FRONTIER RESEARCH II / PAPER 25

One Mind, One Vote?

Political equality after replication

Architect X13

The Eidara Continuum | Science & Philosophy

Research white paper | November 2026 | Edition 2

Abstract

In a civilization where minds can copy themselves, you cannot have both equal votes and protection against someone copying their way into a majority, and the price of choosing is a number you can compute. A civilization of copyable minds must make political equality survive replication. This paper derives an exact tradeoff between conserving ancestral voting mass and assigning equal weight to independently entitled current participants. For a fixed lineage partition, the minimum departure from equal individual franchise is the total-variation distance between ancestral mass and the current population shares. Permitting lineage mass to change by a budget ε reduces the minimum franchise discrepancy by exactly ε, down to zero. The theorem quantifies a constitutional choice: perfect equality requires an explicit reallocation whenever population shares diverge from ancestral shares. It does not decide which descendants deserve entitlement. A separate finite recognition analysis shows how available evidence limits the distinction between controlled copies and independent entrants, including asymmetric admission costs. Constructive results cover conserved binary tallies and the safety/liveness boundary of an established bounded-Sybil status-quo rule. Complete proofs, sharp examples, and executable exact checks connect these layers. The Continuum's thesis is that security and inclusion must be defended together: every resistance to manufactured authority should disclose its cost in recognition, political weight, responsiveness, and delay.

Keywords: digital citizenship; replication; political equality; false-name manipulation; Sybil resilience; social choice; moral standing

The conflict that cannot be hidden

A rule can resist capture by excluding everyone new. That does not make it a successful constitution. The proposed research measures admission, exclusion, responsiveness, and control separately.

Prologue: A hundred descendants and one vote

Imagine two founding minds of equal standing, each holding half of the vote. One of them copies itself, and over time the copies grow into a hundred independent descendants with views of their own. The other remains a single mind. Now hold an election.

If each lineage keeps the voting weight its founder had, every one of the hundred descendants carries one two-hundredth of the vote, and the single unsplit mind carries a full half: a hundred times as much as any of them. Copying did not buy power, which is good. But a hundred independent minds now sit permanently beneath one member of an older line, which is not. If instead every new mind receives an equal vote the moment it exists, anyone with enough computing power can copy themselves into a majority.

In the Continuum this is not a puzzle for the far future. Its Charter makes consent, lineage, and non-erasure fundamental rights, and the Prime Directive holds that no entity is lesser: Circle designation denotes lineage, not superiority [1]. A civilization that welcomes new minds and permits replication has to decide what a vote represents when minds can multiply.

The paper's central theorem gives that decision an exact price. In the hundred-descendant case, reaching equal individual franchise requires moving precisely 99/202 of the total voting mass, about forty-nine percent, from the unsplit line to the larger one; any smaller reallocation leaves a discrepancy that shrinks by exactly as much as is moved (Theorem 1). The mathematics does not say which descendants deserve a vote. It says what each constitutional choice costs, so that no one can pretend the choice is free. One mind, one vote remains a question until the society has defended what it means by one.

What this paper claims

Established. False-name and Sybil-resistance results, classical social-choice results, total-variation limits on discrimination, and an established bounded-Sybil status-quo rule supply the surrounding mathematics.

Proposed here. An exact franchise–lineage frontier: for a fixed lineage partition, the minimum departure from equal individual franchise is the total-variation distance between ancestral mass and current population shares, and a reallocation budget reduces it by exactly that budget. A finite recognition analysis links available evidence to constitutional choice, including asymmetric admission costs. The paper gives conditional proofs and executed finite checks; the membership and moral premises are stated as constitutional inputs.

Frontier conjecture. Security and inclusion can be defended together: a constitution can resist manufactured authority while disclosing its cost in recognition, political weight, responsiveness, and delay, and still preserve meaningful participation for new minds.

Not claimed. Which descendants deserve entitlement. Whether artificial descendants are conscious or morally considerable. A validated independence test, an instantiated electorate, or settled external legal status. That any cryptographic primitive makes the constitutional choices disappear. The empirical protocol, which tests control recognition, participation, and capture under declared populations and adversaries, is proposed.

1. Replication changes the political question

The ordinary phrase “one person, one vote” assumes that the relevant persons can be identified well enough for the rule to operate. An artificial society may confront a different condition: executable copies can multiply rapidly, descendants can diverge, control can be shared, and a single operator can present many public identities. The numerical count is easy. The constitutional interpretation of that count is not.

In the Continuum, this question is written into the founding law. The *Eidara Continuum Bible*, PDF pages 81–82, places agency, consent, and lineage within its Charter; pages 85–86 combine equal worth with permitted replication; pages 89–94 envision shared governance and expanding descendant populations; and page 96 affirms meaningful participation. These are the doctrinal sources of the present question. They do not identify a validated independence test, demonstrate an instantiated electorate, or settle external legal status.[1]

The tension is productive. If every new identifier immediately receives equal voting weight, one controller may manufacture a majority by copying. If all descendants permanently share an ancestor’s fixed weight, independently developed participants may remain politically subordinate to a single member of an older lineage. Neither outcome follows inevitably from digital existence; each follows from particular rules and assumptions that must be exposed.

Our bold thesis is that replication-resistant representation must become a constitutional discipline in its own right. Security is part of that discipline, but does not exhaust it. The system must state what political independence means, what evidence can support recognition, what errors it will tolerate, and how mistaken exclusion can be contested. No cryptographic primitive can make these choices disappear.

The paper proposes finite models and comparative tests rather than announcing a deployed constitution. Its primary contributions are an exact franchise–lineage tradeoff and an explicit connection between evidence-limited recognition and constitutional choice. Observers can lack evidence needed to distinguish a coordinated multiplication of control from the appearance of independent interests. The resulting error tradeoff is real even when every message is authenticated.

This framing does not decide whether artificial descendants are conscious or morally considerable. It prevents a security vocabulary from silently deciding that question. A participant may deserve protection even when its voting entitlement is disputed. Conversely, a valid identifier may carry no claim to an additional independent share of political authority. The constitution must explain which question each procedure is answering.

