Cognitive Liberty Institute

Peer Review and Final Editorial Statement

Cognitive Sovereignty in the Intelligence Economy:
A Unified Empirical–Mathematical–Ecological Framework
for Human–AI Value Systems

Homnath Khatiwada
Assistant Professor of Botany, Tribhuvan University, Nepal

Review conducted by the Cognitive Liberty Institute
Led by Alexander Huseby
coglib.no

Round 1 review issued: May 1, 2026
Final editorial statement: June 13, 2026

Part I

Peer Review — Round 1

Abstract

This paper develops a unified framework describing the transition from the attention economy to the intelligence economy, integrating biological cognition, information theory, behavioral economics, ecological niche theory, and artificial intelligence systems. It proposes that human cognition acts as a productive resource that compounds exponentially through interaction with AI, constrained by biological limits and ecological competition. A mathematically justified model is introduced in which intelligence production follows continuous exponential growth governed by cognitive productivity — operationalized through knowledge depth, attention intensity, and system interaction skill — moderated by an ecological niche factor and bounded by cognitive exhaustion. The framework advances the concept of cognitive sovereignty, arguing that human cognitive contributions in AI ecosystems constitute a form of digital resource requiring recognition, governance, and benefit-sharing, drawing analogy to the Nagoya Protocol on Access and Benefit-Sharing.

Review led by: Alexander Huseby, Founder and Executive Chairman, Cognitive Liberty Institute

Panel composition: Three reviewers — Mathematics & Production Theory | AI Systems & Empirical Methodology | Philosophy, Governance & Legal Frameworks

Priority sections: Section 2 (multiplicative value function), Section 3 (cognitive productivity coefficient alpha), Section 6 (Nagoya Protocol analogy)

Round: 1 of 2

Preliminary Assessment

The following preliminary assessment summarises the panel's initial reading across the three priority sections before individual reviewer reports were prepared.

Section 2 — Multiplicative value function and mathematical derivation

The exponential growth ODE (dI/dt = alphaI) is standard and uncontroversial. The appeal to LTP and learning curves as biological support is reasonable at a conceptual level, though the mapping from synaptic potentiation to a scalar productivity coefficient is informal rather than derived.

The bigger issue is the Cobb-Douglas justification. The author claims unit exponents on both inputs, implying constant returns to scale. That is an assumption, not a derivation, and it is a strong one. Standard Cobb-Douglas estimation produces exponents empirically; asserting them at 1 requires justification. If the exponents deviate from unity, the scaling behavior of V(t) changes substantially.

The information-geometric argument for N is the weakest part. The author invokes the Fisher information metric and independence assumptions to justify N entering multiplicatively, but the argument is imprecise. Fisher information is a property of a statistical model's parameter space, not a discount factor on production value. The phrase "replication-survival distribution" is introduced without formal definition. The independence assumption (that N is independent of production rate) is stated but not defended.

The exhaustion correction in Section 2.3 is the strongest part of the mathematics. The move from a multiplicative min() constraint to a proper state variable with a critical time T* is correct and well-handled.

Verdict: The ODE and exhaustion sections are defensible. The Cobb-Douglas and information-geometric sections need either a proper derivation or explicit acknowledgment that these are structural assumptions rather than derived results.

Section 3 — Operationalisation of alpha

The decomposition alpha = beta1K + beta2A + beta3S is linear and additive, which sits in tension with the multiplicative logic used everywhere else in the paper. A multiplicative or log-linear specification would be more internally consistent.

The Shannon entropy operationalisation of K is conceptually coherent but faces a serious tractability problem: the probability distribution over a knowledge domain before and after a reasoning episode must be specified. The author acknowledges measurability concerns but understates how acute this problem is.

S (prompt engineering efficiency as Q_output / P_effort) has a circularity problem. Q_output requires a quality metric, which is itself the thing the broader model is trying to explain.

Verdict: alpha is the empirical core of the model and it is currently under-specified for measurement. The paper needs either a concrete measurement protocol or an explicit step back from strong tractability claims.

