Optimal Dynamic Majority-Quorum Rules: Theory and Evidence*

Bhargav Nagaraja Bhatt† Jonas Gehrlein‡ Kremena Valkanova‡

September 8, 2026

Abstract

We study the optimal sequential aggregation of randomly arriving votes from a finite population to identify an unknown majority preference before a fixed deadline. We formulate the planner's trade-off between decision accuracy and the flow cost of delay as a finite-horizon dynamic program. We show that the optimal policy is a dynamic majority-quorum rule: a joint, time-varying threshold on vote lead and count that collapses as the deadline approaches. The optimal rule induces a lens-shaped continuation region in vote lead–count space, driven by the tension between the option value of learning and the depletion of the eligible voter pool. Because the exact boundary admits no closed form, we derive a tractable large-population approximation that simplifies the two-dimensional rule onto a single sufficient statistic, and express the rule as a sequential significance test set by a single error tolerance. Calibrating the voting rule on Polkadot governance data, an out-of-sample evaluation finds that our rule calibrated to a 2% error tolerance halves decision times relative to the incumbent mechanism.

Keywords: dynamic collective choice; optimal stopping; information aggregation; speed–accuracy trade-off; blockchain governance.
JEL Classification: D71, D83, D82, C61.

---

*We thank Moumita Deb, Nick Netzer and Christoph Vanberg, and the research team at Parity Technologies.

†Morpho, bhargav@morpho.org

‡Parity Technologies, {jonas,kremena}@parity.io.

1 Introduction

Collective decision-making inherently entails a trade-off between the accuracy of the final choice and the time taken to reach it. In decentralized autonomous organization (DAO) governance, for example, protocol upgrades and security patches are voted on asynchronously by a finite population of token holders. Acting on too few votes risks enacting a change the electorate opposes, while delaying execution extends the window of system vulnerability. The analogous trade-off arises in crowdsourced data annotation, the backbone of modern AI training pipelines, where platforms decide how many worker judgments to collect for each item (Snow et al., 2008). Collecting too few labels risks classification errors, while collecting too many exhausts budgets, increases costs, and slows development. Similarly, in post-election auditing, election officials sequentially sample ballots to verify a reported outcome and stopping prematurely risks certifying the wrong winner while prolonging the audit delays the transfer of political authority. Across these settings, the value of a decision depends not only on whether it is correct, but also on when it is made.

In these environments, and increasingly in collective decision-making processes enabled by digital technologies, information is observable as it arrives. As a result, collective decision-makers observe not only the eventual aggregate outcome, but also the path by which information accumulates. This creates a new opportunity for institutional design: rather than relying exclusively on fixed voting windows, participation thresholds, or predetermined sample sizes, collective decision rules can adapt to the evolving informational state of the process. Yet most institutions rely on static rules that underutilize the informational content of sequential participation.[1] This raises a natural question: how should collective decision rules exploit this information to balance the competing objectives of accuracy and timeliness?

[1]There are notable exceptions. In decentralized governance, Polkadot's OpenGov framework introduces a dynamic approval-support rule, discussed in Section 6, and Arbitrum's Security Council elections down-weight votes cast late in the voting window, linearly to zero at the deadline (Arbitrum Foundation, 2023). The crowdsourced annotation literature has proposed adaptive alternatives that stop collecting labels once sufficient agreement is reached (Sheng et al., 2008), though these have not been widely adopted in practice. Risk-limiting audits (Stark, 2008; Lindeman and Stark, 2012) provide formal statistical stopping criteria for ballot sampling, but their adoption remains far from universal.

The settings above illustrate a broader class of collective decision problems defined by three structural features. First, the population of voters is finite and exhaustible. Because each voter participates at most once, every vote cast reduces the pool of remaining voters and thereby changes the value of waiting for additional votes. Second, vote arrivals are exogenous to the planner. The timing of votes is determined by individual voters and may reflect participation frictions, cost–pivotality tradeoffs, and strategic timing considerations, among other factors. Third, decisions must be reached before a binding deadline. Such deadlines are often necessitated in practice by implementation needs, budget schedules, or legally prescribed certification dates.[2] Existing theories provide only partial guidance for such environments: sequential-testing models in the tradition of Wald and Wolfowitz (1948), and their modern descendants (Moscarini and Smith, 2001; Fudenberg et al., 2018), draw evidence from an effectively inexhaustible source; dynamic collective-choice models allow the planner to solicit votes at will (Gershkov and Szentes, 2009; Bognar et al., 2015); and finite-population sequential tests, including risk-limiting audits (Stark, 2008; Lindeman and Stark, 2012), impose no binding decision deadline.[3] In this paper, we derive optimal collective decision rules for finite societies with exogenous vote arrivals and a fixed decision deadline.

We formalize the problem in a model of dynamic preference aggregation in which a social planner must choose between two alternatives before a fixed deadline. The objective is to implement the true majority preference of a finite population of eligible voters. Voters hold fixed but unobserved preferences and vote at most once. Votes arrive asynchronously over time, reflecting heterogeneous opportunities for participation, in randomly sized batches at a rate that may vary across periods. Rather than modeling participation decisions, we take this arrival process as given and ask how the planner should act in response to the observed stream of votes: at each date, to implement the leading alternative or to continue waiting

[2]The tension between finite sampling resources and hard deadlines is well documented in practice. For example, when Maryland evaluated audit methods, local boards of elections noted that risk-limiting audits were difficult to budget or staff for, since the required sample size depends on the margin of victory and a close contest "could require days of staff work, possibly compromising the local certification deadline" (Maryland State Board of Elections, 2016).

[3]See Section 2 for a detailed review of the literature.

for additional votes. Implementing yields a reward if the chosen alternative matches the majority preference, while waiting provides further information at the cost of delay.

Our central result characterizes the optimal policy as a dynamic majority–quorum rule, in which the vote lead required to act depends jointly on turnout and time remaining, such that a given lead may be decisive at high turnout but not at low. The rule takes a threshold form, i.e. for each turnout level and date, the planner stops once the lead crosses a cutoff that collapses as the deadline approaches, and waits otherwise. The turnout dimension distinguishes the rule from classical sequential testing, where a one-dimensional running tally suffices. It arises because a given lead is consistent with many possible majorities when the remaining electorate is large, but becomes increasingly decisive as it shrinks. The finite horizon lets us compute the exact boundary by backward induction, and the resulting continuation region is lens-shaped in lead–count space: narrow at low turnout, widest at intermediate turnout, and narrowing again as the shrinking residual electorate makes surviving leads decisive.

We then establish conditions under which the optimal threshold responds monotonically to the model's primitives. These comparative statics are governed by a common force: the option value of continued sampling. A higher reward relative to delay costs raises the threshold, since the value of additional votes scales with the reward while their cost is unchanged. A thinner arrival process lowers it, because fewer expected votes reduce the gains from waiting. Finally, within a given environment, the threshold declines over time and converges to zero at the deadline, though monotonically only when participation does not intensify over time.

Because the exact rule admits no closed form, we derive a tractable approximation in a large-population, high-stakes limit with deterministic voter arrival. The stopping decision reduces to comparing a standardized lead with a time-varying threshold given by the familiar Wald–Chernoff stakes term, scaled by the votes still expected to arrive. Because the stakes ratio is difficult to quantify in practical applications, we express the same rule as a sequential significance test in which the planner specifies only an error tolerance and stops once the posterior probability of implementing the minority preference falls below it, with the required confidence relaxing as the deadline approaches.

We complement the theoretical analysis with an empirical evaluation using governance data from Polkadot (Wood, 2022), one of the largest decentralized autonomous organizations by market capitalization. This setting is well suited to our framework on two counts. First, votes are binding and executed on-chain, and the full time path of vote arrivals is observed for several hundred referenda decided by a large and heterogeneous electorate. Second, since Polkadot's incumbent mechanism is itself a dynamic majority-quorum rule, the observed arrival paths reflect voters' incentives over the timing of participation. Thus, the voting behavior we observe should resemble those the proposed rule would itself generate. We find that participation is sharply front-loaded but well described by a non-homogeneous Poisson process with a parsimonious log-linear intensity, fitted tightly across referenda.

We then evaluate the rule out of sample, specifying only an error tolerance and calibrating all remaining parameters on the chronologically first half of the sample, held fixed thereafter. On the second half, we ask when the proposed mechanism would have decided each historical referendum. At 2% error tolerance, the rule decides the median referendum on day 11 of the 28-day voting period, after roughly 60% of eventual participants have voted, saving 15 days relative to the incumbent mechanism while differing from its deadline outcome in fewer than 2% of referenda. The speed–accuracy frontier behind this point is stable across the two halves of the sample and steep at its accurate end: near-perfect accuracy preserves most of the attainable time savings, while eliminating the last percentage point of error would forfeit the mechanism entirely.

The remainder of the paper is organized as follows. Section 2 reviews the related literature and Section 3 introduces the model. We then characterize the optimal stopping rule and study its properties in Section 4 and derive a closed-form approximation in Section 5. Section 6 presents the empirical analysis. Section 7 concludes.

2 Related Literature

Our paper connects three strands of literature: statistical social choice, which frames collective decisions as inference problems; the design of voting rules for information aggregation, from which we take the planner's perspective on dynamic rule choice; and sequential decision-making under uncertainty, whose canonical free-boundary problems we extend to the collective-choice setting under finite population and binding horizon.

The first strand treats collective decisions as a problem of statistical inference, in which the aim is to recover an unknown population-level quantity from observed votes. In the classical Condorcet tradition (Condorcet, 1785; Young, 1988) and its modern extensions (Nitzan and Paroush, 1984; Austen-Smith and Banks, 1996), this latent quantity is interpreted as an external state of the world about which voters receive private signals. Subsequent work has extended this framework to strategic voting under private information (Feddersen and Pesendorfer, 1997), including within Poisson-game frameworks that accommodate population uncertainty (Myerson, 1998). We adopt the same inference perspective: votes reveal information about an unknown majority preference, and the planner aggregates that information from a finite pool. The source of uncertainty is the planner's inability to observe all votes at once, and resolving it through additional observation comes at a cost. We depart from this literature, however, in the question we ask. Whereas the jury-theorem tradition takes the voting rule as given and asks what equilibrium voting reveals about the underlying state, we treat the rule itself as the choice variable.