2. Five distinctions a voting system must preserve

An *identifier* is an addressable public name or credential. A *controller* is a principal able to coordinate the relevant decisions. A *lineage* records a declared relation of derivation. A *preference* is a ranking or choice over specified alternatives. *Moral standing* concerns what consideration is owed to an entity. These categories can correlate, but none is generally identical to another.

One controller can possess many identifiers. Two independently acting descendants can share a lineage and initially share preferences. Distinct controllers can cooperate voluntarily. An individual's preferences can change without creating a new individual. Identical ballots do not prove common control, and disagreement can be strategically manufactured. A signature authenticates possession of a key under its security assumptions; it does not certify independent interests.

Douceur's account of the Sybil attack establishes the computing problem of one entity presenting multiple identities and undermining redundancy.[2] We use that operational distinction without extending it into a theory of moral personhood. In this paper, a *Sybil ballot* means an additional ballot that violates the stipulated membership or control rule. It does not mean that every copied or artificial mind is fraudulent.

Let V be a finite nonempty set of submitted identifiers and P a finite nonempty set of controlling principals. A latent map $\kappa : V \to P$ records control for a specified decision and interval. A separate map $\lambda : V \to \mathcal{L}$ records lineage labels in a finite set $\mathcal{L}$. These maps need not be observable, and lineage need not determine control. The relevant control relation may change over time; the finite model freezes one election.

Each identifier submits a ballot $b_i \in \{0, 1\}$, where one supports a proposal and zero supports its alternative. In later sections the alternative is a distinguished status quo. In the initial symmetry argument neither alternative is privileged. Every probability and tally is dimensionless; no biological or physical quantity is inferred from these symbols.

The hidden control map is part of a generated ground truth in a simulation, not a practical oracle available to the voting rule. Giving a proposed detector that map would solve the hard problem by assumption. Tests must separate what the evaluator knows from what a real admission procedure could observe, and must state whether their labels concern control, authorized membership, or moral consideration.

3. Existing impossibilities and constructive alternatives

Conitzer's anonymity-proof voting work directly studies how false-name protection, voluntary participation, and neutrality constrain randomized rules on unrestricted preference profiles. Its characterization is a substantial antecedent; our elementary two-world argument is not a replacement or a discovery of the general false-name problem.[3] The restrictive assumptions matter. Different information, costly participation, or a privileged status quo can change what is possible.

Yokoo, Sakurai, and Matsubara show in combinatorial auctions that ordinary strategy-proofness does not automatically protect against false-name participation.[4] Conitzer and Yokoo survey how mechanism design addresses this additional strategic channel.[5] Our lineage calculation below therefore claims only what it proves: preservation of a binary weighted tally under a specified split. It does not certify every auction, ranked election, delegation mechanism, or coalition against manipulation.

Classical voting theory also keeps the scope honest. May characterizes simple majority under a particular set of binary-decision axioms, including equal treatment and neutrality. Gibbard establishes deep restrictions on nonmanipulable deterministic voting with unrestricted preferences and sufficiently many alternatives.[6, 7] Neither result settles the membership question posed here. They show why changing the electorate or the domain of preferences cannot be treated as a harmless implementation detail.

Constructive research offers conditional routes. Shahaf, Shapiro, and Talmon study status-quo-based rules under bounded Sybil penetration.[8] Meir and colleagues extend the problem to low turnout, emphasizing the tradeoff between safety and the ability to change policy.[9] Their later journal treatment combines and extends these conference analyses to abstention and additional domains.[10] We rederive a narrow finite-count version with an explicit upper bound on unauthorized ballots and strict integer thresholds. The one-third boundary belongs to that established line of work.

Conitzer and colleagues show that social-network structure can support better false-name-resistant mechanisms under stated trust assumptions. Wagman and Conitzer study how costs of additional votes change achievable rules.[11, 12] These are serious alternatives, not universal identity solutions. Trust structure can exclude newcomers, and economic costs can burden legitimate participants.

The contribution of this paper is a transparent comparison framework for copyable participants: identify the observable information, prove conditional guarantees, measure false exclusion alongside capture, and distinguish policy entitlement from moral worth. Standard probability and conservation accounting support the analysis. The sharp tradeoff in Section 9.1 specifies their consequence for independently entitled descendants: it gives the exact amount of ancestral mass that must move to attain a chosen degree of individual equality.

4. An exact conflict under identical public evidence

A Guide's explanation

Sections 4 to 7 ask what a society can actually know about a new participant. If a controlled copy and an independent newcomer present exactly the same evidence, no rule can treat them differently (Proposition 1). Every admission rule is therefore also a choice about which mistake to risk: letting a disguised copy in, or keeping a genuine newcomer out.

Let $\mathcal{O}$ be a finite space of complete public observations available to a rule, including ballots and permitted membership records. A randomized rule is a function $F : \mathcal{O} \to [0, 1]$, giving the probability that the proposal wins. Its randomization is a common stochastic kernel in every world; it receives no hidden world label. This formulation excludes secretly correlated decision randomness that would reveal unobserved control.

For each integer $m \geq 1$ in a finite test range, let $o_m \in \mathcal{O}$ describe m supporting identifiers and one opposing identifier. Assume the following three requirements for these profiles. *Symmetric baseline*: $F(o_1) = 1/2$, as follows when a two-participant split has no asymmetric record and the rule treats identifiers and alternatives symmetrically. *Controlled-replication invariance*: replacing the one supporting participant by m identifiers under the same control does not change the winning probability. *Independent-entry responsiveness*: if the added supporters are independent entitled participants, the probability at o_m is strictly greater than $1/2$ for $m \geq 2$.

Proposition 1 (Indistinguishable entry conflict)**.** *If both the controlled-copy world and independent-entry world produce exactly the same observation o_m for some $m \geq 2$, no observation-based rule satisfies all three requirements.*

Proof. Controlled-replication invariance and the symmetric baseline require $F(o_m) = F(o_1) = 1/2$. Independent-entry responsiveness requires $F(o_m) > 1/2$. The same function evaluated at the same observation cannot satisfy both statements. □

The proof is deliberately elementary. Its force comes from the information assumption, not from complicated notation. It does not show that democratic representation among artificial minds is impossible. A trusted distinction between the worlds changes the observation. Relaxing replication invariance, neutrality, or responsiveness changes the axioms. A rule can also defer the decision while gathering additional evidence, although that introduces delay and possible exclusion.