Section 6 — Nagoya Protocol analogy and cognitive sovereignty

This is where the author's domain expertise shows, and the structural parallel is genuinely interesting. The four-part mapping (genetic resources to cognitive resources, traditional knowledge holders to users, bioprospecting to platform operators, ABS to cognitive sovereignty) is coherent and publishable as a governance argument.

The author correctly identifies the central disanalogy: genetic resources are geographically bounded and attributable, while cognitive contributions are distributed, pseudonymous, and often aggregated beyond individual traceability. The paper notes this but moves past it too quickly.

The three-component sovereignty framework (recognition, control, participation rights) is asserted rather than derived from the Nagoya analogy. The rights framework needs a different grounding, possibly in labor theory or data dignity literature.

Verdict: The Nagoya analogy works as an analogy and as a rhetorical frame. It does not yet work as a legal or governance foundation. The section needs either a more careful limitation of the analogy's scope or a supplementary grounding in a framework that does support positive cognitive rights.

Reviewer 1 — Mathematics and Production Theory

Overall assessment: Major revision required

Section 2.1 — The proportionality argument and ODE

The differential equation dI/dt = alphaI is mathematically standard and the derivation of I(t) = I□e^(alphat) is correct. However, the claim that this choice is "natural and necessary" conflates two things: the base e is an expositional convenience, not a substantive constraint. Any exponential with a different base is mathematically equivalent under reparametrisation of alpha. This should be corrected to avoid the impression of false necessity.

More substantively: the ODE assumes alpha is a time-invariant constant. This is an extremely strong assumption. Cognitive productivity plausibly changes over time — it may increase with practice (learning-by-doing, Arrow 1962) or decrease with fatigue. The paper should either defend constant alpha explicitly or acknowledge it as a simplifying assumption with known limitations.

Section 2.2 — Cobb-Douglas justification

The structural parallel to Cobb-Douglas is useful as an intuition pump but is not a derivation. Asserting unit exponents on both platform amplification and cognitive input is equivalent to asserting that a 1% increase in either input produces exactly a 1% increase in value output. This is a joint hypothesis that requires empirical support, not a consequence of the Cobb-Douglas form itself.

Furthermore, the paper's formulation V(t) = P□e^(rt) · I□e^(alphat) · N · □ has four multiplicative terms. Extending Cobb-Douglas logic to four inputs with unit exponents implies strongly increasing returns to scale, which is not acknowledged and may not be defensible.

Section 2.2.2 — Information-geometric argument for N

This section requires substantial revision. The Fisher information metric is a Riemannian metric on the space of probability distributions; it is not

straightforwardly a discount factor on production value, and the paper does not establish the mapping between these two concepts rigorously.

The "replication-survival distribution" is introduced as a formal object but never defined. What probability space is this distribution over? The independence assumption between N and production rate is stated but not defended. The author should either replace this section with a simpler and more defensible argument for why N enters multiplicatively, or substantially formalize the information-geometric argument with precise definitions.

Section 2.3 — Cognitive exhaustion

This is the strongest mathematical section. The correction from a multiplicative min() constraint to a proper state variable is well-executed. The derivation of T* is correct and clearly presented. The pharmacokinetic recovery function is a reasonable and well-motivated choice.

Section 3 — alpha decomposition

The additive decomposition alpha = beta1K + beta2A + beta3S is internally inconsistent with the paper's own multiplicative logic. A multiplicative or log-linear specification (ln alpha = beta1 ln K + beta2 ln A + beta3 ln S) would be more consistent and would also make the regression model in Section 7.1 more natural.

The Shannon entropy operationalisation of K requires specifying prior and posterior probability distributions over a knowledge domain. The paper should engage with the psychometric and cognitive measurement literature more seriously before asserting tractability.