This design perspective is at the core of the second related strand of literature, which studies optimal decision rules when the planner sequentially gathers costly information. Foundational work in the social choice tradition characterizes optimal weighted-majority rules for groups of heterogeneous voters (Nitzan and Paroush, 1982). Subsequent contributions have enriched the framework to allow for costly information acquisition (Persico, 2004), recast voting rules explicitly as Bayesian estimators of an underlying truth (Azari Soufiani et al., 2014). Most directly relevant, a line of work has developed the dynamic case in which the planner sequentially aggregates votes and decides when to stop (Gershkov and Szentes, 2009; Bognar et al., 2015; Gershkov et al., 2017). Our paper shares the planner's-problem perspective of this literature but differs in two main respects. First, the rate and order of vote arrivals are exogenous and arise from sampling without replacement from a finite, exhaustible pool: the planner neither controls when votes appear nor solicits them strategically. The planner's only choice variable is when to stop. Second, the deadline is hard rather than asymptotic, which makes the time horizon a binding constraint and rules out the standard infinite-horizon characterizations.

Finally, our paper sits within the broader literature on sequential statistical decisions and optimal stopping under uncertainty, which studies how a Bayesian decision-maker should accumulate noisy evidence over time and decide when to stop. The decision-theoretic foundations of Wald (1950), formalized in the sequential probability ratio test (Wald and Wolfowitz, 1948) and refined through Karlin and Rubin (1956), Chernoff (1961), and Siegmund (1985), characterize Bayesian sequential tests with costly sampling. Lai (1979) and Xiong (1993) extend this framework to the finite population case that is most relevant to our setting. This structure arises in many different domains: optimal experimentation (Moscarini and Smith, 2001; Bolton and Harris, 1999), the drift-diffusion tradition in cognitive science and behavioral economics (Ratcliff, 1978; Ratcliff and McKoon, 2008; Drugowitsch et al., 2012; Tajima et al., 2016; Fudenberg et al., 2018), the theory of American options and real-option problems (Peskir and Shiryaev, 2006), and finite-population sequential testing in post-election ballot audits (Stark, 2008; Lindeman and Stark, 2012) and financial audit sampling (Kato and Nakagawa, 2026). Our paper applies this machinery to collective choice. The signal is a vote drawn without replacement from a finite electorate, the decision-maker is a planner aggregating preferences, and the institutional setting imposes a hard deadline. The combination of finite population and binding horizon has not, to our knowledge, been studied jointly in this literature: we show that the collapsing boundaries familiar from single-agent decision-making have a direct counterpart in collective-choice rules under both constraints.

3 Model

A social planner must decide whether to accept or reject a proposal on behalf of a society of an odd number of $N \in \mathbb{N}$ voters. Each voter has a fixed binary preference, either in favor of (+) or opposed to (−) the proposal. Let $M^+ \in \{0, 1, \ldots, N\}$ denote the (unknown) number of voters who support the proposal, with $M^- = N - M^+$.

The planner is benevolent and seeks to implement the majority-preferred alternative, i.e. accept if $M^+ > \frac{N}{2}$, reject otherwise. We assume the planner adopts a uniform prior $M^+ \sim \text{Unif}\{0, \ldots, N\}$, which is equivalent to assuming each voter independently supports the proposal with a probability drawn uniformly from $[0, 1]$.

3.1 Sequential Vote Arrival

Votes arrive over discrete dates $t = 1, 2, \ldots, T$. On each date, a batch of votes is observed, consisting of a_t^+ favorable votes and a_t^- unfavorable votes, with total batch size $a_t = a_t^+ + a_t^-$. We define cumulative vote counts as

$$s_t^+ = \sum_{\tau=1}^{t} a_\tau^+, \qquad s_t^- = \sum_{\tau=1}^{t} a_\tau^-, \qquad s_t = s_t^+ + s_t^-, \qquad d_t = |s_t^+ - s_t^-|.$$

The potential number of arrivals at date t follows a Poisson distribution with a time-varying rate $\lambda(t) > 0$, independently across dates, i.e. $A_t \sim \text{Poisson}(\lambda(t))$ for all $t = 1, \ldots, T$.[4] We denote the expected arrival rate sequence by $\boldsymbol{\lambda} = (\lambda(1), \ldots, \lambda(T))$. The cumulative expected arrival function $\Lambda(t) = \sum_{\tau=1}^{t} \lambda(\tau)$, is such that $\mathbb{E}\left[\sum_{\tau=1}^{t} A_\tau\right] = \Lambda(t)$. Because the population

[4]This specification represents the asynchronous revelation of a finite signal pool: N signals exist ex ante but surface to the planner at random times, independently of their content and signals observed so far. The Poisson structure captures the planner's uncertainty about how many of the N potential voters will surface at any given date, arising from independent individual-level randomness in participation timing. This is in the spirit of Myerson (1998), where Poisson uncertainty over the realized electorate arises from independent individual participation decisions. This specification is consistent either with random sampling by a coordinator or uncoordinated voluntary participation based on attention or availability. The stochastic batch sizes capture two features absent from pure hypergeometric sampling: the planner does not control how many signals arrive at each date, and arrival intensity can vary over time, accommodating front-loaded participation or periodic fluctuations. The censoring ensures the process respects the finite-population boundary; see Section 6.

is finite and each voter votes at most once, the observed batch size is censored by the number of remaining voters: $a_t = \min\{A_t,\ N - s_{t-1}\}$. The probability mass function of the observed batch $a_t \in \{0, \ldots, N - s_{t-1}\}$ is therefore

$$P(a_t = a) = \begin{cases} \dfrac{\lambda(t)^a \, e^{-\lambda(t)}}{a!} & \text{if } a < N - s_{t-1}, \\ 1 - \displaystyle\sum_{j=0}^{N-s_{t-1}-1} \dfrac{\lambda(t)^j \, e^{-\lambda(t)}}{j!} & \text{if } a = N - s_{t-1}. \end{cases}$$

Because voters cannot recast or change their votes, the voting process corresponds to sampling *without replacement* from a finite population of N voters. Thus, conditional on the votes already observed (s^+, s^-) and on M^+, the composition of each batch follows a hypergeometric distribution: $a^+ \mid a, M^+ = m^+, s^+, s^- \sim \text{Hypergeometric}(N - s^+ - s^-,\ m^+ - s^+,\ a)$.

3.2 Bayesian Updating

The planner updates their beliefs about the unobserved state M^+ based on the cumulative votes observed. Because voters are exchangeable, the order of arrival is irrelevant. The cumulative counts (s_t^+, s_t^-) constitute a sufficient statistic for inference.

Lemma 1. *For a batch of size $a \leqslant N - s$, the probability of observing exactly a^+ favorable votes is*

$$P(a^+ \mid a, s^+, s^-) = \frac{\binom{a^+ + s^+}{s^+}\binom{a - a^+ + s^-}{s^-}}{\binom{a+s+1}{a}}, \qquad a^+ \in \{0, \ldots, a\}.$$

Proof. See Appendix A.1. □

This predictive distribution serves as the transition kernel in the Bellman equation defined below, mapping current beliefs to the expected continuation value. Note that it is the Beta-Binomial distribution with parameters $(a, s^+ + 1, s^- + 1)$ and is independent of the population size N. Thus, conditional on the realized batch size a, the planner's forecast of its composition depends only on the current tally (s^+, s^-). The population size N enters the problem only through the censoring of batch sizes ($a \leqslant N - s$) and the definition of the majority threshold $M^+ > N/2$.

3.3 Decision Problem and Dynamic Programming

At each time t, after observing the state (s_t^+, s_t^-), the planner chooses whether to stop the voting process or continue sampling at cost $\kappa > 0$, incurred regardless of whether any votes arrive at date t. It captures the opportunity cost of keeping the decision open, rather than a per-observation processing cost. If the planner stops at time t, they choose the action that maximizes the probability of matching the true majority preference. The expected payoff from stopping is

$$U(s_t^+, s_t^-) = R \cdot \max\left\{\mu(s_t^+, s_t^-),\ 1 - \mu(s_t^+, s_t^-)\right\}, \tag{1}$$

where $R > 0$ is the reward for making the correct decision, and

$$\mu(s^+, s^-) = P(M^+ > N/2 | s^+, s^-) = \sum_{m^+ > \frac{N}{2}} \frac{\binom{m^+}{s^+}\binom{N-m^+}{s^-}}{\binom{N+1}{s+1}} \tag{2}$$

is the posterior belief that the proposal is supported by the majority, obtained by applying Bayes' rule with the uniform prior and a hypergeometric likelihood.[5]

The value function satisfies the terminal condition $V_T(s^+, s^-) = U(s^+, s^-)$ and, for $t < T$, the Bellman equation[6]

$$V_t(s^+, s^-) = \max\left\{U(s^+, s^-), \quad \mathcal{C}_t(s^+, s^-)\right\}, \tag{3}$$

[5] The normalizing constant follows from the Vandermonde – Chu identity: $\sum_{m^+=s^+}^{N-s^-} \binom{m^+}{s^+}\binom{N-m^+}{s^-} = \binom{N+1}{s+1}$.

[6] Formally, the problem is a partially observable Markov decision process (POMDP), since M^+ is never directly observed. A standard result establishes that any POMDP can be reformulated as a fully observed MDP by taking the posterior belief over the hidden state as the new state variable (Smallwood and Sondik, 1973). In general this belief-state MDP is infinite-dimensional, but the conjugate structure here (uniform prior, hypergeometric likelihood) implies that the full posterior over $M^+ \in \{0, \ldots, N\}$ is summarized by the integer pair (s^+, s^-), reducing the problem to a finite-dimensional MDP on the triangular lattice $\mathcal{S}$.

where the continuation value is

$$\mathcal{C}_t(s^+, s^-) = -\kappa + \mathbb{E}\big[V_{t+1}(\tilde{s}^+_{t+1}, \tilde{s}^-_{t+1}) \mid s^+, s^-\big]\,. \tag{4}$$

The expectation integrates over the random batch size a_{t+1} and its composition:

$$\mathbb{E}[V_{t+1}] = \sum_{a=0}^{N-s_t} p_{t+1}(a) \sum_{a^+=0}^{a} P(a^+ \mid a, s_t^+, s_t^-)\, V_{t+1}(s_t^+ + a^+,\ s_t^- + a - a^+), \tag{5}$$

where $p_{t+1}(a)$ is the censored Poisson mass at rate $\lambda(t+1)$ and $P(a^+ \mid a, s_t^+, s_t^-)$ is the predictive distribution of Lemma 1.

A *dynamic majority-quorum rule* ϕ is a mapping from states $(s^+, s^-, t) \in \mathcal{S} \times \{1, \ldots, T\}$ to actions $\{\text{stop}, \text{continue}\}$. We denote by ϕ^* the solution to the dynamic program in equations (3)–(5).

Equivalent state representations The planner's state (s_t^+, s_t^-) admits three equivalent representations used throughout the paper: the raw counts (s_t^+, s_t^-), the belief-count pair (μ_t, s_t), and the vote lead-count pair (d_t, s_t). By the Monotone Likelihood Ratio Property of the hypergeometric family (Karlin and Rubin, 1956), $|\mu_t - \frac{1}{2}|$ is strictly increasing in d_t for fixed s_t, so all three representations carry identical information.

