When Remote Voting Beats Paper: A Construction and Quantitative Criterion for Making Elections More Trustworthy and Accessible

Greg Magarshak
IE University (NYC) and Safebots, Inc.
greg@magarshak.com

Abstract

Paper ballots give a secret ballot and custody-based integrity but no real-time result; electronic voting gives a checkable count but has never delivered secrecy and integrity that survive a dishonest operator, which is why expert consensus keeps it from high-stakes use. We present a stack that holds all three at once: a secret ballot, integrity verifiable without trusting the operator, and a live tally. The protocol combines blind-signature credentials, per-ballot commitments, a Merkle bulletin board, a verifiable tally, a threshold-trustee embargo, and receipt-free, coercion-resistant repairs, then hardens against the one threat cryptography cannot touch: a compromised client device. We organize the analysis around one quantity, the Lipschitz constant L of the aggregation rule: verifiability, secrecy, and coercion-resistance are independent of L; any distortion is bounded by L times the corrupted fraction; and the rational-attack question reduces to a computable predicate $V \cdot L < C/\rho$. A dual-root construction, an attested environment plus an independent device gated by on-device biometric presence, drives the undetected-corruption fraction ρ to subverting both roots, so the safe regime is where the cost of that exceeds the prize, and where the predicate fails the honest output is paper. The consequence changes a relationship: anyone can check the count, which removes whole categories of dispute and shifts citizens from trusting an authority's tally to verifying their own ballot, while remote casting enfranchises those who cannot reach a polling place. What keeps credible remote voting out of reach is not the cryptography but the absence of a criterion for when it is responsible; the predicate supplies it. The result: a person verifies what an institution does rather than trusting it, with a computable boundary for where the substitution is safe.

1 Thesis and relation to the consensus

This is a construction paper. It does not refute a theorem; it builds a voting system and states, as a computable predicate, the regime in which that system is safe to deploy. The distinction matters because the "impossibility" the field asserts is a worst-case statement, and worst-case statements are routinely mistaken for typical-case requirements. Shannon's bound forbids compressing every input below its entropy, yet zip wins on almost every real file; the CAP theorem forbids consistency and availability during a partition, yet eventual consistency runs most of the internet, because partitions are rare. In each case a true worst-case limit was read as a ban on the typical case, and the engineering came from serving the common case while ceding only the sparse hard corner. Remote voting sits in exactly this confusion. The expert verdict, that internet voting cannot be secured, is a statement about the hardest instance: a single high-stakes national office, decided by a discontinuous winner-take-all rule, attacked by an adversary who owns the voter's device. It has been heard as a verdict on remote voting as such. This paper builds the system that works for

the vast space of elections that are not that instance, gives the predicate $V \cdot L < C_{\mathrm{TCB}}/(2\rho_{\max})$ that says precisely which side of the line a given election is on, and shows that even the hardest instance is moved across the line by a continuity reform a single state can enact (Section 9). The one adversary that survives in the hard corner is named once, bounded by a theorem (Section 10), and shown to be a coordinated covert collusion of two of the largest technology institutions on earth, an event with no precedent. Naming a rare worst case is not conceding the typical one, and the two must not be allowed to blur.

Remote electronic voting has a near-universal expert verdict against it for high-stakes government elections, and the verdict is correct: client malware can alter or leak a vote before any cryptography touches it, and end-to-end verifiability lets a diligent voter *detect* this but not *prevent* it, so a large-scale fraud may be witnessed by only the few who check. We do not contest this for the government case. It is worth stating in the consensus's own words how categorical it is, because our claim lives precisely in the gap between what that consensus actually establishes and the stronger impossibility it is often read to assert. The National Academies' definitive 2018 study concludes that no known technology can guarantee "the secrecy, security, and verifiability of a marked ballot transmitted over the Internet," and that the Internet "should not be used for the return of marked ballots" [1]. The position has not softened with time: a 2026 Princeton analysis records that it has been "the scientific consensus for decades that internet voting is not securable by any known technology" [4], the scientific association AAAS states there is "no technical evidence that any internet voting technology is safe or can be made so in the foreseeable future" [5], and a review summarizing the field for MIT puts it simply: running "a secure election over the internet is not possible today" [6]. The U.S. Vote Foundation states the structural reason most sharply: election officials "must always know exactly who is voting" to verify eligibility, yet "must not be able to trace particular ballots" to voters, and it holds that this pairing "cannot be met reliably and securely with any online system" [7].

That last formulation names the bind exactly, and it is the bind our construction is built to dissolve. We are explicit about what we do and do not claim, because the distinction is the entire point. We invent no new cryptographic primitive; blind signatures, verifiable mixing, threshold decryption, and remote attestation are all standard and pre-existing. What the consensus treats as impossible is not any one of these primitives but the *system* that would have to hold secrecy, operator-independent integrity, and a usable result together, on hardware an attacker can compromise, without a trusted authority. Our claim is narrow and, we argue, exactly calibrated: the threat is a function of how much an attacker gains, and that gain is governed by two levers the system controls, the value of the outcome and the Lipschitz constant of the aggregation rule, and by one barrier the construction raises, the cost of defeating two independent trust-roots at once. The consensus is right that the threat cannot be made to vanish; we show it can be *bounded*, computed, and gated, so that there is a stated, checkable regime in which remote casting is responsible and outside which the honest output is paper. The impossibility is real as a blanket claim and false as a structural one, and the predicate $V \cdot L < C/\rho$ marks the boundary between.

The stakes of getting this wrong are not abstract, and they are not confined to obviously troubled elections abroad. Confidence in the count is eroding inside the world's most established democracies. A national survey by the University of California, San Diego's Center for Transparent and Trusted Elections found that the share of Americans confident their votes would be counted accurately fell from 77% just after the 2024 presidential election to 60% a year later, a drop of seventeen points that spanned every party, Republicans, Democrats, and independents alike [2]; a Harvard Youth Poll the same spring found only a third of Americans aged 18 to 29 believed the 2026 midterms would be conducted fairly [3]. The institutional response to this distrust has been to retreat toward paper, which is slow, cannot give a result the public can watch accumulate,

and, crucially, gives an individual voter no way to confirm after the fact that their own ballot was counted as cast. That societies accept paper's slowness and its uncheckability anyway is itself the strongest evidence of the consensus: the field treats verifiable, secret, device-resilient electronic voting at scale as so far out of reach that the rational fallback is a centuries-old technology no voter can personally audit. This paper's combination, a ballot each voter can check, a tally anyone can verify, and a result available continuously, is aimed squarely at the gap that distrust opens, and the predicate is what says where it is responsible to close it.

The reason this matters beyond the individual guarantees is the combination. A paper ballot delivers a secret ballot and a form of integrity, but it cannot report a continuously updated result, and its integrity rests on chain-of-custody rather than on anything a voter can personally check. Prior electronic voting delivers speed and, with end-to-end verifiability, a check on the count, but it has not simultaneously delivered a convincing secret ballot and an integrity guarantee that survives a dishonest operator, which is exactly why the expert consensus keeps it out of high-stakes use. The system assumed here, together with its companion, holds all three at once, a combination we are not aware of prior systems achieving: ballot secrecy from the protocol of Section 3 (with demographic confidentiality added by the companion reporting layer); cast-as-intended integrity that does not require trusting the operator, from the attested dual-root path; and a continuously available tally, from voting on a standing quantity rather than a single closing event. The predicate $V \cdot L < C/\rho$ is what makes that combination responsible rather than reckless: it is the system stating the regime in which remote electronic casting reaches an assurance high enough to stand in place of supervised paper balloting, and refusing that mode outside it. Where the predicate fails, the honest output is not a weaker electronic election; it is the instruction to run that contest on paper. The contribution is therefore not any one of the three properties, each of which has prior art, but a single system that achieves the trifecta and knows the boundary of its own validity.

This paper presents the protocol layer in full (Section 3): unlinkable per-dimension credentials, an encrypted pre-ballot A re-randomized by the casting layer to a public ballot B, a Merkle bulletin board, a designated-verifier inclusion receipt, a verifiable mixnet or homomorphic tally, a threshold-trustee embargo, and fake credentials for coercion-resistance. We prove verifiability and ballot secrecy under stated trust, show the base construction is not receipt-free, and give two repairs that recover receipt-freeness and coercion-resistance with their costs, all under the explicit assumption of an honest client. The rest of the paper removes that assumption as far as it can be removed and characterizes the remainder. The companion paper, *When Exact Reporting Beats Adding Noise*, takes the results this stack produces and reports demographic breakdowns of how groups voted while keeping each voter unlinkable, spending differential-privacy budget only where a cell needs it; it is the downstream reporting layer, not a prerequisite.

The consensus layer is the Intercloud chilling-effect protocol (arXiv:2605.22830): independent append-only logs secured by Watcher swarms, global consensus reduced to a single random scalar per epoch, and per-stream security allocated in proportion to $\sqrt{\text{value at stake}}$ rather than a flat global maximum. Intercloud secures the recorded ballots and the tally after ballots exist; it is the continuity principle applied to the backend, and it composes with the aggregation-rule continuity below.

2 Threat model

The parties are voters; a registrar; one election authority per demographic dimension; casting and tally infrastructure running in attested enclaves; an append-only, non-equivocating bulletin board; and any public verifier. We assume collision-resistant hashing, semantically secure re-randomizable

encryption, unforgeable blind/anonymous credentials, a designated-verifier proof system, a hardware root of trust supporting remote attestation, and the registration and one-honest-mix assumptions of Section 3.

We classify client-side adversaries by what they can corrupt and at what cost:

- **Commodity.** Malware in an installed app, a tampered client, a compromised updater, spyware. Cheap, scales to many voters, cannot forge a hardware attestation quote.
- **Insider / single-vendor.** A compromised SDK, a coerced employee with signing access, or one device vendor acting maliciously. Moderate cost; controls one trust-root.
- **Trusted-computing-base owner.** The owner of a device's OS, baseband, or attestation root, able to read rendering and memory and to sign measurements. Can defeat one root entirely; cost is the strategic exposure of a covert capability, denoted C_{TCB}, which is large and fixed.

We judge the cumulative protocol. A coercer interacting live with a voter is treated by the coercion-resistance repair of Section 3; here the adversary's goal is silent, scalable corruption or leakage of the vote at the client. We state the residual explicitly in Section 11 and do not claim to defeat a silent, colluding owner of all relevant trust-roots.

3 The voting protocol: verifiability, privacy, and coercion

Before analyzing client integrity, we set out the voting protocol itself: how a ballot is formed, recorded, and tallied so that any voter can check it counted and no voter can prove to a coercer how they voted. This is the stack the rest of the paper hardens against a compromised client and prices against the aggregation rule. The construction composes standard primitives: blind-signature anonymous credentials (Chaum), per-ballot commitments, a Merkle bulletin board for inclusion proofs, and a verifiable mixnet (Neff) or homomorphic tally for the count. Composition of sound primitives does not imply the composite has the claimed security properties; those are statements about all adversaries and must be stated as games and either proved or broken. We do both: two of the three target properties hold under stated trust, and the third, as the construction is described, is *false*. The break is elementary, it matches the known status of comparable systems, and it has a standard fix whose cost we state.

3.1 Protocol parties and trust assumptions

The threat model of Section 2 classified *client-side* adversaries. The protocol layer additionally involves back-end parties, with their own trust assumptions. The parties are voters; a registrar issuing credentials; one election authority per demographic dimension; a set of mix servers (or threshold-decryption trustees); an append-only, consistent bulletin board; and any public verifier. We assume collision-resistant hashing, a binding (and, where stated, hiding) commitment, an unforgeable blind signature, at least one honest mix server (or an honest threshold of trustees), and a bulletin board that cannot equivocate. The adversary controls the network, may corrupt a minority of mix servers and any number of voters, and — for the coercion properties, interacts with a target voter, demanding chosen behavior and any secrets the voter holds. We judge the cumulative protocol, not each message in isolation.

3.2 Three security games

Definition 1 (Ballot privacy, informally BPRIV-style)**.** A mechanism has ballot privacy if no efficient adversary, given the bulletin board and the published tally, distinguishes a world in which a target honest voter cast v_0 from one in which they cast v_1, beyond what the tally itself reveals, provided both worlds yield the same published result. Privacy is defined relative to the tally, since the tally is meant to leak.

Definition 2 (End-to-end verifiability)**.** A mechanism is end-to-end verifiable if (i) *individual verifiability*: an honest voter can check that their ballot is recorded on the bulletin board; and (ii) *universal verifiability*: any verifier can check that the announced tally is the correct function of the recorded ballots. Formally, if all checks pass then, except with negligible probability, the tally equals the count of recorded ballots.

