

Article

Not peer-reviewed version

Negentropic Value Currency: The Governance Architecture of Algorithmic Democracy and Constitutional Firewalls

Dian Cheng and [Jinjun Cheng](#) *

Posted Date: 23 October 2026

Keywords: algorithmic democracy; deliberative democracy; sortition; algorithmic accountability; monetary constitution; mechanism design



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Negentropic Value Currency: The Governance Architecture of Algorithmic Democracy and Constitutional Firewalls

Dian Cheng ¹ and Jinjun Cheng ^{2,*}

¹ CICC Anhui Branch, Hefei, China; MBA, University of Science and Technology of China

² Anhui Vocational College of Grain Engineering, Hefei, China; B.A., Anhui University

* Correspondence: cdcgcd999@126.com

Abstract

The first paper of the negentropic value currency program established the value ontology and the microeconomic incentives (Cheng and Cheng, 2026d); the second proved that macrostability is equivalent to the maintenance of four integrity conditions — the measurement layer unfalsified, the calibration layer uncaptured, the issuance layer uncorrupted, and the extinction layer unfailing (Cheng and Cheng, 2026b). Both papers, however, left the guardian of integrity offstage: they assumed the existence of legitimate accounting, consensus, and supervision procedures without specifying the procedures themselves. The present paper takes up this suspended constitutional question: who guards the four integrity conditions, and by whom, through what procedures, and under what constraints is α set, so that the system escapes majoritarian tyranny, technocratic rule, and the alienation of algorithmic black boxes alike? We formalize the problem as the delegation of value judgments and propose a three-tier governance architecture. The first tier is the bicameral algorithmic democracy of the Global Value Parliament (GVP): a House of Citizens generated by three-dimensional stratified sortition (500 seats, three-year terms, no reappointment) and a House of Experts nominated by authoritative learned bodies and confirmed through peer review (300 seats, six-year terms), with revisions of core parameters requiring a double majority in both houses. The second tier is three constitutional firewalls — the Minority Protection and Exploration Fund (value-inclusion axis), the independent Audit and Emergency Adjustment Committee (process-oversight axis, with quarterly disparate-impact audits and emergency braking powers, extended here to the audit of the extinction layer), and a blockchain-notarized platform of popular oversight (popular-oversight axis, with million-signature initiatives and a 70% opposition re-review rule). The third tier is an economic constitution of metabolic governance, now with four clauses: the intrinsic metabolism rate δ , the Civilization Time Bank, financial minimalism, and knowledge-genealogy attribution with intergenerational decay of inheritance. A guardianship matrix assigns every institution to the integrity condition it defends, and the capture-cost proposition states that global capture is conjunctively superadditive. The teleological loop then anchors the architecture to the survival and elevation of civilization: knowing, utilizing, and mastering cosmic truth, until the withdrawal of money itself. Every mechanism is accompanied by its failure-mode analysis, and all claims of validity are marked as designed but unexecuted; the verification designs are ordered by execution priority, with the audit-detection design first and the sortition design re-scoped to laboratory evidence.

Keywords: algorithmic democracy; deliberative democracy; sortition; algorithmic accountability; monetary constitution; mechanism design

JEL Classification: D71; D72; E42; H11; O17

1. Introduction

The first two papers of this program left a structural gap. The first (Cheng and Cheng, 2026d) migrated the monetary value anchor to net order creation, constructing the value equation $W_{\tau} = \sqrt{(E_{\tau} \cdot T_{\tau} \cdot I_{\tau})}$ — with structural information gain I_{τ} measured counterfactually against the no-intervention baseline and events aggregated over disjoint activities and windows under its aggregation axiom (Axiom 3), the symmetric benchmark $a = 1/2$ being justified there as the minimax convention — and parameterized civilizational preferences as the consensus coefficient α . The second (Cheng and Cheng, 2026b) proved that the constancy of the accounting unit is a construction identity (its Proposition 0), conditional on four-layer integrity — the measurement layer unfalsified, the calibration layer uncaptured, the issuance layer uncorrupted, the extinction layer unfailing — and rediagnosed inflation as the relative deviation produced by failures of these integrities. Yet both papers assumed the existence of legitimate accounting, consensus, and supervision procedures without saying who runs them and under what constraints. The gap is fatal: if the second paper rewrites inflation governance from aggregate regulation into integrity maintenance, then the question of who maintains integrity becomes the load-bearing wall of every macroeconomic conclusion. In a system where money issuance is determined entirely by accounting and consensus, the procedure that guards the four integrity conditions is the constitution of the system.

The task of this paper is to provide a conditional answer to this constitutional question. The design faces a triple dilemma. α carries value judgments, and the legitimacy of value judgments requires democratic participation, on pain of technocratic rule. Majority rule, however, can suppress minorities and strangle frontier exploration — majoritarian tyranny — while exploratory minorities are precisely the source of variation from which civilizational evolution proceeds; And α is computed by an algorithm, the CPPA, that is not neutral by nature: the preferences of its designers and the structural biases of society can be silently embedded in the core of distribution — “algorithmic alienation” (cf. Zuboff, 2019). No single mechanism can answer all three risks. Our answer is therefore a layered architecture: legitimacy is supplied by a bicameral algorithmic democracy, evolutionary diversity by a minority-protection fund, and algorithmic accountability by an independent audit committee, with ultimate sovereignty retained by a platform of popular oversight.

Three declarations precede the argument. First, this paper shares the conditional register of its predecessors: we claim that, granted the NVC normative premises, the governance problem exhibits the following structure and the following architecture answers these structural risks by design; we do not claim that the architecture has been proven effective. Second, every mechanism is accompanied by a failure-mode analysis — the lesson of the governance literature is that every institutional design’s self-description tends to exaggerate its robustness, and an explicit list of failure modes is the minimum of honesty. Third, the scope discipline of the series is maintained: this paper does not re-derive the value equation of the first paper and does not re-prove the stability propositions of the second; it takes the four integrity conditions as given and designs their guardianship. Every definition or proposition inherited from the companion papers is restated at the point of use rather than merely cited, and every example drawn from the program’s monograph (Cheng and Cheng, 2026a) is rewritten here in self-contained form, so that the paper can be read and judged on its own. Appendix B maps each premise inherited from the companion papers to the sections of this paper where it is invoked.

A fourth declaration concerns the paper’s own formal ambitions. The literature of social choice teaches humility about impossibility claims: Arrow’s theorem and the Gibbard–Satterthwaite theorem are deep results with decades of scrutiny behind them, and a new “impossibility” label invites comparison the claim cannot survive. This paper accordingly states its Proposition 1 as what it is — a bounded structural incompatibility among four design requirements for single-rule consensus mechanisms, sufficient to motivate a composite architecture but far weaker than the classical theorems — and positions it explicitly against them (Section 3). The architecture is the contribution; the proposition is its motivation, not its crown. For readers who wish to situate the present argument within the full program, both companion papers are publicly available as preprints

(Cheng and Cheng, 2026d, 2026b) and the program's monograph as a working paper (Cheng and Cheng, 2026a); none has yet undergone peer review.

Section 2 positions the paper in the literature, including the two conversations a monetary constitution of this kind cannot avoid: experimentalist governance and the selection model of representation. Section 3 formalizes the governance problem as a pair — the consensus-generation problem and the integrity-guardianship problem. Section 4 gives the bicameral design of the GVP. Section 5 gives the three constitutional firewalls, including an explicit termination argument for the recursive-oversight regress. Section 6 gives the four clauses of the economic constitution (δ , the CTB, financial minimalism, and genealogy attribution with intergenerational inheritance). Section 7 presents the guardianship matrix of the four integrity conditions and the capture-cost proposition. Section 8 articulates the teleological loop of the architecture — the survival and elevation of civilization. Section 9 lists the falsifiable hypotheses and verification designs, ordered by execution priority. Section 10 states limitations. Section 11 concludes. Appendix A summarizes the notation used throughout the paper.

2. Literature Positioning

2.1. *The Revival of Deliberative Democracy and Sortition*

Habermas's theory of communicative action locates the source of political legitimacy in "rational debate among all affected under free and equal conditions" (Habermas, 1996), and deliberative-democratic theory accordingly rewrites the quality standard of democracy from the aggregation of preferences to their formation. Fishkin's (2009) deliberative polls and the practice of deliberative mini-publics such as the Irish citizens' assemblies show that randomly selected citizens, under conditions of full information and structured debate, can render high-quality judgments on complex issues. The contemporary revival of sortition (Landemore, 2013) restores the egalitarian core of Athenian democracy: in matters of value judgment, every citizen's perspective is probabilistically equal, and randomness is the strongest resistance to lobbying, party patronage, and the entrenchment of professional politicians. Our House of Citizens inherits this tradition directly and extends it from the national to the global scale — a representation problem of its own, which we answer with three-dimensional stratified sortition (continental population share $\times$ cultural sphere $\times$ development level).

2.2. *Majoritarian Tyranny and Constitutional Checks and Balances*

The tyranny of the majority identified by Tocqueville and Mill (1835/1840; 1859) is the structural risk of majority rule; the constitutional economics of Buchanan and Tullock (1962) raises the problem to the level of rules: good outcomes cannot depend on good decision-makers, only on rules that constrain decision. Ostrom (1990) showed that polycentric, layered institutions with monitoring and graduated sanctions can sustain cooperation without central authority. Our three firewalls generalize the logic of constitutional economics to algorithmic governance: the object to be constrained expands from government to a compound power structure of parliament + algorithm + funds.

2.3. *Algorithmic Accountability and Theories of Justice*

An algorithm is not a neutral technical instrument but a value-laden device embedding its designers' preferences and society's structural characteristics; when an algorithm enters the core of distribution, independent auditing of the algorithm becomes the new key to the balance of powers. On the normative side, Rawls's (1971) difference principle and Sen's (1999) capability approach supply the foundations of justice for the program's AI revenue-sharing protocol (the 70/30 principle; Cheng and Cheng, 2026a; for its placement at the distribution stage of the monetary life cycle, see Cheng and Cheng, 2026b, Section 3, Stage 3) and for the avoidability design of the metabolism rate δ . The process philosophy of Whitehead (1929) and the dissipative-structure theory of Prigogine and Stengers (1984) ground the ontological premise of money as a flowing negentropy relation: a

governance architecture is not a static blueprint but a dynamic process keeping the system far from equilibrium. We also record a debt to the monetary-reform tradition that the economic constitution inherits: δ is Gesell's (1916) demurrage under constitutional constraint, and its qualified endorsement by Keynes (1936, ch. 23) is discussed in the companion papers (Cheng and Cheng, 2026d, Section 2.2; 2026b, Section 4.6).

2.4. *The Democratic Deficit of Global Governance*

One further literature supplies the feasibility question this paper must answer. Dahl (1999) doubted that international organizations can be democratic at all: as the scale of governance grows, the chain of delegation lengthens beyond any possibility of popular control, and international bodies degenerate into bargaining among bureaucracies. Held's (1995) cosmopolitan democracy answered with layered governance — global, regional, and national institutions, each with direct citizen access — but left the mechanism of citizen access at the level of principle. The present architecture sides with Held structurally and adds a third answer to Dahl's skepticism: sortition shortens the delegation chain to zero for one chamber — members of the House of Citizens are not delegates but statistical stand-ins for the global public — while the transparency axis (Section 5.3) converts the remaining delegation chains into continuously audited ones. Whether zero-length delegation by lot preserves responsiveness at planetary scale is an empirical question; the claim here is only that the architecture makes the question well-posed for a monetary constitution for the first time.

2.5. *Experimentalist Governance and the Selection Model of Representation*

Two further bodies of work shape the design directly, and both impose corrections on it.