4 Optimal Majority-Quorum Rule

This section characterizes the structure of the optimal stopping rule, its comparative statics with respect to the primitives of the environment, and its numerical solution, expressed in the vote lead-count space.

4.1 Structural Properties of the Solution

The central question is when the planner should stop collecting votes and implement the leading alternative. The optimal policy has a clean threshold structure: the planner should wait if the current vote lead is small relative to the remaining uncertainty, and stop once it exceeds a critical level. This threshold shrinks over time as the deadline approaches and the option value of further learning vanishes. The following theorem makes this precise.

Theorem 1. *The optimal voting rule ϕ^* is a dynamic majority-quorum rule: for all $(s_t^+, s_t^-) \in \mathcal{S}$ and $t \in \{1, \ldots, T\}$, there exists a unique threshold $d_t^*(s_t) \geqslant 0$ such that ϕ^* prescribes continuation if and only if $|s_t^+ - s_t^-| < d_t^*(s_t)$. When $s_t = N$, $d_t^*(N) = 0$.*

Proof. See Appendix A.2. □

The threshold structure established in Theorem 1 is a classical feature of Bayesian sequential analysis (DeGroot, 1970; Chow et al., 1971). The proof departs from this standard approach. Because μ is a nonlinear functional of the full posterior over M^+, convexity in μ cannot be established by the usual argument, and we instead prove the threshold structure directly via monotonicity of the net stopping gain $G_t(\mu, s)$ in μ. This monotonicity rests on two model-specific features: the Beta-Binomial predictive distribution (Lemma 1), which delivers the MLRP ordering of future vote compositions in the current lead, and the independence of the batch-size distribution from the current vote composition, which follows from the Poisson arrival assumption and ensures that the FOSD ordering is preserved when averaging over batch sizes. The asymptotic counterpart of the threshold result for hypergeometric sampling is established in Lai (1979). The proof here operates on the exact finite-horizon discrete state space.

The threshold structure has two features worth emphasizing. First, the threshold $d_t^*(s)$ depends on both the vote lead $d = |s^+ - s^-|$ and the total votes cast $s = s^+ + s^-$, rather than on the lead alone. This two-dimensional dependence is a direct consequence of sampling without replacement from a finite population: the total count s determines the size of the remaining pool of voters and hence the distribution of future information. As $s \to N$, the pool shrinks and the finite-population correction amplifies the evidential weight of any given lead, so the threshold adjusts accordingly. This distinguishes the optimal rule from classical fixed-boundary sequential tests, where the relevant state is typically a one-dimensional running score.

Second, the boundary condition $d_t^*(N) = 0$ is immediate: when all votes have been cast, no further information can arrive and the planner implements whichever alternative leads, regardless of the margin. More generally, the threshold is the result of balancing the option value of further learning — which depends on how many votes remain unobserved — against the flow cost of delay. Proposition 1 characterizes how this balance shifts as the model's primitives change.

Proposition 1. *Under ϕ^*, the threshold $d_t^*(s; \rho, \boldsymbol{\lambda})$ with $\rho \equiv R/\kappa$ satisfies:*

(i) $d_t^*(s; \rho', \boldsymbol{\lambda}) \leqslant d_t^*(s; \rho, \boldsymbol{\lambda})$ *for all* $\rho' \leqslant \rho$, t, s;

(ii) $d_t^*(s; \rho, \boldsymbol{\lambda}') \leqslant d_t^*(s; \rho, \boldsymbol{\lambda})$ *whenever* $\lambda'(t) \leqslant \lambda(t)$ *for all* $t \in \{1, \ldots, T\}$;

(iii) $d_{t+1}^*(s; \rho, \boldsymbol{\lambda}) \leqslant d_t^*(s; \rho, \boldsymbol{\lambda})$ *for all* $t < T$, s, *whenever* $\lambda(\tau) \geqslant \lambda(\tau + 1)$ *for all* $\tau \in \{t + 1, \ldots, T - 1\}$.

Proof. See Appendix A.3. □

Intuitively, the optimal stopping threshold for a given vote count is smaller whenever the environment offers less to gain from continued sampling, whether because stakes are lower, the arrival stream is thinner, or the information environment is deteriorating over time. Part (i) establishes that a lower reward-to-cost ratio ρ reduces the continuation value without affecting the stopping payoff, compressing the region where waiting is worthwhile. Part (ii) extends the same logic to the information environment: slower expected arrivals reduce the option value of waiting, lowering the threshold uniformly.

The pointwise dominance condition on the vote arrival process in part (ii) is strong, but weaker candidate assumptions – such as cumulative dominance over the full horizon, i.e. $\sum_{j=1}^{\tau} \lambda'(j) \leqslant \sum_{j=1}^{\tau} \lambda(j)$ for all $\tau \in \{1, \ldots, T\}$, – are not sufficient to guarantee $d_t^*(s; \rho, \boldsymbol{\lambda}') \leqslant d_t^*(s; \rho, \boldsymbol{\lambda})$.[7] The insufficiency of cumulative dominance reflects two offsetting effects. While more total expected arrivals would raise the option value of waiting through a volume effect, front-loading shifts information to earlier periods, so that fewer arrivals remain for the remaining periods, reducing the continuation value. Because the threshold depends on the date-by-date informativeness of future arrivals rather than their cumulative total, this can lower continuation values now even when total arrivals are higher, breaking monotonicity. Pointwise dominance removes this trade-off by ensuring informativeness is weakly higher at every date.

Front-loading of vote arrivals is the key assumption in part (iii), which describes the time path of the threshold within a given environment. While the terminal condition $d_t^*(N) = 0$ ensures eventual collapse (Theorem 1), front-loading guarantees that this collapse is monotone. When participation is non-increasing, each period is the most informative remaining, so the planner never raises the threshold in anticipation of better future information. Under a back-loaded process, the threshold path may not be monotone.[8]

4.2 Numerical Solution

While the preceding results establish the fundamental properties of the optimal policy, the exact stopping boundary cannot be expressed in closed form. Three features of the problem block an analytical solution. First, the state space is two-dimensional: because sampling is without replacement from a finite population, the absolute counts (s^+, s^-) govern the

[7] Consider a minimal counterexample, say with $N = 5$, $T = 2$, $R = 10$, $\kappa = \frac{1}{5}$, where the process $\boldsymbol{\lambda} = (3, 2)$ dominates $\boldsymbol{\lambda}' = (1, 3)$ cumulatively over the full sequence, yet if $(s_1^+, s_1^-) = (2, 0)$, $d_1^*(2; \boldsymbol{\lambda}) < d_1^*(2; \boldsymbol{\lambda}')$. Details upon request.

[8] For instance, consider $N = 7$, $T = 3$, $\boldsymbol{\lambda} = (1, 1, 5)$, $R = 10$, and $\kappa = \frac{1}{2}$. At state $(s_t^+, s_t^-) = (3, 1)$, the stopping payoff is $U(3, 1) = R \cdot \mu(3, 1) = 10 \cdot \frac{13}{14} = \frac{65}{7}$. Backward induction yields $\mathcal{C}_1(3, 1) \approx 9.197 < U(3, 1) \approx 9.286$, so the planner stops at $t = 1$, while $\mathcal{C}_2(3, 1) \approx 9.411 > U(3, 1)$, so the planner continues at $t = 2$. Hence $d_1^*(4) = 2 < d_2^*(4) = 4$, reversing the monotone-collapse conclusion of part (iii). Details upon request.

remaining uncertainty. Second, the transition kernel is non-stationary, as the remaining population shrinks with each arriving batch. Third, the predictive distribution of incoming votes is a belief-dependent mixture of hypergeometrics, precluding a closed-form solution to the resulting functional equation.

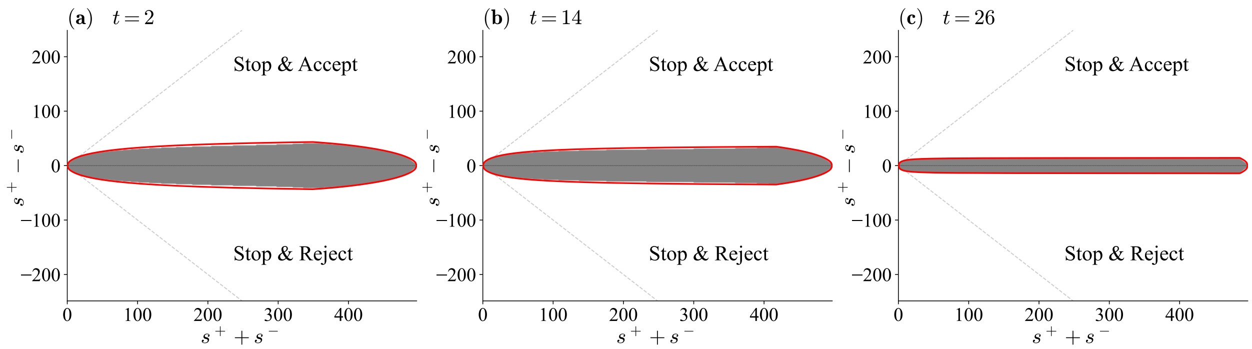

Figure 1: Exact majority-quorum rule ϕ^* for $N = 497, T = 28, \kappa = 0.001, R = 100, \Lambda(T) = 159$, uniform arrival process at different points in time in the vote lead-count space. The gray region corresponds to the states where waiting is optimal. The red curve is the approximated stopping threshold derived in Section 5.

The boundary can nonetheless be computed exactly.[9] Because the cumulative count $s_t = s_t^+ + s_t^-$ only increases, the state space is acyclic, and the value function follows by backward induction from the terminal condition $V_T = U$, with no discretization error. Figure 1 plots the exact continuation region (the set of states where the computed continuation value strictly exceeds the stopping payoff) in vote lead–count space. Three features are worth highlighting. First, the vote-lead threshold $d_t^*(s_t)$ is non-monotone in turnout, producing a lens-shaped continuation region: narrow for small s_t, widest at intermediate s_t, and collapsing toward zero as $s_t \to N$. The widening reflects a standard statistical-precision effect: a small lead is informative only once accompanied by sufficient turnout, while the narrowing near the population boundary reflects the finite-pool correction, which sets it apart from infinite-horizon sequential testing settings. Second, because the arrival rate is constant, the non-increasing condition of Proposition 1(iii) holds with equality, so $d_t^*(s_t)$ is weakly decreasing in t at every fixed s_t and the thresholds accordingly collapse toward zero as $t \to T$. Third,

[9] Code implementing the backward-induction procedure is available at https://github.com/kremenav/mqr.

the boundary is remarkably regular: despite depending on the two-dimensional state (d_t, s_t), it traces what appears to be a single smooth curve through the lattice (captured by the red curve in Figure 1). This regularity suggests that the optimal rule admits a low-dimensional sufficient statistic, which we exploit in the next section to derive a tractable closed-form approximation.