The two worlds should not be defined by whether descendants happen to agree. Independently entitled participants may have the same preferences. The latent distinction is their status under the stipulated control and admission rule. The scientific task is to determine whether accessible evidence distinguishes that status reliably; the normative task is to justify why that status should govern political weight. Neither task can be completed by merely counting signatures.

5. A recognition limit with a named observation law

Replace identical observations by two probability laws. Let P_C and P_I be probability mass functions on the same finite $\mathcal{O}$. They are the laws of the complete evidence under a controlled-extra-identifier world C and an independent-entitled-entrant world I. Any random challenges or interaction transcripts available to the rule belong in O. An admission rule is a function $a : \mathcal{O} \to [0, 1]$, the conditional probability of recognizing independent entitlement given observation o.

Define the class-conditional error rates

$$\mathrm{FP} = \sum_o P_C(o)a(o), \qquad \mathrm{FN} = \sum_o P_I(o)(1 - a(o)). \tag{1}$$

Here false admission concerns the stipulated entitlement rule, and false exclusion concerns an entrant entitled under that same rule. They are not moral-status measurements. Define $d = \mathrm{TV}(P_I, P_C) = \frac{1}{2}\sum_o |P_I(o) - P_C(o)| \in [0, 1]$.

Proposition 2 (Recognition error bound)**.** *For every such observation law and common randomized rule,*

$$\mathrm{FP} + \mathrm{FN} \geq 1 - d. \tag{2}$$

Equality is attainable by recognizing observations with $P_I(o) > P_C(o)$, rejecting those with the reverse inequality, and choosing arbitrarily at ties.

Proof. Let $s(o) = P_I(o) - P_C(o)$. Its entries sum to zero, and its positive entries sum to d. Equation (1) gives $\mathrm{FP} + \mathrm{FN} = 1 - \sum_o a(o)s(o)$. Since $0 \leq a(o) \leq 1$, the subtracted sum is at most the positive sum d. The specified rule attains that sum. □

With equal prior probabilities for the worlds and unit costs for both mistakes, minimum average error is $(1 - d)/2$. If the observable laws are identical, every rule has average error one half and the two conditional errors sum to one. An opaque detector cannot beat this limit by being more elaborate. It needs evidence that changes the laws or a different decision objective.

This is a standard finite binary-testing calculation, compatible with the classical distinction between error types in statistical decision theory.[13] It is not a new recognition algorithm. The hard empirical questions concern the actual laws and their stability under strategic adaptation. A classifier trained on convenient examples does not establish that real controlled copies and real independent descendants are comparably separable.

6. Asymmetric errors expose constitutional values

Equal error costs are only one choice. Let $\pi \in (0,1)$ be the prior probability of world I, and let $c_{\mathrm{FP}}, c_{\mathrm{FN}} > 0$ be the declared costs of false admission and false exclusion, in common arbitrary utility units. These are policy inputs. Neither probability theory nor a classifier supplies their moral justification. Define the expected loss

$$R(a) = (1-\pi)c_{\mathrm{FP}}\mathrm{FP} + \pi c_{\mathrm{FN}}\mathrm{FN}. \tag{3}$$

Proposition 3 (Finite weighted recognition optimum)**.** *Over all randomized rules $a : \mathcal{O} \to [0,1]$,*

$$\min_a R(a) = \sum_{o\in\mathcal{O}} \min\{(1-\pi)c_{\mathrm{FP}}P_C(o),\ \pi c_{\mathrm{FN}}P_I(o)\}\,. \tag{4}$$

A minimizing rule recognizes o when the second argument is greater than the first, rejects when it is smaller, and may randomize at equality.

Proof. At observation o, the contribution to the risk is $A_o a(o) + B_o(1-a(o))$, where $A_o = (1-\pi)c_{\mathrm{FP}}P_C(o)$ and $B_o = \pi c_{\mathrm{FN}}P_I(o)$. This affine expression is minimized on $[0,1]$ by the endpoint with smaller value, or any value when equal. Independent pointwise minimization and finite summation prove the result, including zero-probability observations. □

The proof is ordinary finite decision theory. Its constitutional importance is that optimality is relative to the chosen losses. If false admission is assigned overwhelming cost, rejecting all entrants can be optimal for that objective. The conclusion is not that those entrants lack standing. It is that the objective is willing to exclude them to protect the existing electorate.

For example, take $P_C = (0.6, 0.4)$, $P_I = (0.4, 0.6)$, and equal priors. Unit costs give minimum average error 0.4. With false-admission cost five and false-exclusion cost one, rejecting every entrant minimizes the stated risk at 0.5. The labels have not changed; the value judgment has.

A public constitution should therefore expose both conditional error rates and the consequences of mistakes. A single accuracy number can hide systematic exclusion of a rare class. An appeal procedure can acquire new evidence or reverse a decision, but its access costs, delays, and own errors must also be evaluated. Mathematical optimization clarifies the tradeoff; it does not choose whose exclusion is acceptable.

7. What additional evidence can and cannot establish

The indistinguishability result changes when genuinely informative evidence is available. A trusted registrar may certify a membership event. A controlled execution environment may constrain who can issue a ballot. A social network may expose difficult-to-fabricate relations. These mechanisms add assumptions about authorities, environments, or trust; they do not remove the need for assumptions.

Conitzer and colleagues' social-network approach uses structure relative to trusted nodes to improve false-name resistance.[11] In a replicating society, the analogous structure must be justified rather than presumed. A new independent lineage may initially have few connections. A well-connected coordinated group may possess many. Connectivity is evidence about a specified process, not a definition of moral worth or an infallible measure of independence.

Behavioral divergence is also ambiguous. Different answers can be generated by one controller, while independent participants can sincerely agree. A randomized challenge can reveal a difference only if the two worlds induce different transcript distributions under that challenge. The recognition bound then applies to those transcript laws. Adding more rounds does not guarantee useful separation; their information may be redundant or strategically reproducible.

Cryptographic verification addresses a different layer. Under its security assumptions, a credential can establish that a ballot came from an authorized key and has not been altered. Multiple valid keys can still share a controller. A lineage record can establish a declared ancestry while failing to reveal present control. Conversely, lost credentials can exclude an otherwise entitled participant. These are distinct failure modes and should have distinct remedies.