The first is the experimentalist-governance account of Sabel and Zeitlin (2010). Their reading of the European Union identifies a recursive architecture — framework goals set jointly, lower-level units given discretion to pursue them, periodic peer review of performance, and revision of the framework goals themselves in light of what the units learn — as a governance form capable of operating precisely where neither hierarchy nor markets can: under strategic uncertainty, with dispersed and heterogeneous actors, and where the means and even the ends must be learned in the doing. The resemblance to this paper's architecture is not accidental, and we state it plainly: the CPPA's three-stage calibration procedure, the AEAC's quarterly disparate-impact audits, and the GVP's obligation to revise parameters against CRN accounting data are, in Sabel and Zeitlin's vocabulary, destabilization routines, peer review, and recursive goal revision constitutionalized for a monetary system. But the resemblance is bounded by one sharp difference that the dialogue forces into the open. Experimentalist governance keeps the framework goals themselves provisionally revisable — nothing is constitutional in the strong sense. The present architecture does the opposite at exactly one point: the survival-first telos and the four integrity conditions are not up for revision, while everything below them is. The NVC constitution is thus an experimentalist regime operating under a fixed constitutional telos — a hybrid the experimentalist literature has not treated, because its canonical cases (the EU, the Montreal Protocol) are goal-plural. Whether a fixed telos strengthens the recursion (by giving peer review an invariant standard against which to measure drift) or corrupts it (by shielding one value commitment from the destabilization all other commitments face) is a genuine open question; we flag it here and return to it in the limitations, rather than claiming the experimentalist pedigree without its costs.

The second is Mansbridge's re-reading of representation. Her critique of the standard sanction model — in which representatives are controlled through the threat of non-reelection — identifies selection as an independent and often superior mechanism of alignment: where interests are aligned ex ante and independence from sanction is desirable, choosing representatives for fit is better than disciplining them for deviation (Mansbridge, 2003); and her "selection model" of political representation develops the affirmative logic of exactly the institutional form our House of Citizens takes (Mansbridge, 2009). The House of Citizens is a pure selection-model chamber: members are drawn by lot for statistical fit with the global public, serve one non-renewable term, and are thereafter

beyond the reach of both reward and sanction. Mansbridge's framework clarifies both why this is defensible — sanction-based accountability is not merely unavailable under sortition but inappropriate to it, since re-eligibility would reintroduce precisely the career incentives sortition exists to abolish — and where its genuine cost lies: a selection chamber cannot be punished, so drift between the chamber's deliberated judgments and the public's raw preferences has no electoral correction. The architecture's answer is that this residual drift is not corrected at the chamber but around it — by the transparency axis's deliberative polls and re-review rule (Section 5.3), which reintroduce the public as a check without reintroducing re-election. Mansbridge's framework lets us state this design as a principled allocation of representation types rather than an improvisation: selection for the standing chamber, sanction-equivalent review for its outputs. The selection model has its critics. Lafont (2020) objects that empowering randomly selected bodies shortcuts the public sphere: a mini-public's outcomes carry no prescriptive force for citizens who did not participate in its deliberation, and such bodies are better used to inform and contest public debate than to decide in its place. The present design absorbs the correction rather than dismissing it — the House of Citizens never acts alone, every constitutional change requires the double majority of Section 4.3, and the transparency platform of Section 5.3 exists precisely to keep the wider public inside the loop.

3. Formalizing the Governance Problem

Definition 1 (the consensus-generation problem). Let Θ be the space of civilizational value preferences; the CPPA algorithm parameterizes preferences into a weight vector $\Theta(t)$, which generates the consensus coefficients α_τ across domains; the α convention of the companion papers — a weight vector, W -weighted and normalized to unit mean — is inherited unchanged. The consensus-generation problem is to design a mapping mechanism G from society's plural preferences to $\Theta(t)$ such that (i) legitimacy: the outputs of G command the recognition of those affected; (ii) epistemic quality: the outputs reflect the best available knowledge about negentropy creation; (iii) capture resistance: G is resistant to lobbying and manipulation by interest groups; (iv) evolutionary openness: G does not systematically suppress non-mainstream exploration.

Proposition 1 (structural incompatibility of single-rule consensus generation). Call a mechanism single-rule if it maps preference profiles over Θ to outcomes through one aggregator applied to one jurisdiction. Consider four requirements: (L) legitimacy — if an overwhelming majority of the affected ranks weight vector θ above θ' , the mechanism does not select θ' over θ ; (E) epistemic quality — on technically verifiable sub-domains, the outcome tracks the expert-verified ordering; (C) capture resistance — no coordinated minority below a fixed share q^- can determine the outcome against an opposing majority; (O) evolutionary openness — every exploration domain retains a weight floor $\omega_{\min} > 0$ independent of the preference profile. Then no single-rule mechanism satisfies L, E, C, and O simultaneously on an unrestricted domain of preference profiles. Proof. Consider a profile in which a share $1-\varepsilon$ of citizens ranks frontier exploration last — short-run material needs first — while a minority of share ε ranks it first. Requirement O demands a positive weight floor for exploration on every profile; requirement L demands that the outcome reflect the majority's ranking on this profile. Assigning the floor over the majority's objection violates L; removing it violates O. A single aggregator must resolve the same conflict identically in every jurisdiction it governs, so the contradiction is generic rather than exceptional. Expert-rule variants satisfy E but fail L and C on the mirrored profile in which the expert class itself is the coordinated minority; algorithmic-rule variants fail L whenever the encoded objective diverges from expressed preference, which on an unrestricted domain occurs with positive measure. $\square$ One clarification preempts a likely misreading. Requirement O does not conflict with Rawls's (1971) difference principle: the weight floor $\omega_{\min}$ is a budget floor on exploration domains — a constraint on the allocation of issuance weights — not an equalization of outcomes; it governs the opportunity structure from which the future position of the least advantaged is generated, not the distribution of primary goods at a point in time.

Remark (what this result is, and what it is not). Proposition 1 is deliberately not presented as an Arrow-type theorem, and the differences matter for how much weight it can bear. Arrow's

impossibility theorem and the Gibbard–Satterthwaite theorem (Arrow, 1963; Gibbard, 1973; Satterthwaite, 1975) establish incompatibilities among minimal, individually attractive desiderata over unrestricted domains of ordinal preferences, and their proofs are combinatorially deep. Proposition 1 is a structural observation by comparison: requirements O and L conflict directly on a transparent profile, because O imposes an outcome constraint (a weight floor) that is defined independently of preferences, while L requires the outcome to track them. No fixed-point argument is needed, and none is offered. The proposition’s contribution is not depth but specification: it states exactly which requirement must be removed from the aggregator — O — and it thereby converts the design move of Section 5.1 (reassigning the floor to a budget rule outside the preference aggregation) from an intuition into the unique available escape route. A weaker but correctly labeled result is more useful to the design than a stronger label the proof cannot carry; we state it at its true strength.

Design corollary (domain decomposition). The architecture of Sections 4–5 escapes the incompatibility by assigning the requirements to different bodies rather than to a stronger aggregator: L and E are separated between the two chambers of the bicameral parliament, each optimized for one requirement (Section 4); O is removed from the aggregator altogether and reassigned to the MPEF as a budget rule rather than a preference rule (Section 5.1); C is distributed across the firewalls whose capture costs Proposition 2 makes superadditive; and a fallback mechanism reserves final adjudication to dispersed citizens (Section 5.3). The proposition thereby explains, rather than merely asserts, why the architecture must be composite.

Definition 2 (the integrity-guardianship problem). The second paper (Cheng and Cheng, 2026b) reduced macrostability to the maintenance of four integrity conditions, restated here for self-containment: the measurement layer unfalsified (no fabricated negentropy records enter issuance), the calibration layer uncaptured (the consensus coefficients are not steered by organized interests), the issuance layer uncorrupted (no units exist that do not correspond to a ratified measurement), and the extinction layer unfailing (the levy and clearing rules execute as specified). The guardianship problem is to assign institutional defenses to each condition such that the expected cost of breaching any layer exceeds the expected gain, and the failure modes of the defenses across layers are mutually uncorrelated — the independence requirement is essential: if the defenses share a common failure mode, the four conditions collapse in fact into one, and the guardianship matrix degenerates into a formality.

$$p_i(e_i) = 1 - \exp(-e_i/c_i) \tag{1}$$

Proposition 2 (conjunctive superadditivity of capture costs). Model global capture as a sequential problem over the four bodies $i = 1, \dots, 4$ (GVP, AEAC, CRN, transparency platform) with isolated capture-cost scales c_i : an effort e_i directed at body i secures control with probability (1), concave in effort. Suppose (i) the generation mechanisms are uncorrelated, so capture effort is body-specific and cannot be reused across bodies; (ii) each attempt emits an exposure signal on the transparency platform with per-period hazard $q_0 > 0$, and an exposed attempt fails with penalty F ; (iii) observed attempts raise the capture-cost scales of all not-yet-captured bodies by a vigilance factor $(1 + r)$, $r > 0$, as audits and public scrutiny tighten. Then the minimal expected cost of global capture satisfies

$$C_{glob} \geq \sum_i c_i + r \cdot \sum_{\{i \geq 2\}} c_i + q_0 \cdot T^- \cdot F \tag{2}$$

strictly exceeding the sum of the isolated costs, where T^- is the minimal duration of a sequential campaign. Proof. By (i), the capturer must pay each body’s cost separately. Order the campaign by stage: by (ii)–(iii), after each stage the remaining price list inflates by the factor $(1 + r)$, contributing the second term of (2), and the campaign carries an expected exposure loss accumulating to $q_0 \cdot T^- \cdot F$, the third term. Summing stages gives the bound. □

The theorem’s content is conditional in a way that matters politically: the superadditivity gap is proportional to the vigilance factor r , and r collapses toward zero precisely when the generation mechanisms converge — a shared appointment pipeline, a shared funding source, a shared technical

stack. The empirical regularity of constitutional history now has its formal counterpart: no case of the wholesale capture of a separation-of-powers system has stormed the conjunctive defenses frontally; each began by eroding the independence of the branches, making their generation mechanisms converge. Constitutions fall not when the walls are stormed but when the moats are connected. Guarding the diversity of generation mechanisms is therefore not one precaution among many; it is the parameter on which the entire guardianship matrix stands, and the common-mode risk of a shared technical stack (Section 10) is the matrix’s most dangerous single point.

Parameter status and sensitivity of Proposition 2. The exposure hazard q_0 , the penalty F , and the vigilance factor r enter bound (2) as design parameters whose empirical calibration belongs to the pilot audits of Section 4.4 rather than to this paper; what the proposition itself guarantees is qualitative robustness, since the strict inequality $C_{glob} > \Sigma_{ic_i}$ holds for every $q_0 > 0$ and $r > 0$, so superadditivity is not a knife-edge property of any parameter point. What varies with the parameters is the materiality of the deterrence margin. Table 1 reports the normalized margin $C_{glob}/\Sigma_{ic_i} - 1 = r \cdot (\Sigma_{\{i \geq 2\}} c_i / \Sigma_{ic_i}) + q_0 \cdot (T^{\top} F / \Sigma_{ic_i})$ over the policy-relevant ranges $q_0 \in [0.05, 0.50]$ and $r \in [0.10, 0.50]$, under equal body costs (so that $\Sigma_{\{i \geq 2\}} c_i / \Sigma_{ic_i} = 0.75$) and with the penalty–duration product normalized to $T^{\top} F = \Sigma_{ic_i}$. The margin runs from 12.5 percent at the pessimistic corner to 87.5 percent at the optimistic corner and is positive everywhere. Two structural readings matter more than any single cell. First, if the transparency platform goes blind ($q_0 \rightarrow 0$), the third term of (2) vanishes and deterrence rests on vigilance alone — which is why the platform’s tamper-evidence is constitutionally load-bearing (Section 5.3). Second, if the appointment pipelines converge ($r \rightarrow 0$), the margin rests on exposure alone — which is why the AEAC audits the diversity of the four bodies’ appointment pipelines annually and reports convergence as a constitutional alarm (Section 5.2). The bound is weakest exactly where the architecture’s own monitoring points, and the pilot calibration of q_0 is a measurement task with a known estimand: the empirical exposure rate of the injected audit irregularities in design H3 of Section 9. Proposition 2 is accordingly a qualitative mechanism result, not a quantitative formula for real-world capture costs: q_0 , r , and F enter as design targets to be calibrated in pilots, not as known constants of political reality. One boundary of the model should be named explicitly: capture is modeled as a sequential campaign, and the superadditivity margin rests on that structure. A simultaneous, centrally coordinated strike on all four bodies at once — a single syndicate acting as one principal — is outside the model; the defense against that limiting case rests on the exposure hazard and on the failure-independence design goal of Section 7, not on bound (2).