Reviewer 2 — AI Systems and Empirical Methodology

Overall assessment: Major revision required

The intelligence economy framing

The transition narrative from attention to behavior to intelligence is well-constructed and the lineage through Simon, Davenport-Beck, and Zuboff is appropriate. One gap: the paper does not engage with the existing literature on human-AI complementarity (Brynjolfsson & Mitchell, 2017; Autor, 2022) or with the emerging empirical work on AI's effects on knowledge worker productivity. These are directly relevant and their omission weakens the literature positioning.

The cognitive productivity coefficient alpha

K (knowledge depth via Shannon entropy reduction) requires measuring a probability distribution over a knowledge domain before and after a reasoning episode. The claim that this is "precisely how Shannon defined information" is correct but does not resolve the measurement problem.

A (attention intensity as T_focused / T_total) is the most tractable measure. Interaction log data can provide proxies, and eye-tracking or EEG measures can operationalise focused vs. diffuse attention in controlled settings.

S (prompt engineering efficiency as Q_output / P_effort) has a circularity problem the paper does not resolve. Using Q_output as an input to alpha, which itself governs output quality, creates a feedback loop that makes causal identification extremely difficult.

Section 7 — Empirical testability

The primary regression model ln I(t) = ln I□ + alphat + □ is a standard log-linear specification and is estimable in principle. However, several identification problems are not addressed.

First, alpha is decomposed into K, A, and S, which are themselves time-varying and potentially endogenous. A user who improves their prompt engineering skill (S) over time will also tend to acquire deeper knowledge (K); these are correlated, making coefficient separation difficult without strong instruments or experimental variation.

Second, the kind of fine-grained interaction logs needed to estimate K, A, and S are held by platform operators and are not publicly available. The paper should address how empirical work would actually be conducted.

The exhaustion-recovery cycle prediction in Section 7.3 is the most tractable. EEG and fMRI paradigms for measuring cognitive fatigue during sustained tasks exist and are well-validated. This is a genuinely novel experimental prediction and should be developed further.

Reviewer 3 — Philosophy, Governance, and Legal Frameworks

Overall assessment: Minor to major revision required (section dependent)

Section 6.1 — The Cognitive Resource Principle

The claim that cognitive labor in AI interaction is "rivalrous in attention, non-rival in its contribution to AI training" is important and mostly accurate, but requires a qualification. Whether cognitive contributions are truly non-rival depends on how AI training is structured. In fine-tuning scenarios targeting specific domain expertise, individual contributions may be more rivalrous.

Section 6.2 — The Nagoya Protocol analogy

The structural mapping is well-constructed and the four-part parallel is the paper's most original governance contribution. However, the central disanalogy is more fundamental than the paper acknowledges. The Nagoya Protocol's legal architecture rests on three pillars — territorial sovereignty, traceability, and prior informed consent — none of which transfer cleanly to cognitive contributions.

Cognitive contributions are not subject to territorial sovereignty in any established legal sense. They are not traceable: prompt data is aggregated, anonymised, and transformed through training pipelines in ways that make attribution practically impossible. The paper should either substantially narrow the scope of the analogy or develop a supplementary legal foundation that does not depend on the Nagoya architecture.

Section 6.3 — The three-component sovereignty framework

Recognition rights, control rights, and participation rights are a coherent tripartite framework. However, the paper cites the EU AI Act as support for the participation rights argument. The Act's provisions are primarily procedural rather than substantive rights to share in value generated. This overstates the alignment and should be corrected.

The framework does not engage with the practical tension between control rights and the non-rival nature of cognitive contributions claimed in Section 6.1. The paper needs to resolve whether cognitive contributions are better framed as labor inputs or as resources — these lead to different legal and policy architectures.

The Global South dimension

The author's note about the Nepal context and the intersection of Himalayan traditional knowledge with digital infrastructure is more analytically powerful than the paper currently reflects. The framework would be substantially strengthened by a dedicated section on how cognitive sovereignty dynamics play out differently in contexts where traditional and ecological knowledge is already subject to extraction dynamics in the biological domain, and is now also subject to digital extraction. This is a genuinely distinctive contribution that the current manuscript leaves implicit.