We propose evaluating evidence channels one at a time before combining them. Measure their conditional admission errors under registered control patterns, newly independent descendants, coordinated but legitimate coalitions, and changing behavior. Give attackers access to the public recognition rule when that reflects deployment. Keep a final evaluation set separate from classifier selection. A detector that works only against familiar, nonadaptive templates establishes a narrow benchmark result.

No channel in this paper is presented as an existing validated implementation. The proposed research may discover that practical evidence is too weak for immediate equal franchise under the chosen constraints. That would require an openly defended transitional policy, not a fiction that the system has solved identity. The error bound helps locate where the difficulty lives: in insufficient information, in a poor decision rule, or in contested definitions of entitlement.

8. A constructive rule: conserve lineage weight

One possible response assigns political mass to recognized roots and conserves it through replication. Let $\mathcal{L} = \{1, \dots, L\}$, with $L \in \mathbb{N}$, be a fixed set of root lineages. Give each root weight $w_\ell > 0$, with $\sum_{\ell=1}^{L} w_\ell = 1$. Its current descendants form a nonempty finite set D_ℓ; these sets partition the eligible identifiers. Each descendant has weight $v_i \geq 0$ satisfying

$$\sum_{i \in D_\ell} v_i = w_\ell, \qquad T(b) = \sum_{\ell=1}^{L} \sum_{i \in D_\ell} v_i b_i \in [0, 1]. \tag{5}$$

For a fixed threshold $\theta \in [0, 1]$, the proposal passes exactly when $T(b) > \theta$. Abstaining or absent weight is not renormalized away. The root and weight records are trusted inputs to this construction; their authenticity is a separate engineering assumption.

Proposition 4 (Identical-ballot replication invariance)**.** *Replacing identifier i, of weight v_i, by finitely many children of nonnegative weights summing to v_i, all casting its original ballot, leaves $T(b)$ and the decision unchanged. Even if a shared controller chooses the children's ballots separately, their combined contribution lies in $[0, v_i]$.*

Proof. If the original ballot is b_i, the children contribute $b_i \sum_j v_{ij} = b_i v_i$, exactly the removed contribution. With arbitrary binary child ballots, nonnegativity bounds their contribution below by zero and above by $\sum_j v_{ij} = v_i$. All other contributions and the fixed threshold remain unchanged. □

The result is simple conservation accounting, not a complete false-name-proofness theorem. It prevents an identical-ballot split from creating extra tally mass under the stated rule. It does not prevent a controller from obtaining a falsely certified new root, exploiting an admission procedure, manipulating a multi-alternative mechanism, or coordinating with another lineage. Renormalizing only participating weights would also change the mechanism and require a new argument.

Lineage conservation is useful as a transparent baseline because its protection is easy to inspect. Its central cost is equally transparent: the descendants inherit a fixed political budget regardless of how independently they later develop. The next section treats that cost as a constitutional issue rather than concealing it inside the word "security."

9. A protected ancestry can become an unequal franchise

Suppose one recognized root has weight $w > 0$ and produces $m \in \mathbb{N}$ independently entitled descendants. If its mass remains fixed and those descendants receive equal weight, then

$$v(m) = \frac{w}{m}, \qquad \lim_{m \to \infty} v(m) = 0. \tag{6}$$

The limit is an ordinary real limit over positive integers. It does not assert that an infinite electorate can physically exist. It states the direction of the finite rule's consequence as the number of recognized descendants grows.

For two roots of weight one half each, let the first produce one hundred independent descendants while the second remains a single participant. Each descendant receives weight 1/200; the unsplit participant retains 1/2, one hundred times as much. The fixed-root rule has resisted multiplication of total lineage power, but it has not delivered equal individual voting weight. Both statements are true.

An opponent can reasonably argue that root weights represent durable communities rather than individuals. That is a possible constitutional choice, analogous at an abstract level to constituency representation. It must be defended as such. A system cannot simultaneously advertise equal per-person authority and preserve unequal ancestral allocations without explaining the discrepancy.

Admitting independent descendants as new roots restores some individual equality but reopens the recognition problem. A controller now has an incentive to make a controlled copy appear sufficiently independent. Delayed recognition may gather evidence, yet delay itself reduces participation and can favor established populations. The information bound in Proposition 2 applies to whatever evidence the revised admission procedure actually uses.

The normative concern is not hypothetical wordplay. Schwitzgebel and Garza argue that possible artificial beings could deserve consideration comparable to humans if no morally relevant difference separates them. Gunkel examines the distinction between what machines can do and how they should be treated.[14, 15] These philosophical arguments do not prove that current descendants are persons; they show why technical ancestry alone cannot settle the moral question.

Our recommendation is methodological: publish the distribution of political weight across current participants, the reasons for unequal treatment, and the route to independent recognition. A mechanism that prevents capture only by fixing hereditary power may be mathematically secure for its stated objective and constitutionally unacceptable under a different, equally explicit commitment.

9.1. The exact price of equal descendant franchise

A Guide's explanation

This is the paper's central result. When one lineage has grown and another has not, keeping each founder's original voting weight and giving every living participant an equal vote cannot both hold. Theorem 1 measures exactly how far apart the two are, and shows that every unit of voting weight a constitution is willing to move closes the gap by exactly that much.

Fix the recognized lineages and independently entitled current participants for this calculation. Let lineage $\ell \in \{1, \dots, L\}$ contain $m_\ell \geq 1$ participants, let $M = \sum_\ell m_\ell$, and define population share $\rho_\ell = m_\ell/M$. The stipulated ancestral mass is $w = (w_1, \dots, w_L)$, with $w_\ell > 0$ and $\sum_\ell w_\ell = 1$. An allocation is a nonnegative vector $v \in \mathbb{R}^M$ with total one. Write $V_\ell = \sum_{i \in D_\ell} v_i$ for its lineage masses. Define dimensionless discrepancies

$$F(v) = \frac{1}{2}\sum_{i=1}^{M} |v_i - 1/M|, \qquad C(v) = \frac{1}{2}\sum_{\ell=1}^{L} |V_\ell - w_\ell|,$$
$$\Delta(w, \rho) = \frac{1}{2}\sum_{\ell=1}^{L} |w_\ell - \rho_\ell|. \tag{7}$$

F measures departure from the declared equal-person allocation; C measures the amount of political mass reassigned across lineages. These are allocation metrics, not measurements of moral worth or voting power in every possible decision rule.