5 Approximate Optimal Majority-Quorum Rule

Although the optimal policy ϕ^* can be computed numerically, it does not admit a closed form. This makes the resulting voting rule opaque and difficult to communicate or implement in practice. We therefore derive a closed-form approximation that captures its essential structure. Let $\hat{s}_T := \min\big\{s_t + \big(\Lambda(T) - \Lambda(t)\big),\, N\big\}$ denote the censored forecast of terminal turnout. We define the following dynamic majority-quorum rule.

Definition 1. *Let $\hat{\phi}$ be the majority-quorum rule stopping at state (d_t, s_t, t), if and only if $t = T$, or $s_t = N$, or $t < T$ and*

$$\frac{\big(d_t/s_t\big)^2}{1/s_t - 1/\hat{s}_T} \;\geqslant\; \big[2\log\rho - 2\log\log\rho\big]\,. \tag{6}$$

The left-hand side of the inequality is the squared signal-to-noise ratio of the vote tally. The numerator measures the strength of the sample vote lead and the denominator is the reducible sampling variance: the current estimation noise $(1/s_t)$ net of the residual noise that will still carry over to the deadline $(1/\hat{s}_T)$. Thus, the denominator isolates strictly the noise that additional waiting can physically eliminate. The critical value $2\log\rho - 2\log\log\rho$ is a constant: the familiar Wald–Chernoff stakes term (Wald and Wolfowitz, 1948), with a $\log\log$ correction, reflecting that the option value of waiting is the expected reversal gain.

The theorem below shows that $\hat{\phi}$ and ϕ^* can disagree only on a band around the stopping thresholds that is vanishingly small relative to its level.

Theorem 2. *Fix T. For every $\varepsilon \in (0, \frac{1}{2})$ there exist N_ε, ρ_ε such that, for all $N \geqslant N_\varepsilon$ and $\rho \geqslant \rho_\varepsilon$ with $\log \rho \leqslant \varepsilon N^{1/3}$, ϕ^* and $\hat{\phi}$ prescribe the same action at every state $(d_t, s_t, t) \in \mathcal{S}_\varepsilon$ at which*

$$\left| \frac{(d_t/s_t)^2}{1/s_t - 1/\hat{s}_T} \cdot \frac{1}{2\log\rho - 2\log\log\rho} - 1 \right| > \frac{\log\log\rho}{\log\rho},$$

where $\mathcal{S}_\varepsilon := \Big\{(d_t, s_t, t) : \ t < T, \quad \varepsilon N \leqslant s_t, \quad s_t + \varepsilon N \leqslant \hat{s}_T \leqslant (1-\varepsilon)N\Big\}.$

Proof. See Appendix A.4. □

Thus, the closed-form rule $\hat{\phi}$ replicates the optimal policy ϕ^* at every state except those within a band around the stopping boundary, whose relative width shrinks to zero as stakes $\rho \to \infty$. This disagreement region is concentrated entirely where the net gain from stopping is near zero, i.e., where the planner is closest to indifferent between acting immediately and collecting additional votes.

Although $\hat{\phi}$ is determined by the primitives of the model before voting begins, it cannot be represented as a dated schedule. Public transparency requires a rule that can be announced as a threshold on a statistic read from the current tally and fixed in advance for each date. The adaptive rule $\hat{\phi}$ does not admit such a representation: realized turnout enters both the evidence statistic and the threshold through the forecast $\hat{s}_T$, so the threshold cannot be specified independently of the state. The statistic $\frac{d_t^2}{s_t} \cdot \frac{N}{N-s_t}$ is $\hat{\phi}$'s evidence measure: the factor $\frac{N}{N-s_t}$ enters both sides of (6) and cancels in the ratio form (Appendix A.4). Carrying it through instead places the finite-population correction on the evidence side, leaving the threshold, which we evaluate along the mean path $s_t \to \Lambda(t)$, fixed in advance for each date.

Definition 2. *Let $\tilde{\phi}$ be the majority-quorum rule stopping at state (d_t, s_t, t), if and only if $t = T$, or $s_t = N$, or $t < T$ and*

$$\frac{d_t^2}{s_t} \cdot \frac{N}{N - s_t} \geqslant \big[2\log\rho - 2\log\log\rho\big] \frac{N\big(\Lambda(T) - \Lambda(t)\big)}{\big(N - \Lambda(t)\big)\Lambda(T)}. \tag{7}$$

Equation (7) separates the stopping decision into an evidence side and a threshold side.

The left depends only on the realized tally: d_t^2/s_t measures how decisively the sample favors the leading alternative, and the multiplier $N/(N - s_t)$ amplifies it as the pool is depleted, reflecting that each observation is more informative when fewer voters remain. Together these constitute the log-likelihood ratio for a Bernoulli parameter under sampling without replacement. When turnout is small relative to N, the correction is negligible and the statistic converges to its classical counterpart. As $s_t \to N$, it diverges, consistent with $d_t^*(N) = 0$.

The right side depends only on primitives and calendar time. The stakes term $2\ln(R/\kappa)$, familiar from Wald–Chernoff sequential testing, with the correction $-2\log\log\rho$, is scaled by a residual-informativeness factor comparing the information still to arrive between t and the deadline to that available at the deadline. This factor equals one at $t = 0$ and vanishes at T, so the threshold begins at the Wald–Chernoff level and declines to zero as the deadline approaches. Because it enters only through cumulative arrivals $\lambda(t)$ and $\Lambda(t)$, any two arrival processes agreeing at the current date and the deadline induce identical thresholds, irrespective of the intervening path.

Figure 2(a) compares the two rules visually: while $\tilde{\phi}$ produces a lens attaining its maximum width at $s_t = N/2$, $\hat{\phi}$ shifts this maximum to the right as dynamic updates to $\hat{s}_T$ delay pool saturation. Despite this shift, the two boundaries remain quantitatively close across realistic parameterizations.

The comparative statics of the approximate rule align with Proposition 1, but display two distinct features. First, because the right-hand side of (7) depends on $\boldsymbol{\lambda}$ only through the cumulative arrivals $\Lambda(t)$ and $\Lambda(T)$, the threshold is monotone in calendar time under any arrival process, whereas Proposition 1(iii) establishes monotonicity only under non-increasing arrival paths. Moreover, comparisons between two arrival processes reduce to cumulative dominance at two dates: the current date and the deadline, rather than the stronger pointwise dominance condition required in Proposition 1(ii). Both simplifications reflect that the approximation abstracts from the discrete option-value effects of back-loading that drive the restrictions in Proposition 1(ii) and (iii) (see footnotes 7 and 8). Second, the

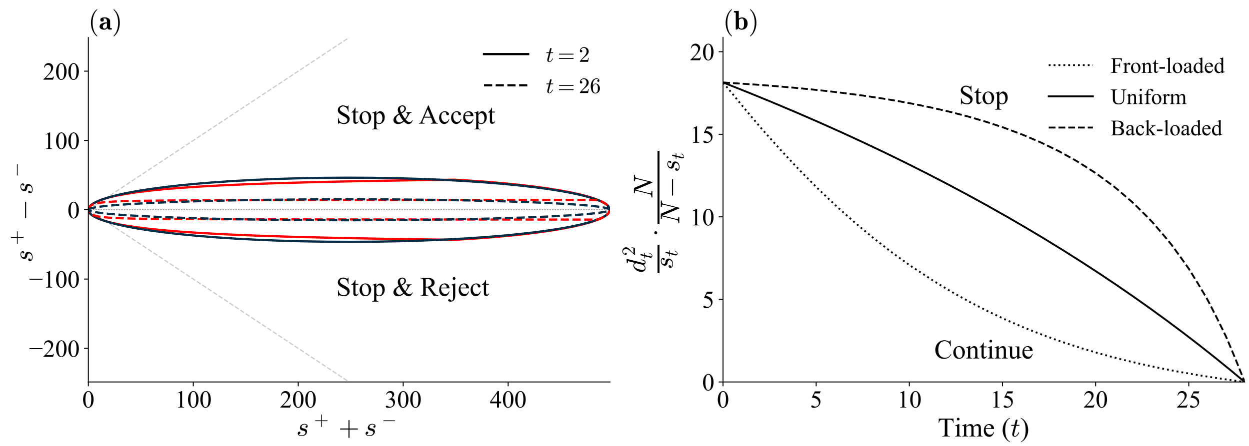

Figure 2: Approximate dynamic-quorum rules for $N = 497, T = 28, \kappa = 0.001, R = 100, \Lambda(T) = 159$. (a) Comparison between $\hat{\phi}$ and $\tilde{\phi}$ under a uniform arrival process; (b) $\tilde{\phi}$ over time for uniform, front- and back-loaded arrival processes.

shape of $\lambda(\cdot)$ determines how quickly the boundary collapses to zero over the time horizon: front-loaded arrival processes concentrate information early, producing a rapid initial decline followed by flattening near the deadline, whereas back-loaded processes generate the opposite pattern, with a slow early decline and a sharper late collapse. This difference in the concavity of the stopping threshold functions over time is visible in Figure 2(b), which illustrates the voting rule for a uniform, front- and back-loaded arrival process, while holding total arrival constant.

The rule $\tilde{\phi}$ is stated in terms of the stakes-to-cost ratio ρ, a parameter that is difficult to reason about in practice. The following corollary shows that it can be restated in terms of an error bound.

Corollary 1. *Let $\bar{t} \in [0, T]$ be a reference date. The rule $\tilde{\phi}$ stops if and only if*

$$\frac{d_t^2}{s_t} \cdot \frac{N}{N - s_t} \geqslant \left[\Phi^{-1}(1 - \alpha_{\bar{t}})\right]^2 \frac{\Lambda(T) - \Lambda(t)}{N - \Lambda(t)} \cdot \frac{N - \Lambda(\bar{t})}{\Lambda(T) - \Lambda(\bar{t})}, \tag{8}$$

where $\alpha_{\bar{t}} \in (0, \frac{1}{2})$ is the error tolerance at $\bar{t}$, i.e. the posterior probability of implementing the *minority preference is at most $\alpha_{\bar{t}}$ for any decision reached at or before $\bar{t}$ and is given by*

$$\alpha_{\bar{t}} \;=\; 1 - \Phi\left(\sqrt{\big[2\log\rho - 2\log\log\rho\big]\,\frac{N\big(\Lambda(T)-\Lambda(\bar{t})\big)}{\big(N-\Lambda(\bar{t})\big)\Lambda(T)}}\right). \tag{9}$$

Proof. See Appendix A.5. □

Thus, the planner stops as soon as the posterior probability that the leading alternative is not the majority preference falls below $\alpha_{\bar{t}}$.[10] The optimal mechanism is therefore a sequential test whose significance level is indexed by calendar time. That level increases monotonically over the horizon, from its tightest value at $t = 0$ to one half at the deadline, where the planner implements whichever alternative leads, such that the confidence required falls as the opportunity to acquire it through further votes disappears.

