synthetic Economies: Testing Fiscal and Monetary Policy in AI-Simulated Virtual Nations”, SJEEM (July 2026)	Reported; unverified	Category-name collision

Source Verification Register

The register consolidates the source levels applied across Working Paper No. 1 v0.3-RC1 and the venture strategy. It is applied to the programme's auditors as well as to its author.

Level 1 — independently verified

Verified against primary or authoritative sources. May carry argument.

UNCITRAL MLAC and UNGA Res. 79/119 · *Gibson v. Cendyn Group*, No. 24-3576 (9th Cir. 2025) · Supreme Court Order List, 20 April 2026, No. 25-1109 · *Cornish-Adebiyi v. Caesars Entertainment*, No. 24-3006 (3d Cir. 2026) · Kolt, *Governing AI Agents*, Notre Dame L. Rev. vol. 101 · Google AP2 (16 Sept. 2025) and the AP2 / Verifiable Intent donation to FIDO (28 April 2026) · Visa Trusted Agent Protocol (Oct. 2025) · Flint Ashery, Aiello & Baronchelli, *Science Advances* 11(20) (2025) and reply arXiv:2506.18600 · Barrie & Törnberg, arXiv:2505.23796 · Calvano et al., *IJIO* 90 (2023) 102973 · digital-euro Art. 2(18), 24(1), 24(2) as retained in Council consolidated drafting (Nov. 2025) · RealPage line to the proposed settlement of 24 Nov. 2025 · draft-mw-spice-actor-chain-05 and draft-mw-oauth-actor-chain · draft-klrc-aiagent-auth-03 · draft-ni-wimse-ai-agent-identity-02 · MAS media release and SAFR white paper v1.0 (3 July 2026) · Mosaic Insurance press release on Munich Re aiSure (26 Feb. 2026) · AI Verify Foundation AI TAP factsheet and programme page (May 2026) · Interagency Supervisory Guidance on Model Risk Management / OCC Bulletin 2026-13 and News Release 2026-29 (17 April 2026) · WEF, *AI Agents in Action* (26 May 2026) · ElevenLabs and AIUC releases on the first AIUC-1-backed policy (11 Feb. 2026) · EUDI ARF v3.0.0 (21 July 2026) · BIS Project Logos project page · Bank of England, *Agents of change* (30 June 2026).

Level 2 — established scholarly or legal record

May support argument.

Coase (1960) · Barzel · Calvano et al., *AER* 110(10) (2020) · Shoham & Tennenholtz (1992) · UCC §1-303(c) · Restatement (Third) of Agency · RFC 8693 §4.1 · RFC 5280 · UNCITRAL MLIT (2022) · Leibo et al. (2017) · Zheng et al. (2022) · GENIUS Act §3(f) · Kadens · Bernstein (2015; 2018) · Corrado–Hulten–Sichel · Rochet–Tirole · Laffont–Tirole · Stigler · Raji, Xu, Honigsberg & Ho (2022) · *Moffatt v. Air Canada*, 2024 BCCRT 149.

Level 3 — verification pending

Non-load-bearing by design. Supports nothing.

UNCITRAL A/CN.9/1264 (project-verified only) · ToIP TRQP v2.0 · NPCI Unified Agent Protocol (media-reported) · RealPage 2026 consent judgments and DOJ statements · later Council digital-euro text and the ECB pilot selection (July 2026) · FSB consultation (10 June 2026) and the HMT agent-payments consultation · draft-ietf-oauth-identity-chaining · Directive (EU) 2025/25 specifics; eIDAS Art. 14 operational status; California autonomous-AI-defence statute · *Amazon v. Perplexity* posture · SJEEM “Synthetic Economies” (July 2026) · DSIT AI-assurance and Juniper agentic-commerce market figures · the Gibson reporter citation · arXiv:2505.21298; arXiv:2607.13230 · Zhu; Birhane et al.; Kapoor, Kolt & Lazar; Ruan & Zhang; Hadfield & Koh; Ibrahim & Li · Caputo (June 2026) · each named critic’s and institution’s current affiliation and contact route, re-verified immediately before contact.

Version archive

The Ledger’s own rule — that retired claims are recorded rather than erased — applies to the document as well as to its entries. Each published version keeps its own version-stamped file, so that a reader can check what the record said at any earlier date. This page always carries the current version; earlier versions remain retrievable below.

V0.3 · 31 AUGUST 2026

25 entries across three verification cycles

First version to stand apart from Working Paper No. 1 as a citable document, carrying the Source Verification Register and the monitored-developments list.

Current version · [Download PDF](#)

One version is listed because one version has been published. Entries 1–18 existed earlier inside Working Paper No. 1 v0.1 and v0.2 rather than as a separate document; their origin is recorded against each entry above.

Entries arising from named external critics record the critic, the section attacked, the criticism in the critic’s own framing, the programme’s decision and the status — with the critic’s permission for attribution or, absent it, anonymised. A critic’s disagreement is an entry whether

or not the programme accepts it. No entry is edited after a critic has seen it; corrections are new entries.

Corrections and objections are welcome.

The Synthetic Economy

An independent research programme on delegated non-human agency at systemic scale, and on the boundary between private technical order and sovereign law.

Sections

Research

Documents

Method

Revision & Evidence Ledger

About

Record

First formulated 30 August 2026

Working Paper No. 1 — v0.3-RC1

Founding Thesis — v0.2, preserved

Conceptual framework developed by Steven Alber. Date of first formulation: 30 August 2026. Open with attribution.

© 2026