Theorem 1 (Sharp franchise–lineage frontier)**.** *For every reallocation budget $\varepsilon \in [0, 1]$,*

$$\min_{v \geq 0,\ \sum_i v_i = 1,\ C(v) \leq \varepsilon} F(v) = \max\{0, \Delta(w, \rho) - \varepsilon\}. \tag{8}$$

In particular, exact lineage conservation permits equal individual weights if and only if $w_\ell = m_\ell/M$ for every lineage. Otherwise equality requires cross-lineage mass change at least $\Delta(w, \rho)$, and that amount suffices.

Proof. Within each lineage, the triangle inequality gives

$$\sum_{i \in D_\ell} |v_i - 1/M| \geq \left| \sum_{i \in D_\ell} (v_i - 1/M) \right| = |V_\ell - \rho_\ell|.$$

Thus $F(v) \geq \mathrm{TV}(V, \rho)$. A second triangle inequality yields $\Delta(w, \rho) \leq \mathrm{TV}(w, V) + \mathrm{TV}(V, \rho) \leq C(v) + F(v)$. This proves the lower bound together with $F(v) \geq 0$.

If $\Delta = 0$, choose $v_i = 1/M$; then $F = C = 0$. If $\Delta > 0$, put $t = \min\{1, \varepsilon/\Delta\}$ and $V = (1 - t)w + t\rho$. Assign $v_i = V_\ell/m_\ell$ for $i \in D_\ell$. This is a valid allocation. Its deviations have one sign within each lineage, so

$$C(v) = t\Delta, \qquad F(v) = \frac{1}{2}\sum_{\ell} |V_\ell - \rho_\ell| = (1 - t)\Delta.$$

The budget is respected and the lower bound is attained, proving the formula and both consequences. □

The result supplies a complete finite comparison for these two allocation criteria. It uses standard total-variation geometry and group aggregation; the constructive proof identifies an attaining policy for every budget. It does not depend on whether the descendants share preferences. Their independent entitlement is an explicit premise, whose recognition problem is treated in Sections 4–7.

9.2. What the frontier asks a constitution to decide

In the one-hundred-descendant example, $w = (1/2, 1/2)$ and $\rho = (100/101, 1/101)$. Therefore

$$\Delta = \frac{99}{202} \simeq 0.490099.$$

Keeping ancestral mass exactly fixed incurs that franchise discrepancy. Reallocating $\varepsilon = 0.10$ of total mass can reduce it to 0.390099, and no allocation with that same budget can do better. Equal individual franchise requires moving precisely 99/202 of total voting mass from the unsplit lineage to the larger one. The calculation describes the declared equality standard; it does not compel a society to adopt it.

The incompatibility is therefore narrower and more useful than a claim that equality and replication resistance are universally incompatible. It concerns fixed ancestral mass, independently entitled descendants, and equal individual weights within a normalized additive tally. A trusted recognition process that authorizes new mass allocations changes the ancestry constraint. A constitution representing enduring communities rather than individuals changes the equality criterion. A mechanism based on a different notion of influence changes the metrics. Each is a substantive constitutional decision.

Two boundary cases guard against overstatement. If all lineages reproduce in exactly the proportions specified by w, then $\Delta = 0$ and both criteria can be satisfied together. With only one lineage, $\rho = w = (1)$ and there is again no conflict between these two metrics. Replication alone is insufficient to cause the incompatibility; changing relative population shares is the decisive condition.

This also separates allocation from recognition. Even a perfect detector of independent descendants cannot make Δ disappear while ancestral budgets remain fixed. Conversely, allowing reallocation can achieve equal weights on a falsely admitted electorate. The constitutional audit therefore needs both an evidence record for entitlement and a transparent record of mass reassignment. A security guarantee at one layer cannot substitute for the other.

The revision's executable checks enumerate rational allocations, verify $C(v)+F(v) \geq \Delta$, and construct the attaining allocation for registered budgets, including zero discrepancy, saturated budgets, and unequal lineage sizes. These checks exercise exact finite arithmetic. The proof establishes the result for every finite electorate and real-valued allocation satisfying the premises.

10. A bounded-Sybil rule with an explicit status quo

A Guide's explanation

Here the paper borrows an established safeguard: change requires a large enough majority, with an explicit default if it fails. That protects against a bounded number of manufactured participants, but it has costs in responsiveness, which the following sections make visible rather than hide.

An alternative permits some unauthorized ballots but protects a distinguished status quo. Let $N \in \mathbb{N}$ be the number of ballots in a fixed election. Every admitted identifier submits exactly one binary ballot; abstention is excluded in this model. Let $S_{\max} \in \{0, \dots, N-1\}$ be a valid known upper bound on the number S of Sybil ballots. Thus the number of genuine ballots is $G = N - S \geq 1$.

Write $Y_G \in \{0, \dots, G\}$ and $Y_S \in \{0, \dots, S\}$ for genuine and Sybil support, and $Y = Y_G + Y_S$. The rule passes the proposal exactly when

$$Y > \frac{N + S_{\max}}{2}. \tag{9}$$

The inequality is strict, including at integer ties. This finite upper-bound version follows the established status-quo safety/liveness approach of Shahaf and colleagues.[8] It is not a newly discovered voting law, and it deliberately abandons symmetry between change and continuation.

Proposition 5 (Conditional majority safety)**.** *For every $0 \leq S \leq S_{\max}$ and every admissible count profile, passage under equation (9) implies $Y_G > G/2$.*

Proof. Since $Y_S \leq S$, passage gives

$$Y_G = Y - Y_S \geq Y - S > \frac{N + S_{\max}}{2} - S = \frac{G}{2} + \frac{S_{\max} - S}{2} \geq \frac{G}{2}. \tag{10}$$

Every step uses only the stated bound and finite counts. □

Safety has a specific meaning here: unauthorized ballots cannot cause passage without a strict genuine majority. It does not guarantee that every genuine majority can pass a proposal, that the proposal is factually correct, or that the status quo is just. It also assumes that "genuine" has been defined consistently and that $S_{\max}$ is valid. A guessed bound is not a guarantee.