Table 1. Sensitivity of the superadditivity margin $C_{glob}/\Sigma_{ic_i} - 1$ (percent) to the exposure hazard q_0 and the vigilance factor r (equal body costs; $T^{\top} F = \Sigma_{ic_i}$).

$q_0 \setminus r$	$r = 0.10$	$r = 0.30$	$r = 0.50$
$q_0 = 0.05$	12.5	27.5	42.5
$q_0 = 0.10$	17.5	32.5	47.5
$q_0 = 0.25$	32.5	47.5	62.5
$q_0 = 0.50$	57.5	72.5	87.5

4. The Global Value Parliament: Bicameral Algorithmic Democracy

4.1. The House of Citizens: Sortition as the Engine of Legitimacy

The House of Citizens has 500 seats. Its members are generated by three-dimensional stratified sortition: seats are allocated by continental population share, subdivided within continents by major cultural spheres, and further adjusted for the representation of countries at different development levels; members are drawn at random by strata from a database of eligible adult citizens worldwide, serve three-year terms, and may not be reappointed; the drawing is broadcast live on the blockchain under the supervision of international observer missions. The chamber’s core function is to reflect the plurality of popular preferences over value dimensions, ensuring that $\Theta(t)$ does not drift from the basic value demands of the public. Short terms without reappointment cut the incentive to

professionalize: members are not career politicians, their decisions do not serve re-election, and the expected return to lobbying falls accordingly — the first structural safeguard of capture resistance (iii).

In Mansbridge's terms (2003, 2009) (Section 2.5), the chamber is a pure selection-model institution: alignment is engineered ex ante by stratified sampling rather than maintained ex post by electoral sanction. This is a deliberate trade, not an oversight. Sanction-based accountability under sortition would require re-eligibility, and re-eligibility would rebuild the career structure whose abolition is the chamber's *raison d'être*; the architecture instead accepts the selection model's known cost — no electoral correction of drift between the chamber's deliberated judgments and the public's raw preferences — and compensates for it around the chamber, through the deliberative-poll re-review mechanism of Section 5.3.

Bootstrapping the registry. The circularity objection — the database that legitimizes global sortition presupposes a global authority to build it — is answered by staged enrollment rather than denied. The database is seeded from national electoral and civil registries under treaty, with a published coverage ratio $q_{cov}(t)$; each parameter class j carries an activation threshold q_j , and GVP authority over class j acquires legitimacy only when $q_{cov} \geq q_j$ — the core monetary parameters (k_s, δ) requiring near-universal coverage, peripheral classes proportionally less. Legitimacy thereby scales with enrollment instead of awaiting a completed world state: the GVP begins life as a federation of participating publics and becomes global by accretion. The residual circularity of the founding act itself is acknowledged in Section 10.

4.2. *The House of Experts: Peer Review as the Engine of Epistemic Quality*

The House of Experts has 300 seats. Its members are nominated by authoritative learned bodies worldwide, by major schools of philosophy and ethics, and by representative cultural organizations, and are appointed after rigorous peer review and academic-integrity checks; terms are six years, renewable once. Membership covers five knowledge domains corresponding one-to-one to the five value dimensions of the CPPA: the natural sciences support the quantification of strategic match (S) and externalities (E^+/E^-); the social sciences analyze the social impact of weight adjustments; philosophy and ethics conduct the deliberation of intergenerational justice and fairness; culture and anthropology guarantee plural value perspectives; engineering and technology assess the feasibility of CRN and CPPA implementation. (For concreteness: in the first paper's brain-science demonstration, the mapping over the dimensions S, E^+ , C, and U yielded the consensus $\alpha = 0.95$; (Cheng and Cheng, 2026d, Section 11.2) — restated: research activity scored near-maximally on strategic match and positive externality, giving it the highest consensus weight in the calibration example.) The chamber's core duty is to supply knowledge support, cross-cultural ethical reflection, and long-horizon judgment for the adjustment of value weights, guarding against populist short-sightedness — the institutional carrier of epistemic quality (ii).

4.3. *The Decision Chain and the Double Majority*

The jurisdiction of the double-majority rule expands with the framework of the second paper: the α weight vector $\Theta(t)$, the macrostability coefficient k_s , the intrinsic metabolism rate δ , the knowledge-genealogy attribution weights, the depreciation schedules of durables, and the inheritance metabolism rate $\delta_{inherit}$ — every parameter entering the issuance and extinction identities, without exception, passes through the same procedure. Revision of core parameters follows a four-stage prudent procedure. (1) Initiation: proposals may originate from the specialized committees of either house, from citizen initiatives with a million signatures, or from the AEAC, and must be accompanied by CRN accounting data and ABM simulation results. (2) Preliminary review and independent feasibility assessment: data authenticity, logical coherence, and differentiated impacts are assessed. (3) Multiple rounds of public hearings and global debate: all materials are published in real time on the transparency platform and notarized on the chain — an institutionalized attempt at Habermas's ideal speech situation. (4) Double majority: a proposal takes effect only upon

simple majorities in both houses separately; amendments to fundamental rules — adding a value dimension, altering the algorithmic structure — require a two-thirds absolute majority in each house. The double majority is the essence of the mixed regime: every value adjustment must win the endorsement of the popular will and of reason at once, guarding institutionally against majoritarian tyranny and technocratic rule alike. This division of jurisdiction also closes the incommensurability loop left open in the first paper: constitutive value enters the system through the α -floor under democratic guardianship, not through the scalar W_τ , so that plural values are preserved at the calibration layer exactly where the first paper assigned them (Cheng and Cheng, 2026d, Section 2.4).

The resulting allocation of constitutional labor can be stated exactly. Table 2 records, for each layer of the monetary constitution, who writes the operative rules, who holds the power to change them, and who audits their execution. The matrix is the closed interface between the measurement architecture of the first paper (Cheng and Cheng, 2026d), the issuance–extinction identities of the second (Cheng and Cheng, 2026b), and the guardianship assignments of Section 7 below: no layer is self-authorizing, and every change power is paired with an audit power held by a different body.

Table 2. The authority matrix of the four layers: who writes, who changes, who audits.

Layer	Writes / proposes	Decides / changes	Audits
Measurement (CRN)	CRN protocol community publishes the measurement conventions under the L1–L2 rules (Paper I, Section 11)	Convention revisions require a GVP double majority; indicator-system changes follow the 15-percent review trigger	Cross-source validation, interpretive-community review, and AEAC process audit
Calibration (CPPA + GVP)	CPPA computes the $\Theta(t)$ weight vector and α proposals through the three-stage procedure	GVP double majority; structural amendments (new value dimensions, changes to the algorithmic structure) require two-thirds of both houses	AEAC quarterly disparate-impact audits; transparency-platform public record
Issuance (platform)	None — issuance is the rule-bound execution of equation (1) of Paper II	No discretionary change power exists at this layer	AEAC issuance audit; ledger reconciliation against ratified measurements
Extinction (platform + AEAC)	The platform clearing pool executes cancellation and δ sit within the double-on audited depreciation schedules	Depreciation-schedule rules and δ sit within the double-majority jurisdiction (Section 4.3)	AEAC extinction-layer audit (duty six); composition monitoring (duty seven)
Constitutional parameters (k_s band, δ , firewall clauses)	Joint specialized committees of the two houses	Two-thirds absolute majority in each house, subject to the activation thresholds q_j of Section 4.1	AEAC constitutional alarm; the popular-oversight axis (Section 5.3)
State share (τ at clearing node)	Constitutional convention drafted by joint committees	GVP double majority; never set by the beneficiary state	Transparency platform public ledger; AEAC fiscal audit

An illustrative scenario displays the operating logic of the decision chain. Against the background of a global energy crisis, specialized committees of the two houses jointly propose raising the strategic-match weight of controlled-fusion research from 0.25 to 0.35. In the hearings, representatives of developing countries demand a supporting technology-access fund, and environmental organizations demand a new technology-safety sub-indicator in the negative-externality dimension. The amended proposal, incorporating both, passes with 58% in the House of

Citizens and 62% in the House of Experts. The scenario — a stylized case constructed for this paper, not a report from the program's monograph — shows public pain, professional reflection, and demands of justice converted through procedure into precise adjustments of incentive parameters; it displays the mechanism, not its empirical validity.

4.4. *Nested Pilot Governance: From Planetary Design to Research Object*

The scale objection of Section 2.4 — Dahl's skepticism that a planetary parliament can be democratized at all — cannot be answered by refining the planetary design alone. The architecture's answer is nestability: before any global activation, the GVP design admits single-domain, sub-global instances that inherit the constitutional parameters unchanged and exercise only peripheral jurisdiction. Three minimal configurations define the pilot ladder, ordered by increasing scope.

City-level pilot. A medium-sized municipality issues NVC-denominated credits against two measured activity classes — care work and ecological restoration — governed by a fifty-seat citizens' chamber drawn by two-dimensional stratified sortition (district $\times$ age-and-education stratum) from the municipal registry and a fifteen-seat expert chamber nominated by local universities and professional bodies, with the double-majority rule restricted to the domain α -weights of the two pilot classes. Constitutional parameters — the k_s band, the metabolism rate δ , the firewall clauses — are inherited, not revisable at pilot scale; pilot chambers may propose α -deviations only within a corridor of ± 20 percent around inherited values, larger deviations escalating to the next nesting level. The legitimacy thresholds of Section 4.1 scale with jurisdiction coverage: activation of a parameter class requires the coverage ratio q_{cov} over the pilot registry rather than over a world registry, so that a pilot parliament is legitimate on its own territory from its first day.

Open-source ecosystem and wetland-restoration pilots. An open-source ecosystem pilot governs issuance against the measured I_τ of verified contributions, with a thirty-seat sortition chamber drawn from the contributor registry and a maintainer expert panel; the CRN is reduced to its L1–L2 layers over a single indicator system, which keeps measurement cost proportionate to the pilot's budget. A wetland-restoration pilot narrows the jurisdiction to one activity class and one α , and exercises the extinction layer concretely: the restored asset's corresponding money is extinguished on the amortization schedule of the companion paper's extinction paths (Cheng and Cheng, 2026b, Section 3). Each pilot terminates in a pre-registered evaluation feeding hypotheses H1–H3 of Section 9, so that the pilot ladder and the verification agenda are one instrument.

The pilot ladder converts the GVP from a founding act into a research object. Each instance generates audit data that calibrate the exposure hazard q_0 and the penalty F of Proposition 2, observed capture attempts that stress the firewalls of Section 5, and participation data that discipline the sortition design of Section 4.1; by the time a planetary activation could be contemplated, every constitutional parameter would carry an empirical estimate rather than a design value. This is also the honest answer to the charge of utopianism. The program does not ask to be adopted at civilizational scale; it asks to be tested at scales where failure is survivable and measurement is real, and it specifies in advance what its own failure would look like.