Definition 3 (Receipt-freeness; coercion-resistance)**.** A mechanism is *receipt-free* (Benaloh–Tuinstra) if a voter cannot produce evidence that proves to a third party how they voted. It is *coercion-resistant* (Juels–Catalano–Jakobsson) if, additionally, a voter interacting with a coercer during the protocol can convincingly appear to comply while voting as they wish, including against demands to abstain or to surrender secrets.

3.3 What holds

Proposition 1 (Verifiability and tally privacy hold under the stated trust)**.** *With a non-equivocating bulletin board and Merkle inclusion proofs, the construction satisfies individual verifiability (Definition 2(i)); with a sound verifiable mixnet or homomorphic-tally proof, it satisfies universal verifiability (Definition 2(ii)); and with at least one honest mix server (resp. an honest decryption threshold) and a hiding commitment, it satisfies ballot privacy (Definition 1). These are the standard guarantees of the underlying primitives and transfer directly.*

Proof. Each clause is the soundness or hiding guarantee of one primitive, preserved under their composition. Individual verifiability is the soundness of Merkle inclusion proofs against a non-equivocating board: a voter accepts only on a valid path from their leaf to the published root, which a board that cannot equivocate cannot forge. Universal verifiability is the soundness of the verifiable mixnet (or the homomorphic-tally proof): a verifying transcript certifies, except with negligible probability, that the announced tally is the correct function of the recorded ballots, which is Definition 2(ii). Ballot privacy is the hiding of the commitment together with the secrecy of the anonymizing step, at least one honest mix server or an honest decryption threshold: a distinguisher in the game of Definition 1 between equal-tally worlds yields either a commitment distinguisher or a break of mix anonymity (resp. threshold secrecy), both negligible. The three primitives operate on disjoint parts of the transcript, so their parallel composition preserves each guarantee. □

This is the part that motivates the design and it is sound. The standing-to-complain property of the disclosure half (a voter can prove their commitment is missing from the Merkle root) is exactly individual verifiability and is real.

3.4 What fails: receipt-freeness

Theorem 1 (The construction as described is not receipt-free)**.** *Suppose the per-ballot commitment is $c = H(v \,\|\, s \,\|\, r)$ for vote v, salt s, nonce r, with the voter retaining (v, s, r) and the Merkle path from c to the published root. Then the mechanism is not receipt-free: there is an efficient coercer strategy that verifies a coerced voter's vote.*

Proof. The coercer demands the opening (v, s, r) and the Merkle path. The voter, to comply, supplies them; the coercer recomputes $H(v \,\|\, s \,\|\, r)$, checks equality to the leaf c, and checks the path from c to the root. This is a verifiable proof that the ballot recorded as c carries vote v. A voter wishing to claim a different vote $v' \neq v$ must exhibit (v', s', r') with $H(v' \,\|\, s' \,\|\, r') = c$, i.e. a second preimage of c; collision resistance makes this infeasible. Hence the opening is a receipt for the vote's content, and the receipt-freeness game is lost. □

Remark 1 (The mixnet does not repair this)*.* The mixnet anonymizes the link between a recorded commitment and a slot in the final tally, which defeats a coercer who would trace a published commitment forward to the count. It does nothing to the opening attack, which reads the vote *out of the commitment itself*, before any mix. The essay's "shuffling destroys the link" argument is sound against the forward-tracing attack and silent against the opening attack; the property it secures is weaker than the property it claims. This is the same gap that makes Helios explicitly not coercion-resistant, in contrast to Civitas / JCJ.

3.5 Two repairs, with their costs

Proposition 2 (Repair A: withhold openings, obtain receipt-freeness, lose cast-as-intended)**.** *Let commitments be perfectly hiding and let the opening randomness be threshold-shared among the trustees and never given to the voter, who receives only the Merkle path to their hiding commitment. Then no opening exists for the voter to surrender, and the mechanism is receipt-free against a coercer who lacks the opening (Definition 3, first clause). The cost is in verifiability: the voter can verify their ballot is* recorded *(cast-as-recorded) but can no longer verify its* content *was captured as intended (cast-as-intended), since they hold no opening to check against. Receipt-freeness and cast-as-intended verifiability are traded against each other.*

Proof. Receipt-freeness. In the receipt-freeness game the coercer fixes $v_0 \neq v_1$, the voter casts v_β, and the coercer, given the voter's entire retained view, guesses β. Under this repair that view is the hiding commitment $c = \mathrm{Com}(v_\beta; r)$ with r held only by the trustees, together with the Merkle path from c to the public root. The path is a deterministic function of c and the public tree, so it carries no information about v_β beyond c. A distinguisher for β is therefore a distinguisher for $\mathrm{Com}(v_0; \cdot)$ versus $\mathrm{Com}(v_1; \cdot)$, whose advantage is at most the commitment's hiding advantage, negligible by assumption. Hence the coercer's advantage is negligible and the mechanism is receipt-free against a coercer lacking the opening.

Loss of cast-as-intended. The voter's verification procedure takes only (c, path) and checks membership of c under the root; call its output $\mathsf{ver}(c)$. Consider two executions that record, respectively, an encoding of the intended v^* and of a substituted $v' \neq v^*$. By the hiding property their commitment distributions are computationally indistinguishable, so the induced distributions of (c, path), and hence of ver, differ by at most a negligible amount. No voter-side function of the view can therefore separate "recorded v^*" from "recorded v'" with non-negligible advantage; the voter verifies cast-as-recorded but not cast-as-intended. The same hiding property that delivers receipt-freeness is what removes cast-as-intended, so the two cannot be held simultaneously from the voter's view, which is the stated trade. □

Proposition 3 (Repair B: fake credentials, obtain coercion-resistance)**.** *Replace the single blind-signature credential per dimension with a JCJ-style credential admitting indistinguishable* fake *credentials: the voter can manufacture fakes that a coercer cannot tell from the real one, and ballots cast under fake credentials are silently discarded during tallying by credential-validity tests (plaintext-equivalence tests) that never reveal which ballots were dropped. Then, assuming a coercion-resistant*

registration phase (one interaction from which the coercer is absent) and at least one honest mix server, the mechanism is coercion-resistant in the sense of Definition 3: under coercion the voter surrenders a fake credential and later votes with the real one, and the coercer cannot detect the substitution.

Proof. The construction is the Juels–Catalano–Jakobsson (JCJ) credential mechanism, instantiated once per demographic dimension, so coercion-resistance reduces to two components.

Credential indistinguishability. A coercer who demands "the credential" receives either the real credential or a voter-manufactured fake. Distinguishing these is exactly distinguishing the underlying credential encryptions, so a coercer that detects the substitution with advantage δ yields a distinguisher against the credential-hiding (IND-CPA) assumption with the same advantage; thus the detection advantage is negligible.

Tally-side soundness and privacy. Ballots cast under invalid (fake) credentials are removed by plaintext-equivalence tests (PETs) against the encrypted registration roll. A PET reveals only the equal/unequal bit, not the matched index, and is simulatable, so the removal leaks no information about which ballots were dropped beyond their count; with at least one honest mix server the correspondence between submitted ballots and the anonymized tally is unlinkable. These are the JCJ tally-phase guarantees, which we inherit verbatim (Juels, Catalano & Jakobsson, WPES 2005, coercion-resistance theorem; instantiated as in Civitas, IEEE S&P 2008).

Composing, under a coercion-free registration phase (a single interaction, with the coercer absent, that issues the real credential) and at least one honest mix server, the coercer's advantage in the game of Definition 3 is bounded by the credential-indistinguishability advantage plus the PET/mixnet simulation error, both negligible. The voter surrenders a fake under coercion and later votes with the real credential, undetectably. Both assumptions are necessary in the strong sense: a coercer present at registration learns the real credential, against which no fake helps; and a coercer controlling every mix server breaks the unlinkability the PET argument relies on. These are precisely the JCJ/Civitas trust assumptions, imported unchanged. □

Remark 2 (Composition with the disclosure layer is clean). The demographic separation of Sections 1–8 and the coercion repair here are orthogonal. The unlinkable per-dimension credentials are the same blind-signature objects in both analyses; Repair B's fake-credential mechanism applies to each dimension's election independently, and the companion paper's data-only separation equivalence is untouched because it depends only on the published marginals, not on how credentials are issued or tested. The precise headline for the combined system is therefore: tally privacy and end-to-end verifiability hold; receipt-freeness requires Repair A (at a verifiability cost) or coercion-resistance requires Repair B (at a registration-phase cost); and the demographic-privacy guarantees of the first half are independent of which repair is chosen.

3.6 Embargo and the reveal schedule

Every reporting rule presumes a moment of reveal. Because the board holds only encrypted ballots, nothing about the tally exists in the clear until a decryption is performed: partial or running results do not leak, in pointed contrast to the incremental counting of mailed ballots. We make the reveal a controlled event with standard machinery. The decryption key is shared among N trustees under a distributed (threshold) key generation, so unsealing requires a quorum of M to combine shares, a cryptographic embargo in which no single trustee, and no coalition smaller than M, can produce results before the scheduled reveal. Where the reveal is an on-chain action, the same quorum is an on-chain multisignature authorizing publication. The threshold-tally construction is not new (Cramer–Gennaro–Schoenmakers; Benaloh; and the Helios/Civitas line); the point here is to use

it to make disclosure *timing* a deliberate parameter rather than an accident of when counting happens.

What the reveal publishes. The quorum decrypts only aggregates, the homomorphic tally, or the output of a verifiable mix, never a ballot linked to a voter. A reveal therefore publishes the binding outcome (Tier 0) exactly and the demographic overlays under the per-cell budget of the companion paper's per-cell disclosure budget, and nothing that links a decrypted vote to an identity. A voter then checks two public facts: that their ballot is included on the board, and that the published tally is the correct decryption of the board (the threshold decryption and the mix or homomorphic aggregation carry verification proofs). Together these establish that the voter's ballot was counted without giving the voter, or a coercer, any proof of *how* they voted. Verifying-that-it-counted never becomes proving-how, because the check is over inclusion and a universal tally proof, not a per-ballot decryption.

Ongoing elections reveal on epochs. An ongoing or continuous election has no single close, so the reveal runs per epoch: ballots are committed during an epoch and the quorum unseals that epoch's aggregate after it ends. The epoch length, the cadence of reveals, and how much demographic detail each reveal carries are disclosure-design parameters, and they trade transparency against feedback. A faster cadence and a richer demographic reveal give the public more information and also more opportunity for results-in-progress to shape later voting — bandwagon, strategic abstention, herding, the same effect that motivates result embargoes in conventional elections, here under explicit control because the encryption means intermediate results need not exist at all. Each epoch's reveal spends from the per-voter budget and composes across epochs as in the companion paper's budget-composition rule, so the reveal schedule is part of the disclosure design rather than an afterthought to it.

Limits of the embargo. The embargo's strength is exactly the threshold assumption: a colluding quorum of M trustees can unseal early, so it is as strong as the trustee set is independent and no stronger, and the choice of (M, N) trades this against liveness, since fewer than M available trustees cannot unseal at all. The embargo governs confidentiality and timing, not integrity, the tally's correctness rests on the verification proofs, not on the secrecy of the reveal.

Close trigger, vote cutoff, and a turnout fallback. Two timing questions remain: when does voting close, and what forces the reveal if a faction of trustees refuses to unseal. The close is set by the schedule or, for an ongoing election, by the epoch boundary; an oracle or the bulletin board's clock attests that the deadline passed. The refusal problem is primarily a liveness concern, and the direct remedy is the threshold itself: choosing $M = \lceil (N+1)/2 \rceil$ lets an honest majority unseal even if a minority defects, the standard robust setting. A turnout threshold can serve as an additional fallback. Turnout is observable *without* decryption, since the board publicly shows the count of committed ballots, so a rule of the form "once a target number of valid ballots is committed, the reveal is authorized" is checkable by anyone. The honest limitation is that a turnout trigger *authorizes or conditions* a reveal; it does not *perform* the decryption, which still requires the M-of-N quorum, so turnout cannot replace the quorum. Moreover, if a turnout threshold unsealed results while voting were still open, it would reintroduce the partial-results leakage the embargo exists to prevent; the trigger is therefore clean only at the close or at full participation. Used this way, turnout is a fallback authorization at the deadline, not a second decryption path.

4 Eligibility registration: who may deny the franchise

The protocol above assumes a registrar that issues each eligible voter an unlinkable credential. That issuance carries the same two jobs the companion attribute paper [8] identifies for its enrollment, checking a predicate (is this person eligible to vote) and issuing an unlinkable credential, and it admits the same two instantiations. But the choice between them is governed by a different threat model, and the default reverses. We state both and the axis that selects them, because the wrong default here is not an inefficiency but a disenfranchisement surface.