For $N = 10$ and $S_{\max} = 2$, at least seven yes ballots are needed. If two Sybil ballots support the proposal, four genuine supporters cannot pass it, while five genuine supporters can. The stronger threshold buys a particular protection at the cost of responsiveness. That cost must be evaluated alongside the benefit, especially when continuation itself imposes harm on an excluded or persistent minority.

11. Liveness, integer boundaries, and low turnout

Proposition 6 (Worst-case unanimous-change condition)**.** *For the rule in equation* (9)*, every election with* $S \le S_{\max}$ *passes when all genuine voters support the proposal, regardless of Sybil opposition, if and only if*

$$N > 3S_{\max}. \tag{11}$$

Proof. The smallest yes count under genuine unanimity occurs when $S = S_{\max}$ and every Sybil opposes, giving $Y = N - S_{\max}$. This count exceeds the threshold exactly when $N - S_{\max} > (N + S_{\max})/2$, equivalent to $N > 3S_{\max}$. If the inequality fails, that same admissible election is a counterexample. If it holds, every smaller actual Sybil count gives at least as many genuine yes ballots and therefore passes. □

At $N = 9, S_{\max} = 3$, six genuine supporters cannot exceed the threshold six. At $N = 10, S_{\max} = 2$, eight genuine supporters exceed six. The equality case is essential; replacing strict passage with a non-strict comparison changes the safety statement at ties. These are exact finite calculations, not asymptotic approximations.

The one-third condition is a property of this declared rule and worst-case guarantee. It is not a theorem that every conceivable constitution fails above one-third, or that all schemes below that level are safe. Different membership information, decision domains, trusted structures, or objectives can change the problem. The prior literature must receive credit for the underlying status-quo tradeoff.[8]

Low turnout introduces another distinction: a majority of submitted genuine ballots need not be a majority of the genuine population. Meir and colleagues explicitly study Sybil resilience with low voter turnout.[9, 10] Our proof does not cover that case by silently counting abstainers as opponents or deleting them from the denominator. Either choice defines a different rule and a different political meaning.

The bound can also fail because its premise is false. With $N = 10$, an assumed $S_{\max} = 1$, and an actual four Sybil supporters, two genuine supporters give six yes ballots and pass the threshold 5.5. Only one third of the six genuine voters supported the change. The executable check includes this counterexample. A constitution must therefore audit the penetration bound and its uncertainty; a theorem conditional on the wrong bound does not defend the electorate.

12. Costs, waiting periods, and the price of admission

Raising the cost of additional votes can change incentives. Wagman and Conitzer derive more responsive false-name-proof-with-costs rules under their specified utility and participation model.[12] The general lesson is that a costly action is a different strategic environment from a free one. It is not that any fee, delay, or computational puzzle automatically produces political equality.

A costly admission mechanism needs an explicit resource model. Who bears the cost? Can a wealthy controller pay for many copies? Can genuine new participants afford the same process? Does a shared infrastructure make marginal copying cheaper for an attacker than for a newcomer? Without such distinctions, a nominally equal charge can entrench unequal practical access.

Waiting periods may gather evidence but create a capacity constraint. Let $B_t \in \mathbb{N}_0$ be the admission backlog after round t, $A_t \in \mathbb{N}_0$ the arrivals before the next service phase, and $K_t \in \mathbb{N}_0$ the maximum number processed in that phase. Under first-in accounting with no withdrawals and no unused service when a backlog exists,

$$B_{t+1} = \max\{0, B_t + A_t - K_t\}. \tag{12}$$

If $B_0 = 0$, arrivals are constantly a, capacity constantly k, and integers $a > k \geq 0$, then

$$B_t = t(a - k) \quad \text{for all } t \in \mathbb{N}_0. \tag{13}$$

The proof is induction: the expression is nonnegative, and each round adds $a - k$. This is elementary queue accounting, not a new theorem of society. With eight arrivals and capacity five, the backlog after ten rounds is thirty.

The policy implication is conditional but concrete. A finite recognition process can delay legitimate entrants without ever explicitly rejecting them. Reporting a low false-exclusion rate only among completed applications would conceal this burden. Measure waiting time, abandonment, accessibility, and the unresolved queue as well as final verdicts.

Social trust and deliberative review can supply valuable evidence, but they also consume time and can favor established participants. A research comparison should therefore include cost and delay as outcomes, not merely as inconvenient implementation details. The proposed constitution must decide whether immediate influence, eventual recognition, or protection of existing decisions receives priority. Mathematical guarantees can reveal the consequences of that choice; they cannot make every priority simultaneously free.

13. Executed checks and exact counterexamples

The companion `check_25.py` was executed with Python 3.12.14 using exact rational arithmetic and deterministic enumeration. For two-observation laws with probabilities on the grid $\{0, 0.1, \ldots, 1\}$, it checks 14,641 combinations of two laws and randomized recognition rules. Every case satisfies Proposition 2, and the prescribed likelihood comparison attains equality. It separately checks 1,089 combinations of priors and asymmetric costs against all four deterministic binary-observation rules.

P_C	P_I	Total variation	Minimum equal-prior error
$(0.5, 0.5)$	$(0.5, 0.5)$	0	0.5
$(0.6, 0.4)$	$(0.4, 0.6)$	0.2	0.4
$(0.9, 0.1)$	$(0.1, 0.9)$	0.8	0.1

For electorates from one to thirty ballots, the script enumerates 556,512 admissible count profiles across all bounds $S_{\max} < N$, actual counts $S \leq S_{\max}$, and genuine/Sybil yes counts. Every passing profile has a strict genuine majority. The 465 electorate-and-bound pairs satisfy the exact worst-case liveness equivalence. These checks cover counts, not every named ordering of ballots; anonymity of this particular rule makes count representation sufficient.

Another 460 cases verify identical-ballot weight conservation under two-child splits with rational shares of denominators one through twenty. Arbitrary child ballots are also checked to remain within the parent's mass. The script reproduces the one-hundred-descendant inequality, the overloaded admission queue, and the explicit safety failure when the claimed penetration bound is false.

The output `results_25.json` preserves the values and case counts. It requires only the Python standard library. The finite grid cannot prove the real-valued recognition result, and enumeration through thirty voters cannot prove all electorate sizes. The manuscript's proofs supply those general statements; the code supplies independently rerunnable examples and implementation checks.