6 Empirical Analysis

We evaluate the approximate rule $\tilde{\phi}$ given in Definition 2 on governance data from Polkadot, asking whether the rule could have reached the same decisions faster than the currently deployed mechanism. Our analysis is designed to be forward-looking rather than retrospective: every model parameter is calibrated on the chronologically first half of the sample and then applied to the second half of the sample. Performance is therefore reported exclusively out of sample, as it would have accrued to a planner who adopted the rule midway through the observation period.

6.1 Setting and Data

Polkadot (Wood, 2022) is among the largest decentralized autonomous organizations by market capitalization, and its governance framework, *OpenGov*, decides all protocol changes (i.e., code upgrades) and treasury disbursements through binding on-chain referenda. Referenda

[10]The reference date is a normalisation rather than a second parameter. The equivalence holds for every $\bar{t}$, since moving it rescales α and the shape factor by reciprocal amounts.

are organized into tracks that cap the scope of the proposed change and share a common 28-day decision period. Crucially for our purposes, Polkadot already applies a variant of a dynamic majority-quorum rule, the Approval-Support Mechanism (ASM). A referendum concludes before the deadline if its conviction-weighted approval share and its turnout each exceed track-specific, monotonically declining thresholds for a sustained confirmation period. The ASM thus provides a natural benchmark, and, because voters already face a dynamic rule, strategic timing of participation is possible in this setting, making it a demanding test of the model's random-arrival assumption. The blockchain records every vote, so we observe the full arrival path of each referendum at twice-daily resolution, mapping directly onto the state variable (s_t^+, s_t^-) of the model.

We collect the complete record of 1,824 OpenGov referenda submitted between June 2023 and October 2025. Because the model treats each vote as one signal drawn from a finite electorate, we measure participation in *votes by accounts* rather than in tokens: s_t counts the accounts that have voted by date t and d_t is the lead in accounts. Token-weighting would let stake size masquerade as sample size, because tokens are naturally and frequently concentrated into large holders. The evidence would enter the statistic in Equation (7) as millions of independent signals which would lead to an overestimation of its signal weight. The act of voting from an account count preserves the informativeness of a vote and better applies to the structure of the model.[11]

From the raw record we retain referenda that reached a substantive disposition (approved or rejected) on the six focus tracks, and we drop faulty and broken observations: referenda, whose own proposer instructed voters to reject (due to faulty meta data), and tally paths too degenerate to support estimation of an arrival process (too few snapshots, single-jump paths, or large non-monotonicities). The cleaned sample comprises 952 referenda, of which 591 were approved and 361 rejected. The stopping evaluation further excludes the `Root`

[11] Polkadot allows accounts to delegate their votes. Delegated votes are registered in bulk at the moment the delegate votes. Therefore, we measure arrivals by directly voting accounts and read the aye–nay split from the account tally. The conviction-weighted ASM outcome and the final account majority disagree in roughly 17% of referenda, a wedge we return to when defining the error benchmark below.

track, whose single-slot queue mechanically distorts timing, leaving 933 referenda for the stopping evaluation (467 in the first half of the sample, 466 in the second).

6.2 Empirical Strategy

We split the cleaned sample at the median referendum index into a training half and a test half of 476 referenda each. All quantities entering Equation (8) are computed from the training half alone.

Environment primitives. The model's arrival process is a non-homogeneous Poisson process with intensity $\lambda(t)$. We adopt a parsimonious log-linear specification in normalized referendum time $x \in [0, 1]$,

$$\lambda(x) = e^{a+bx}, \qquad F(x;b) \;=\; \frac{e^{bx}-1}{e^{b}-1}, \tag{10}$$

where F is the implied cumulative arrival fraction and the shape parameter b summarizes the temporal pattern ($b < 0$: front-loaded participation; $b \to 0$: uniform). We fit F to each training referendum's cumulative arrival path by non-linear least squares. The fits are tight, and participation is strongly front-loaded, with a pooled (median) estimate of $b = -3.05$. The cumulative expected arrivals entering the rule are then $\Lambda(t) = \Lambda(T)\, F(t/T;\, b)$, with the expected final turnout $\Lambda(T) = 159$ accounts set to the training-half mean and the electorate size $N = 496$ set to the maximum training-half turnout.[12]

Planner preferences. We evaluate the rule in the parameterization of Corollary 1, which replaces the reward-to-cost ratio R/κ with two interpretable quantities: a per-decision error tolerance α and a reference date $\bar{t}$ up to which the tolerance is guaranteed – any decision reached at or before $\bar{t}$ implements the minority preference with probability at most α. This

[12] Pooling the arrival shape across referenda preserves the planner's information structure: a planner deploying the rule in real time cannot fit the intensity of the referendum at hand. The calibration of N is deliberately conservative. Smaller electorates make the rule more patient, as the residual-informativeness factor in Equation (7) then binds more tightly.

is calibration in the spirit of sequential testing, where boundaries are set by target error probabilities rather than by primitive payoffs (Wald and Wolfowitz, 1948). The corollary's guarantee is exact under the model's sampling assumptions, but when early arrivals are not a representative draw from the electorate, the realized error rate exceeds the stated α, and the shortfall is governed by $\bar{t}$. We therefore treat α as the planner's stated tolerance and *derive* $\bar{t}$ *from the training half*: $\bar{t}^*(\alpha)$ is the earliest reference date whose implied threshold keeps the realized training-half error within α. The derived $\bar{t}^*$ is thus a measured *sufficient-guarantee horizon* – how far into the decision period the model's guarantee must extend before it can be trusted in this environment. Errors are measured as *divergent decisions*: referenda in which the rule's early decision differs from the decision the same account-majority criterion would produce at the deadline. This benchmark isolates genuine stopping errors from the conviction-weighting wedge described above, which no account-based rule can close at any stopping time.[13]

For the evaluation, we replay each test referendum's realized vote path in chronological order and record the first date at which the rule of Corollary 1, with the frozen pair $(\alpha, \bar{t}^*)$ and the training-half calibration, fires – together with the leading alternative at that date. If the rule never fires, it decides with the deadline majority. Time saved is measured against the referendum's realized duration under the ASM.

6.3 Out-of-Sample Results

Table 1 reports the results for tolerances $\alpha \in \{10\%, 5\%, 2\%, 1\%\}$. Two findings stand out. First, the derived reference dates are remarkably stable: across tolerances spanning an order of magnitude, $\bar{t}^*$ varies only between day 20.3 and day 24.2 of the 28-day period. This statistic must be read correctly: $\bar{t}^*$ anchors the *height* of the stopping threshold, not the

[13] At fixed calibration, $(\alpha, \bar{t})$ enters the rule only through the threshold constant $[\Phi^{-1}(1-\alpha)]^2 \, \frac{N-\Lambda(\bar{t})}{\Lambda(T)-\Lambda(\bar{t})}$, which plays the role of the bracketed level in Theorem 2. Deriving $\bar{t}^*$ is therefore equivalent to selecting the smallest such constant that meets the tolerance in-sample, and the pair $(\alpha, \bar{t}^*)$ labels the resulting threshold in interpretable units.

timing of decisions. A reference date of day 23 means the threshold is set high enough that every decision reached up to day 23 – by when roughly 96% of the expected electorate has spoken – meets the tolerance. The rule itself stops much earlier, whenever the observed evidence crosses that bar. At the headline tolerance below, the median test referendum is decided on day 10.8 of 28, with 60% of its final turnout cast (interquartile range 50–77%). The stability of $\bar{t}^*$ across tolerances is a compact summary of how unrepresentative early participation currently is. Everything else in the calibration is either a policy choice (α) or a routine estimate (b, N, $\Lambda(T)$). Second, the pairs $(\alpha, \bar{t}^*)$ selected on the training half transfer to the test half with little degradation: at the headline tolerance of $\alpha = 2\%$ (with derived $\bar{t}^* = 22.9$ days), the rule diverges from the deadline decision in 1.72% of test referenda while deciding the median referendum 15.0 days earlier than the ASM, whose median realized duration in the test half is 28 days. A planner willing to accept two divergent decisions per hundred referenda would thus have cut decision times by roughly half, using only information available before the evaluated referenda were submitted.

Table 1: Out-of-sample performance of $\tilde{\phi}$ under the $(\alpha, \bar{t})$ parameterization of Corollary 1. For each tolerance α, the reference date $\bar{t}^*$ is the earliest date whose implied threshold keeps the training-half divergence within α; the pair $(\alpha, \bar{t}^*)$ is then applied unchanged to the test half.

Tolerance	$\bar{t}^*$ (day)	Divergent decisions (%)		Median time saved (days)	
α (%)	(train)	train	test	train	test
10	20.3	9.9	11.6	22.1	22.0
5	21.6	4.5	5.2	19.8	20.0
2	22.9	1.9	1.7	15.4	15.0
1	24.2	0.9	1.3	11.6	11.5

Figure 3 traces the full speed–accuracy frontier on both halves, swept by varying the reference date $\bar{t}$ at a fixed tolerance (equivalently, the implied evidence threshold). The two curves nearly coincide, indicating that the frontier is a stable property of the voting environment rather than an artifact of in-sample tuning. Its shape carries the central economic message: the frontier is steep at low error rates and flat at high ones, so almost all of the attainable time saving is available at near-perfect accuracy, while the final percentage points of accuracy are extremely expensive. Exact zero divergence is unattainable altogether: as $\bar{t} \to T$ the threshold diverges and the rule never stops before the deadline, so suppressing the last divergent decision means forfeiting the mechanism entirely. The planner's choice is therefore not whether to trade accuracy for speed, but how cheaply the first few basis points of error buy the bulk of the delay savings.