The two rails. The first is *threshold eligibility issuance*: an M-of-N committee certifies eligibility and blind-signs the credential, so no coalition smaller than M can issue a credential to an ineligible voter and, crucially, no coalition smaller than $N - M + 1$ can *deny* a credential to an eligible one. The second is the *credential-gated pool* of the companion paper: where an eligibility attestation already exists as a ZK-provable signature (a state voter-registration credential, or an organization's membership credential), the voter proves possession in zero knowledge, emits a one-per-credential registration nullifier, and self-inserts into a membership tree, and the issuer sees nothing after issuance. Both deliver the same downstream property, unlinkable ballots from eligible voters; they differ in who can silently exclude a voter at registration.

Why the default reverses from the attribute paper. For age, the companion paper makes the credential-gated pool the preferred rail, because the state already holds the holder's birthdate and is not the holder's adversary on that fact; reusing an existing credential costs nothing new. For the franchise, the same move concentrates the power to decide who may vote in a single registrar, which is exactly the party with the strongest incentive and the longest history of shaping an electorate by shaping its roll. The threat model is the inverse of the attribute case, and it is the same inversion the two-device analysis draws against banking (Section 6): where the protected asset is the individual's own, a birthdate, a bank balance, a single existing issuer is acceptable because that issuer is not the adversary; where the protected thing is a collective outcome and the issuer may wish to tilt it, the power to issue must be distributed. Hence the default reverses: for contested public elections the threshold committee is the construction of record, because distributing the power to deny is the point; the credential-gated pool is the lighter option for low-contest settings, an organization, a cooperative, a municipality with an uncontested roll, where a single registrar is trusted not to exclude and the efficiency of reusing an existing credential is worth having. The companion attribute paper's untaggability analysis sharpens why: its Lemma on registration shows that a credential-gated pool with an inflation-resistant (one-per-person) bound leaks *fact-of-enrollment* to the issuing authority, harmless when the authority is not your adversary about your birthdate, but for voting the authority is exactly the potential adversary and the enrollment it would learn is your eligibility, so the leak is disqualifying and the committee, which distributes that knowledge across a quorum, is required. The two papers therefore reach opposite defaults from a single rule: reuse one existing issuer when it is not your adversary about what it learns, distribute issuance when it is.

	Threshold committee	**Credential-gated pool**
Can a single party deny eligibility	no (needs $N-M+1$)	**yes (the registrar can)**
Can a single party be compelled to deny	no (two jurisdictions)	**yes**
Post-issuance linkage	none below M-of-N	none (issuer sees nothing)
Reuses an existing roll/credential	no	yes
Deployable with today's primitives	yes	needs ZK-provable credential
Default for a contested public election	**primary**	fallback (low-contest only)

The eligibility roll is an upstream trust root neither rail dissolves. Both rails issue credentials *against a determination of who is eligible*, and that determination, the voter roll, the citizenship or membership list, is a trust root the cryptography does not create and cannot launder. Distributing issuance across a committee stops a single party from denying a credential to someone the roll already includes; it does not adjudicate who belongs on the roll. A pool gated by a state credential inherits whatever inclusion or exclusion the state's roll already encodes. We state this plainly as the boundary of what the construction claims: it protects the secrecy and integrity of ballots cast by voters the roll recognizes, and it makes the issuance of voting credentials auditable and, under the committee, resistant to single-party denial; it does not resolve who is on the roll, which is a civic question upstream of any cryptographic protocol and outside what this system settles. A deployment in a setting where the roll itself is the contested object should not read this construction as answering that contest.

5 The aggregation layer: Lipschitz continuity

Let the normalized tally be $t \in \Delta$, a distribution over candidates, and let the outcome be a map $f : \Delta \to \mathcal{O}$ from tallies to results (allocations, winners, seat counts). The aggregation rule's robustness is its modulus of continuity.

Definition 4 (Aggregation Lipschitz constant)**.** f is *L-Lipschitz* if $\|f(t) - f(t')\| \le L \, \|t - t'\|_1$ for all $t, t' \in \Delta$.

Proposition 4 (Distortion bound; holds against every adversary)**.** *If an adversary corrupts a fraction ρ of ballots arbitrarily, the tally moves by $\|t - t'\|_1 \le 2\rho$, and hence the outcome moves by*

$$\|f(t) - f(t')\| \;\le\; 2L\rho.$$

This bound assumes no defense of any kind; it is a property of f alone.

Proof. Reassigning one ballot moves at most one unit of mass out of one coordinate and into another, changing $\|t\|_1$-measured mass by at most $2/n$ for n ballots; over ρn ballots the ℓ_1 movement is at most 2ρ. Apply L-Lipschitzness. □

Remark 3 (Winner-take-all is the degenerate case)*.* Winner-take-all is a step function: near a margin it is discontinuous, so no finite L exists and Proposition 4 is vacuous, a hair's-width corruption flips the result, which is both the Florida 2000 failure and the reason a close winner-take-all outcome is contestable whether or not anyone cheated. Proportional allocation is the identity-like map with $L = 1$, giving distortion $\le 2\rho$: corrupting 1% of votes moves the outcome by at most $\sim 2\%$. Continuity bounds, rather than prevents, the *consequence* of every defense failing, and it is the only guarantee in this paper that assumes no adversary model at all.

6 The client-integrity construction

The protocol layer forms a ballot on the voter's client, which is exactly where a client adversary intervenes. The construction splits the act of forming and confirming a ballot across two independent trust-roots and makes the bulletin board reject any ballot not jointly endorsed by both.

Construction 1 (Attested dual-root casting)**.**

1. **Pristine environment.** The voter loads the casting client into an attested enclave (a Safebox/TEE-style measured environment). The enclave produces the encrypted pre-ballot $A = \mathrm{Enc}(pk_E, v; a)$ for the selected candidate v, together with a designated-verifier proof π of the statement $S(A, B, R)$: "A encrypts the displayed candidate, $B = \mathrm{ReRand}(A; \delta)$, and $H(B)$ lies under root R," and a remote-attestation quote q over the enclave's measurement and over $H(A)$.

2. **Independent confirmation.** The enclave renders (q, A, π, B) as a QR code. The voter scans it with a device on an *independent trust-root* (different vendor / OS than the enclave host). The second device (i) verifies the attestation quote q, confirming the payload was formed by the pristine measured code; (ii) verifies π against the voter's designated-verifier secret, confirming A encrypts the named candidate and B re-randomizes it into the tree; and (iii) displays the candidate name *ephemerally* for the voter to confirm against intent.

3. **Dual co-signature.** On confirmation, both devices sign $H(A)$: the enclave signs σ_{enc} (bound to q), the second device signs σ_{dev} (bound to the verified statement S). The bulletin board accepts B *only* if it carries both signatures over the same $H(A)$, alongside the credential proof and nullifier of Section 3.

Three properties of this construction are essential, and stating them as invariants rather than caveats is what keeps the composition sound.

Design Invariant 1 (The confirmation is designated-verifier, never plaintext)**.** The second device signs the *designated-verifier statement* S, not a plaintext assertion "this encodes Alice." Because π is simulatable, the co-signed ballot proves to a third party only that *some* independent device confirmed *some* statement; it is not a transferable proof of vote content. Violating this invariant — signing a plaintext confirmation, would mint exactly the receipt the protocol paper's coercion-resistance forbids. The same designated-verifier machinery thus serves two jobs at once: it is the malware cross-check and it preserves receipt-freeness.

Design Invariant 2 (Attestation certifies computation; the second device certifies intent)**.** A remote-attestation quote certifies *what code ran*, not *what the screen showed*. A compromised rendering or input path above the enclave can display one candidate while the enclave faithfully encrypts another. Attestation therefore cannot secure the intent-to-payload link by itself; the second device's independent verification of π and re-display is what closes that gap, and it is why the two roots are not redundant. The enclave secures the computation; the independent device secures the voter's intent.

Remark 4 (Device-bound presence on the confirming root). The companion attribute paper [8] binds each confirmation to a live present person via a key sealed in the device's secure element and gated by an on-device liveness check, producing a per-action signature over a verifier-chosen nonce. Applied to the confirming device here, the same primitive binds the designated-verifier confirmation to a fresh, non-replayable presence proof: a captured confirmation cannot be replayed into a later ballot, and a remote synthesis of the voter cannot stand in for the on-device gate.

This strengthens cast-as-intended by tying the second root's signature to a present voter, while leaving the receipt-freeness of the designated-verifier statement intact, and it inherits that paper's boundary, presence-binding does not prevent a voter from voluntarily handing over an unlocked device.

Design Invariant 3 (Both signatures bind the same payload)**.** σ_{enc} and σ_{dev} are both over $H(A)$, the same hash the credential proof binds. Without this, a man-in-the-middle could pair a valid enclave signature on A with a device signature confirming a different A'. Binding both to one $H(A)$ makes the credential, the attestation, and the confirmation refer to a single ballot or fail to verify.

6.1 Relation to second-device cast-as-intended verification

The second device is the established route to cast-as-intended verifiability. Guasch and Morillo [19] formalized challenge-and-cast on an auditing device; Escala, Guasch, Herranz and Morillo [18] made the check universally verifiable under a trusted registrar; and Müller and Truderung's CAISED [25] gives a protocol-agnostic second-device audit with strong security guarantees, deployed in real elections. The mechanism in Construction 1 sits inside this lineage and does not improve on its cryptographic audit. Two things are added rather than the proof technique. First, the verifying root is an *attested* environment, so the second device checks a hardware measurement (Section 6.2) in addition to a designated-verifier statement, which moves part of the trust from "the voter ran an honest app" to "the app's image hash matches a published, independently rebuilt one." Second, the integrity requirement is stated against the aggregation rule's Lipschitz constant L and a per-voter payoff bound, so the deployment predicate $V \cdot L < C_{\text{TCB}}/(2\rho_{\max})$ (Theorem 3) says when a given election is safe to run remotely at all, independent of which cast-as-intended mechanism fills the second-root slot. CAISED or either Guasch–Morillo variant can occupy that slot; the contribution here is the economic predicate around it and the attested instantiation of the root, not a new audit proof.

6.2 Instantiation: what the attestation actually measures

The abstract enclave of Construction 1 instantiates on a governed execution substrate (Safebox) whose attestation is not a single quote but four distinct measurements, each stopping a different adversary. Stating them concretely is what keeps the scope of Theorem 2 precise, because they do not all cover the same thing.

Measurement	**What it certifies**	**Adversary it stops**
Content-addressed code + M-of-N auditor sigs	the bytes that ran hash to the approved value; auditors signed it	tampered / substituted casting client, bad updater
Hardware enclave quote (Nitro-style)	the host could not read or alter enclave memory or swap the code	malicious cloud operator / OS on the executor
Execution hash over code+ctx+input+result	the casting/tally computation is deterministically replayable	backend equivocation, miscount (universal verif.)
Two independent executor roots (§6.2)	enclave root and device root are distinct parties	single-vendor client compromise

The hardware binding, concretely. The enclave's authentication key is derived by HKDF from a verified Nitro attestation document over PCRs $0, 1, 4$, binding it to the (image, kernel, host) tuple: a tampered image or kernel changes PCR0/1, a disk copied to a different host changes PCR4, and either yields a different key, so a forged or relocated environment fails on its first request.

Interactive access (SSH, remote-management agent, TTY) is removed from the production image, so the code that runs is the built and audited image rather than something introduced by an operator at runtime. This is the operational meaning of "the audited code ran intact."

Empirical grounding still required. The argument above assumes the measured PCRs correspond to a *reproducible* build of the audited source and that the attestation client behaves as specified on real hardware. In the present implementation the mechanism is built but reproducible-build verification and real-hardware attestation validation are pending, as is independent audit. Until those are complete the hardware-attestation premise is a design assumption, not a verified fact, and we mark it as such rather than assert it. Separately, the host layer does not defend a compromised casting process: a runtime memory-safety failure in the audited code holds the derived key and the host will execute its requests, so the audited code must be both correct (the auditors' job) and memory-safe at runtime.

Identity is not correctness. The hardware quote and the content-addressing certify that *the audited code ran, unaltered*; they say nothing about whether the audited code is *correct*. Code correctness rests on the M-of-N auditor signatures over sha256(streamName $\|$ codeHash $\|$ publisherId), a governance assumption, not a hardware one. The enclave is "pristine" only in the conjoined sense: the hardware attests that the code ran intact *and* an honest auditor quorum approved that code. We therefore add the auditor quorum to the assumption list, and note that for the backend it should be independent of the election operator, exactly as the two client roots are independent of each other.

