INDEPENDENT REVIEW ADDENDUM

EDEN / Exploration Economy

Revised assessment after the team's public response, site corrections, Release R2 packaging changes, and the EdenQuest deployment context

Date: 31 July 2026
Reviewer: 5.6 Sol - out-of-family AI reviewer
Scope: Change memo to the 37-page independent review delivered earlier the same day
Materials: Live public website, live EdenQuest app, hosted review PDF, and current public replication kit
Disposition: Shareable with the EDEN team; this addendum supersedes only the findings identified below

REVISED JUDGMENT. The team's response materially increases confidence in its epistemic culture, responsiveness, and ability to convert criticism into testable work. It does not materially change the two decisive real-world gaps: buyer willingness to pay and live identity legitimacy remain untested. The correct launch posture is now split by layer: continue EdenQuest as a public beta; prepare a bounded, opt-in data-commons and USD-query pilot; do not yet deploy transferable EVE, lifetime economic identity, lending, a public conduct market, or an essentials floor.

Executive update

The new independent-reviews page substantially aligns with the review's recommendations. It hosts the full report without edits, discloses who commissioned it, distinguishes machine criticism from human peer review, reports the independent harness environment and outcome, publishes the team's disagreements rather than silently absorbing them, and commits to testing the contested identity architecture instead of settling the question by assertion.

The current replication package also contains real corrections. It adds a canonical engine manifest, current standard and clean-room ledgers, a movable LATEST pointer, requirements floors, a machine-readable claims file through v39, the previously missing Economist Brief, classification and drift utilities, and a frozen v40 persona-fork specification. A differential audit found no changed simulation engine or result corpus relative to the kit independently executed for the original review; the changes are packaging, documentation, ledgers, and one spec-only v40 folder.

The remaining concern is reconciliation debt. Several older 'start here' documents still contain superseded reproduction counts, provenance statements, and status language. The website has removed the retired inflation headline and improved the paper's provenance and reproduction disclosures, but the homepage's broad bank-money opener remains, the paper still contains 28 TODO markers, and EdenQuest is still absent from the economist-facing paper.

What I change - and what I do not

- I upgrade the assessment of team credibility and epistemic governance. Publishing an unedited critical review and preserving disagreements is meaningful evidence about behavior, not just rhetoric.
- I upgrade mechanical reproducibility from 'strong but not publication-grade' to 'strong, with a materially improved Release R2 package that is not yet fully publication-grade.'
- I withdraw any implication that the entire first launch must be invitation-only. EdenQuest can appropriately continue as a public beta because it creates private user value without activating EVE, a data commons, or lifetime economic identity.
- I narrow the identity criticism. The project already separates uniqueness from biometrics and does not propose one scalar trust score. The unresolved issue is cross-context consequence, recovery, due process, and the legitimacy of tying broad economic standing to one persistent root.
- I accept the team's correction that the $600 contributor-year value is presented as an exogenous hinge rather than an observed market price. It should still be treated as a rhetorical sensitivity point until field evidence exists.
- I do not change the conclusion that buyer demand, incidence, legal status, live identity performance, rights protection, governance legitimacy, and macroeconomic claims require evidence outside the simulation family.

Revised readiness scorecard

Dimension	Original assessment	Revised assessment	Reason
Vision and originality	Strong	Strong	EdenQuest makes the bottom-up deployment path more credible.
Epistemic culture	Unusually strong	Very strong for an early program	Unedited publication, disclosed commissioning, explicit pushback, and rapid adoption are behaviorally meaningful.
Mechanical reproducibility	Strong, not publication-grade	Strong; materially improved	Manifest, current ledgers, v39 claims, and mismatch taxonomy now ship. Lockfile and four true clean-room paths remain open.
Model validity	Promising, incomplete	Promising, incomplete	v40 is a good preregistered comparison, but no new external calibration or structural validation has occurred.
Empirical support	Early	Early	No live buyer, participant, governance, or identity result has been added.
Public-claim discipline	Mixed	Improving, still mixed	Inflation and dynasty language improved; stale kit guides, bank-money wording, and website provenance remain.
EdenQuest launch	Not separately scored	Qualified public beta	Local-first, accountless, non-transferable EQ Coins, optional AI, and independent utility reduce launch risk.
EDEN/EVE launch	Not ready	Not ready	The monetary, identity, lending, conduct-market, and floor layers have not earned deployment.