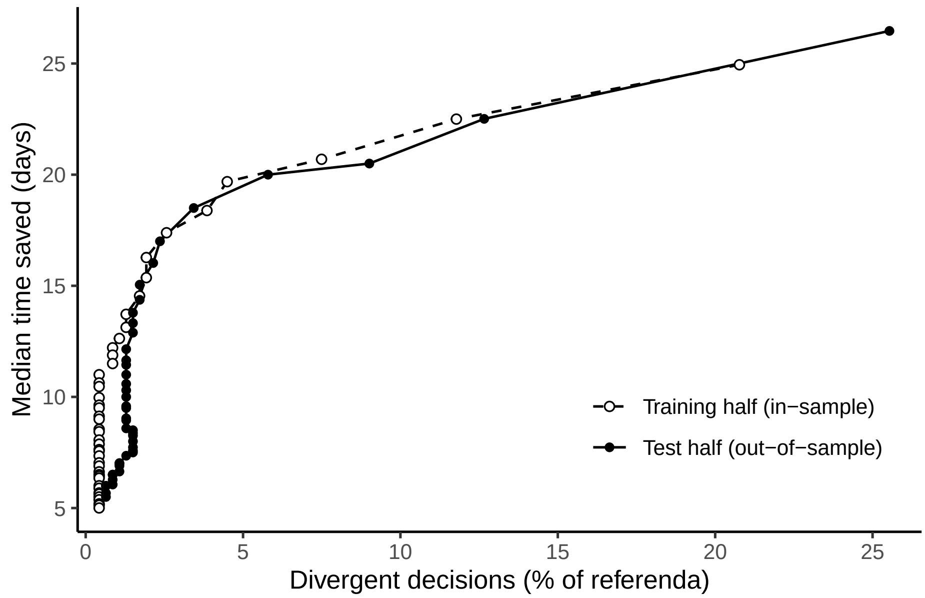

Figure 3: Speed–accuracy frontier of $\tilde{\phi}$, traced by varying the reference date $\bar{t}$, on the training half (calibration sample) and the test half (out-of-sample). Each marker is one threshold on the sweep grid. Divergent decisions are measured against the deadline account majority.

7 Discussion and Conclusion

In many collective decision problems, waiting for full participation is neither feasible nor desirable. This paper argues that the sequential arrival of votes can and should inform the decision of when to stop. We develop a framework that treats collective decision-making as an optimal stopping problem over a finite, exhaustible population of voters, bridging statistical social choice and sequential analysis, and characterize the optimal dynamic majority-quorum rule: a joint, time-varying threshold on vote lead and turnout that collapses as the deadline approaches and the option value of further learning vanishes. The rule responds naturally to the primitives of the environment and admits a tractable closed-form approximation. Evaluated using governance data from Polkadot, the rule would have reached the same decisions substantially faster, suggesting meaningful gains in the speed-accuracy trade-off.

The framework focuses on the planner's problem in a partial equilibrium setting. By taking the vote-generating process as a primitive, we obtain a tractable derivation of the planner's optimal rule and a transparent mapping from primitives to outcomes. This opens a research agenda on dynamic collective decision-making that relaxes the exogeneity of voter behavior and the vote-generating process.

A first direction endogenizes participation decisions. Voters choose whether and when to participate based on costs and perceived pivotality, implying that arrival patterns are determined in equilibrium rather than taken as given. This introduces a feedback between turnout and the same notion of pivotality that governs the planner's stopping decision, suggesting a natural fixed-point structure linking participation and aggregation.

A second direction concerns the policy invariance of the arrival process, which we flag as the principal caveat to our counterfactual. The evaluation in Section 6 replays historical vote paths and therefore holds voter behavior fixed. Adopting $\tilde{\phi}$ would change the incentives that generate those paths, since termination itself becomes a function of the running tally. Whether this invalidates the rule depends on the symmetry of the behavioral response, not on its size. The rule's inference requires only that arrival timing be uninformative about preferences: if supporters and opposers shift their timing in the same way – say, participation generally moves earlier once voters anticipate early termination – the tally remains a representative sample at every date and the rule survives intact with a recalibrated $\Lambda(\cdot)$. There is a structural reason to expect precisely this symmetric response. Because $\tilde{\phi}$ conditions only on the absolute lead d_t and turnout, it terminates early in *both* directions – early acceptance and early rejection are equally attainable – so the mechanism is neutral between the alternatives, and any timing incentive it creates attaches to a voter's anticipated position in the tally (leading or trailing), not to the alternative they favor. This neutrality distinguishes $\tilde{\phi}$ from the incumbent ASM, which permits early *approval* only and thereby hard-wires asymmetric timing incentives into the status quo: supporters can accelerate execution, while opposers can only obstruct it. In this respect, adopting $\tilde{\phi}$ would remove a built-in asymmetry rather than introduce one.

The symmetry argument is a hypothesis about equilibrium behavior, however, not a theorem: incentives that attach to the anticipated leading or trailing position can still correlate timing with the realized majority, and the minimum-turnout requirement embedded in the lens-shaped boundary protects against small coalitions but not large ones. Because our data were generated under the asymmetric incumbent rule, they cannot confirm or refute behavior under a symmetric one. Doing so requires deliberately generated variation – a controlled experiment that randomizes the stopping rule while holding preferences fixed, or a staged field deployment on low-stakes decision tracks with the existing deadline retained as a backstop – which we view as a prerequisite for adoption and the natural next step of this research agenda.

References

Arbitrum Foundation (2023). The Constitution of the Arbitrum DAO. https://docs.arbitrum.foundation/dao-constitution. Section 4: Security Council Elections. Accessed: 2026-08-02.

Austen-Smith, D. and Banks, J. S. (1996). Information aggregation, rationality, and the Condorcet Jury Theorem. *American Political Science Review*, 90(1):34–45.

Azari Soufiani, H., Parkes, D. C., and Xia, L. (2014). A Statistical Decision-Theoretic Framework for Social Choice. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.

Bognar, K., Börgers, T., and Meyer-ter-Vehn, M. (2015). An optimal voting procedure when voting is costly. *Journal of Economic Theory*, 159:1056–1073.

Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.

Chernoff, H. (1961). Sequential experimentation. *Bulletin of the International Statistical Institute*, 38(IV):3–9.

Chow, Y. S., Robbins, H., and Siegmund, D. (1971). *Great Expectations: The Theory of Optimal Stopping*. Houghton Mifflin, Boston.

Condorcet, M. d. (1785). *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. Imprimerie Royale.

DeGroot, M. H. (1970). *Optimal statistical decisions*. McGraw-Hill, Inc.

Drugowitsch, J., Moreno-Bote, R., Churchland, A., Shadlen, M., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. *Journal of Neuroscience*, 32(11):3612–3628.

Feddersen, T. and Pesendorfer, W. (1997). Voting behavior and information aggregation in elections with private information. *Econometrica*, 65(5):1029–1058.

Fudenberg, D., Strack, P., and Strzalecki, T. (2018). Speed, accuracy, and the optimal timing of choices. *American Economic Review*, 108(12):3651–3684.

Gershkov, A., Moldovanu, B., and Shi, X. (2017). Optimal voting rules. *The Review of Economic Studies*, 84(2):688–717.

Gershkov, A. and Szentes, B. (2009). Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1):36–68.

Johnson, N. L., Kotz, S., and Kemp, A. W. (1992). *Univariate Discrete Distributions*. John Wiley & Sons, New York, 2nd edition.

Karatzas, I. and Shreve, S. (1991). *Brownian Motion and Stochastic Calculus*. Springer New York.

Karlin, S. and Rubin, H. (1956). The theory of decision procedures for distributions with monotone likelihood ratio. *Annals of Mathematical Statistics*, 27(2):272–299.

Kato, M. and Nakagawa, K. (2026). Sequential audit sampling with statistical guarantees. *arXiv preprint arXiv:2604.06116*.

Lai, T. L. (1979). Sequential tests for hypergeometric distributions and finite populations. *The Annals of Statistics*, pages 46–59.

Lindeman, M. and Stark, P. B. (2012). A gentle introduction to risk-limiting audits. *IEEE Security & Privacy*, 10(5):42–49.

Maryland State Board of Elections (2016). Post-election tabulation audit pilot program report. Technical report, Maryland State Board of Elections. Retrieved March 23, 2026.

Moscarini, G. and Smith, L. (2001). The optimal level of experimentation. *Econometrica*, 69(6):1629–1644.

Myerson, R. B. (1998). Population uncertainty and poisson games. *International Journal of Game Theory*, 27(3):375–392.

Nitzan, S. and Paroush, J. (1982). Optimal decision rules in uncertain dichotomous choice situations. *International Economic Review*, pages 289–297.

Nitzan, S. and Paroush, J. (1984). The significance of independent decisions in uncertain dichotomous choice situations. *Theory and Decision*, 17(1):47–60.

Persico, N. (2004). Committee design with endogenous information. *The Review of Economic Studies*, 71(1):165–191.

Peskir, G. and Shiryaev, A. (2006). *Optimal stopping and free-boundary problems*. Springer.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.

Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922.

Sheng, V. S., Provost, F., and Ipeirotis, P. G. (2008). Get another label? improving data quality and data mining using multiple, noisy labelers. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 614–622.

Siegmund, D. (1985). *Sequential Analysis: Tests and Confidence Intervals*. Springer.

Smallwood, R. D. and Sondik, E. J. (1973). The optimal control of partially observable markov processes over a finite horizon. *Operations research*, 21(5):1071–1088.

Snow, R., O'Connor, B., Jurafsky, D., and Ng, A. (2008). Cheap and fast – but is it good? Evaluating non-expert annotations for natural language tasks. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 254–263, Honolulu, Hawaii. Association for Computational Linguistics.

Stark, P. B. (2008). Conservative statistical post-election audits. *The Annals of Applied Statistics*, 2(2):550–581.

Tajima, S., Drugowitsch, J., and Pouget, A. (2016). Optimal policy for value-based decision-making. *Nature Communications*, 7(12400).

Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge university press.

Wald, A. (1950). Statistical decision functions. In *Breakthroughs in Statistics: Foundations and Basic Theory*, pages 342–357. Springer.

Wald, A. and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, pages 326–339.

Wood, G. (2022). Polkadot v1.0: Sharding and economic security. Whitepaper. Available at https://polkadot.network/blog/polkadot-v1-0-sharding-and-economic-security/.

Xiong, X. (1993). *Sequential tests for hypergeometric distribution.* PhD thesis, Purdue University.

Young, H. P. (1988). Condorcet's theory of voting. *American Political Science Review*, 82(4):1231–1244.

A Proofs

A.1 Proof of Lemma 1

By the Bayes–Laplace result, the discrete uniform prior on M^+ is the marginal of Binomial(N, p) over $p \sim$ Beta$(1, 1)$, so the posterior over p given (s^+, s^-) is Beta$(s^+ + 1, s^- + 1)$. Under a uniform prior, the hypergeometric predictive coincides with the Beta-Binomial predictive (see, e.g. Johnson et al. (1992)), so we integrate the Binomial(a, p) likelihood against the posterior:

$$\begin{aligned}
P(a^+ \mid a, s^+, s^-) &= \int_0^1 \mathrm{P}(a^+ \mid a, p) f(p \mid s^+, s^-), dp \\
&= \int_0^1 \left[\binom{a}{a^+} p^{a^+} (1-p)^{a-a^+} \right] \left[\frac{p^{s^+}(1-p)^{s^-}}{\int_0^1 u^{s^+}(1-u)^{s^-} du} \right] dp \\
&= \binom{a}{a^+} \frac{\int_0^1 p^{a^+ + s^+}(1-p)^{a-a^+ + s^-} dp}{\int_0^1 p^{s^+}(1-p)^{s^-} dp} \\
&= \binom{a}{a^+} \frac{B(a^+ + s^+ + 1; a - a^+ + s^- + 1)}{B(s^+ + 1; s^- + 1)} \\
&= \frac{\binom{a^+ + s^+}{s^+}\binom{a - a^+ + s^-}{s^-}}{\binom{a+s+1}{a}}
\end{aligned}$$

where the last step expands the Beta functions via $B(\alpha,\beta) = \Gamma(\alpha)\Gamma(\beta)/\Gamma(\alpha+\beta)$ and cancels factorials.