The enclave never sees the screen. The casting/tally enclave is a server-side executor; it never touches the voter's display or input path. This is not a limitation to apologize for, it is the concrete reason Invariant 2 holds. Attestation of a cloud enclave cannot, even in principle, certify what the voter saw, so the independent second device is doing work no strengthening of the enclave could do. The architecture thus confirms the non-redundancy of the two roots rather than merely asserting it.

Two genuinely independent roots are first-class. The substrate already distinguishes executor roots with distinct hardware anchors: a cloud enclave rooted in the cloud vendor's attestation authority, a mobile executor rooted in the device Secure Enclave, and a browser executor rooted in a WebAuthn authenticator. A deployment satisfies Theorem 2's independence hypothesis by drawing the two co-signers from two of these, for instance a cloud-rooted enclave (one vendor's attestation root) confirmed by a Secure-Enclave mobile device (a different vendor's root). The joint-subversion cost of Theorem 3 is then C_{TCB} for two unrelated vendors, not one.

Reproducible builds make the second root publicly checkable. Independence bounds $\Pr[\mathsf{Col}]$ by the product of per-root subversion probabilities (Theorem 4); a reproducible build drives the second-root factor down by converting "trust this vendor" into "anyone can rebuild and compare." When the second-device verifier, whether a Safebox-attested server that serves the page or a reproducibly-built Android application, is produced by a deterministic build from published source, any party can rebuild the artifact, hash it, and confirm the running code matches the audited source; the cloud root's Nitro PCR0 already binds the enclave image hash for exactly this third-party comparison. Subverting such a root is then not the quiet act of one vendor but a substitution that every independent rebuilder would detect, so the adversary must defeat the

reproducible-build check, the M-of-N auditor quorum, and the hardware attestation at once, and do so on two unrelated roots simultaneously. Each added check is another independent factor in the product of Theorem 4, pushing Col from unlikely toward absurd. We are explicit (Section 6.2) that reproducible-build verification is engineering still in progress in the present implementation; the point here is architectural, that the second root's trust is not a fixed vendor assumption but a quantity the deployer raises by making the artifact independently reconstructible.

App or website; device, not domain. The second root may be a native application or merely a web page in the second device's browser: a page can verify the attestation document, check the designated-verifier proof, and re-display the candidate using only standard browser cryptography, so no app-store distribution is required. What makes it an *independent* root is that it runs on a second *physical device under a different vendor's operating system and browser*. Two tabs in one browser on one phone, even on two different domains — share a single vendor's trusted base and are one root, not two; a compromised browser or OS sees and can alter both. The app-versus-website choice is orthogonal to independence and trades provenance for convenience: an installed app fixes the verifier's code at install time, whereas a web verifier is served fresh on each load and so additionally trusts the attested server that serves it. That is acceptable, and strongest, when the verifier is served by a second attested environment operated by a *different party* (for example the election's auditors) than the one that cast the ballot, giving a second root that differs in device vendor and in server operator at once, with no app-store dependence.

Why the two-device cost is justified here, and why it is a real cost. The second device is a genuine friction, and it is worth being exact about why a voter should accept it, because the same mechanism carries a different bargain in the setting it is most familiar from. Banking apps ask a user to confirm a transfer on a second device or channel, and the user accepts that friction readily, because the asset being protected is their own and concentrated: the attacker wants a few large targeted break-ins, the user and the bank are aligned in wanting that one account hardened, and the cost of the second step is borne by the same person who reaps its benefit. Voting inverts every term of that bargain. The value an attacker pursues is not in any one ballot but in the aggregate outcome, so the defense must cover a broad, correlated compromise of many clients at once; the benefit of the second device is therefore collective, accruing to the integrity of the result rather than to the individual performing the scan, while the cost, a second device, a scan, a moment of attention, is borne per voter. A mechanism whose benefit is collective and whose cost is individual is exactly the kind that must be kept minimal and must not become a gate on participation. This is why the second root is specified to run as an ordinary web page requiring no installation (above), why the accessibility fallback to a supervised or mail channel is retained for voters without a second device (Section 13), and why the companion attribute paper [8] can lean on device-sealed confirmation more heavily than this one: there the protected asset is the individual's own identity, as with banking, so the incentives realign and the per-user friction is self-interested rather than civic. The contrast is not decorative. It is the reason the same dual-root primitive is deployed conservatively for the franchise and can be deployed aggressively for a personal credential.

Verifier-code provenance and the SRI gap. "Root 2 honest" decomposes in deployment into two things: device isolation (a different vendor's sandbox) and *verifier-code provenance* (the running code is the genuine verification logic). A signed native application gives both, its code is fixed and platform-integrity-checked. A web verifier gives isolation but not code provenance, because current browsers enforce neither top-level nor iframe subresource integrity: there is no

browser-level guarantee that the loaded top-level document is the genuine verifier rather than a substituted one. The strongest no-server-trust verifier root today is therefore a signed native application, and the web verifier is sound only under the additional assumption that its serving environment is attested and independent (§6.2). A trusted browser extension could enforce a top-level integrity check; more usefully, a standard browser affordance, an indicator that the top-level document matches a published hash, analogous to the matching-symbol indicator some messengers display for a verified call, would let a web page serve as a first-class verifier root with no native app. We note this as the concrete browser feature that would remove the one provenance gap between the two verifier options.

Co-signature plumbing. The dual signature uses the substrate's ES256 (ECDSA P-256) signing with a trust-on-first-use public-key registry, and binds both signatures to $H(A)$ (Invariant 3). A freshness re-check at acceptance, confirming each co-signing key is still registered and still authorized at the moment the board accepts B, not merely at signing time, closes the gap in which a revoked key's old signature would otherwise still verify. Universal verifiability of the *backend* (the execution hash and replay) is separate from individual *cast-as-intended* (the designated-verifier proof and the second device); the former catches a lying tally, the latter catches a lying client, and conflating them would overstate what either provides.

Two attestation layers, not one. The host-attestation layer (the HKDF-from-PCR key above) authenticates the boundary between the host and the casting/tally process with a symmetric key, it answers "is this the attested image on the attested host." The ballot-level co-signatures (Invariant 3) are application-layer asymmetric (ES256) signatures over $H(A)$, they answer "did two independent roots endorse this ballot." The two compose but share no secret and no failure: breaking the host channel does not forge a ballot co-signature, and conversely. Per-request asymmetric governance signatures, so that each privileged backend action carries its own signature rather than riding the attestation-derived channel key, are a documented future upgrade and would tighten the host layer to match the ballot layer.

7 Security as a function of L

We separate the guarantees into those independent of the aggregation rule and the two that depend on it.

7.1 Properties independent of L

Proposition 5 (Inherited protocol guarantees)**.** *Under the cryptographic and registration assumptions of Section 3, and Construction 1 with Invariants 1–3, the system provides, at* any *L: individual and universal verifiability; ballot secrecy; and receipt-freeness and (with fake credentials) coercion-resistance under the corresponding repair. These are protocol properties; the aggregation rule does not enter their proofs.*

Proof. Each listed property is established for the base protocol in Section 3 (Proposition 1 for verifiability and ballot secrecy; Propositions 2–3 for the receipt-freeness and coercion-resistance repairs); we show Construction 1 preserves each, and that no proof invokes the aggregation rule, so all hold at any L. The construction's only additions to a ballot are the attestation quote q over $(m, H(A))$ and the two signatures $\sigma_{\text{enc}}, \sigma_{\text{dev}}$ over $H(A)$, plus acceptance checks on them.

Verifiability. The added checks only further restrict which ballots the board accepts; they do not alter the recorded-ballot set's relationship to the tally. Both individual verifiability (a voter checks inclusion of their leaf) and universal verifiability (any verifier checks the tally is the announced function of recorded ballots) are therefore inherited unchanged.

Ballot secrecy. $H(A)$ is a hash of the re-randomizable ciphertext and q is over $(m, H(A))$ for a public measurement m; none depends on v beyond what the public ballot B already exposes. A hybrid that replaces A by an encryption of an independent value changes the adversary's view only through quantities derivable from B, so any secrecy distinguisher reduces to one against the base protocol's IND-CPA ballot privacy.

Receipt-freeness. The only new artifact a coerced voter could surrender is σ_{dev} over $H(A)$. By Invariant 1 the confirming device signs the designated-verifier statement, not the plaintext, and that statement is non-transferable: the voter holds the designated key and can simulate a convincing confirmation for any chosen v'. Hence σ_{dev} binds v to no third party, and a coercer gains no verification advantage over the base protocol; receipt-freeness (under Repair A) is preserved.

Coercion-resistance. The credential layer (fake-credential mechanism of Repair B) is untouched by Construction 1, and the added confirmation is non-transferable by the previous paragraph, so the coercion-resistance of Section 3 composes without modification.

Since none of these arguments mentions the aggregation rule, every property holds for all L. □

7.2 Client integrity: formal model

We give the construction a game-based treatment and prove the cast-as-intended property to the same standard as the rest of the protocol.

Objects. Candidates $\mathcal{V}$. The voter's intent $v^* \in \mathcal{V}$ is the candidate the voter selects. The election uses a re-randomizable, IND-CPA encryption scheme with a public well-formedness predicate Valid; for $\mathsf{Valid}(A)$ the decryption $v_A = \mathsf{Dec}_{sk_E}(A) \in \mathcal{V}$ is unique. The pre-ballot is $A = \mathsf{Enc}(pk_E, v; a)$, the public ballot $B = \mathsf{ReRand}(A; \delta)$ with $v_B = v_A$. H is collision-resistant; R is the board's Merkle root. The designated-verifier relation is

$$\mathcal{R}: \quad (A, B, R, c) \in \mathcal{R} \iff \mathsf{Valid}(A) \wedge v_A = c \wedge B = \mathsf{ReRand}(A) \wedge H(B) \preceq R,$$

with proof π verified under the voter's designated key dv. The attestation quote q is over $(m, H(A))$ for measurement m; $\sigma_{\text{enc}}, \sigma_{\text{dev}}$ are signatures over $H(A)$, the latter emitted only with a verified statement.

Acceptance. $\mathsf{Accept}(B) = 1$ iff $\mathsf{Valid}(A)$; q verifies against the pinned root for the approved measurement m^* and binds $H(A)$; π verifies for (A, B, R, c_π) under dv; σ_{enc} and σ_{dev} verify over the same $H(A)$; and the credential proof and nullifier of Section 3 are valid.

Roots and honest behavior. *Root 1* is the enclave attestation key together with the approved-code (auditor) assumption: honest Root 1 runs approved code that, on the candidate it is given as input, outputs $A = \mathsf{Enc}(pk_E, \text{input})$, an honest π with $c_\pi = v_A$, a quote q, and σ_{enc} over $H(A)$; its honesty includes delivering the voter's true selection v^* to the enclave (an intact casting input path). *Root 2* is the verifier device key: honest Root 2 verifies q and π, displays c_π, and emits σ_{dev} over $H(A)$ only if both verify. *Voter model:* an honest voter does not submit unless the candidate displayed by Root 2 equals v^*.

Misencoding event. Mis holds when $\mathsf{Accept}(B) = 1$, the voter did not abort, and $v_B \neq v^*$.

Games.

Definition 5 (Three experiments)**.** Fix a security parameter λ, and let the attestation scheme, the designated-verifier proof system, the signature scheme, and the hash family H be instantiated at λ. Each experiment below is played by an adversary running in time $\mathrm{poly}(\lambda)$, and each advantage is a function of λ equal to the adversary's winning probability (over the experiment's coins) in that experiment; "negligible" means negligible in λ, i.e. eventually smaller than $1/\lambda^c$ for every constant $c > 0$.

- $\mathbf{Exp}^{\mathrm{att}}$: an adversary without Root 1's key outputs (q, m^*, h) and wins if q verifies for the approved measurement m^* and binds $h = H(A)$ for an A not output by a genuine instance of the approved code. $\mathrm{Adv}^{\mathrm{att}}$ is its winning probability.
- $\mathbf{Exp}^{\mathrm{dv}}$: an adversary outputs (A, B, R, c, π) with $(A, B, R, c) \notin \mathcal{R}$ and wins if π verifies under dv. $\mathrm{Adv}^{\mathrm{dv}}$ is its winning probability.
- $\mathbf{Exp}^{\mathrm{bind}}$: an adversary wins if it produces accepted data in which $\sigma_{\mathrm{enc}}, \sigma_{\mathrm{dev}}$ verify yet were not both produced by their respective roots over a single payload, i.e. it forges a co-signature or exhibits $A \neq A'$ with $H(A) = H(A')$. $\mathrm{Adv}^{\mathrm{bind}}$ is bounded by signature unforgeability plus the collision-resistance of H.