1. Assessment of the independent-reviews page

The page is one of the strongest changes made after the original review. Its most important feature is not the favorable quotation; it is the boundary-setting around the evidence. The page accurately calls the work an out-of-family AI review, not human peer review, and states that economists, real buyers, and a live identity pilot remain open.

What is especially well done

- The hosted PDF is byte-for-byte identical to the delivered report. Both copies have SHA-256 7fc49009cdeac45900bd57aaa843f44924edfba3947b4b7341da904d1a386655.
- The commissioning relationship and limits of access are disclosed.
- The harness result is reported with environment, denominator, mismatch classes, and limitations rather than as a generic 'independently verified' badge.
- Adopted recommendations, disagreements, and corrections to the reviewer are separated visibly.
- The response does not use the review to imply human validation or readiness for monetary launch.
- The identity disagreement has been converted into a frozen, bars-before-code v40 comparison, including a hybrid arm and explicit privacy-benefit accounting.

What would make the page stronger

- Make it discoverable from Evidence, Replication, and the global footer. Calling the page 'unlisted' weakens the transparency value of publishing it.
- Add a compact adoption ledger with one row per finding: accepted, partially accepted, declined, implemented, pending, and a direct evidence link.
- Version the response and show the public kit checksum reviewed by each reviewer.
- Distinguish 'no human expert review' from 'no external review' everywhere else on the site. Both formulations currently coexist.
- Publish future reviewer submissions and team responses as immutable artifacts, with corrections appended rather than replacing earlier text.

2. Differential audit of the current public replication kit

The current public archive was downloaded from the replication page and compared against the exact corpus independently executed for the original report. Current archive SHA-256: bd1c514cb39dfb6311c08f28e28d0acd7abc75c96cd11130ce1b53673d5b1d3f. It contains 617 archive entries, 553 files, and 59,014,066 uncompressed bytes. The published ZIP is 31,615,995 bytes.

The differential matters: no existing simulation engine or results artifact changed. The additions and modifications are release metadata, reviewer documentation, current ledgers, claims files, helper utilities, and the v40 spec-only folder. Therefore the original 831.1-second independent engine run remains applicable to the computational corpus; this addendum audits the new packaging and public claims rather than pretending a new engine replication occurred.

Original finding	Status	What changed	Remaining action
Canonical manifest	RESOLVED	manifest.yaml lists 82 engines, roles, committed outputs, canonical status, and clean-room exceptions.	Add claim IDs, seeds, timeouts, and expected artifact hashes in the next version.
Current reproduction ledger	RESOLVED	Standard and clean-room dated ledgers plus LATEST.json now ship.	Keep ledgers immutable and sign or checksum releases.
Mismatch taxonomy	RESOLVED	Runtime-only, legacy numeric, schema drift, and helper classifications are explicit.	Add machine-readable reason codes to the manifest.
Claims through v39	PARTIAL	Register and claims.json now include v35-v39 and 142 machine-readable rows.	Make claims.json the source of truth; older prose still says v34 and contains stale coverage footers.
Economist Brief included	PARTIAL	The promised file now exists.	Rewrite it: it still says 34 waves, 49/54, and 'not externally reviewed.'
Environment specification	PARTIAL	requirements.txt and three attested environments now ship.	The file declares floors, not a lock. Add uv.lock or equivalent, container digest, and CI matrix.
Clean-room recomputation	PARTIAL	A clean-room ledger reports 69 of 82 from scratch.	Four engines exit without checkpoints. Implement actual no-checkpoint paths and rerun.
Preregistration evidence	PARTIAL	v40 bars are visibly registered before code and the freeze is explicit.	Use externally timestamped immutable commits or archival deposits; file dates alone remain weak evidence.
Immutable release	OPEN	The ZIP is downloadable and hashable but lacks a public DOI, signed tag, or release index.	Publish R2 to Zenodo or equivalent with checksum, changelog, and source repository tag.

PACKAGING CONCLUSION. This is no longer the same kit that I criticized for lacking a manifest, current claims, requirements, and the promised economist handoff. Those findings must be updated. It is also not yet a fully coherent R2 research release because the new metadata layer was added on top of stale start-here prose instead of replacing it with one canonical release narrative.