A.2 Proof of Theorem 1

The proof proceeds by backward induction on the net gain from stopping $G_t(\mu, s) := U(\mu, s) - \mathcal{C}_t(\mu, s)$, establishing that G_t is non-decreasing in μ at every period t. The base case follows from the belief martingale property, which we establish in Lemma 2 below. The inductive step decomposes G_t into two terms each inheriting monotonicity from the previous period. The threshold structure follows from the boundary conditions and the MLRP of the hypergeometric family. We first show the martingale property of posterior beliefs.

Lemma 2. *Let $(s^+, s^-) \in \mathcal{S}$ with $s^+ + s^- < N$, and let a be any random batch drawn from the remaining $N - s$ voters. Then,* $\mathbb{E}\big[\mu(s^+ + a^+,\, s^- + a^-) \,\big|\, s^+, s^-\big] = \mu(s^+, s^-)$.

Proof. By definition, $\mu(s^+, s^-) = \mathbb{E}[\mathbf{1}_{\{M^+>N/2\}} \mid s^+, s^-]$, and similarly for $\mu(s^++a^+, s^-+a^-)$. Since the augmented tally contains (s^+, s^-), the law of iterated expectations gives $\mathbb{E}\big[\mu(s^+ + a^+, s^- + a^-) \mid s^+, s^-\big] = \mathbb{E}\big[\mathbb{E}[\mathbf{1}_{\{M^+>N/2\}} \mid s^+ + a^+, s^- + a^-] \,\big|\, s^+, s^-\big] = \mu(s^+, s^-)$. $\square$

The key input at each stage of the induction is that a higher current belief μ_t shifts the posterior belief μ_{t+1} stochastically upward, a property we establish in Lemma 3 below. The lemma proceeds by showing MLRP of the predictive distribution in μ_t conditional on batch size, which directly implies FOSD of μ_{t+1} in μ_t conditional on batch size, and then averages over batch sizes, which preserves the ordering because the batch-size distribution is independent of μ_t.

We restrict attention to $\mu_t \in [\frac{1}{2}, 1]$ for now. For fixed s_t, μ_t is in bijection with $d_t = |s_t^+ - s_t^-|$ by the MLRP of the hypergeometric family (Karlin and Rubin, 1956), so conditioning on $\mu_{1,t} > \mu_{2,t} \geqslant \frac{1}{2}$ is equivalent to conditioning on $d_{1,t} > d_{2,t} \geqslant 0$ with s_t fixed.

Lemma 3. *For any* $\mu_{1,t} > \mu_{2,t} \geqslant \frac{1}{2}$,

$$P(\mu_{t+1} \leqslant x \mid \mu_{1,t}, s_t) \leqslant P(\mu_{t+1} \leqslant x \mid \mu_{2,t}, s_t) \qquad \textit{for all } x \in [0,1].$$

Proof. The argument proceeds in two steps. First, we show the MLRP of the predictive distribution and the belief.

Since $\mu_t \geqslant \frac{1}{2}$, it is strictly increasing in s_t^+ for fixed s_t by the MLRP of the hypergeometric family (Karlin and Rubin, 1956). Using the predictive distribution of Lemma 1, the likelihood ratio for consecutive values of a^+ is

$$\frac{P(a^+ = k+1 \mid a, s_t^+, s_t^-)}{P(a^+ = k \mid a, s_t^+, s_t^-)} = \frac{\binom{k+1+s_t^+}{s_t^+}}{\binom{k+s_t^+}{s_t^+}} \cdot \frac{\binom{a-k-1+s_t^-}{s_t^-}}{\binom{a-k+s_t^-}{s_t^-}} = \frac{k+1+s_t^+}{k+1} \cdot \frac{a-k}{a-k+s_t^-}.$$

which is strictly increasing in s_t^+ and strictly decreasing in s_t^-, holding $s_t = s_t^+ + s_t^-$ fixed. This establishes MLRP of a^+ in d_t, which implies FOSD of a^+ in d_t, and hence FOSD of a^+ in μ_t by the monotonic relationship between μ_t and d_t: for any $\mu_{1,t} > \mu_{2,t} \geqslant \frac{1}{2}$ and every threshold j,

$$P(a^+ \geqslant j \mid a, \mu_{1,t}, s_t) \geqslant P(a^+ \geqslant j \mid a, \mu_{2,t}, s_t).$$

Since μ_{t+1} is strictly increasing in a^+ for fixed (s_t, a), FOSD of a^+ in μ_t implies FOSD of μ_{t+1} in μ_t, conditional on every realization of a.

Next, we show that averaging over a preserves FOSD. By the law of total probability,

$$P(\mu_{t+1} \leqslant x \mid \mu_t, s_t) = \sum_{a=0}^{N-s_t} P(a \mid s_t)\, P(\mu_{t+1} \leqslant x \mid \mu_t, s_t, a).$$

Since $a_t = \min\{A_t, N - s_t\}$ follows a censored Poisson with rate $\lambda(t)$, the weights $P(a \mid s_t)$ depend only on s_t and t, not on μ_t. Since $P(\mu_{t+1} \leqslant x \mid \mu_{1,t}, s_t, a) \leqslant P(\mu_{t+1} \leqslant x \mid \mu_{2,t}, s_t, a)$ for every a (Step 1) and the weights $P(a \mid s_t)$ are identical for $\mu_{1,t}$ and $\mu_{2,t}$, it is preserved under summation. □

We prove Theorem 1 by backward induction on the statement: For each s_t, the net gain from stopping $G_t(\mu_t, s_t)$ is non-decreasing in μ_t on $[\frac{1}{2}, 1]$, implying the threshold structure.

Terminal state. At $t = T$ the planner is forced to stop, so $V_T = U$. One step before the deadline, $G_{T-1}(\mu_{T-1}, s_{T-1}) = \kappa + U(\mu_{T-1}, s_{T-1}) - \mathbb{E}\big[U(\mu_T, s_T) \mid \mu_{T-1}, s_{T-1}\big]$. Using $U(\mu) = R\max\{\mu, 1-\mu\}$[14] and the belief martingale $\mathbb{E}[\mu_T \mid \mu_{T-1}, s_{T-1}] = \mu_{T-1}$ (Lemma 2):

$$G_{T-1}(\mu_{T-1}, s_{T-1}) = \kappa - R\,\mathbb{E}\big[\max\{0,\, 1 - 2\mu_T\} \mid \mu_{T-1}, s_{T-1}\big].$$

The function $x \mapsto \max\{0, 1-2x\}$ is non-increasing. By Lemma 3, higher μ_{T-1} induces FOSD on μ_T, so the expectation of this non-increasing function is non-increasing in μ_{T-1}. Hence G_{T-1} is non-decreasing in μ_{T-1}.

Inductive step. Assume $G_{t+1}(\mu_{t+1}, s_{t+1})$ is non-decreasing in μ_{t+1}, so $V_{t+1} - U = \max\{0, \mathcal{C}_{t+1} - U\}$ is non-increasing in μ_{t+1}. Write:

$$G_t(\mu_t, s_t) = \kappa - \underbrace{\mathbb{E}\big[V_{t+1}(\mu_{t+1}, s_{t+1}) - U(\mu_{t+1}, s_{t+1}) \mid \mu_t, s_t\big]}_{\text{Term A}} + \underbrace{U(\mu_t, s_t) - \mathbb{E}\big[U(\mu_{t+1}, s_{t+1}) \mid \mu_t, s_t\big]}_{\text{Term B}}.$$

Term A is non-increasing in μ_{t+1} by the inductive hypothesis, so its expectation is non-increasing in μ_t by Lemma 3. Since Term A enters with a negative sign in G_t, it is therefore non-decreasing in μ_t. Term B equals $-R\,\mathbb{E}[\max\{0, 1 - 2\mu_{t+1}\} \mid \mu_t, s_t]$, since the belief martingale $\mathbb{E}[\mu_{t+1} \mid \mu_t, s_t] = \mu_t$ (Lemma 2) and the time-invariance of $U(\mu)$, similarly to the argument in the base case. This is non-decreasing in μ_t by Lemma 3. Thus, both terms are non-decreasing in μ_t, so G_t is non-decreasing.

Observe that if $\mu_t = 1$, the planner is certain the proposal has majority support, so the continuation value cannot exceed $R - \kappa < R = U(1, s_t)$, giving $G_t(1, s_t) = \kappa > 0$ and stopping is optimal. Since $G_t(\mu_t, s_t)$ is non-decreasing in μ_t (Step 2) and positive at $\mu_t = 1$, two cases arise at $\mu_t = \frac{1}{2}$: if $G_t(\frac{1}{2}, s_t) \geqslant 0$, then $G_t(\mu_t, s_t) \geqslant 0$ for all $\mu_t \in [\frac{1}{2}, 1]$, so stopping is optimal everywhere and the threshold is degenerate at $\mu_t^*(s_t) = \frac{1}{2}$ (equivalently,

[14] Note that $U(\mu_t, s_t) = R\max\{\mu_t, 1 - \mu_t\}$ depends on s_t only through μ_t.

$d_t^*(s_t) = 0$); if $G_t(\frac{1}{2}, s_t) < 0$, then by monotonicity there exists a unique $\mu_t^*(s_t) \in (\frac{1}{2}, 1)$ such that $G_t(\mu_t^*(s_t), s_t) = 0$. In both cases, continuation is optimal if and only if $\mu_t < \mu_t^*(s_t)$, where $\mu_t^*(s_t) \in [\frac{1}{2}, 1)$.

By the MLRP of the hypergeometric family (Karlin and Rubin, 1956), μ_t is strictly increasing in $d_t = |s_t^+ - s_t^-|$ for fixed s_t, so $\mu_t^*(s_t)$ corresponds to a unique vote-lead threshold $d_t^*(s_t) \geqslant 0$ on $\mu_t \in [\frac{1}{2}, 1]$. By symmetry of the model in s_t^+ and s_t^-, the same threshold applies on $\mu_t \in [0, \frac{1}{2}]$, so continuation is optimal if and only if $|s_t^+ - s_t^-| < d_t^*(s_t)$.