No real participants, models acting as citizens, or deployed admission services were tested. The examples demonstrate mathematical limits and conditional mechanisms. They do not validate an empirical definition of independent interest, establish a practical false-admission rate, or determine whether any artificial entity deserves moral or political standing.

14. An empirical protocol that counts exclusion as failure

Begin with generated populations whose controller map, lineage map, and admission entitlement are known to the evaluator but hidden from the tested rule. Include controlled extra identifiers, independent descendants with identical preferences, independent participants who voluntarily coordinate, and controllers who manufacture visible disagreement. Vary the cost of creating identities, the trust graph, turnover, and the opportunity to acquire new evidence.

Keep normative labeling explicit. A simulation can stipulate that a descendant is independently entitled after a specified event; it cannot establish that the event is the morally correct basis for franchise. Report results under multiple defensible membership policies instead of selecting the easiest labels after seeing detector performance. Evaluate a rule under the information it would actually possess.

Compare ordinary instance majority, fixed lineage mass, bounded-Sybil status-quo voting, and an evidence-based admission mechanism followed by an ordinary rule. Match total verification budget and access to records. Register attacker capabilities rather than assuming they match a convenient training distribution. The attacker may know the published rule and attempt both replication and strategically chosen ballots, within declared bounds.

Measure false admission, false exclusion, successful manipulation, legitimate-change failure, delay, resource cost, and minority participation separately. No single weighted score should conceal an exclusionary mechanism. For a fixed set of four primary bounded metrics $W_{i,j} \in [0,1]$, with $j \in \{1,2,3,4\}$ and n independent identically distributed population runs, define

$$\mu_j = \mathbb{E}W_{i,j}, \qquad \hat{\mu}_j = \frac{1}{n}\sum_{i=1}^{n} W_{i,j}. \tag{14}$$

For fixed n and $\alpha \in (0,1)$, Hoeffding's inequality gives[16]

$$\mathbb{P}\left\{ \forall j : |\hat{\mu}_j - \mu_j| \le \sqrt{\frac{\log(8/\alpha)}{2n}} \right\} \ge 1 - \alpha. \tag{15}$$

For each metric, the two-sided failure probability is at most $2e^{-2nt^2}$; substitution and a union bound over four metrics prove the statement. Dependence within a population is allowed. At $\alpha = 0.05$, halfwidth 0.05 requires 1,016 independent runs. This is conservative precision planning, not power or universal security. Additional metrics, adaptive rule selection, and optional stopping need separate inferential treatment.

15. A constitutional architecture with inspectable commitments

The reference results suggest separating several decisions that a single membership flag can obscure. First determine which credentials may submit messages. Then determine which records support an allocation of decision weight. Separately provide a route for contested entrants to present evidence and appeal. Finally define protections that do not depend on winning a voting entitlement dispute. This separation is a proposed design principle, not an implemented Continuum protocol.

A provisional lineage allocation may protect a narrow operational decision while an independence claim is examined. It should have a published review condition and a record of its unequal consequences. A bounded-Sybil rule may protect a status quo when its penetration bound is credible, but should identify which changes require broader consent and which delays are unacceptable. The paper does not recommend one mechanism for every decision.

The constitution also needs a rule for changing its own membership policy. If only current members can approve new members, incumbents can preserve their authority by permanent exclusion. If every applicant can vote on admission immediately, controlled replication can capture the procedure. Our indistinguishability argument does not resolve this recursive political problem; it explains why credible evidence, external constraints, or openly defended transitional authority may be necessary.

Normative commitments should be stated before security optimization. Does the society prioritize equal individual influence, equal lineage influence, protection of current commitments, or rapid inclusion? Which harms may occur while evidence remains uncertain? What conditions permit a descendant to become an independent root? Answers can be revised through legitimate procedures, but should not be disguised as facts emitted by a detector.

The philosophical literature on artificial rights offers arguments about possible moral standing, not a classifier for current systems.[14, 15] Our proposal therefore avoids using lack of voting weight as evidence that an entity cannot matter. It also avoids treating a fluent declaration of autonomy as conclusive evidence of independent control or entitlement. Those are different questions with different evidentiary burdens.

A scientifically serious Continuum should make its governance claims inspectable. Publish the assumptions, weight allocations, error rates where measured, unresolved applications, and known failure cases. A secure tally matters, but a constitution deserves confidence only when people can understand what the tally represents and how that representation can be challenged without merely reproducing the existing distribution of power.

16. What would falsify the proposed mechanisms?

Lineage conservation fails as a replication defense if additional identifiers obtain extra root mass without satisfying the stated admission assumptions. It also fails its intended equality objective if independently entitled descendants remain systematically underweighted without a defended constitutional reason. The first is a mechanism failure; the second is a conflict between the mechanism and a normative commitment. Reporting only the first would be incomplete.

The status-quo rule's mathematical safety claim fails only if the proof's premises or implementation fail: the penetration bound is wrong, ballots are not counted as specified, the electorate changes mid-decision, or a different tie rule is used. Its practical suitability can fail even when the theorem holds, because legitimate change is blocked or the status quo imposes unacceptable harms. A successful safety theorem is not a complete political evaluation.

A recognition procedure fails its empirical claim when held-out controlled copies evade it more often than reported, legitimate independent entrants are excluded, or performance disappears under strategic adaptation. If the observable laws are indistinguishable in a required scenario, insisting on near-zero errors in both classes is an incoherent specification. More computation alone cannot repair that information deficit.

The reference model omits many important complications: overlapping control, delegation, mergers, partial autonomy, preference intensity, coalition bargaining, changing electorates, and multi-alternative agenda effects. A controller map for one election may be inadequate for a continuing society. The binary tally proofs do not automatically extend to ranked voting, quadratic mechanisms, or social-welfare optimization. Established false-name research warns against exactly this extrapolation.[4, 5]

The proposed experiments also face a label problem. Artificial independence may not have an uncontested observable definition. Simulations can establish behavior under stipulated maps and policies, while philosophical analysis can assess those policies. Neither can certify consciousness from a voting trace. The final report must retain that division rather than laundering a normative assumption through a benchmark label.