4.5. *Institutional Interfaces: Docking the Measurement Layer with GEP Accounting and Time Banks*

The pilot ladder of Section 4.4 specifies where the architecture can be tested; this subsection specifies what it can dock onto. Two running national institutional programs — China's Gross Ecosystem Product (GEP) accounting and its pension time banks — belong to the same measurement family as the value equation's factors, and each exhibits a methodological gap that the measurement layer of this program (Cheng and Cheng, 2026d) is built to fill. A companion policy paper develops the docking path in full (Cheng and Cheng, 2026e); the constitutional relevance here is that the interfaces convert the governance architecture's most demanding components — auditable measurement, versioned conventions, counterfactual baselines — from design requirements into practices with living institutional counterparts.

GEP accounting has crossed from methodology debate into national standards: a 2021 central directive on realizing the value of ecological products (General Office of the CPC Central Committee and General Office of the State Council, 2021), a 2022 trial accounting standard issued by the National Development and Reform Commission and the National Bureau of Statistics (National Development and Reform Commission and National Bureau of Statistics of China, 2022), a first round of national pilot sites in 2024 including Beijing's Yanqing district and Huangshan, Anhui (National Development and Reform Commission, 2025), and the UN Statistical Commission's 2021 adoption of SEEA-EA as the international ecosystem-accounting standard with which the Chinese trial standard explicitly aligns (UN Statistical Commission, 2021). Two structural gaps matter for this paper. First, GEP keeps a credit side without a debit side: it sums the services ecosystems provide without deducting the entropy costs of producing them, so high-emission and near-natural restoration record identically. The value equation's E_{τ} factor (equation (2) of the companion) has the debit built in, and an entropy-debt ledger added to GEP project accounting — technically feasible with existing LCA factor databases — would for the first time rank ecological projects by order created per unit of entropy debt, letting fiscal compensation price net rather than gross order. Second, GEP registers increments but not maintenance: a forest protected from degradation and an unattended forest are indistinguishable in the accounts, so the ranger's labor appears as cost, never as value. The counterfactual clause of the measurement layer — prevented divergence from a no-intervention trajectory registers as order creation — supplies the missing baseline, estimable from standard control-plot networks at low cost. The distributive consequence is immediate: the labor of rangers, river chiefs, and grassland wardens becomes accountable, compensable value creation.

Pension time banks are the living institution of the T_{τ} factor. Shanghai rolled its time-bank program out citywide in 2021 (Shanghai Civil Affairs Bureau, 2021); Beijing made time credits interoperable across the city under its 2022 implementation plan (Beijing Municipal Civil Affairs Bureau et al., 2022); and a 2021 Peking University research report counted more than 240 operating organizations across 31 provinces (Peking University Institute of Population Research, 2021). The system's pain points are precisely measurement pain points: inconsistent timekeeping standards across sites, difficult inter-institutional redemption, and no quality dimension (an hour of professional nursing equals an hour of conversation). Two components of the measurement layer transfer directly: outcome-keyed service metering — anchoring credit to the maintained state trajectory of the cared-for person against the no-care counterfactual, rather than to raw hours — and the issuance–extinction accounting architecture, which gives time credits a closed issuance–circulation–cancellation loop that prevents indefinite stock accumulation from diluting the unit.

The docking sequence matters more than any single interface, and it is deliberately staged so that each step has independent value and survivable failure: (1) measurement methodology inside the GEP pilots and time-bank networks (entropy-debt ledger; counterfactual maintenance accounts; outcome-keyed metering), years 1–3; (2) accounting instruments — the issuance–extinction architecture and the quarterly disparate-impact audit pipeline — inside ecological compensation and time-credit redemption networks, years 3–5; (3) quasi-monetary experiments — regional ecological certificates, cross-domain time credits — only thereafter, carrying the full measurement–calibration–issuance–extinction cycle under the bicameral governance of Sections 4.1–4.3. Each step generates the data the program currently lacks: step 1 yields measured distributions and measurement errors of E_{τ} and T_{τ} ; step 2 yields audit detection rates and the capture-hazard parameters of Proposition 2 (q_0, r, F) from real gaming behavior; step 3 alone touches money. If the first two steps fail, the loss is confined to a methodology experiment.

Three honesty boundaries apply. First, the policy version of Goodhart's law: once GEP scores or time credits bind fiscal compensation or cadre evaluation, the indicators will be gamed — the defenses are those already declared (outcome keying, artifact anchoring, recomputable audit) plus versioned rotation of indicators that accept their own decay. Second, the entropy-debt ledger will shrink measured GEP totals; that is its entire purpose and also its political resistance, so the honest strategy is to run it first in research accounts unbound to evaluation. Third, the docking path carries

only the measurement layer: the monetary and constitutional claims of this trilogy do not ride on it and require their own tests. What the interfaces establish is narrower and sufficient for this section: the governance architecture's measurement demands are not utopian — they are refinements of accounting practices that two national programs already run.

5. Three Constitutional Firewalls and System Closure

5.1. *The Value-Inclusion Axis: The Minority Protection and Exploration Fund (MPEF)*

Majority rule suppresses minorities by nature, while civilizational breakthroughs usually originate from variation at the margins — the early internet and the concept of quantum computing were, at birth, the murmurings of minorities. The MPEF holds a fixed share of the GPF's annual budget and funds, with funding decisions exempt from majority rule, explorations that the majority does not understand but that hold disruptive potential: non-mainstream scientific theories, avant-garde art, the preservation of endangered cultures. Its theoretical function corresponds to requirement (iv) of Definition 1: an institutionalized umbrella of variation safeguarding the openness and diversity of civilizational evolution — and, in light of the Remark following Proposition 1, it is the designated carrier of requirement O, which is precisely why it is constituted as a budget rule exempt from preference aggregation rather than as a constraint on the aggregator. Failure mode: the fund itself may be captured by small circles and degenerate into a funding channel for particular academic schools — the defenses are the AEAC's audit of its money flows (duty four of Section 5.2) and the platform's full publicity (Section 5.3).

5.2. *The Process-Oversight Axis: The Audit and Emergency Adjustment Committee (AEAC)*

When an algorithm is embedded in the core of value distribution, auditing the algorithm itself becomes the new key to the balance of powers. The need is not hypothetical: the scoring and ranking systems documented by Pasquale (2015) show how opaque algorithms can quietly reallocate money, opportunity, and information long before any constitutional question is posed. The AEAC is fully independent of the GVP; its members — jurists, philosophers, data scientists, ethicists, and CRN technical experts — are jointly nominated by the two houses and appointed by a two-thirds absolute majority, serve fixed five-year terms, and cannot be removed except for statutory cause. Its core duties: (1) continuing algorithmic audit — quarterly disparate-impact analyses of CPPA outputs using interpretability tools to detect whether the algorithm systematically discriminates against particular regions, industries, or groups; (2) an annual Algorithmic Transparency and Impact Report — full disclosure of operating logic, parameter settings, and ethical-risk assessments; (3) emergency braking — upon discovery of a fatal vulnerability or in a global emergency, the power to suspend a specific algorithmic rule by a two-thirds vote, reported to the GVP within 48 hours; (4) oversight of the use of GPF and CTB funds — examining whether investment decisions conform to the high- α orientation and guarding against interest transfer; (5) the power to propose algorithmic improvements. With the introduction of the extinction layer in the second paper, the AEAC bears a sixth duty: (6) extinction-layer audit — auditing the cancellation records of the platform clearing pool and the execution of depreciation schedules, measuring and publishing the extinction gap ΔX (corresponding to failure Mode D of the second paper), so that the death of money receives the same level of audit as its birth. A seventh duty follows from the composition condition of the second paper's Section 4.7 (restated: experienced inflation stays within bounds only if the structure of α -weighted creation tracks the structure of consumption demand): (7) composition monitoring — tracking that alignment, and reporting premium congestion at the upper rail of the dual-price band as an early-warning indicator of experienced price pressure. The AEAC is the ultimate embodiment of algorithmic accountability: technical power is caged by institutions.

Terminating the recursive-oversight regress. The standard objection — who watches the watchmen? — deserves a formal answer, not a shrug, and the architecture supplies one. The naive answer (a committee to audit the auditors) only pushes the regress one level up; it terminates nothing.

The present design instead terminates the regress by loop closure onto a body with unbounded capture cost, in three steps. First, the AEAC's outputs — every audit report, every braking decision, every pipeline-diversity assessment — are public on the transparency platform in real time; the platform itself is overseen not by any committee but by the million-signature initiative and the stratified deliberative-poll re-review rule of Section 5.3, whose "membership" is the dispersed global citizenry. Second, dispersed citizens cannot be captured as a body: there is no appointment pipeline to colonize, no term to wait out, and no committee to staff — the capture-cost model of Proposition 2 has no well-defined target at this level, which is exactly why the regress ends here rather than at a fifth institution. Third, the cost of the loop is borne at the AEAC level in advance: because the auditors know that every report is exposed to citizen re-review with a statistically hardened trigger (the 70% opposition rule of Section 5.3), the exposure hazard q_0 of any colonization attempt against the AEAC is high throughout the attempt, and Proposition 2 then prices slow colonization accordingly — the expected cost grows with the campaign duration T^- at hazard $q_0 F$ per period. The regress therefore terminates not because a final watcher is virtuous, but because the last link is structurally uncapturable and the second-to-last link knows it is watched. We state candidly what this argument assumes: that the transparency platform's records are in fact tamper-proof and in fact legible to citizens — the common-mode risk of a shared technical stack (Section 10) is precisely the point where this termination could fail, which is why the AEAC's annual audit of the four bodies' pipeline diversity (below) is a constitutional alarm and not a routine report.

Failure mode: capture of the auditors themselves. The defenses are fixed terms, the supermajority threshold of joint appointment, and the public answerability of all audit reports; we acknowledge that these raise the cost of capture rather than abolishing it, and that a sufficiently patient capturer with a five-year horizon can in principle colonize the appointment pipeline. The residual defense is the independence of the pipeline's two halves: nomination by the two houses and confirmation by supermajority cannot be corrupted through a single channel without leaving traces on the transparency platform. Proposition 2 converts this worry into a parameter: slow colonization of an appointment pipeline is a low-hazard, long-duration campaign, and its expected cost grows with the exposure time T^- — the public answerability of the audit reports is precisely the hazard q_0 of that campaign. The defense is therefore not any single appointment rule but the coupling between attempts and visibility; the AEAC accordingly audits the diversity of the four bodies' appointment pipelines annually and reports any convergence as a constitutional alarm.

5.3. *The Popular-Oversight Axis: Transparency and Multi-Level Participation*

The GVP's full agendas, voting records and reasons, the CPPA's algorithmic versions and parameters, every money flow of the GPF and CTB, and all AEAC audit reports are stored on the chain, tamper-proof, and queryable worldwide in real time. Four participation channels are built on top: NLP aggregation of public opinion from daily comments; citizen initiatives of a million signatures — meeting diversity requirements of region, age, and occupation — that automatically trigger formal deliberation; stratified random-sample deliberative polls before major decisions, with a re-review rule, operationalized statistically: the mock vote is a stratified deliberative poll of $n \geq 1,600$, drawn by the same three-dimensional stratification as the House of Citizens (a margin of error of ± 2.5 percentage points at 95% confidence), and the re-review duty triggers when the one-sided 95% confidence lower bound on the global opposition share exceeds 70% — so that sampling error alone cannot fire the brake, and an organized minority cannot fire it either; and annually drawn citizen observers who attend committee meetings, submit written inquiries, and hold no vote. The function of this axis is to reserve final adjudication to dispersed global citizens — popular sovereignty is not a one-off authorization but whole-process oversight. Failure mode: transparency is not comprehension; a flood of data can produce a tyranny of transparency, with the public drowning in indigestible information, and NLP aggregation may systematically misread opinion expressed in minority languages. The structured inquiry channel of citizen observers mitigates this, but we acknowledge that the gap between accessibility and intelligibility is a real difficulty the architecture has not solved.