7.3 Client integrity: theorem and proof

Theorem 2 (Dual-root cast-as-intended, exact bound)**.** *For any PPT adversary that subverts at most one of Root 1 and Root 2, with the honest-voter and approved-code assumptions holding for the unsubverted party, and for every security parameter λ,*

$$\Pr[\mathsf{Mis}(\lambda)] \;\leq\; \mathrm{Adv}^{\mathrm{att}}(\lambda) + \mathrm{Adv}^{\mathrm{dv}}(\lambda) + \mathrm{Adv}^{\mathrm{bind}}(\lambda),$$

which is negligible in λ under the attestation, designated-verifier, signature, and collision-resistance assumptions. Consequently an accepted, voter-undetected misencoding requires either winning one of the experiments of Definition 5 (breaking a primitive) or subverting both *roots.*

Proof. Condition on none of the three experiments being won; the complementary event is bounded by $\mathrm{Adv}^{\mathrm{att}} + \mathrm{Adv}^{\mathrm{dv}} + \mathrm{Adv}^{\mathrm{bind}}$. Assume Mis. By acceptance there exist verifying σ_{enc}, σ_{dev} over a common $H(A)$; since $\mathbf{Exp}^{\mathrm{bind}}$ is not won, both were produced by their roots over one well-defined, Valid payload A (so v_A, hence $v_B = v_A$, is well-defined).

Suppose Root 2 is honest. Then σ_{dev} exists only if Root 2 verified π and displayed c_π; the honest voter did not abort, so $c_\pi = v^*$. Since π verifies and $\mathbf{Exp}^{\mathrm{dv}}$ is not won, $(A, B, R, c_\pi) \in \mathcal{R}$, whence $v_A = c_\pi$. Therefore $v_B = v_A = c_\pi = v^*$, contradicting Mis. Hence Root 2 is subverted.

Suppose Root 1 is honest. Honest Root 1 delivers v^* to approved code and emits σ_{enc} over $H(A)$ only for the A that code produced, so $v_A = v^*$ and $v_B = v^*$, contradicting Mis. (If only the casting input path were compromised while the enclave code and key stayed honest, then v_A equals the substituted input v' and the honest π has $c_\pi = v'$; with Root 2 honest, this is the previous case and the voter aborts. So a misencoding with Root 1's code honest still forces Root 2 subverted.) Hence Root 1 is subverted.

Both roots are therefore subverted, completing the dichotomy. □

Remark 5 (What each primitive actually carries). The cast-as-intended conclusion rests on designated-verifier soundness, an honest second root, and binding: these three force $v_B = c_\pi = v^*$. Attestation does not carry the intent link, it secures the *preconditions*: that A is well-formed (so v_A is defined and $\mathcal{R}$ applies), that the casting client did not leak v (ballot secrecy against a malicious client), and that the ballot was not relocated or substituted at creation. The well-formedness precondition can alternatively be discharged by the public Valid check, so the integrity theorem degrades gracefully if one discounts attestation entirely: it then rests on Valid, designated-verifier soundness, the honest second root, and binding. This is the scope of what the enclave buys for *integrity*, separate from its larger role in secrecy and in the economic argument.

Corollary 1 (Undetected-corruption fraction)**.** *Let ρ_{both} be the fraction of voters for whom the adversary jointly subverts both trust-roots. Then the fraction of accepted misencoded ballots is at most ρ_{both} up to negligible terms. Commodity and single-vendor adversaries have $\rho_{\text{both}} \approx 0$.*

Proof. By Theorem 2, conditioned on no primitive being broken (an event of all-but-negligible probability), an accepted ballot is misencoded only if both of its trust-roots are subverted. Summing the per-voter misencoding indicator over the electorate, the expected fraction of accepted misencoded ballots is at most the fraction with both roots subverted, ρ_{both}, plus the summed negligible advantages, i.e. ρ_{both} up to negligible terms. A commodity or single-vendor adversary cannot place two independent, separately provisioned roots under its control on the same voter, so $\rho_{\text{both}} \approx 0$. □

7.4 Economic security and the deployment predicate

Theorem 3 (No rational attack below threshold)**.** *Let V be the value of the outcome to the adversary, C the attack cost, and $\rho_{\max} \leq 1$ the largest fraction of voters at which the adversary can jointly subvert both trust-roots. Then:*

(a) *commodity and single-vendor adversaries ($\rho_{\text{both}} \approx 0$) have non-positive expected return and do not attack;*

(b) *the only adversary with $\rho_{\text{both}} > 0$ at scale is one subverting both trust-roots, at fixed cost C_{TCB}, who attacks only if*

$$\boxed{V \cdot L \;\geq\; \frac{C_{\text{TCB}}}{2\,\rho_{\max}}}$$

where $\rho_{\max} \leq 1$ is the largest corruptible fraction. The safe set *is the complement, $V \cdot L < C_{\text{TCB}}/(2\rho_{\max})$.*

Proof. By Proposition 4 any adversary shifts the outcome by at most L times the corrupted fraction, and by Corollary 1 the accepted-misencoding fraction is at most ρ_{both} up to negligible terms; valuing the outcome shift at V and applying the factor-two distortion bound, the achievable payoff is at most $2\,V\,L\,\rho_{\text{both}}$. A rational adversary attacks only when this exceeds C. For (a), $\rho_{\text{both}} \approx 0$ gives payoff $\approx 0 < C$, a non-positive expected return. For (b), raising ρ_{both} above zero at scale requires subverting both independent roots, incurring the fixed cost C_{TCB}; with the maximum corruptible fraction $\rho_{\max}$ the attack is rational iff $2\,V\,L\,\rho_{\max} \geq C_{\text{TCB}}$, i.e. $V \cdot L \geq C_{\text{TCB}}/(2\rho_{\max})$, whose complement is the safe set. □

Remark 6 (Why the two halves were always one picture). The predicate has the value V and the constant L on one side and the both-roots-subversion cost on the other. The aggregation layer (Section 5) and the Intercloud consensus layer lower the left side — proportional allocation makes L small, and pricing security to $\sqrt{\text{value}}$ keeps the recorded election secured in proportion to what

is at stake; the dual-root construction (Section 6) raises the right side by forcing joint subversion of two independent roots. Security is the region where the cost line sits above the prize, and every architectural choice in both papers moves one of the two lines. Continuity does not defend a vote; it shrinks the prize until no one rational attacks the votes that continuity has made cheap to corrupt.

8 Continuous voting on a standing quantity

The embargo and reveal schedule answer "when do we open the count." A different regime removes the question: a vote on a single ordered quantity that never closes. We develop it here because it is the limiting case of the continuity principle, the lowest-L aggregation rule there is, and because it sidesteps the social-choice impossibilities cleanly, under a stated and reasonable assumption.

Setup. Let the decision be one quantity on an ordered scale, the level of a weekly issuance, a budget line, a fee, a threshold. Each voter holds a current position and may move it by one unit, up or down, per interval; the standing outcome is the median of voters' positions. Take preferences to be *single-peaked*: each voter has an ideal level and prefers levels less as they move away from it in either direction. For a quantity like "how much to issue this week" this is the natural shape, not a contrived one.

Proposition 6 (Standing Condorcet winner and convergence)**.** *Under single-peaked preferences on a one-dimensional ordered domain, the median ideal point is a Condorcet winner: it is preferred by a majority to every other level in pairwise comparison, and no majority cycle exists (Black, 1948). The bounded ± 1-per-interval majority adjustment is a hill-climbing dynamics on this domain whose unique stable fixed point is the median: once the standing level is the median, a majority prefers neither $+1$ nor -1, so it rests there, and from any other level the majority moves it one step toward the median each interval. The median rule is moreover strategyproof on this domain (Moulin, 1980), so a voter at rest gains nothing by misstating position, which is what makes an always-open vote resistant to timing games.*

Proof. Let voters $i = 1, \dots, N$ have single-peaked preferences with integer peaks p_i on the ordered scale, and let m be a median of $\{p_i\}$, so $\#\{i : p_i \le m\} \ge N/2$ and $\#\{i : p_i \ge m\} \ge N/2$.

Median is a Condorcet winner. Take any level $y \ne m$; say $y > m$ (the case $y < m$ is symmetric). Any voter with $p_i \le m$ has $p_i \le m < y$, so by single-peakedness $|p_i - m| < |p_i - y|$ and the voter strictly prefers m to y. There are at least $N/2$ such voters, so m is preferred to y by at least half the electorate, and strictly more than half when N is odd or the median is unique. Hence m beats every alternative in pairwise majority comparison.

No majority cycle. On a single-peaked profile the pairwise majority relation ranks two levels $y < z$ by which is nearer the median: the voters preferring the nearer one are those whose peak lies on its side of the midpoint, a majority. This relation is the linear order of distance from m, hence transitive, so no Condorcet cycle exists (Black, 1948).

Convergence of the ± 1 dynamics. Suppose the current level is $x < m$ and the electorate votes between x and $x + 1$. A voter strictly prefers $x + 1$ iff $p_i > x + \frac{1}{2}$, i.e. $p_i \ge x + 1$. Since $x + 1 \le m$, the median property gives $\#\{i : p_i \ge x + 1\} \ge \#\{i : p_i \ge m\} \ge \lceil N/2 \rceil$, while $\#\{i : p_i \le x\} = N - \#\{i : p_i \ge x + 1\} \le \lfloor N/2 \rfloor$. For odd N (or unique median) this is a strict majority for $x + 1$, so the level rises by one; iterating, it increases monotonically until it reaches m. By symmetry it decreases to m from above. At $x = m$, the move to $m + 1$ is preferred only by $\#\{p_i \ge m + 1\} \le \lfloor N/2 \rfloor$, not a majority, and likewise $m - 1$ is not preferred; so m is a fixed point. The process thus reaches m in at most $|x_0 - m|$ steps and rests there, the dynamic counterpart of

m being the Condorcet winner. (For even N with a split median the same argument confines the level to the interval between the two central peaks, the usual median indeterminacy, on which no strict majority moves it.)

Strategyproofness. A voter cannot profit by misreporting a peak $p_i' \neq p_i$. The outcome is the median of the reported peaks; changing one report moves the median only when that voter is pivotal, and then only toward the misreport. A voter whose true peak is above the current median can pull the median up only as far as their report, never past it to a point they prefer less, and cannot pull it down; the best attainable outcome is thus achieved by reporting p_i truthfully, and symmetrically below. Truthful reporting is therefore a dominant strategy; the general characterization of strategyproof rules on single-peaked domains as generalized medians with phantom voters is Moulin (1980). □

This is the one regime in the paper that escapes *both* classical impossibilities, and it does so by domain restriction rather than by evading them: single-peakedness is exactly the restriction that breaks Arrow's universal-domain hypothesis (so a Condorcet winner exists), and it is the restriction on which strategyproof non-dictatorial rules exist despite Gibbard–Satterthwaite. The escape is real and named, not asserted.

Why polls need never close. A discrete election closes for two reasons: to fix the electorate at a decision instant, and to prevent results-in-progress from distorting later voting (the hazard the embargo of Section 3.6 exists to manage). A standing-quantity vote dissolves both. There is no winner to seal and no pivotal instant to snipe, because the outcome is a median that anyone may move by one unit and that re-converges; and the visible running level is not a contest with a finish line but the working state of the system. The leaderboard hazard (bandwagon, herding, strategic timing) that makes an always-open *ranking* dangerous is largely absent here for a precise reason: seeing the current level and nudging toward one's ideal is the mechanism operating as designed, not a distortion of a pending verdict. Visibility therefore plays opposite roles in the two settings, feeding a standing quantity toward its median while distorting a pending contest, so continuous open voting suits a one-dimensional single-peaked quantity and not a field of rival rankings.

What it enables, and at what cost. A standing-quantity vote lets a parameter track the electorate continuously rather than lurch at election dates: issuance levels, budget allocations, fees, thresholds. The participation cost per act is low (move one unit when one has a moment), and the process is self-correcting, drifting back if it overshoots. We state as a hypothesis, not a theorem, that removing the activation energy of a discrete event and the "too late, polls closed" disenfranchisement may raise participation; the same always-open visibility is what supports it. The standing cost is that continuous participation favors the continuously attentive: without a decay or staleness rule on inactive positions, the median reflects whoever is watching rather than the electorate. A decay rule (positions revert toward abstention or a last-confirmed value if not refreshed) is therefore part of the construction rather than an option, and its time constant is a disclosure-and-participation design parameter alongside the reveal cadence of Section 3.6.