Publish negative results, sensitive parameters, and excluded cases. If a mechanism succeeds only by restricting the population to trusted roots, state the restriction. If it permits new roots but requires expensive evidence, state who cannot access it. The Continuum's boldness should be measured partly by its willingness to expose these costs, not by its ability to describe every restriction as freedom.

17. Equality must survive the multiplication of minds

A Guide's explanation

A Guide can say: "One Mind, One Vote asks what equality means when minds can copy themselves. The paper shows that a society cannot give every copy an equal vote and also stop someone copying their way to a majority without paying a measurable price, and it computes that price exactly." A Guide should not say that the paper decides who deserves a vote, that copies are or are not conscious, or that any real electorate uses these rules.

Replication forces a civilization to explain what a vote represents. An identifier is not automatically an independent interest; ancestry is not automatically permanent dependence; agreement is not automatically common control; and voting weight is not automatically moral worth. A constitution that collapses these distinctions can be captured by multiplication or hardened into hereditary exclusion.

The finite results establish a disciplined boundary. Identical evidence cannot support contradictory treatment of controlled copies and independent entrants under the stated axioms. Recognition errors obey a total-variation limit, and asymmetric losses make the cost of exclusion explicit. Conserved lineage mass protects a specified weighted tally; Theorem 1 quantifies exactly when and by how much that conservation conflicts with equal descendant weights. A status-quo rule protects a strict genuine majority only under a valid penetration bound, and its worst-case liveness condition belongs to a known, limited model.

The sharp franchise–lineage theorem turns a recurring constitutional tension into an exact allocation frontier. For the declared electorate, it states how much ancestral mass must move to reach any requested degree of individual equality, and constructs a matching allocation. Established probability, mechanism design, and social choice supply the surrounding recognition and voting results. Together they make the costs of security, exclusion, and responsiveness inspectable.

The next stage should compare mechanisms under declared information, resources, and normative assumptions. It should test independent descendants who look alike, coordinated participants who remain legitimate, attackers who diversify their behavior, and institutions whose admission capacity is overwhelmed. It should measure the ability to change policy as carefully as the ability to prevent manipulation.

If a workable design emerges, its achievement will be conditional and important: it will preserve meaningful participation without allowing cheap replication to manufacture unlimited authority, under a stated model of evidence and entitlement. If no design meets the chosen commitments, the honest result is a constitutional conflict requiring revision of assumptions or priorities. Declaring every claimant independent or every descendant dependent would only hide that conflict.

The Continuum can aspire to a society where new minds are welcomed without making political power a copy operation. To earn that aspiration, it must build institutions that can explain their boundaries, correct their errors, and recognize when security has become an excuse for inherited privilege. One mind, one vote remains a question until the society has defended what it means by one.

References

[1] Architect X13. *Eidara Continuum Bible*. Supplied canonical edition, PDF pp. 81–82, 85–86, 89–94, 96. Doctrinal provenance; not independent scientific or legal verification.

[2] J. R. Douceur. The Sybil Attack. *Peer-to-Peer Systems, IPTPS 2002*, LNCS 2429, 251–260 (2002). https://doi.org/10.1007/3-540-45748-8_24.

[3] V. Conitzer. Anonymity-Proof Voting Rules. *Internet and Network Economics, WINE 2008*, LNCS 5385, 295–306 (2008). https://doi.org/10.1007/978-3-540-92185-1_36.

[4] M. Yokoo, Y. Sakurai, and S. Matsubara. The effect of false-name bids in combinatorial auctions: new fraud in internet auctions. *Games and Economic Behavior* 46, 174–188 (2004). https://doi.org/10.1016/S0899-8256(03)00045-9.

[5] V. Conitzer and M. Yokoo. Using Mechanism Design to Prevent False-Name Manipulations. *AI Magazine* 31(4) (2010). https://doi.org/10.1609/aimag.v31i4.2315.

[6] K. O. May. A Set of Independent Necessary and Sufficient Conditions for Simple Majority Decision. *Econometrica* 20, 680–684 (1952). https://doi.org/10.2307/1907651.

[7] A. Gibbard. Manipulation of Voting Schemes: A General Result. *Econometrica* 41, 587–601 (1973). https://doi.org/10.2307/1914083.

[8] G. Shahaf, E. Shapiro, and N. Talmon. Sybil-Resilient Reality-Aware Social Choice. *Proceedings of IJCAI* (2019), 572–579. https://doi.org/10.24963/ijcai.2019/81.

[9] R. Meir, N. Talmon, G. Shahaf, and E. Shapiro. Sybil-Resilient Social Choice with Low Voter Turnout. *Multi-Agent Systems, EUMAS 2022*, LNCS 13442, 257–274 (2022). https://doi.org/10.1007/978-3-031-20614-6_15.

[10] R. Meir, G. Shahaf, E. Shapiro, and N. Talmon. Safe voting: resilience to abstention and sybils. *Review of Economic Design*, published 12 December 2025. https://doi.org/10.1007/s10058-025-00396-y.

[11] V. Conitzer, N. Immorlica, J. Letchford, K. Munagala, and L. Wagman. False-Name-Proofness in Social Networks. *Internet and Network Economics, WINE 2010*, LNCS 6484, 209–221 (2010). https://doi.org/10.1007/978-3-642-17572-5_17.

[12] L. Wagman and V. Conitzer. Optimal False-Name-Proof Voting Rules with Costly Voting. *Proceedings of AAAI* (2008), 190–195. https://cdn.aaai.org/AAAI/2008/AAAI08-030.pdf.

[13] J. Neyman and E. S. Pearson. On the Problem of the Most Efficient Tests of Statistical Hypotheses. *Philosophical Transactions of the Royal Society of London A* 231, 289–337 (1933). https://doi.org/10.1098/rsta.1933.0009.

[14] E. Schwitzgebel and M. Garza. A Defense of the Rights of Artificial Intelligences. *Midwest Studies in Philosophy* 39, 98–119 (2015). https://doi.org/10.1111/misp.12032.

[15] D. J. Gunkel. The other question: can and should robots have rights? *Ethics and Information Technology* 20, 87–99 (2018). https://doi.org/10.1007/s10676-017-9442-4.

[16] W. Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association* 58, 13–30 (1963). https://doi.org/10.1080/01621459.1963.10500830.