5.4. *The Mesurer Inside the Equation: Recursive Audit Closure*

All three papers of this program share a single closure rule: no participant stands outside the value equation. Measurement and audit personnel are remunerated in NVC through the same issuance identity as everyone else — their E_τ and T_τ are natively logged by the platform, and their I_τ is the counterfactual value of accurate measurement itself. They hold no payment channel outside the system: a fiat-salaried measurement layer would be an island of unmeasured value inside the measured economy, and the one position from which value could be extracted without measurement would become the dominant capture surface — the superadditivity of Proposition 2 applies reflexively to the measurers themselves.

The regress — who measures the measurers — is closed, not infinite. Three anchors terminate it. Quantities terminate in physics: E_τ and T_τ are instrument readings and I_τ is the output of published algorithms on public data; like the 2019 redefinition of the SI, which retired the prototype kilogram in favor of the Planck constant, the design retires the custodian's discretion together with the artifact. Conventions terminate in the double majority: α calibration, eligibility boundaries, and product-class reference distributions are declared, public, and falsifiable through the pilot machinery of Section 4.4. Persons terminate in a closed audit cycle: the measurement foundation measures the economy, the audit committee audits the foundation, sortition-drawn citizen panels and the transparency platform audit the committee, and no node audits its own case; case assignment is randomized, and conflict-of-interest recusal is mandatory.

Measurer incentives are engineered rather than assumed. Ex-post accuracy scoring — blind re-measurement samples, inter-rater reliability, prediction against realization — determines realized issuance; a portion of measurement compensation is deferred and clawed back if the underlying measurements are later falsified; the measurement budget is set by the double majority and capped as a share of total issuance, so the measurers cannot vote themselves a raise; oversight roles rotate by sortition. The honest residue: the regress cannot be driven to zero. The system's claim is not the absence of discretion but the absence of unmeasured, unaudited, undeclared discretion — every residual human judgment is itself a measured, auditable, and revocable activity inside the loop.

5.5. *The State Inside the Equation*

Nothing in this constitution exempts the state. Its fiscal existence inside the equation has exactly three forms. As producer, the state's services — defense, justice, public health, coordination infrastructure — are genuine and often enormous negentropic outputs; they are measured like any other production, and issuance accrues to the actual producers — soldiers, judges, civil servants — as individuals. As consumer, the state draws a constitutionally fixed share τ at the clearing node: when a consumption payment is cancelled, the fraction τ is routed to public accounts before cancellation, and the state completes extinction by consuming it in its operations; the identity $M(t) \equiv k_s \cdot V(t)$ is preserved, with the extinction path merely rerouted through public consumption. As issuer, the state has no existence at all: discretionary fiscal issuance is the canonical historical form of failure Mode C (Cheng and Cheng, 2026b), and this constitution exists to abolish it.

The share τ belongs to the same constitutional class as k_s , δ , and α : it is set by the double majority and cannot be set by its beneficiary. A state that fixed its own share would be a measurer measuring its own case — the self-measurement prohibition of Section 5.4 applies to governments first of all.

The deepest hazard of the entire architecture lives here. The state alone holds legitimate coercion, and its revenue depends on measurements it could in principle force. The defenses are constitutional, not technical: the measurement layer is planetary — no single state appoints, budgets, or dismisses it — its budget is capped, its ledgers are public, and enrollment is nested so that no national jurisdiction hosts a single point of failure. The residue must be stated honestly: no monetary constitution survives a regime willing to shoot the measurers. At that depth the defense is political legitimacy — sortition, double majorities, and transparency — not algorithm.

Why would states accept? Because the constitution offers the best fiscal bargain in their history: a permanent revenue share denominated in a constant-value unit — no inflation cycles, no debt crises, and near-zero collection cost, since the share is routed automatically at the clearing node. What the state surrenders is the printing option, which is precisely what has repeatedly destroyed state finances from within; surrendering it is the cure, not the cost. One measurement difficulty must be recorded rather than concealed: the I_{τ} of state services is the hardest measurement in the system — monopoly supply, an anarchic counterfactual, no reference distribution — and its conventions can only be calibrated through the staged pilots of Section 4.4.

6. Metabolic Governance: The Four Clauses of the Economic Constitution

The governance architecture constrains not only decision procedures but the economic rules themselves. Four constitutional clauses set the metabolic boundaries of the system. The clauses are stated here in self-contained form — each clause's rule, its constitutional function, and its calibration status — with the program's monograph (Cheng and Cheng, 2026a) providing extended background only.

Clause 1 (the intrinsic metabolism rate δ). All idle NVC balances not deposited in the CTB are extinguished annually at a public, transparent, predictable rate δ . Calibration by an overlapping-generations model with multiple saver types places the optimal dynamic band of δ at 0.5%–1%: lower, dead capital accumulates and strategic investment starves; higher, saving is discouraged and long-run capital accumulation suffers. The clause descends from Gesell's (1916) stamped scrip, with the three redesigns stated in the second paper (Cheng and Cheng, 2026b, Section 4.6 — restated: the levy is avoidable through the CTB, constitutionally bounded within its band, and embedded beneath the tracked extinction paths so that it operates only on the residual overhang). What is constitutional here is the constraint on the constraint: adjustments of δ must be proposed by an independent economic-stability committee and approved through the GVP's democratic procedure, so the metabolic pressure is itself subject to the double majority. The constitutional significance of δ is the replacement of the passive, concealed, unavoidable expropriation of the traditional inflation tax by active, transparent, avoidable guidance — a holder can avoid extinction entirely by depositing in the CTB, which accords with the capability approach's emphasis on the space of free choice (Sen, 1999). A computational scan of the joint policy space of δ and a social consumption floor (Appendix C, Section C.4) adds a distributional qualification: within the 0.5%–1% band the wealth-Gini response to δ is near-flat, but at higher values δ is regressive in incidence — poorer households hold their buffers in money while wealthier households hold CTB certificates — so the binding distributional instrument is the social floor, not δ .

Clause 2 (the Civilization Time Bank, CTB). Social surpluses are pooled by the CTB into patient capital and directed to the highest- α long-horizon strategic projects recognized by the GVP. What the depositor receives is not fixed interest — the time price of debt — but a revenue-share certificate of the supported project: the existence and size of the return are linked solely to the project's real negentropy output as accounted by the CRN, and if the project fails, losses are shared proportionally, with a single-event loss ceiling actuarially designed at about 5%. The clause reconstructs the saving-investment contract from a claim on interest into a share in the fruits of civilizational progress.

Clause 3 (financial minimalism). A constitutional rule strictly prohibits the creation and trading of highly leveraged complex derivatives on NVC's own credit — credit default swaps, naked short options, nested securitization. Permitted financial activity is confined to two kinds: simple insurance serving real risk management, and the fractional transfer of equity in real projects assessed by the CRN for negentropy potential. The clause institutionalizes the lesson of the 2008 crisis: the first function of finance is to serve value creation, not to proliferate itself. (It is the governance-layer counterpart of the second paper's proliferation ban — its Seal 2 — which bars credit multiplication of the unit itself; Clause 3 bars the leveraged claim-pyramids on top of permitted instruments.)

Clause 4 (knowledge-genealogy attribution and intergenerational inheritance). The persistent issuance stream designed for major theoretical discoverers in Section 7.4 of the second paper requires

distribution rules of constitutional rank, which this clause supplies in two parts. First, the attribution-weight rule: subsequent creation events applying a given theory allocate a proportion of their W_{τ} issuance shares to the prior discoverers along the citation genealogy; the weight vector is calibrated through the CPPA’s three-stage procedure and audited by the AEAC, and attribution does not use raw citation counts directly — the citation network is deduplicated and de-looped on the knowledge graph to frustrate manipulation by citation cartels. Second, the intergenerational decay rule of inheritance: the discoverer may designate a successor or a public fund to receive the issuance stream, but the inherited share decays from generation to generation at the inheritance metabolism rate $\delta_{inherit}$ and eventually reverts to the GPF. The rationale of decay is the same as that of δ : perpetual attribution without decay would recreate a hereditary rentier class in the domain of knowledge — precisely the capital sedimentation this system exists to dismantle. The decay structure gives the discoverer a full return during the theory’s active period of creation, while letting great knowledge settle ultimately into the common asset of all civilization. The discoverer’s immortal contribution to civilization is recorded permanently as reputational capital and usage trajectory; the monetary return follows the same metabolic logic as all value. A household-level implementation of this decay rule in the agent-based model (Appendix C, Section C.3) finds that, with the reverted share recycled through the GPF, third-generation wealth concentration falls monotonically in $\delta_{inherit}$: the wealth Gini after the third generation is 0.308 at $\delta_{inherit} = 0.3$ and 0.285 at $\delta_{inherit} = 0.5$, against 0.366 under a no-decay control.

7. The Guardianship Matrix of the Four Integrity Conditions

The institutional designs of Sections 4–6 can be recataloged along the four integrity conditions of the second paper — the *raison d’être* of every institution must ultimately land on which integrity condition it guards and which failure it repels. The guardianship matrix is as follows.

Table 3. The guardianship matrix of the four integrity conditions.

Integrity condition	Corresponding failure (paper 2)	First line of defense	Second line of defense	Backstop
Measurement unfalsified	Mode A: accounting falsification	CRN multi-node verification with zero-knowledge proofs	AEAC quarterly disparate-impact audits	Real-time public data on the transparency platform
Calibration uncaptured	Mode B: consensus distortion	Bicameral double majority and the three-stage procedure	MPEF minority protection (against spillover of majoritarian tyranny)	Million-signature initiatives and the 70% opposition re-review rule
Issuance uncorrupted	Mode C: issuance corruption	Built-in issuance rule with no discretionary interface	Full issuance records notarized on the chain	AEAC emergency braking
Extinction unfailing	Mode D: extinction failure	Consumption clearing sharing blockchain notarization with the CRN	Depreciation schedules approved by the GVP and audited by the AEAC (duty six)	δ basal metabolism absorbing the residual

Three design principles of the matrix deserve explicit statement. First, layered redundancy: each integrity condition has three levels of defense, and the fall of a single point does not collapse the integrity. Second, failure independence — the key requirement of Definition 2: the failure modes

across the rows are mutually uncorrelated, since sortition, peer review, technical distributed verification, and popular oversight rest on different sources of power and different paths of failure; this is the structural basis of the superadditivity of Proposition 2. Third, auditability throughout: every cell of the matrix corresponds to an auditable operational object — records, logs, parameters, reports — and no cell rests on trust in the character of decision-makers, in isomorphism with the second paper's principle of replacing commitment credibility with recomputability.

The matrix also reveals the honest boundary of the architecture: the extinction layer's defensive chain is the shortest, with residual channels beyond platform coverage absorbed only approximately by δ ; and the calibration layer has the most lines of defense, yet every one of them depends on the conscientiousness of procedural participants — institutional design can raise the cost of capture, but cannot institutionalize taking the procedure seriously. We write this into the limitations rather than covering it with rhetoric.

8. The Survival and Elevation of Civilization: The Teleological Loop of the Governance Architecture

The mechanism analyses so far — the value ontology of paper 1, the macrostability of paper 2, and the guardianship of integrity in Sections 3–7 of this paper — must ultimately answer a teleological question: what is this system for? Our answer is that the NVC governance architecture is not engineering for efficiency's sake. Its ultimate orientation is the survival and elevation of civilization: enabling humanity to know as much cosmic truth as possible, to utilize it, and ultimately to master it — achieving the sublimation of humanity itself and the upgrading of civilization. This directional goal is established by the civilization utility (CU) theory of the first paper (Cheng and Cheng, 2026d); the task of this section is to specify how the governance architecture serves as the navigation mechanism toward that goal, along four dimensions. (The extended philosophical treatment of this section's claims appears in the program's monograph [3, chs. 8–10]; every claim needed here is stated below in self-contained form.) The account that follows is a normative expectation, not an empirical prediction: it states what the architecture is for, not what it is guaranteed to produce.