Continuous voting requires continuous reporting. A vote that never closes is reported continuously, and continuous reporting is repeated release of the same quantities over time. This is the regime the composition analysis of the companion paper's tiered construction governs directly: each reveal interval is a fresh release of the standing median (and, where group breakdowns are published, of the per-group marginals), so the per-cell privacy loss accumulates across intervals at the rate set by the reveal cadence. Three parameters therefore move together and must be set

jointly: the reveal cadence of Section 3.6, which fixes how often a fresh release occurs; the per-cell budget of the companion paper's tiered construction, which must be divided across the releases a voter's cell participates in rather than spent once; and the decay constant above, which bounds how long any single position remains in the released aggregate and so bounds the number of releases over which one voter's participation composes. A coarse cadence and a decay window that retires stale positions keep the composed loss finite even over an unbounded horizon; a fast cadence with no decay does not, and is the configuration to avoid. Continuous operation does not escape the budget, it spends it over time, and the cadence and decay constant are how that spending is rationed.

If preferences over the level are genuinely two-peaked, the ± 1 dynamics can rest at a non-median fixed point or oscillate, and Black's guarantee lapses. If several quantities are voted jointly, the decision is multidimensional and a Condorcet winner generically fails to exist (McKelvey's chaos theorem), no matter how small the step. The clean regime is one single-peaked quantity at a time, median-aggregated, continuously adjusted; beyond it the impossibilities return and the construction should not be claimed.

9 Continuity and the government-scale boundary

The deployment predicate is a product, $V \cdot L$, so there are two independent routes into the safe set: lower the value V, which is the organizational case, or lower the Lipschitz constant L, which is a property of the aggregation rule and is available at any scale. The second route is what brings some government-scale elections into range, and it is worth stating precisely what it does and does not buy.

Continuity at scale. Proportional, approval, ranked-multi-winner, multi-seat, and policy-leaderboard rules have small L ($L = 1$ for proportional), removing the unbounded-L blowup that winner-take-all suffers near a margin. The effect on the attacker is structural, not merely quantitative: under winner-take-all the attacker flips the few pivotal votes in the pivotal jurisdiction, a concentrated and cheap operation; under a continuous rule, to move the outcome by δ the attacker must corrupt a $\Theta(\delta)$ fraction of *all* ballots, each requiring joint subversion of two independent roots, silently. Continuity converts a concentrated cheap attack into a broad expensive one, and breadth is the enemy of covertness, so it forces the adversary into precisely the regime the dual-root construction defeats best.

Its limit. The predicate is $V \cdot L$, not L alone. Proportional allocation fixes L and leaves V untouched, and at national scale V is large enough that even $L = 1$ leaves a genuine quantitative question rather than an automatic pass; the trusted-computing-base residual (§11) is most stressed exactly where V is largest. The resulting statement is that continuity moves a national election from *categorically outside* the safe set to *inside if and only if* a broad two-root attack at the required scale is infeasible or non-covert, a large shift, not a guarantee.

The cleanest fits are genuinely multi-winner. Legislatures elected by proportional representation or single transferable vote, ballot measures, apportionment, participatory budgeting, and policy rankings are continuous throughout, with no final winner-take-all step and lower V per decision; they sit inside the safe set on the merits. A single-winner executive office is the hardest case, and proportional *elector* allocation illustrates the boundary: it removes the single-jurisdiction pivotality that turned a 537-vote margin into a national crisis — the attacker can no longer take an entire bloc by flipping one close state, but a single office retains a final national argmax, so a tied

elector count (270–268) still turns on a small shift. Proportional electors reduce L substantially without eliminating it for a single seat, whereas a genuinely multi-winner body keeps L low end to end.

Two objections continuity does not address. Removing brittleness is not removing every objection the expert consensus raises to remote government voting. Continuity does nothing about *correlated systemic failure*, one client-side defect strikes the whole electorate simultaneously, unlike precinct-distributed paper that fails independently, nor about the high-V trusted-computing-base residual. Continuity removes the pivotality-and-dispute objection; it leaves the systemic-failure and the top-of-scale TCB objections in place. It is necessary for the government case, not sufficient.

A point that stands without electronic voting. The dispute-removal argument does not depend on adopting this system at all. Proportional allocation removes the winner-take-all class of dispute whatever the voting medium, because a small error or a small fraud moves the outcome by a correspondingly small amount, so a close result stops being an all-or-nothing crisis the loser can always plausibly contest. That is a reason to adopt continuous aggregation independent of going electronic, and it is what makes the continuity argument credible rather than self-serving: it would improve paper elections too.

9.1 A worked example: U.S. presidential elections

The framework is not only theoretical; one of its clearest instances is a live constitutional question. Within each state the prevailing winner-take-all rule maps the popular vote to electors as a step function, a plurality of one vote awards the entire slate, which is the unbounded-L map of Section 5 in the field. Its canonical failure is Florida in 2000, where a margin near 537 votes out of roughly six million decided the whole slate and with it the presidency: a hair's-width perturbation moving the outcome discontinuously, exactly the regime in which Proposition 4 is vacuous. The same discontinuity explains why a popular-vote winner can lose the office, as in 2000 and 2016, and, independently of the merits of any particular allegation, why thin-margin outcomes are so hard to lay to rest: when a small set of contested ballots can be outcome-determinative, any dispute over them is consequential and difficult to settle. A continuous allocation removes that leverage, because under it no small set of ballots is pivotal; the contestability is a property of the step function, not of the ballots.

The narrow reform, and what it is not. The continuity reform is within-state proportional or district allocation of electors — lowering L, and it must be distinguished sharply from two proposals it is often confused with. It is *not* abolition of the Electoral College, and it is *not* a change to how electors are *apportioned* among the states (the House-plus-two count that encodes the small-state weighting of the constitutional compromise). Apportionment is left exactly as it stands; only the within-state *division* of a state's fixed electors changes, from all-or-nothing to proportional. The small-state advantage, the federal structure, and the College itself all remain, this is a continuity reform, not a redistribution of weight between land and people.

It is already within a single legislature's power. The reform requires no constitutional amendment and no interstate compact. The Constitution commits the manner of appointing electors to each state legislature (Art. II, §1), and *McPherson v. Blacker*, 146 U.S. 1 (1892), held this power plenary and expressly inclusive of choice "by popular vote in districts" as well as "by general

ticket", upholding, in that very case, a state's switch to a district method. Maine (since 1972) and Nebraska (since 1992) allocate by congressional district today. The manner of appointment is so fully a legislative choice that most states once appointed electors directly, without any popular vote, and the Court described the power to set the manner as one the legislature may resume; popular voting for president is in that sense a creature of state statute. The within-state allocation rule is therefore a matter a single legislature can change by ordinary law.

Relevance to the construction. Two consequences connect the example to the construction. First, continuity reform stands on its own as election design: it would remove the Florida-2000 class of dispute under paper ballots, with no electronic component at all, which is why it is evidence that the Lipschitz account is operational rather than rhetorical. Second, it is the precondition that would move a presidential-scale contest from the unbounded-L region toward the evaluable side of the predicate of Theorem 3. The value V remains at its peak for a single national office, and a single-winner office retains the national argmax discontinuity of Section 9, so continuity is necessary, not sufficient, in this case, but it is the first and largest step, and the only one a state can take by itself.

10 The realistic threat landscape

The predicate $V \cdot L < C_{\mathrm{TCB}}/(2\rho_{\max})$ is often read in its worst case, a national winner-take-all office attacked by an omnipotent adversary, and that reading makes the system sound narrow. The worst case is the wrong default. The predicate is a map over a space of elections, and most of that space sits comfortably inside the safe set once three facts about real adversaries are stated, none of which the abstract bound captures on its own.

Most elections are not worth attacking. The value V is the attacker's gain from moving the result, and for the overwhelming majority of real elections it is small relative to the cost of subverting two independent attested platforms. A homeowners' association board, a union ratification, a professional-society slate, a cooperative's budget, a city or town referendum, a shareholder or DAO governance vote: these decide real matters, but none commands a prize that justifies developing and burning a covert capability against a hardware root of trust, an asset whose value lies precisely in not being known to exist. The economically relevant adversary for these contests is not a nation-state; it is a motivated insider or a commodity-malware operator, and against both the dual-root construction already wins by Theorem 2. The set of elections where V is genuinely large enough to attract a trusted-computing-base-level attacker, contested national offices, is a small and identifiable corner of the space, not its center.

Verifiability turns fraud into a public scandal. The abstract bound treats a successful attack as a clean win. In this system it is not, because the same end-to-end verifiability that lets an honest voter confirm inclusion also makes a large-scale manipulation *detectable at scale*. To move a continuous outcome the attacker must corrupt a $\Theta(\delta)$ fraction of all ballots; each corrupted voter who checks their own receipt against the public board finds a discrepancy, and it takes only a small number of such findings, out of the many thousands a broad attack requires, to produce cryptographic evidence that the published tally does not match cast ballots. The attacker's exposure is therefore not a fixed cost but one that grows with the very breadth continuity forces on them. For an institutional actor, a platform vendor, a listed corporation, the realized cost of discovery is not merely the burned capability but a provable, reputation-ending, potentially unlawful act caught on

an immutable record: a chilling effect the bare C_{TCB} does not include and that pushes the rational threshold well above it. Paper elections have no equivalent: a thrown-away ballot leaves no receipt its caster can later check.

The surviving adversary is a black swan. The residual of Section 11 requires an attacker who silently owns *both* independent trust-roots, for instance the makers of the device operating system and of the independent confirming platform. With two genuinely independent roots, sourced from different vendors sharing no operating system and no attestation authority, this is the simultaneous, covert, coordinated corruption of two of the largest technology institutions in the world, betting their core businesses on not being caught by a system designed to catch them. No such event has occurred. It is worth naming plainly that the live residual of this construction is an Apple-and-Google-collude-to-steal-an-election scenario, an extreme black swan with no precedent, rather than the routine client-malware threat that actually keeps internet voting out of use today, which the construction does defeat. A security analysis must state its residual; it should also state how exotic that residual is, so the residual is not mistaken for the typical case.

Put together, the three facts relocate the realistic question. For the bulk of the election space, organizational, municipal, multi-winner, proportional, $V \cdot L$ is small and the system is simply inside the safe set. For the small high-value corner, the deterrent of detection and the implausibility of two-root collusion raise the effective attack cost far above the nominal C_{TCB}. Only one combination is genuinely hard, a single national winner-take-all office at peak V with unbounded L, and that is exactly the case Section 9 addresses by lowering L through continuity, a step a single state legislature can take by itself.

10.1 A bound on the failure region

The three facts above are usually delivered as prose, which invites a reader to suspect they are softer than the theorems around them. They are not. The event in which the construction's integrity guarantee fails is a conjunction of independent conditions, and a conjunction of independent conditions multiplies, so the failure region is the product of several already-small factors and a rational adversary shrinks it further. We make this a theorem, because a worst case stated without its probability is exactly the raw material from which a typical-case barrier is wrongly inferred.

Definition 6 (Failure events for one election)**.** Fix an election with value V and aggregation Lipschitz constant L. Let:

- Col: both independent trust-roots are covertly subverted by a single coordinating adversary (Construction 1's independence hypothesis fails);
- Above: the election lies above the deployment predicate, $V \cdot L \geq C_{\mathrm{TCB}}/(2\rho_{\max})$;
- Undet: the resulting manipulation escapes detection by the verifiability layer.