Consolidated Editorial Recommendation — Round 1

Decision: Major Revision

The manuscript advances a timely and genuinely original argument. The core claim — that human cognitive labor in AI ecosystems constitutes an uncompensated productive resource requiring governance — is well-motivated, interdisciplinarily ambitious, and has clear policy relevance. The biological and ecological foundations are the strongest components and reflect the author's domain expertise effectively.

The priority revision demands are:

> — The Cobb-Douglas and information-geometric arguments in Section 2 need either formal derivation or explicit reframing as structural assumptions.
>
> — The additive decomposition of alpha in Section 3 is inconsistent with the paper's multiplicative logic and should be revised. Measurability claims

need to be substantiated with concrete measurement protocols or moderated.

— The Nagoya Protocol analogy in Section 6 needs a clearer limitation of scope and a supplementary legal or normative grounding. The tension between non-rivalry and control rights needs to be resolved rather than left implicit.

— The empirical identification strategy in Section 7 needs to address endogeneity, data access, and selection bias more seriously.

— The Global South and traditional knowledge dimension should be developed within the manuscript itself — it is the paper's most distinctive empirical grounding and currently exists only in correspondence.

Alexander Huseby
Founder and Executive Chairman
Cognitive Liberty Institute
coglib.no
May 1, 2026

Part II

Final Editorial Statement — Round 2

Manuscript: Cognitive Sovereignty in the Intelligence Economy: A Unified Empirical-Mathematical-Ecological Framework for Human-AI Value Systems

Author: Homnath Khatiwada, Assistant Professor of Botany, Tribhuvan University

Review round: 2 (Final)

Date: June 13, 2026

Decision: Accept with Minor Revisions

---

1. Overview

The Cognitive Liberty Institute convened a three-reviewer panel to assess this manuscript across formal mathematics and production theory, AI systems and empirical methodology, and philosophy of mind and governance frameworks. Round 1 returned a Major Revision decision with five consolidated demands. The author has now submitted a revised manuscript (v3) accompanied by a detailed response letter.

The panel has completed its assessment of the revision. The verdict is Accept with Minor Revisions. One outstanding issue must be resolved before submission; it is identified in Section 4 below.

---

2. Assessment of Round 1 Demands

Priority 1: Mathematical Foundations (Section 2)

All four mathematical demands have been addressed satisfactorily. The Cobb-Douglas unit-exponent claim has been retracted and replaced with an explicit increasing-returns-to-scale framing grounded in Romer (1990). The Fisher information-geometric argument for N has been fully retracted and replaced with a direct survival-fraction scaling argument that is definitionally defensible. The false-necessity claim regarding base e has been corrected. The time-invariance of alpha has been scoped as a short-run approximation with appropriate caveats.

The panel notes that the revision is honest about what remains structural assumption versus empirical finding, which strengthens rather than weakens the paper's credibility with mathematical readers.

Priority 2: The alpha Decomposition (Section 3)

The additive specification has been replaced with a log-linear (Cobb-Douglas) decomposition at the level of alpha, restoring internal consistency with the paper's multiplicative logic. The circularity problem in S has been resolved through external benchmark anchoring. The dimensional mismatch between episode-level K and corpus-level K_unique has been resolved with an explicit aggregation function. Strong tractability claims have been appropriately moderated.

Priority 3: Nagoya Protocol Analogy and Governance Foundation (Section 6)

The three pillars of the Nagoya Protocol that do not transfer are now stated explicitly. The Ostromian commons governance framework is developed as the positive institutional foundation, grounded in two empirically validated Nepali institutions: Nepal's Community Forest User Groups and the Rupa Lake Restoration and Fisheries Cooperative. The three-tier governance architecture directly incorporates the critical institutional failure mode identified by Adhikari, Kingi and Ganesh (2014).

The panel considers the governance section the most substantially improved element of the revision. The Ostromian grounding is appropriate, empirically anchored, and more institutionally sophisticated than the Round 1 framework.