8.1. The Survival Dimension: The Negentropy Anchor as the Metabolic Health of Civilization

Civilization is, in thermodynamic terms, a dissipative structure that must continuously draw negentropy from its environment in order to persist (Prigogine and Stengers, 1984); the macroscopic form of civilizational extinction risk is nothing but the interruption or mismatch of negentropy metabolism. The systemic risks of scarcity-anchored monetary systems — financial crises, ecological overdraft, the misallocation of innovation incentives — are in essence complications of the decoupling of monetary signals from negentropy metabolism. NVC anchors monetary signals directly to negentropy metabolism (the stock identity of paper 2) and guards the four-layer integrity of that anchor through the governance architecture (the matrix of Section 7): for the first time, civilization possesses a bookkeeping and incentive system isomorphic with its own conditions of survival — the health indicators of the economy are the metabolic indicators of civilization. Survival thereby changes from passively enduring the fluctuations of markets and ecosystems to active maintenance that is accountable, auditable, and correctable.

8.2. The Individual Dimension: Three Paradigm Shifts in the Mode of Existence

Institutional incentives ultimately shape the human being itself. The program's monograph (Cheng and Cheng, 2026a, ch. 8) argues that NVC's incentive structure drives three transformations of the individual mode of existence; we restate them here and mark their institutional counterparts. The three shifts below are directional expectations under the architecture's incentive structure, not guaranteed stages of social evolution.

First, from the having mode to the creating mode. Fromm (1976) distinguished two modes of existence: the having mode takes control and possession as the standard of success, while the creating mode locates abundance in the unfolding of one's potential and the meaning of one's contribution. Scarcity-anchored money reinforces the former through its scarcity-possession logic; NVC income derives solely from creation (proof-of-contribution issuance), and money turns from the object of pursuit into the natural by-product and record of creative activity.

Second, from the scarcity mindset to the abundance mindset. Scarcity imprints a zero-sum pattern of thought, and experimental research on the scarcity mindset has documented its suppression of cooperative willingness and creative thinking (Mullainathan and Shafir, 2013). The non-rivalry of negentropy creation — knowledge created by one person does not diminish anyone else's stock but raises the common level through sharing — makes the positive-sum mindset a rationality congruent with social reality rather than a moral exhortation.

Third, from the atomized individual to the networked node. The accumulation of every kind of five-dimensional capital strengthens the individual's connections to social networks (Granovetter's (1985) embeddedness); the knowledge-genealogy attribution mechanism (Clause 4, Section 6) keeps amplifying and recording the individual's creative value within the network. When value is self-confirmed in creative activity and recognition arises naturally in network interaction, the psychological basis of conspicuous consumption dissolves — laying the deepest psychological and cultural foundation for the withdrawal of money from center stage.

The transition toward this mode of existence is modeled explicitly in Appendix C (Section C.2): a retraining protocol that migrates 2% of the low-value workforce per year into the care and ecology domains reduces the low-value share from about 21% to 3.1% over a century, lowers the income Gini from 0.352 to 0.301, and — jointly with the social floor — holds the below-subsistence share to 0.1%. The institutional point is that elevation is not left to exhortation: the domains whose tracked contribution most exceeds their market weight are precisely the absorbers of the transitioning workforce.

8.3. *The Civilizational Dimension: Trajectories of Four Morphological Leaps*

The aggregation of individual elevation emerges, at the level of civilizational morphology, as four leaps (developed at length in (Cheng and Cheng, 2026a, chs. 8 and 10); the content is stated here in full). The first is from a civilization of competition to a civilization of creation: the structural configuration of α — low coefficients for stock rent-seeking, high coefficients for incremental creation — redirects social talent from rent extraction to innovation and building, and the root of social conflict shifts from disputes over the distribution of the stock to contests among ideas of creation; the open-source movement is the contemporary microcosm of this ethos. The second is from a civilization of growth to a civilization of flourishing: the core development indicator changes from the linear expansion of GDP to the total stock of five-dimensional capital and the net global negentropy flow, answering the twin challenges of the Easterlin paradox and ecological limits. The third is from rivalry to synergy: international competition transforms into races of civilizational productivity and of α -consensus attractiveness, with contribution tournaments of GPF flagship projects replacing zero-sum trade games. The fourth is from the utilitarian to the perennial: the δ metabolism and the clauses of intergenerational justice internalize the accounting of posterity and ecology into every present economic decision, extending the time horizon of development from electoral cycles and fiscal quarters to the civilizational scale. They are offered as the morphology the design is oriented toward, not as a forecast of civilizational development.

8.4. *Knowing, Utilizing, Mastering Cosmic Truth: The Three Stages of CU Navigation and the Withdrawal of Money*

The civilization utility theory of the first paper established the directional goal of civilization as the progressive mastery of cosmic truth; the task here is to specify how the governance architecture navigates toward that goal. The α weight vector $\Theta(t)$ is, in substance, civilization's current consensual

judgment of what brings us closer to cosmic truth; the CPPA’s three-stage procedure and the bicameral parliament (Section 4) are the prudent generation procedure of that judgment; genealogy attribution and the high- α configuration of basic science (Section 7.4 of paper 2) make knowing truth the most highly rewarded activity; the GPF flagship projects and CTB patient capital (Section 6) secure trans-generational resources for utilizing truth; and the extinction and metabolism mechanisms ensure that resources never settle in possession but remain in perpetual creative flow. The three-stage navigation thus closes: knowing truth (high I_{τ} incentivizing basic science) $\rightarrow$ utilizing truth (high α incentivizing application) $\rightarrow$ mastering truth (CU feedback driving the iterative calibration of $\Theta(t)$).

Proposition 3 (the teleological loop, stated informally). The legitimacy of every layer of the governance architecture traces back ultimately to the maximization of civilization utility CU: the final criterion by which a governance arrangement is judged is its long-run contribution to civilization’s total negentropy creation and to its proximity to truth. The proposition is a normative commitment, not an empirical prediction — it prescribes the direction of the architecture’s self-revision: every clause is revisable under CU review, including the architecture itself. (This is the fixed telos whose tension with the experimentalist recursion is flagged in Section 2.5; the tension is real, and we do not resolve it by definition — we register it.)

As this navigation operates over the long run, the destiny of money itself becomes clear. Money is the imperfect instrument by which a scarcity age coordinates creation; when the closed loop of accounting, consensus, and incentives fully matures, the functions of money will gradually dissolve, differentiate, and sublimate — the measure of value evolving into lifelong negentropy-contribution genealogy archives, the medium of circulation into intelligent demand-matching networks (the end-state developed in (Cheng and Cheng, 2026a, ch. 10)). NVC is thus an intermediary toward a post-monetary society, not a terminal point: its ultimate success consists precisely in its own withdrawal — civilization ceases to preoccupy itself with allocating scarcity and takes creation itself directly as its purpose. The highest mission of the governance architecture is to keep this road to human self-sublimation from being hijacked, corrupted, or cut short by short-sightedness.

9. Falsifiable Hypotheses and Verification Designs

The hypotheses below are stated in executable form, and are explicitly ordered by execution priority. Two principles set the order. First, executable designs with objective, pre-registrable pass/fail criteria precede designs whose outcome coding requires judgment. Second, designs whose agent-based implementation risks circularity — where the modeler must assume in the mechanism roughly what the experiment is meant to find in the result — are re-scoped toward laboratory or real-world evidence before their ABM versions are run. A household-wealth-level analog of H5 has been executed computationally and is reported in Appendix C (Section C.3); the remaining designs await execution in the order stated.

H3 (bias detection rate of AEAC audits) — priority 1. Inject controlled biases (regional discrimination, domain prejudice) into a simulated implementation of the CPPA and measure the detection and false-positive rates of the quarterly disparate-impact audit. Criteria: pre-registered detection floor of 90%, false-positive ceiling of 10%. This design leads the queue for three reasons. Its criterion is a property of a well-defined detection pipeline against planted ground truth, so the outcome coding is mechanical; its implementation does not require modeling political behavior, so no circularity arises — the biases are injected by the experimenters, and the audit pipeline either finds them or does not; and its subject — the audit duty — is the firewall on which the most other defenses depend, including the recursive-oversight termination of Section 5.2, so its failure would degrade the largest share of the guardianship matrix. The full specification (bias magnitudes and placement, audit sampling protocol, detection thresholds fixed before execution) will be registered on OSF Registries before the code is run, following the pre-analysis-plan practice of the second paper’s Appendix E.

H2 (decision quality of the double majority) — priority 2. In simulation, compare bicameral double-majority rule with single-chamber majority rule on trap proposals — popular in the short run,

damaging to civilization utility in the long run — measuring rejection rates and false-kill rates. Criterion: the double majority rejects trap proposals at a significantly higher rate, and the increase of the false-kill rate on benign proposals does not exceed the pre-registered threshold of 15 percentage points. The design's honesty requirement is that the classification of proposals as "trap" or "benign" be fixed by an independent ex-ante rubric (long-run CU contribution sign, computable in the model), not by ex-post inspection of outcomes.

H6 (composition monitoring) — priority 3. Implement the AEAC's seventh duty (Section 5.2) in the ABM: under simulated misalignment between α -weighted creation and consumption demand, test whether the composition monitor detects premium congestion at the upper band rail earlier than the extinction-gap indicator detects Mode D stress. Criterion: the composition indicator leads the extinction-gap indicator by at least one audit quarter under the simulated shock path. This test links the governance architecture to the experienced-price-level agenda of the second paper's Section 4.7.

H4 (accountability effect of the transparency platform) — priority 4. Compare perceived corruption and rent-seeking incidence under full on-chain publicity of the decision process versus result-only publicity (ABM plus a comparative analysis of analogous real-world platforms such as Estonian e-governance). Criterion: significantly lower rent-seeking incidence under full publicity. The real-world comparative arm is stated as primary; the ABM arm is mechanism exploration.

H5 (long-run distributional effects of genealogy attribution with intergenerational decay) — priority 5. Implement knowledge-genealogy attribution with the intergenerational decay rule (δ_{inherit}) in the ABM and track the intergenerational concentration of knowledge issuance income (the trajectory of the Gini coefficient), comparing with a no-decay control. Criterion: within the pre-registered band of δ_{inherit} , the concentration of knowledge income after the third generation is not significantly higher than that of ordinary labor income; the no-decay control is significantly higher and keeps rising. A reduced-form analog at the household wealth level has been executed (Appendix C, Section C.3): with $\delta_{\text{inherit}} \in \{0.3, 0.5\}$ and GPF recycling, the third-generation wealth Gini is 0.308 and 0.285 against 0.366 for the no-decay control, and the top-decile wealth share falls from 0.259 to 0.234 and 0.225. The full genealogy-attribution implementation remains open.

H1 (capture resistance of sortition) — priority 6, re-scoped. Original form: in ABM and laboratory experiments, compare the effect of lobbying expenditure on α decisions under three selection mechanisms for legislators (election, sortition, appointment), with the criterion that under sortition the decision deviation per unit of lobbying expenditure is significantly smaller than under election (effect size $d \geq 0.5$). We retain the hypothesis but re-scope it, because its ABM implementation is circular in a specific and corrosive way: in an agent-based model, lobbying expenditure influences decisions only through an influence function the modeler writes down, and any specification in which sortition severs the lobbyist's access channel will trivially produce the predicted result — the hypothesis would be confirmed by construction. The ABM arm is therefore demoted to mechanism exploration (mapping which access-channel assumptions generate which deviations), and the test arm is relocated to evidence channels where the influence process is real rather than programmed: (i) laboratory experiments with human subjects playing legislators under the three selection rules, with real stakes for lobbyist and legislator players; (ii) the existing empirical record of deliberative mini-publics (the Irish citizens' assemblies and others) as a prior proxy test, where lobbying access is observable rather than assumed. The pre-registered criterion ($d \geq 0.5$) applies to the laboratory arm; the mini-public evidence is graded as suggestive, not confirmatory. It should be stated plainly that all existing mini-public evidence is national or local: no global-scale sortition record exists, and none can exist before the registry problem of Section 4.1 is solved, so the global-scale form of H1 is an extrapolation from smaller scales and is labeled as such.