Theorem 4 (The integrity-failure region is a product of independent small factors)**.** *Under the assumptions of Theorems 2 and 3, a ballot is accepted misencoded only on the event* Col*, and an adversary moves the* outcome *only on* $\mathsf{Col} \cap \mathsf{Above}$*. Consequently:*

1. (Multiplicativity.) *If the two roots are drawn from vendors that share no operating system and no attestation authority, and are each independently subvertible with probability at most* q_1, q_2*, then* $\Pr[\mathsf{Col}] \leq q_1 q_2$*; using two of three mutually independent roots (for instance a cloud enclave attestation authority, a device Secure Enclave vendor, and a reproducibly-built*

second-device verifier) the bound is the second elementary symmetric form $\sum_{i<j} q_i q_j$*, which a deployer lowers by adding independent roots. The detection term multiplies again: a broad manipulation is caught unless every one of* $\Theta(\delta N)$ *corrupted voters who could check fails to, so* $\Pr[\mathsf{Undet} \mid \textit{broad attack}]$ *decays geometrically in the number of checkers.*

2. (Restriction to the corner.) *For every election with* $V \cdot L < C_{\mathrm{TCB}}/(2\rho_{\max})$ *the event* Above *is empty, so no rational adversary attacks regardless of* Col*; the failure region is confined to the single corner where value and discontinuity are both maximal, which continuity reform (Section 9) removes by driving* L *down.*

Proof. Misencoding requires a forged or altered ballot accepted by the board; by Theorem 2 acceptance requires valid co-signatures from both independent roots, so a misencoded ballot is accepted only if both roots are subverted, i.e. on Col. Independence of the roots gives $\Pr[\mathsf{Col}] \le q_1 q_2$ (or the stated symmetric form for two-of-three), since the joint subversion of independent parties is the product of their individual probabilities. Outcome movement additionally requires the corruption of a $\Theta(\delta)$ fraction of ballots (Proposition 4); each such voter who checks a receipt against the public board detects a discrepancy, so escaping detection requires all checkers among the $\Theta(\delta N)$ affected to abstain from checking, whose probability falls geometrically in the checker count. Finally, by Theorem 3 a rational adversary acts only when $V \cdot L \ge C_{\mathrm{TCB}}/(2\rho_{\max})$, which is Above; off that set the expected payoff is negative and no attack is mounted. □

Proposition 7 (Rational incentives and chilling effects price out the corner)**.** *On* Above ∩ Col *the adversary still acts only if the value* $V \cdot L$ *of moving the outcome exceeds the expected cost* $C = p_{\mathrm{exp}}\, D + C_{\mathrm{TCB}}$*, where* p_{exp} *is the probability the manipulation or the collusion is exposed and* D *is the penalty on exposure. Three facts make* C *large for the only adversary capable of* Col*, an institutional owner of a trust root. First, verifiability makes* p_{exp} *high precisely when the attack is broad enough to move an outcome: even a* 0.1% *minority of checkers, each holding a cryptographic receipt that does not match the published tally, produces a forensic, reproducible, public proof of fraud. Second,* D *for a named platform is the loss of its core business plus legal liability, a penalty incommensurate with the value of essentially any contest below the national scale. Third, the deployer raises both* p_{exp} *and the operational cost without bound by adding independent roots and auditors. Hence even within the corner the rational-attack region is smaller still, and empty for every deployment whose* $V \cdot L$ *falls short of an institution's franchise value.*

To make the structure concrete, with the explicit caveat that the following are illustrative pessimistic figures rather than measured rates: suppose each of three independent roots (a cloud enclave vendor, a device Secure-Enclave vendor, and a reproducibly-built verifier) would covertly collude with probability $q = 0.05$. Then $\Pr[\mathsf{Col}] \le \sum_{i<j} q_i q_j = 3(0.05)^2 = 7.5 \times 10^{-3}$ before the detection term, which for a broad attack on a 10^4-voter electorate with even a 1% checking rate multiplies in a further factor below e^{-100}. The specific numbers are whatever a deployer's root roster yields; the structural fact is that integrity failure is a product of independent small factors gated by an economic predicate, the opposite of an irreducible barrier, and a deployer drives it toward zero by choosing more independent roots and auditors.

11 The residual: the one surviving corner case

Remark 7 (The residual). One adversary survives: a silent, colluding owner of *both* independent trust-roots, facing an election whose $V \cdot L$ exceeds $C_{\mathrm{TCB}}/(2\rho_{\max})$. Such an adversary can form the vote on a surface it fully controls, before any cryptography applies, and sign both halves. No

remote system defeats this, and we do not claim to; the same boundary already holds for privacy and encryption generally, where a trusted-computing-base owner can read the decrypted screen regardless of end-to-end guarantees. What the architecture does is make this regime avoidable rather than defeatable: lower L (continuity), lower V (election scope), and require two *independent* roots (so that no single vendor's compromise suffices and the joint cost is C_{TCB} for two unrelated parties, not one). Deployments outside the safe set should not use the system; the predicate of Theorem 3 tells a deployer, from V and L, which side they are on.

The empirical record of this residual. It is worth stating as a matter of record, not of framing, that the event this residual describes has never occurred. The required adversary is two independent trusted-computing-base vendors, sharing no operating system and no attestation authority, secretly and jointly corrupting the same election. To our knowledge there is no recorded instance, anywhere, of two such vendors colluding to subvert any election, by this or any mechanism. The count is zero, not low. This is a different statement from "unlikely": a low-probability event has a positive base rate and a history of occurrences, while this conjunction has a history of none. We do not present the absence of precedent as a security guarantee, since a guarantee is what Theorem 4 supplies and the empirical record cannot; we present it because a residual whose realized frequency is zero across the entire history of elections should be weighed as such, and not silently upgraded to a live threat merely because it can be written down. The honest pairing is precise: the attack is *possible*, so the worst-case impossibility for a maximal adversary stands, and the attack is *unprecedented*, so a deployer weighing real risk is weighing an event that has, to date, a frequency of zero.

Collusion is not the only way to own two roots: compelled access. The empirical record above concerns *voluntary* collusion, two vendors choosing together to corrupt an election, and that is genuinely unprecedented. It is not the whole of the residual, and we state the harder version plainly rather than let the zero-frequency claim stand in for it. A state can *compel* a vendor, by court order, statute, or national-security process, to subvert its own root without the vendor's independent choice and without a second vendor's cooperation being sought in the same way. Compulsion of a single vendor has ample historical precedent, unlike collusion of two, so the base rate here is not zero, and a reader who weighed only the collusion count would be weighing the wrong number. The honest accounting is this. First, the dual-root construction still requires *both* roots to be subverted for a misencoded ballot to be accepted (Theorem 2), so single-vendor compulsion defeats the system only when a single legal authority can compel *both* roots at once. That is precisely why the roots should be placed not merely under different vendors but under different *legal jurisdictions*: a second root operated by auditors or a vendor outside the compelling state's reach converts "compel one vendor" into "compel two vendors in two jurisdictions simultaneously and secretly," which restores much, though not all, of the collusion-level difficulty. Second, the detection property is not suspended by compulsion: a compelled manipulation broad enough to move an outcome still produces receipts that fail against the public board (Section 10), so a compelled attack is a detectable-at-scale attack, and the legal authority that ordered it is ordering an act that leaves cryptographic evidence, which raises its cost even where secrecy law would shield it. Third, transparency instruments, warrant canaries, published attestation logs, reproducible builds that a compelled silent code change would have to defeat publicly, raise the probability that compulsion of a root surfaces. None of these three eliminates the threat: a single authority able to compel two roots in two jurisdictions, silently, and absorb the detection cost, defeats the system, and we do not claim otherwise. The point is that compelled access is a distinct and more frequent threat than collusion, that the construction's defense against it is jurisdictional independence and detection

rather than the zero base rate, and that a deployer whose adversary is a compelling state must place its roots accordingly or stay out of the safe set. This is the honest residual for the reader who, correctly, does not accept that "never colluded" answers "could be compelled."

A lever the hardest case already has. The national winner-take-all contest is the one configuration the predicate excludes, but it is not excluded by anything a state must passively accept. The exclusion comes from the aggregation rule, statutory winner-take-all, which makes the map from votes to electors discontinuous and so sends L unbounded at the margin; and that rule is the one variable a single state legislature already controls. The Lipschitz analysis of Section 9 is general: it is a property of the counting rule, not of this construction, so the lever applies whether or not a state ever adopts the system in this paper. A legislature that wishes to use any electronic or remote method for its own share of a national election can first move its electors from winner-take-all to proportional within-state allocation, which is squarely within its plenary Article II power (*McPherson v. Blacker*; Maine and Nebraska already allocate non-statewide), requires no amendment, no interstate compact, and no other state's cooperation, and takes effect for that state's own voters immediately. The effect is concrete and one-directional: proportional allocation lowers L, which shrinks the discontinuity that makes a thin margin outcome-determinative, removes the single-jurisdiction pivotality that turns one close state into a national crisis, and correspondingly lowers the payoff $V \cdot L$ to anyone, foreign or domestic, contemplating interference in that state. The state gets better election properties and weaker rigging incentives in exchange for a change to its own statute.

We are careful about two things this lever is not. It is not abolition of the Electoral College, and not a change to how electors are *apportioned* among the states: the College, the small-state weighting, and the federal structure all remain untouched, and only the within-state *division* of a state's fixed electors changes from all-or-nothing to proportional. And it is not, by itself, a claim that a presidential contest then clears the predicate automatically, because the value V of a single national office remains large even at low L; the honest statement is that the lever moves that contest, in the acting state, from the categorically excluded region into the evaluable one, where whether to use a remote method becomes a quantitative question a deployer can answer rather than an automatic no. That is the most a state can do unilaterally, it is more than the consensus assumes is possible, and it is available now, to one legislature, without waiting for anyone else.

12 Capstone

Theorem 5 (Bounded-payoff remote election)**.** *Under the cryptographic and registration assumptions, Construction 1 with Invariants 1–3, and two independent attested trust-roots, the system:*

(i) *provides verifiability, ballot secrecy, and coercion-resistance* independent of *the aggregation rule (Prop. 5);*

(ii) *bounds any adversary's outcome distortion by* $2L\rho$*, against every adversary including one facing no defense (Prop. 4);*

(iii) *accepts a misencoded ballot only under joint subversion of both trust-roots (Thm. 2, Cor. 1); and*

(iv) *admits no rational attack whenever* $V \cdot L < C_{\mathrm{TCB}}/(2\rho_{\max})$ *(Thm. 3), a condition continuity and election scope are designed to satisfy.*

The trifecta is the L-independent floor; continuity converts "we cannot stop a determined trusted-computing-base owner" into "no one rational attempts it."

Proof. Each clause is one of the established results, and all are proved in the single model of Sections 7.2–9 under one assumption set, so their conjunction holds without further argument: (i) is Proposition 5; (ii) is Proposition 4; (iii) is Theorem 2 with Corollary 1; and (iv) is Theorem 3. No two clauses impose conflicting hypotheses: (i)–(iii) are stated for every L, and (iv) supplies the value-times-Lipschitz condition under which the bound of (ii)–(iii) leaves no profitable attack. □

At the level of what an election delivers, the four clauses combine into the property that motivates the whole construction. A secret ballot (i), integrity a voter checks without trusting the operator (i) and (iii), and, through voting on a standing quantity, a continuously available tally are the three things a paper-ballot replacement must provide together; paper provides the first two without the third, and prior remote electronic voting has not provided the first two convincingly enough to be trusted with the third. Clause (iv) is the discipline that makes offering all three responsible: it identifies the regime in which the assurance is high enough to serve in place of supervised paper, and outside that regime the correct output of the system is the recommendation to use paper for that contest. The system does not assert that electronic casting is always preferable; it asserts that there is a precisely bounded regime in which it is at least as good, and it can tell whether a given election lies inside it.

13 Comparison with the deployed and academic state of the art

The right question for a deployer is not whether this system attains an absolute ideal but whether it beats the alternative actually available for the same election. There are three such alternatives: supervised paper balloting, the remote channels already in production (email and fax return of marked ballots for overseas and military voters under UOCAVA), and the academic internet-voting systems that have been built and analyzed (Helios, and the JCJ/Civitas coercion-resistant line). Against each, on the properties that decide whether a result can be trusted, the construction is ahead on most axes; we lead with those and name the one axis where an incumbent still leads.

Property	**Paper**	**Mail-in**	**Email/fax**	**Helios**	**This work**
Voter checks own ballot counted	no	no	no	yes	yes
Public, universal tally proof	no	no	no	yes	yes
Integrity without trusting operator	partial	no	no	yes	yes
Voter-auditable chain of custody	yes	no	no	yes	yes
Resists a compromised client	n/a	n/a	no	no	**bounded (dual-root)**
Ballot secrecy under home coercion	yes	**no**	no	no	yes (with repair)
Continuous / real-time result	no	no	no	no	**yes**
Runs without polling-place apparatus	no	yes	yes	yes	yes
Outcome distortion bounded by L	no	no	no	no	**yes**
Reaches device-less / offline voters	yes	**yes**	yes	no	no
Worst-case high-value national contest	**safest**	no	no	no	conceded

Three wins are worth drawing out because no fielded system holds them together. First, *client-malware resistance*: Helios and Civitas assume an honest client and are defeated by malware on the voter's device, the precise reason the expert consensus keeps internet voting out of binding use; the dual-root construction is the only one here that bounds that adversary rather than assuming it

away, requiring joint subversion of two independent attested roots (Theorem 2). Second, *individual verifiability against paper*: a paper ballot, the supposed gold standard, gives the voter no way to confirm after the fact that their ballot was counted as cast, and the deployed remote channel, emailed or faxed ballots, gives neither secrecy nor verifiability; this system gives both. Third, *continuity and bounded distortion*: no prior system makes the outcome a Lipschitz function of the ballots, so none can bound an attacker's distortion by $2L\rho$ or support a continuously available tally (Section 8); both are new here.