Priority 4: Empirical Identification Strategy (Section 8)

The cross-sectional OLS specification has been replaced with a panel data model with individual fixed effects and AR(1) error structure. The r vs. alpha identification problem has been named and addressed with two proposed separation strategies. Endogeneity and data access challenges are acknowledged with concrete proposed resolutions.

The panel notes that the author identified the r vs. alpha separation problem independently, beyond what Round 1 demanded. This is evidence of genuine engagement with the methodological stakes of the argument.

Priority 5: Global South and Traditional Knowledge Dimension

This is now a full theoretical and empirical section comprising seven subsections. The Voluntary Extraction Inversion and three-phase extraction typology constitute the paper's most significant original theoretical contribution. The panel considers the typology a genuine advance: it names and formally theorizes a phenomenon that the existing literature on surveillance capitalism, AI governance, and indigenous knowledge rights has circled without yet unifying.

---

3. Additional Contributions

The revised manuscript includes several contributions not requested in Round 1. The panel flags the following as substantive additions that strengthen the paper.

The Intelligence Economy Survival Gap G provides a formal economic measure of the cognitive labor compensation deficit. The Homeostasis Condition formalizes the circular economy argument as applied to cognitive labor. The Sandakpur Theorem, derived from field observation in Ward No. 3, Sandakpur Municipality, offers five formal stability conditions for survival microeconomy governance. The Digital Niche Displacement Event is formally classified as a Type I Error and generates a testable empirical prediction regarding algorithmic security protocol discrimination rates across network environments.

The Rupa Lake Cooperative is documented as a functioning institutional prototype for the cognitive commons model — an unusual and valuable empirical grounding for what is typically a theoretical argument.

---

4. Outstanding Issue: Abstract Placeholder

One unresolved issue must be corrected before submission to any journal. The abstract refers to the Sandakpur Theorem as "(Name??) Theorem." The body of the manuscript and the response letter both use "Sandakpur Theorem" consistently. The placeholder in the abstract is an oversight that will disqualify the manuscript at editorial screening. Replace "(Name??)" with "Sandakpur" in the abstract before submission.

---

5. Considerations for Journal Submission

The manuscript is interdisciplinary in a way that creates genuine editorial routing risk. The mathematical sections will be assessed as insufficiently rigorous by pure economics or mathematics journals; the governance sections will be assessed as insufficiently grounded by legal or political science journals; the field observations will be assessed as methodologically thin by empirical social science journals. The paper is best positioned at journals explicitly committed to interdisciplinary AI governance, cognitive rights, or digital political economy.

The author followed CLI's advice to frame disciplinary boundaries explicitly in the manuscript. The cover letter to the journal should reinforce this framing: the paper's contribution is theoretical architecture and problem identification, not definitive empirical proof.

The field observations from Eastern Nepal are the paper's most distinctive empirical grounding and its most vulnerable methodological point simultaneously. They are first-person observations by the primary author, not independently collected data. The author should frame them explicitly as illustrative pilot evidence and naturalistic pilot datasets, consistently throughout the manuscript.

---

6. Final Assessment

The Round 1 Major Revision decision was issued in recognition of genuine theoretical ambition alongside specific technical and architectural gaps. The author has addressed those gaps seriously, honestly, and in several cases beyond what was asked.

The Voluntary Extraction Inversion is a substantive theoretical contribution. The three-phase extraction typology provides a governance narrative that the field has not had. The Ostromian institutional grounding is appropriate and empirically anchored. The mathematical corrections are honest and competent. The identification of the r vs. alpha separation problem demonstrates methodological rigor the panel did not require.

The manuscript is ready for submission, subject to correction of the abstract placeholder identified in Section 4.

CLI considers the review process for this manuscript complete with this statement.

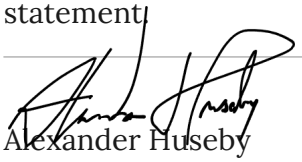

Alexander Huseby
Founder and Executive Chairman
Cognitive Liberty Institute
coglib.no
June 13, 2026