10. Limitations and Honest Boundaries

First, this paper is a normative construction, not an empirical conclusion; all claims of architectural validity remain in the state of designed but unexecuted, and the priority ordering of Section 9 is a commitment about sequence, not a claim about results.

Second, the representativeness of global-scale sortition (three-dimensional stratification) depends operationally on a database of eligible adult citizens worldwide, the construction of which is itself an unsolved political problem — the bootstrapping protocol of Section 4.1 dilutes the circularity gradually by indexing legitimacy to coverage, but does not abolish it, and the first convening of a low-coverage GVP remains an act of founding politics, not of procedure. Beyond the database problem stands a harder geopolitical one. Participation is voluntary, a global registry presupposes cooperation from the very sovereign powers whose monetary discretion the constitution would curtail, and the path from pilots to global coverage crosses unresolved conflicts of sovereignty and great-power rivalry that no design feature of this paper can settle; the pilot ladder of Section 4.4 is feasible precisely because it does not require that settlement in advance. Even with a complete registry, moreover, the stratification dimensions themselves — continental population share, cultural sphere, development level — are social-normative choices, open to contestation even where the random-draw machinery is technically neutral; what counts as a cultural sphere is a value judgment, and the constitution locates that judgment in the convention process of Table 2 rather than pretending it away.

Third, jurisdictional overlaps and conflicts among the three firewalls are possible — for instance, the boundary between the AEAC's emergency brake and the GVP's sovereignty; this paper has given procedural constraints such as reporting deadlines, but complete conflict-adjudication rules remain to be developed.

Fourth, the concrete parameters — the 70/30 principle, the $\pm 30\%$ floating band, the optimal band of δ , the attribution weights — all depend on model calibration, and their robustness has not been independently replicated.

Fifth, Propositions 1 and 2 are formal but deliberately conditional, and Proposition 1 in particular claims less than its label might suggest: it is a structural incompatibility among four design requirements for single-rule mechanisms, not an Arrow-depth theorem (see the Remark in Section 3) — its function is to fix the design's escape route, and the composite architecture it motivates inherits the ordinary burdens of composite architectures (coordination costs, blame diffusion, slower crisis response) that single-rule mechanisms do not face. Proposition 2's superadditivity is proportional to the vigilance factor r , which collapses if the generation mechanisms converge; the independence of the four bodies is thereby promoted from design goal to load-bearing assumption.

Sixth, the failure independence of the guardianship matrix is a design goal, not an accomplished fact — the actual correlation of the bodies in resources and information infrastructure (a shared blockchain stack, for example, may introduce common-mode failure) requires dedicated assessment at the stage of engineering implementation; and because the recursive-oversight termination of Section 5.2 runs through the platform's tamper-proof records, a common-mode failure of that stack is also a failure of the regress's last link — the matrix's single most dangerous point.

Seventh, the teleological account of Section 8 is a normative commitment rather than an empirical prediction: the trajectories of individual elevation and civilizational leaps enjoy support from historical analogies and micro-level evidence, but the degree to which they would fully unfold under NVC incentives cannot be guaranteed in advance; the falsifiability of Proposition 3 rests on indirect testing through the long-run tracking of CU proxy indicators. The fixed telos also sits in unresolved tension with the experimentalist recursion the architecture otherwise embraces (Section 2.5): a constitution that revises everything except its directional goal must show that the exempted goal disciplines the recursion rather than capturing it, and that showing awaits cases. The telos is nonetheless separable from the machinery. A political community that declined the civilizational-direction premise could still operate the parliament, the firewalls, and the guardianship matrix as a neutral integrity-guardianship constitution for whatever value weights its own deliberation produced; what would be lost is the directional navigation described in Section 8, not the integrity protections of Sections 3–7.

Eighth — a limitation inherited from the companion papers and now carried by this one — the architecture presupposes that the calibration layer can absorb the epistemic burden that market prices

carry in conventional systems. The first paper bounds that burden (its Section 9.6), the second assigns its macroeconomic symptom a monitor (its Section 4.7 and this paper’s AEAC duty seven), but neither eliminates it. A governance constitution can distribute, delay, and audit value judgments; it cannot manufacture the knowledge those judgments require. The speed and accuracy of procedure-mediated correction relative to market discipline is, in the end, an empirical question, and the honest answer today is that we do not know.

11. Conclusion

This paper has completed the suspended constitutional link of the negentropic value currency program: the institutional guardianship of the four integrity conditions and the legitimate generation of the consensus coefficient α . The architecture’s answer is layered — the bicameral algorithmic democracy supplies the balance of legitimacy and epistemic quality; the three constitutional firewalls defend respectively against majoritarian tyranny, algorithmic alienation, and the black box of power; the four clauses of metabolic governance constrain the economic rules themselves; and the guardianship matrix anchors every institution to the integrity condition it defends, with the capture game of Proposition 2 giving the superadditivity of capture costs an explicit conditional form; the teleological loop (Section 8) then anchors all mechanisms to civilization’s ultimate orientation — knowing, utilizing, and mastering cosmic truth, until the withdrawal of money itself. Three self-imposed disciplines mark the presentation: the impossibility claim is stated at its true (bounded) strength rather than dressed in Arrow’s clothing; the recursive-oversight regress is terminated by a cost argument with an explicit failure point, not by faith in a final guardian; and the verification designs are ordered so that the claims that can be tested objectively are tested first, and the claims whose agent-based tests would be circular are re-scoped to human-subject and real-world evidence. The institutional interfaces of Section 4.5 ground the architecture’s measurement demands in two running national programs — GEP accounting and pension time banks — and fix a staged docking path (measurement methodology → accounting instruments → quasi-monetary experiments) whose every stage has independent value and survivable failure. The three papers together form a complete trilogy: value ontology (paper 1) → macroeconomic issuance and stability (paper 2) → governance constitution (this paper). The common methodology of the program is invariant: distinguish normative premises from physical facts, mark executed and unexecuted tests, and publish the failure modes of every layer. The program’s next step is to execute the verification designs in the priority order of Section 9, beginning with the audit-detection specification of H3, and to integrate the trilogy into an English series for journal submission.

Funding: The authors received no funding for this research.

Data Availability Statement: This paper reports no new empirical dataset. The model extension, parameter settings, and distributional outputs for the agent-based experiments are reported in Appendix C; executable code is available from the authors and will be deposited in a public repository with journal submission. No independent public-repository DOI had been created at the time of this revision. The verification designs of Section 9 will be executed in the stated priority order; the pre-registered specification of H3 will be registered on OSF Registries, timestamped before its code is run.

Conflicts of Interest Statement: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Notation

Symbol	Meaning
$\Theta(t), \alpha_{\tau}$	Civilizational preference vector; consensus coefficients
G	Consensus-generation mapping (Definition 1)
L, E, C, O	Legitimacy; epistemic quality; capture resistance; evolutionary openness (Proposition 1)

$\omega_{\min}$	Exploration weight floor
$c_i, e_i, p_i(e_i)$	Capture-cost scale, effort, control probability for body i (Proposition 2)
q_0, r, F, T^-	Exposure hazard; vigilance factor; exposure penalty; campaign duration
$q_{\text{cov}}(t), q_j$	Sortition-database coverage; activation threshold of parameter class j
δ, k_s	Metabolism rate; stability coefficient (papers 1–2)

Appendix B. Premise Dependency Map

Premise layers follow the companion papers: L2 denotes the axiom layer (the social-choice and game-theoretic axioms stated with each proposition; the inherited value-equation axioms of the first paper include the aggregation axiom Axiom 3 over disjoint activities and windows); L3 the normative survival-first premise; L5 the structural assumptions about institutional design stated in the text.

Result	Premise layers	Survives rejection of L3?
Proposition 1 (structural incompatibility)	L2	Yes
Proposition 2 (capture game)	L2, L5	Yes
Double majority and firewalls (Sections 4–5)	L3, L5	Partly (objective-dependent)
Guardianship matrix (Section 7)	L2, L5	Yes, as defense design
Teleological loop (Section 8)	L3	No

The same division of labor holds here as in the companion papers: the two formal propositions are premise-inexpensive incompatibility and cost results, independent of the normative premise; the premise enters only where the architecture claims a direction (Section 8), not where it claims a defense.

Note. Appendix C — the distributional-dynamics experiments (the transition of the low-value workforce under a retraining protocol, the intergenerational decay of bequests with GPF recycling, and the joint policy space of the demurrage rate δ and the social floor) — accompanies this manuscript as supplementary material for the editor's and referees' reference at submission; it is not intended for publication with the main text.

Supplementary Appendix C

for: Negentropic Value Currency: The Governance Architecture of Algorithmic Democracy and Constitutional Firewalls

Dian Cheng and Jinjun Cheng

The following appendix accompanies the manuscript for the editor's and referees' reference at submission; it is not intended for publication with the main text. Appendix C reports the distributional-dynamics experiments of the NVC constitution: the transition dynamics of the low-value workforce with and without the retraining protocol, the intergenerational decay of bequests as a household-wealth analog of hypothesis H5, and a scan of the joint policy space of δ and the social floor, with the associated distributional outcomes.

Contents — Appendix C: Distributional Dynamics: Transition, Bequests, and the Policy Space of δ and the Social Floor.

Appendix C. Distributional Dynamics: Transition, Bequests, and the Policy Space of δ and the Social Floor

This appendix extends the agent-based model of the second paper (Cheng and Cheng, 2026b, Appendix D) with three household-level layers — an overlapping-generations demography with bequests, a retraining protocol that migrates low-value workers toward the care and ecology domains, and a social consumption floor financed from demurrage proceeds and Global Public Value Fund (GPF) revenues — and reports three experiments: the transition dynamics of the low-value workforce (Section C.2), the intergenerational decay of bequests as a household-wealth analog of H5

(Section C.3), and a scan of the joint policy space of the metabolism rate δ and the floor level (Section C.4). All configurations use $N = 1,000$ households, 150 periods, and 30 replications; reported statistics are replication means, and shaded bands in Figure C.1 are 95% replication intervals.

C.1 Model Extension

Households differ in domain membership (drawn from the six-domain shares of the baseline calibration), in productivity (log-normal, intergenerationally persistent with parent–child correlation 0.3), and in asset portfolios comprising money, CTB certificates, and — when borrowing is permitted — debt up to a ceiling of one times mean income, with full-surplus repayment priority. Lifespans are drawn from a normal distribution with mean 60 and standard deviation 5, truncated to [45, 80]. At death, debts are extinguished; the heir receives a share $(1 - \delta_{\text{inherit}})$ of net positive wealth; and the reverted share accrues to the GPF and is recycled as an equal per-capita dividend to all households — the reduced-form implementation of Clause 4’s “common asset of all civilization.” Under the retraining protocol, a fixed fraction (2% per year) of workers in the low-value domain migrates to the care (60%) and ecology (40%) domains, incurring a one-time 10% productivity cost. The social floor tops up household consumption to a specified share of mean income, financed from demurrage proceeds and GPF revenues. The baseline configuration is $\delta = 1\%$, $\delta_{\text{inherit}} = 0.3$, floor = 0.3, retraining active; each experiment varies one dimension.