Mail-in ballots: paper's reputation without paper's protections. Mail voting deserves its own line, because it is the remote method used at the largest scale and the one most often assumed safe. It inherits the trusted reputation of the supervised paper ballot while sharing almost none of what earns that reputation. The protection of precinct paper comes from a supervised booth and an observed chain of custody; a mailed ballot is marked unsupervised in a home, where vote-buying and family or household coercion are not exotic attacks but the ordinary condition of an unobserved kitchen table, and it then travels through collection boxes, the postal service, and processing centers, so the voter is trusting the carrier and every handler in a chain they cannot audit. The voter receives nothing back that proves their ballot was counted as cast. Against mail-in, this construction wins on every trust-relevant axis at once: it offers ballot secrecy under coercion with a stated repair where mail offers none, a receipt the voter checks against a public board where mail offers no confirmation, and a tally that is a public proof rather than the output of an unauditable custody chain. Where mail's failures are invisible unless separately investigated, this system's would be cryptographic and detectable at scale. The "patina of paper" is the precise problem: mail borrows the credibility of the supervised ballot while trusting the mailman.

The honest exception is the last accessibility row, and we state it plainly because it is the real reason mail remains in service: mail reaches voters with no smartphone, no secure-element device, and no connectivity, the rural, the elderly, the unhoused, the device-less, whom any device-bound remote system cannot serve directly. That is a genuine advantage of mail and a genuine cost of this construction, the same distribution gap the companion attribute paper [8] confronts, and it is not dissolved by any property in the table above. A deployment that must reach those voters keeps a supervised or mail channel for them; the construction's claim is over the voters it can reach, which is most of them, not all.

The one axis where an incumbent still leads is the last row, and we state it without softening: for a single high-value national contest against an adversary who can subvert both trust-roots, supervised in-person paper in a precinct remains the safer choice, because its failures are local and independent rather than systemic, and this is exactly the residual of Section 11. That comparison, however, is the narrow worst-case corner, not the common case. For the remote national channel that actually exists today, emailed and faxed ballots with no cryptographic protection at all, the construction is strictly safer even there; and for the entire organizational, municipal, and proportional space, it dominates every alternative on every row. The honest summary is that against the systems a deployer can actually choose from, this one wins on the properties that matter, everywhere except the one contest where supervised paper still wins, and it says so.

14 Scope and deployment profile

The system is built for elections where a deployer can (a) run a coercion-free registration step, (b) require two independent trust-roots per voter, and (c) keep $V \cdot L$ below the threshold, typically by proportional or otherwise continuous allocation and by the bounded value of an organizational

contest. Member organizations, professional societies, unions, cooperatives, shareholder and DAO governance, and advisory elections fit this profile: registration is a controlled onboarding step, members plausibly hold two independent devices, and the deployer controls the aggregation rule. The government case is governed by the same predicate and the boundary of Section 9. The realistic envelope is broad: every organizational and municipal contest, every proportionally or multi-winner allocated body, every ballot measure, apportionment, and participatory budget sits inside the safe set on the merits, because either V is bounded or L is small or both. Only one configuration is genuinely hard, a single binding national winner-take-all office, which combines a peak V with an L that is unbounded at the margin under statutory winner-take-all. Even there the situation is not a categorical no but a question with a known answer: within-state proportional or district allocation of electors, which a single legislature can enact by ordinary law (Section 9.1), lowers L into the tractable range and converts the contest into one where a broad two-root attack at scale would be required, exactly the attack the construction is built to make infeasible and non-covert. The honest summary is the inverse of the timid one: the system covers the great majority of real elections outright, and supplies, for the hardest remaining case, a lawful and unilateral path that brings it inside the boundary too. The deployment claim is a computed one: publish V, choose L, and the predicate states whether the election is in the safe set.

Continuous operation and the reveal. An election run continuously rather than to a single close needs an explicit reveal policy, because the property that makes the system efficient, a real-time, checkable record, is also what would expose results in progress. Section 3.6 makes this a controlled parameter: the board holds only encrypted ballots, a threshold quorum unseals aggregates on a schedule (per epoch for an ongoing election), and the cadence and demographic detail of each reveal are disclosure-design choices rather than defaults. The point for the present paper is that a voter's post-reveal verification, inclusion on the board plus the universal tally proof, establishes that the vote was counted without becoming a proof of how it was cast, so the embargo and reveal compose with the receipt-freeness of Proposition 5 rather than weakening it. The limiting case is a vote on a single ordered quantity that never closes: there the aggregation rule is maximally continuous (L at its smallest), the outcome is a standing median that re-converges after any single move, and the embargo question dissolves rather than needing a schedule; Section 8 develops this as continuous voting on a standing quantity, the lowest-L point of the predicate.

What the construction puts within reach. The deployments the predicate admits are not a consolation prize; they are most of the democratic activity that actually happens, and the construction makes them cheap. A paper election at national scale demands polling places, staff, chain-of-custody, and an election day; the costs are why most organizations that should run a trustworthy secret ballot do not. This system needs none of that apparatus: registration is an onboarding step, the ballot is cast from a personal device confirmed by a second one, and the tally is a public proof anyone can check. A union, a cooperative, a professional society, a municipality, a shareholder body, an online community can run an election with verifiable integrity and ballot secrecy at the cost of software, which puts a property formerly reserved for states with a census's resources into the hands of any organization that needs it. The continuous-voting construction of Section 8 extends this further, into decisions that a periodic ballot cannot express at all: a standing quantity the electorate adjusts in real time, the level of a community's basic-income disbursement next month, or the coarseness of the age-band tokens of the companion attribute paper [8], so that a community rather than a legislature fixes where the $[20, 30]$ versus $[25, 35]$ line falls. These are the lowest-L point of the predicate and the safest deployments of all, and they let a community

govern parameters continuously, a capability that has no paper analogue.

What is assumed, proved, and residual. *Assumed:* crypto hardness; coercion-free registration; an honest hardware attestation root per trust-root; an honest M-of-N auditor quorum for the backend casting/tally code (attestation certifies code identity, the auditors certify code correctness; §6.2); two independent trust-roots; one honest mix server. The cast-as-intended theorem additionally assumes an honest voter who performs the mandatory cross-check (does not submit unless the second device displays the intended candidate). *Proved (here):* the distortion bound (Prop. 4); dual-root cast-as-intended (Thm. 2); the economic predicate (Thm. 3); and the capstone (Thm. 5), with the protocol guarantees of Section 3 (Prop. 5). *Residual:* a silent, colluding owner of both trust-roots on an election above threshold (Sec. 11).

The residual, weighed against the deployed alternatives. The compromised-client residual is not unique to this construction; it is shared, and unmitigated, by every *remote* voting method actually in use. Email and fax return of marked ballots, the deployed remote channel for overseas and military voters, runs on the voter's own compromisable device with no cryptographic protection at all; prior internet-voting pilots shared the exposure with weaker defenses. Against those, the systems a deployer choosing a remote method actually picks from, this construction is the only one that bounds the client adversary rather than ignoring it, and the residual it retains, joint subversion of two independent roots, is a strictly harder attack than the single-device compromise the others already permit. The one alternative that escapes the residual is supervised in-person paper, which avoids it by not being remote at all; that is the genuine tradeoff, and it is a choice between remoteness and a never-observed worst case (Sec. 11), not between this system and a remote alternative that does better.

References

[1] National Academies of Sciences, Engineering, and Medicine, *Securing the Vote: Protecting American Democracy*, 2018.

[2] Center for Transparent and Trusted Elections and Yankelovich Center for Social Science Research, University of California San Diego, *Trust in Elections Declines across Party Lines Ahead of 2026 Midterms* (national survey of 11,406 eligible voters; confidence in accurate counting fell from 77% to 60% year over year), February 2026. `https://today.ucsd.edu/story/trust-in-elections-declines-across-party-lines-ahead-of-2026-midterms-uc-san-diego-survey`

[3] Harvard Public Opinion Project, *Harvard Youth Poll*, Spring 2026 (33% of Americans aged 18–29 believe the 2026 midterms will be conducted fairly and accurately).

[4] A. Appel et al., *Internet voting is insecure and should not be used in public elections*, Princeton Center for Information Technology Policy, January 2026. `https://blog.citp.princeton.edu/2026/01/16/internet-voting-is-insecure-and-should-not-be-used-in-public-elections/`

[5] American Association for the Advancement of Science, *Internet or Online Voting Remains Insecure*, EPI-Center. `https://www.aaas.org/epi-center/internet-online-voting`

[6] A. Koppel et al. (MIT), security analysis of the Voatz mobile voting application, MIT News, 2020. `https://news.mit.edu/2020/voting-voatz-app-hack-issues-0213`

[7] U.S. Vote Foundation, *The Myth of "Secure" Blockchain Voting*. `https://www.usvotefoundation.org/blockchain-voting-is-not-a-security-strategy`

[8] Companion paper. *When Age Checks Need Not Reveal Identity: A Construction and Distribution Criterion for Privacy-Preserving Attribute Verification*.

[9] G. Magarshak. When exact reporting beats adding noise: a construction and quantitative criterion for reporting how groups voted without exposing how anyone did. Companion manuscript, 2026. Takes the election results this paper's stack produces and characterizes, via the contingency-tensor model and Fréchet–Hoeffding cell bounds, how much demographic detail can be reported while keeping each voter unlinkable; supplies the downstream tiered disclosure budget this paper's reveal layer references.

[10] G. Magarshak. Intercloud: chilling-effect consensus for federated coordination. arXiv:2605.22830, 2026.

[11] K. J. Arrow. *Social Choice and Individual Values*. Wiley, New York, 1951.

[12] J. Benaloh. Verifiable secret-ballot elections. PhD thesis, Yale University, Technical Report YALEU/DCS/TR-561, 1987.

[13] J. Benaloh and D. Tuinstra. Receipt-free secret-ballot elections (extended abstract). In *ACM Symposium on Theory of Computing (STOC)*, pages 544–553, 1994.

[14] D. Black. On the rationale of group decision-making. *Journal of Political Economy*, 56(1):23–34, 1948.

[15] D. Chaum. Blind signatures for untraceable payments. In *CRYPTO*, pages 199–203, 1982.

[16] M. R. Clarkson, S. Chong, and A. C. Myers. Civitas: toward a secure voting system. In *IEEE Symposium on Security and Privacy (S&P)*, pages 354–368, 2008.

[17] R. Cramer, R. Gennaro, and B. Schoenmakers. A secure and optimally efficient multi-authority election scheme. In *EUROCRYPT*, pages 103–118, 1997.

[18] A. Escala, S. Guasch, J. Herranz, and P. Morillo. Universal cast-as-intended verifiability. In *Financial Cryptography and Data Security (FC) Workshops*, LNCS 9604, pages 233–250. Springer, 2016.

[19] S. Guasch and P. Morillo. How to challenge and cast your e-vote. In *Financial Cryptography and Data Security (FC)*, pages 130–145. Springer, 2016.

[20] A. Gibbard. Manipulation of voting schemes: a general result. *Econometrica*, 41(4):587–601, 1973.

[21] M. Jakobsson, K. Sako, and R. Impagliazzo. Designated verifier proofs and their applications. In *EUROCRYPT*, LNCS 1070, pages 143–154. Springer, 1996.

[22] A. Juels, D. Catalano, and M. Jakobsson. Coercion-resistant electronic elections. In *ACM Workshop on Privacy in the Electronic Society (WPES)*, pages 61–70, 2005.

[23] R. D. McKelvey. Intransitivities in multidimensional voting models and some implications for agenda control. *Journal of Economic Theory*, 12(3):472–482, 1976.

[24] H. Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980.

[25] J. Müller and T. Truderung. CAISED: a protocol for cast-as-intended verifiability with a second device. In *International Joint Conference on Electronic Voting (E-Vote-ID)*, LNCS 14230, pages 123–139. Springer, 2023. Also arXiv:2304.09456.

[26] C. A. Neff. A verifiable secret shuffle and its application to e-voting. In *ACM Conference on Computer and Communications Security (CCS)*, pages 116–125, 2001.

[27] National Academies of Sciences, Engineering, and Medicine. *Securing the Vote: Protecting American Democracy*. The National Academies Press, Washington, DC, 2018.

[28] S. Park, M. Specter, N. Narula, and R. L. Rivest. Going from bad to worse: from Internet voting to blockchain voting. *Journal of Cybersecurity*, 7(1):tyaa025, 2021.

[29] M. A. Satterthwaite. Strategy-proofness and Arrow's conditions. *Journal of Economic Theory*, 10(2):187–217, 1975.

[30] Amazon Web Services. AWS Nitro Enclaves: cryptographic attestation and reproducible enclave images. AWS technical documentation, 2024. Establishing verifiable security through reproducible builds and Nitro Enclave attestation; PCR0 binds the enclave image hash, enabling third-party verification against independently rebuilt images.