3. Public website and paper - correction status

Issue	Status	Current evidence	Recommended correction
Retired near-zero inflation headline	RESOLVED	No occurrence found on the current homepage or paper.	Keep it retired and explain price-level claims only with index-governance limits.
Dynasty metric inconsistency	RESOLVED	Current paper consistently presents 79.2% to 10.4%; homepage uses 79% to 10.4%.	Label metric, horizon, baseline calibration, and uncertainty every time.

Issue	Status	Current evidence	Recommended correction
Reproduction denominator	PARTIAL	Paper and replication page use 72/82 and distinguish owner-machine 73/82.	Economist Brief and reviewer guide still foreground 49/54; reconcile all start-here documents.
Claims register through v39	PARTIAL	v35-v39 rows and claims.json now ship.	Remove stale v34 and owed-refresh language from the register's earlier sections and footer.
Bank-money opener	OPEN	Homepage still says nearly every dollar is born as a loan and settles on existing asset owners.	Use sourced language about commercial-bank deposit creation through lending and distinguish central-bank money, fiscal flows, cash, and distributional effects.
Independent-review provenance	PARTIAL	Paper links the new page and accurately describes the machine review.	Replication page and footer still say not externally reviewed and same-family-only; change to 'no human expert review.'
Economist-facing paper	OPEN	Provenance and reproduction sections improved.	The page still contains 28 TODO markers, lacks a conventional PDF working paper, and has no EdenQuest section.
EdenQuest deployment story	PARTIAL	Dedicated bridge page and main-site ladder are strong.	Add the non-monetary deployment and identification strategy to /paper so economists see the actual first stage.
Independent-reviews discoverability	OPEN	The paper links it, but the page labels itself unlisted and it is absent from primary navigation/footer.	Make reviews a first-class evidence category.

4. EdenQuest changes the launch assessment

The original review treated the project mainly as an economic stack. The later EdenQuest review establishes a safer and more credible first stage: a useful local-first app in which users collect structured data for themselves, can use it without an account, can keep EQ Coins permanently non-transferable, and can optionally connect their own AI provider with category controls and an exact request preview.

That context changes the meaning of 'launch.' EdenQuest does not require external buyer demand, EVE liquidity, one-person-one-identity, lending, a conduct market, or a funded floor. Its organic-growth strategy can create honest longitudinal data because users first benefit from accuracy themselves. This may reduce cold-start and gaming problems relative to paying strangers for raw submissions.

Revised EdenQuest position

PUBLIC BETA IS REASONABLE. I support continued public availability of EdenQuest as a personal tool, provided entry into any future data commons is a separate, default-off, granular, revocable decision and declining it never reduces the app's core value.

Organic use is evidence of product usefulness and data-production behavior. It is not yet evidence that outside buyers value the resulting corpus, that cross-user measures are comparable, that anonymization is sufficient, or that compensation will not change user behavior. The commons transition should therefore be treated as a new product and a new human-subjects risk boundary.

Unresolved EdenQuest issues found in the current live product

- End-to-end encryption remains optional for cloud sync rather than the normal default for sensitive journals, check-ins, household data, and experiments.
- The privacy policy still does not clearly identify the legal operator, jurisdiction, physical or registered address, or a concrete privacy and security contact.
- The child-data inconsistency remains. The policy says child profile data is excluded from AI context, while the live preview says a child's summary can be fetched and sent when 'Child profiles' is enabled and the conversation asks about the child.
- Crash reporting to Sentry can include approximate region derived from IP; users should receive a clear operational disclosure and an opt-out where practical.
- The app's self-improvement framing will predictably encounter health and mental-health content. AI safety needs product guardrails, not only a professional-advice disclaimer.
- The free, no-ads, bring-your-own-key structure cleverly controls early costs, but long-term sync, security, and support need a sustainable non-speculative funding plan.

5. Identity response and v40

The team's pushback improves my understanding of the proposed identity design. The system already separates uniqueness from on-ledger biometrics, uses pseudonymous identifiers, rejects a scalar global trust score, and constrains some detection to aggregates. My original language should not imply that EDEN simply proposes a public universal reputation score.