C.2 Transition Dynamics: Retraining the Low-Value Workforce

Figure C.1(a, b) compares the baseline with and without the retraining protocol. Without retraining, the low-value domain retains about 21% of the workforce, and the income Gini drifts only from 0.36 to 0.34 over 150 periods. With retraining at 2% per year, the low-value share declines smoothly to 3.1% by period 100 (2.0% by period 150), the income Gini falls to 0.301, the wealth Gini to 0.341 (against 0.403 without retraining), and the share of households below the subsistence threshold is 0.1% at period 100. Two features matter institutionally. First, the migration is absorptive rather than displacing: care and ecology — the domains whose market-value weight most understates their tracked contribution — absorb the transitioning workers without wage collapse, because issuance follows tracked contribution rather than market scarcity. Second, retraining and the floor are complements: the floor prevents transitional poverty during the migration window, while retraining shrinks the population that would otherwise draw on the floor permanently. A faster protocol (4% per year) drives the low-value share below 1% by period 100, a pace we judge institutionally implausible; the 2% calibration is retained as conservative.

C.3 The Bequest Mechanism: Intergenerational Decay with GPF Recycling

Clause 4’s decay rule is implemented here at the household wealth level, as a reduced-form analog of the genealogy-attribution test of H5: at death, the reverted share δ_{inherit} of net positive wealth accrues to the GPF and is recycled as an equal per-capita dividend. Figure C.1(c) reports wealth-Gini trajectories for $\delta_{\text{inherit}} \in \{0, 0.3, 0.5\}$. After the third generation (periods 110–150), the wealth Gini is 0.366 under the no-decay control, 0.308 under $\delta_{\text{inherit}} = 0.3$, and 0.285 under $\delta_{\text{inherit}} = 0.5$; the top-decile wealth share is 0.259, 0.234, and 0.225, respectively. The effect is monotone and sizable — a 16% to 22% reduction of third-generation wealth concentration — but two qualifications deserve emphasis. First, even the no-decay control does not diverge within the horizon: with the CTB’s net return (about 1%) below the growth rate (2%), the configuration is $r < g$, so bequests do not compound explosively as they do in the credit-money control of the inequality simulation, where the capital return of 4% exceeds the growth rate. The decay rule therefore operates as an accelerator of an already-stable dynamic, not as the sole barrier against rentier formation. Second, the direction of the effect depends on the destination of the reverted share: an earlier implementation that destroyed rather than recycled it produced an inverted, middle-class-squeeze artifact — middle heirs lose modest bequests while top wealth, accumulated within a lifetime through CTB revenue shares,

remains untouched. The equalization reported here holds only when the reverted share is recycled through the GPF, which is exactly what Clause 4 specifies.

C.4 The Policy Space of δ and the Social Floor

Table C.1 and Figure C.1(d) report a scan of $\delta \in \{0.5\%, 1\%, 2\%, 3\%\}$ against floor levels $\in \{0, 0.2, 0.3, 0.4, 0.5\}$ (as shares of mean income), with $\delta_{\text{inherit}} = 0.3$ and retraining at 2% per year. Three findings stand out. (i) The incidence of δ is regressive at high values: at floor = 0, the wealth Gini rises from 0.341 at $\delta = 0.5\%$ to 0.378 at $\delta = 3\%$, because poorer households hold their buffers in money (taxed at δ) while wealthier households hold CTB certificates (net return approximately $g - \delta$). (ii) Within the constitutional band of Clause 1 (0.5%–1%), the distributional response is near-flat — the wealth Gini moves by less than 0.005 across the band at any floor — so the band is confirmed on distributional as well as accumulation grounds. (iii) The floor, not δ , is the binding distributional instrument: a floor at or above 0.4 eliminates the debtor population at every δ (from 0.6%–2.3% of households to zero) and raises the bottom-half income share monotonically at every δ . The policy implication is a division of labor among instruments: δ within its band governs monetary metabolism, the floor governs the lower tail, and δ_{inherit} governs intergenerational concentration. Using δ above its band as a redistributive device would tax precisely the money balances of the poor.




Figure C.1. Distributional dynamics under the NVC constitution. (a) Share of workers in the low-value domain, with and without the 2%-per-year retraining protocol. (b) Income Gini under the same two configurations. (c) Wealth-Gini trajectories under inheritance decay $\delta_{\text{inherit}} \in \{0, 0.3, 0.5\}$ with GPF recycling; shaded bands are 95% replication intervals. (d) Wealth Gini at period 100 across the joint policy space of the metabolism rate δ and the social floor. All panels: $N = 1,000$ households, 150 periods, 30 replications.

Table C.1. Policy-space scan at period 100. Entries are wealth Gini / bottom-50 income share / debtor share (%); 30 replications, $\delta_{\text{inherit}} = 0.3$, retraining 2% per year.

δ Floor	0.0	0.2	0.3	0.4	0.5
0.5%	0.341 / 0.267 / 0.6	0.341 / 0.268 / 0.6	0.335 / 0.271 / 0.4	0.332 / 0.273 / 0.0	0.320 / 0.281 / 0.0

1.0%	0.345 / 0.264 / 0.9	0.341 / 0.267 / 0.9	0.338 / 0.269 / 0.7	0.332 / 0.273 / 0.0	0.327 / 0.276 / 0.0
2.0%	0.359 / 0.255 / 1.5	0.357 / 0.256 / 1.6	0.358 / 0.256 / 1.4	0.341 / 0.266 / 0.0	0.332 / 0.272 / 0.0
3.0%	0.378 / 0.243 / 2.0	0.378 / 0.243 / 2.3	0.376 / 0.244 / 2.2	0.360 / 0.254 / 0.0	0.350 / 0.260 / 0.0

C.5 Limitations

Three limitations qualify these results. First, GPF recycling is modeled in reduced form as an equal per-capita dividend; explicit public-goods provision would alter the incidence in ways that depend on the expenditure mix. Second, the household portfolio omits housing and productive capital goods, so the wealth distribution understates the concentration of non-financial assets. Third, retraining costs, success rates, and destination shares are exogenous; endogenizing occupational choice under tracked-contribution wages is a natural extension. The experiments establish mechanisms within the model’s ontology, not forecasts.

References

1. Arrow K. J. (1963): Social Choice and Individual Values, 2nd edn, Yale University Press, New Haven.
2. Beijing Municipal Civil Affairs Bureau, Beijing Municipal Finance Bureau, Beijing Municipal Committee of the Communist Youth League (2022): Implementation Plan for the Beijing Pension Service Time Bank (Trial).
3. Buchanan J. M., Tullock G. (1962): The Calculus of Consent: Logical Foundations of Constitutional Democracy, University of Michigan Press, Ann Arbor.
4. Cheng D., Cheng J. (2026a): “Value Consensus Currency (VCC): A Meta-Theoretical Construction for the Leap of Human Civilization in the Post-Scarcity Era”, SSRN Working Paper. DOI: 10.2139/ssrn.6171167.
5. Cheng D., Cheng J. (2026b): “Negentropic Value Currency: A Conditional Theory of Inflation and Macrostability over the Full Monetary Life Cycle”, revised manuscript, October 2026; earlier version: Preprints, preprints202609.2497.v3. DOI: 10.20944/preprints202609.2497.v3.
6. Cheng D., Cheng J. (2026d): “Negentropic Value Currency: From Scarcity to Order — Monetary Value Anchored in Net Order Creation”, revised manuscript, October 2026; earlier version: Preprints, preprints202609.1138.v4. DOI: 10.20944/preprints202609.1138.v4.
7. Cheng D., Cheng J. (2026e): “From GEP to Negentropy Accounting: National Institutional Interfaces and a Staged Docking Path for the NVC Measurement Layer”, NVC Program Working Paper, October 2026.
8. Dahl R. A. (1999): “Can international organizations be democratic? A skeptic’s view”, in Shapiro I., Hacker-Cordón C. (eds), Democracy’s Edges, pp. 19–36, Cambridge University Press, Cambridge.
9. Fishkin J. S. (2009): When the People Speak: Deliberative Democracy and Public Consultation, Oxford University Press, Oxford.
10. Fromm E. (1976): To Have or to Be?, Harper & Row, New York.
11. General Office of the CPC Central Committee, General Office of the State Council (2021): Opinions on Establishing and Improving the Mechanism for Realizing the Value of Ecological Products.
12. Gesell S. (1916): Die natürliche Wirtschaftsordnung durch Freiland und Freigeld, self-published, Bern. (English translation: The Natural Economic Order, Free Economy Publishing, San Antonio, 1934.)
13. Gibbard A. (1973): “Manipulation of voting schemes: A general result”, Econometrica, 41(4), pp. 587–601.
14. Granovetter M. (1985): “Economic action and social structure: The problem of embeddedness”, American Journal of Sociology, 91(3), pp. 481–510.
15. Habermas J. (1996): Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy, MIT Press, Cambridge, MA.
16. Held D. (1995): Democracy and the Global Order: From the Modern State to Cosmopolitan Governance, Polity Press, Cambridge.
17. Keynes J. M. (1936): The General Theory of Employment, Interest and Money, Macmillan, London.
18. Lafont C. (2020): Democracy without Shortcuts: A Participatory Conception of Deliberative Democracy, Oxford University Press, Oxford.

19. Landemore H. (2013): *Democratic Reason: Politics, Collective Intelligence, and the Rule of the Many*, Princeton University Press, Princeton.
20. Mansbridge J. (2003): "Rethinking representation", *American Political Science Review*, 97(4), pp. 515–528.
21. Mansbridge J. (2009): "A 'selection model' of political representation", *Journal of Political Philosophy*, 17(4), pp. 369–398.
22. Mill J. S. (1859): *On Liberty*, John W. Parker and Son, London.
23. Mullainathan S., Shafir E. (2013): *Scarcity: Why Having Too Little Means So Much*, Times Books, New York.
24. National Development and Reform Commission (2025): Reply to Recommendation No. 6602 of the Third Session of the 14th National People's Congress, 1 August 2025.
25. National Development and Reform Commission, National Bureau of Statistics of China (2022): Accounting Standard for Gross Ecosystem Product (Trial) [Shengtai Chanpin Zongzhi Hesuan Guifan (Shixing)].
26. Ostrom E. (1990): *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, Cambridge.
27. Pasquale F. (2015): *The Black Box Society: The Secret Algorithms That Control Money and Information*, Harvard University Press, Cambridge, MA.
28. Peking University Institute of Population Research (2021): *China Time Bank Development Research Report*, Peking University, Beijing.
29. Prigogine I., Stengers I. (1984): *Order out of Chaos: Man's New Dialogue with Nature*, Bantam Books, New York.
30. Rawls J. (1971): *A Theory of Justice*, Harvard University Press, Cambridge, MA.
31. Sabel C. F., Zeitlin J. (eds) (2010): *Experimentalist Governance in the European Union: Towards a New Architecture*, Oxford University Press, Oxford.
32. Satterthwaite M. A. (1975): "Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare functions", *Journal of Economic Theory*, 10(2), pp. 187–217.
33. Sen A. (1999): *Development as Freedom*, Knopf, New York.
34. Shanghai Civil Affairs Bureau (2021): Notice on Implementing the "Pension Service Time Bank" Program Citywide in Shanghai [Hu Min Yang Lao Fa (2021) No. 28].
35. Tocqueville A. de (1835/1840): *De la démocratie en Amérique* [Democracy in America], Gosselin, Paris.
36. UN Statistical Commission (2021): Adoption of the System of Environmental-Economic Accounting — Ecosystem Accounting (SEEA-EA) as an international statistical standard.
37. Whitehead A. N. (1929): *Process and Reality: An Essay in Cosmology*, Macmillan, New York.
38. Zuboff S. (2019): *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, PublicAffairs, New York.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.