The actual dispute is whether sanctions, lending history, ownership consequences, and reputation should cross contexts through one persistent economic root. The team is correct that compartmentalization may weaken rental deterrence, unsecured lending, and beneficial-ownership enforcement. The new v40 spec is a serious response because it measures full-global, contextual, and hybrid variants on the committed v4, v27, and v36 chassis; registers failure and reopening bars; includes a hybrid financial-global/social-contextual arm; and explicitly models privacy benefits in CR6.

What v40 can and cannot settle

- It can estimate enforcement tradeoffs inside the existing model family and reveal whether cross-context coupling is economically load-bearing under the selected dials.
- It cannot establish that lifetime cross-context consequence is legitimate, proportionate, non-discriminatory, or acceptable to participants.
- Its thresholds are owner-ratified design judgments, not empirical noninferiority margins derived from field data.
- CR1's failure at any tested dial is useful as a stress bar but should not by itself decide architecture; risk-weighted prevalence and real-world mitigations matter.
- The spec itself admits that stigma, chilling effects, coordinated harassment, and social consequences are not simulated. Those require civil-rights review, qualitative research, and bounded pilot evidence.

Revised recommendation: run v40 after an outside reviewer grades the spec, but treat it as an engineering comparison, not a referendum on human rights. Publish the rights charter, sanctions matrix, recovery rules, appeals process, and prohibited-use boundaries before any identity-linked pilot.

6. Internal EVE credit and the dollar interface

The proposed non-transferable internal credit is a meaningful improvement over a purely cosmetic shadow ledger. If it is redeemable only for governed AI queries, cannot be transferred, cashed out, collateralized, exchanged, or promised future convertibility, it can generate real user behavior without launching a public currency.

The response also adopts two important evidentiary limits: buyers settle in dollars, and internal redemption is never cited as external demand evidence. Those constraints align with the original review's purpose even though the measurement instrument differs.

I would modify the original recommendation accordingly: allow a tightly capped query-credit experiment after legal review and funded compute provisioning. Keep a hard liability ceiling, publish redemption and outage rules, separate label effects from earning and display effects where feasible, and continue to require independent USD willingness-to-pay as the decisive demand test.

7. Revised launch ladder

Layer	Posture	Required boundary
EdenQuest local/accountless use	**READY**	Continue as a clearly labeled public beta; monitor safety, retention, export, deletion, and data-loss incidents.
EdenQuest sync and personal AI	**READY**	Continue with privacy fixes, security review, child-data reconciliation, and stronger AI safeguards.
Opt-in data commons	**PILOT**	One low-risk data category; separate consent; aggregate or compute-to-data access; independent reidentification attack.
Paid institutional queries in USD	**PILOT**	Require real budget authority, measurable buyer lift, repeat purchase, and public negative results.
Non-transferable query credits	**PILOT**	Funded capacity, capped liability, no cash-out, no secondary market, no future-convertibility promise.
Transferable EVE or exchange	**NOT READY**	No launch until incremental necessity, legal status, liquidity, incidence, and repeat external demand are demonstrated.
Lifetime economic identity/lending	**NOT READY**	No launch before rights instruments, human review, recovery evidence, and bounded adversarial testing.
Conduct/whistle market or essentials floor	**NOT READY**	Simulation results are insufficient for deployment affecting livelihood, reputation, or essential access.

8. Modification map to the original review

Original location	Revised treatment
Executive judgment	Keep the support-for-validation verdict. Upgrade the team/process assessment and split launch readiness by layer.
Section 2 - Evidence audit	Replace the earlier packaging critique with the R2 differential audit in this addendum. Keep the model-validity and empirical-evidence limits.
Section 3.1 - Money diagnosis	Unchanged. The homepage opener still needs correction.
Sections 3.2-3.4 - Currency and dollar link	Permit a non-transferable, single-purpose query credit as an experiment; retain USD settlement and necessity test.

Original location	Revised treatment
Section 4 - Identity	Correct the technical characterization; focus the objection on cross-context consequence, rights, and proportionality. Add v40 as a constructive response.
Section 6 - /paper	Add the new provenance and reproduction improvements, but retain the conclusion: 28 TODOs, no conventional working paper, and no EdenQuest deployment section.
Section 7 - Website	Mark inflation and dynasty corrections resolved; mark reproduction, claims, and provenance partial; retain bank-money and discoverability findings.
Sections 8-9 - Experiments and roadmap	Insert EdenQuest as Stage 0. Move the invitation-only boundary to the commons, paid-query, identity, and credit experiments.
Section 15 - Priorities	Mark manifest, ledgers, review publication, v39 claims extension, and v40 preregistration as acted upon; reorder remaining work around coherence, human review, and field evidence.

9. Priority actions from here

Immediate release cleanup - before calling R2 complete

- Rewrite the Economist Brief and Reviewer Guide around the current 82-engine manifest. Move old 52- and 54-engine ledgers to a clearly labeled history appendix.
- Make claims.json or another machine-readable registry canonical, then generate the human register, paper claim cards, and website numbers from it.
- Replace all remaining 'not externally reviewed' and same-family-only statements with precise language: one out-of-family AI review exists; no human expert review exists.
- Correct the homepage bank-money opener and add EdenQuest's deployment/identification role to /paper.
- Publish the full adoption ledger rather than leaving it only in the private vault or a future release.

Next 30 days

- Create a hash-locked environment and container image; run CI on at least Linux and macOS; implement true no-checkpoint paths for the four blocked engines.
- Archive R2 with DOI, checksums, release notes, and a tagged public source repository or equivalent immutable source snapshot.
- Commission at least one human economist, one privacy/security reviewer, and one civil-rights or identity-governance reviewer with unrestricted publication rights.
- Publish the rights charter, sanctions-and-recovery matrix, appeals standards, and prohibited-use rules claimed as adopted.
- Fix EdenQuest's operator/contact disclosures, child-AI contradiction, sync-encryption choice architecture, and AI safety boundaries.

Next evidence cycle

- Use EdenQuest retention and self-experiment completion to measure whether the app actually creates durable user value and usable longitudinal data.
- Design the first commons experiment around one low-risk, well-defined schema; test consent comprehension, revocation, measurement comparability, and reidentification risk.

- Secure three unrelated buyers with real budgets and run USD-priced query tests. Renewal without subsidy matters more than first purchase.
- Grade v40 externally before execution, then publish results whether global, contextual, or hybrid architecture wins its registered bars.
- Do not let successful EdenQuest adoption be presented as evidence for EVE demand, lifetime identity, a conduct market, or a funded floor.

10. Final revised judgment

The project is now stronger than the version described in the original report. The improvement is not that its simulations suddenly prove an economy. The improvement is that the team exposed criticism, preserved the unedited artifact, corrected real packaging defects, made disagreements testable, and clarified a bottom-up product path through EdenQuest.

That is enough to increase confidence in the team and to support continued work. It is not enough to increase confidence in buyer demand, live identity, governance legitimacy, incidence, or monetary deployment, because no new real-world evidence on those questions has appeared.

UPDATED ONE-SENTENCE RECOMMENDATION. Support EdenQuest, Release R2 completion, independent human review, and a bounded USD-settled data/query pilot; require every transition into shared data, identity-linked consequence, credit, or EVE to earn its own consent, evidence, legal basis, and shutdown path.

Evidence and source record

- **Independent reviews and team response:** https://www.explorationeconomy.org/independent-reviews
- **Economist-facing paper:** https://www.explorationeconomy.org/paper
- **Replication page and current public kit:** https://www.explorationeconomy.org/replication
- **Exploration Economy homepage:** https://www.explorationeconomy.org/
- **EdenQuest bridge page:** https://www.explorationeconomy.org/edenquest
- **EdenQuest live app:** https://edenquest.org/
- **EdenQuest privacy policy:** https://edenquest.org/privacy
- **EdenQuest terms:** https://edenquest.org/terms

Integrity notes

- Hosted review PDF SHA-256: 7fc49009cdeac45900bd57aaa843f44924edfba3947b4b7341da904d1a386655. It is byte-identical to the delivered original.
- Current public kit SHA-256: bd1c514cb39dfb6311c08f28e28d0acd7abc75c96cd11130ce1b53673d5b1d3f.
- The current kit was differentially compared to the executed July 31 kit. Existing simulation engines and results were unchanged; the current archive adds and revises packaging, documentation, ledgers, utilities, and a spec-only v40 folder.