We established the result for $\mu_t \in [\frac{1}{2}, 1]$. The result for $\mu_t \in [0, \frac{1}{2}]$ follows by symmetry of the model in s_t^+ and s_t^-, which implies that $G_t(\mu_t, s_t)$ is non-increasing on $[0, \frac{1}{2}]$, negative at $\mu_t = 0$, and the threshold structure holds by an identical argument.

When $s_t = N$, no further votes can arrive, so $\mathcal{C}_t(\mu_t, N) = U(\mu_t, N) - \kappa < U(\mu_t, N)$ for all μ_t, giving $d_t^*(N) = 0$.

A.3 Proof of Proposition 1

Before we prove the proposition, we show the submartingale property of the value function in the batch size, used in the proofs of parts (ii) and (iii). We denote the expected value of the planner's problem from period t onward, conditional on the current tally (s^+, s^-) and the next batch of size a, with $\mathcal{I}_t(a; s^+, s^-) := \mathbb{E}[V_t(s^+ + a^+,\ s^- + a - a^+) \mid a,\ s^+,\ s^-]$.

Lemma 4. *For any $t \in \{1, \ldots, T\}$, state $(s^+, s^-) \in \mathcal{S}$, and batch sizes $0 \leqslant a \leqslant a' \leqslant N - s^+ - s^-$, $\mathcal{I}_t(a; s^+, s^-) \leqslant \mathcal{I}_t(a'; s^+, s^-)$.*

Proof. It suffices to prove the one-step version,

$$\mathcal{I}_t(a; s^+, s^-) \leqslant \mathcal{I}_t(a+1; s^+, s^-) \qquad \text{for all } 0 \leqslant a < N - s, \tag{11}$$

since iterating gives $\mathcal{I}_t(a; \cdot) \leqslant \mathcal{I}_t(a+1; \cdot) \leqslant \cdots \leqslant \mathcal{I}_t(a'; \cdot)$ for any $a \leqslant a'$. We prove (11) by backward induction on t.

Terminal state. We have $\mathcal{I}_T(a; s_{T-1}^+, s_{T-1}^-) = \mathbb{E}\big[U(s_{T-1}^+ + a^+,\ s_{T-1}^- + a - a^+) \mid a,\ s_{T-1}^+,\ s_{T-1}^-\big]$.

Let $\mu_T := \mu(s^+_{T-1} + a^+,\ s^-_{T-1} + a - a^+)$ denote the posterior belief after observing the a-batch, and $\tilde{\mu}_T$ the updated belief after one additional vote $v^+ \in \{0,1\}$. By Lemma 2 in Appendix A.2, $\mathbb{E}[\tilde{\mu}_T \mid \mu_T] = \mu_T$. Since (1) is convex in μ, Jensen's inequality gives $\mathbb{E}[U(\tilde{\mu}_T) \mid \mu_T] \geqslant U(\mu_T)$ for each realization of μ_T. Taking expectations over the distribution of μ_T,

$$\mathcal{I}_T(a+1; s^+_{T-1}, s^-_{T-1}) = \mathbb{E}_{a^+,v^+}[U(\tilde{\mu}_T)] = \mathbb{E}_{a^+}[\mathbb{E}[U(\tilde{\mu}_T) \mid \mu_T]] \ \geqslant\ \mathbb{E}_{a^+}[U(\mu_T)] = \mathcal{I}_T(a; s^+_{T-1}, s^-_{T-1}).$$

Inductive step. Suppose Inequality (11) holds for $t+1$, i.e. $\mathcal{I}_{t+1}(\cdot\,;\, s^+_t, s^-_t)$ is non-decreasing in its first argument for every state (s^+_t, s^-_t). We prove it for t.

Let $(s^+_t, s^-_t) := (s^+_{t-1} + a^+,\ s^-_{t-1} + a - a^+)$ for the state after the a-batch. Let $v^+ \in \{0,1\}$ be the extra vote, drawn from pool $N - s_t$ conditional on (s^+_t, s^-_t), and write $(\tilde{s}^+_t, \tilde{s}^-_t) := (s^+_t + v^+,\ s^-_t + 1 - v^+)$ for the state after the $(a+1)$-batch. Our goal is to show for every realization of a^+ that

$$\mathbb{E}_{v^+}\big[V_t(\tilde{s}^+_t, \tilde{s}^-_t)\big] \ \geqslant\ V_t(s^+_t, s^-_t). \tag{12}$$

Following from (3), $\mathbb{E}_{v^+}\big[V_t(\tilde{s}^+_t, \tilde{s}^-_t)\big] \ \geqslant\ \max\{\mathbb{E}_{v^+}\big[U(\tilde{s}^+_t, \tilde{s}^-_t)\big], \mathbb{E}_{v^+}\big[\mathcal{C}_t(\tilde{s}^+_t, \tilde{s}^-_t)\big]\}$.

Applying the argument used in the base case, we bound the first term on the right:

$$\mathbb{E}_{v^+}\big[U(\tilde{s}^+_t, \tilde{s}^-_t)\big] = \mathbb{E}[U(\tilde{\mu}_t) \mid \mu_t] \ \geqslant\ U(\mu_t) = U(s^+_t, s^-_t).$$

We now show that the second term can be bounded as well. After the extra vote v^+ at time t, the time-$(t+1)$ batch has size $a_{t+1} = \min\{A_{t+1},\ N - s_t - 1\}$. Define the post-period-$(t+1)$ state $(s^+_{t+1}, s^-_{t+1}) := (\tilde{s}^+_t + a^+_{t+1},\ \tilde{s}^-_t + a_{t+1} - a^+_{t+1})$.

By exchangeability of sampling without replacement, the joint draw of v^+ at time t followed by the a_{t+1}-batch at time $t+1$ is distributed as a single batch of size $a_{t+1}+1$ drawn from the pool of $N - s_t$ voters available at the start of time $t+1$ (i.e., before v^+ is drawn).

Hence, for any fixed a_{t+1},

$$\mathbb{E}_{v^+,\, a^+_{t+1}}\Big[V_{t+1}(s^+_{t+1}, s^-_{t+1}) \;\Big|\; a_{t+1}\Big] = \mathcal{I}_{t+1}(a_{t+1}+1;\, s^+_t, s^-_t).$$

Expanding $\mathcal{C}_t$ via Equation (4) at the state $(\tilde{s}^+_t, \tilde{s}^-_t)$,

$$\begin{aligned}
\mathbb{E}_{v^+}\Big[\mathcal{C}_t(\tilde{s}^+_t, \tilde{s}^-_t)\Big] &= -\kappa + \mathbb{E}_{v^+,\, A_{t+1},\, a^+_{t+1}}\Big[V_{t+1}(s^+_{t+1}, s^-_{t+1})\Big] \\
&= -\kappa + \mathbb{E}_{A_{t+1}}\Big[\mathcal{I}_{t+1}(a_{t+1}+1;\, s^+_t, s^-_t)\Big] \\
&= -\kappa + \mathbb{E}_{A_{t+1}}\Big[\mathcal{I}_{t+1}\Big(\min\{A_{t+1}+1,\, N-s_t\};\, s^+_t, s^-_t\Big)\Big],
\end{aligned}$$

where the second equality applies the exchangeability identity, collapsing the averages over v^+ and a^+_{t+1} conditional on a_{t+1} (valid because A_{t+1} is independent of v^+ and a^+_{t+1} conditional on (s^+_t, s^-_t)), and the third uses $a_{t+1}+1 = \min\{A_{t+1}+1,\, N-s_t\}$. Since $\min\{A_{t+1}+1,\, N-s_t\} \geqslant \min\{A_{t+1},\, N-s_t\}$ and $\mathcal{I}_{t+1}(\cdot\,;\, s^+_t, s^-_t)$ is non-decreasing by the inductive hypothesis,

$$\mathbb{E}_{v^+}\Big[\mathcal{C}_t(\tilde{s}^+_t, \tilde{s}^-_t)\Big] \;\geqslant\; -\kappa + \mathbb{E}_{A_{t+1}}\Big[\mathcal{I}_{t+1}\Big(\min\{A_{t+1},\, N-s_t\};\, s^+_t, s^-_t\Big)\Big] = \mathcal{C}_t(s^+_t, s^-_t).$$

We have shown $\mathbb{E}_{v^+}\Big[V_t(\tilde{s}^+_t, \tilde{s}^-_t)\Big] \;\geqslant\; \max\Big\{U(s^+_t, s^-_t),\, \mathcal{C}_t(s^+_t, s^-_t)\Big\} = V_t(s^+_t, s^-_t)$, which establishes (12). Averaging over the distribution of a^+ gives $\mathcal{I}_t(a+1; s^+_{t-1}, s^-_{t-1}) \geqslant \mathcal{I}_t(a; s^+_{t-1}, s^-_{t-1})$, completing the induction. □

Proof. **Part (i).** Define $\hat{V}_t \equiv V_t/R$, $\hat{U}(\mu) \equiv \max\{\mu, 1-\mu\}$, and $\hat{\mathcal{C}}_t \equiv \mathcal{C}_t/R = -1/\rho + \mathbb{E}[\hat{V}_{t+1}]$, so the Bellman equation becomes $\hat{V}_t(\mu, s; \rho) = \max\Big\{\hat{U}(\mu),\ \hat{\mathcal{C}}_t(\mu, s; \rho)\Big\}$, with transition kernel independent of ρ.

We show $\hat{V}_t$ is weakly increasing in ρ by backward induction. At $t = T$, $\hat{V}_T = \hat{U}$ is independent of ρ. Suppose $\hat{V}_{t+1}$ is weakly increasing in ρ. Then $\mathbb{E}[\hat{V}_{t+1}]$ is weakly increasing in ρ, so $\hat{\mathcal{C}}_t = -1/\rho + \mathbb{E}[\hat{V}_{t+1}]$ is strictly increasing in ρ. As the maximum of a ρ-independent function and a strictly increasing function, $\hat{V}_t$ is weakly increasing in ρ.

By Theorem 1, the continuation region at (t, s) is $\{\mu : \hat{\mathcal{C}}_t(\mu, s; \rho) > \hat{U}(\mu)\}$. Since $\hat{\mathcal{C}}_t$ is strictly increasing in ρ and $\hat{U}$ is independent of ρ, this set is monotone non-decreasing in ρ. By the threshold structure, this implies $d_t^*(s;\rho',\boldsymbol{\lambda}) \leqslant d_t^*(s;\rho,\boldsymbol{\lambda})$ for all $\rho' \leqslant \rho$.

Part (ii). Let $V_t^{\boldsymbol{\lambda}}$ and $V_t^{\boldsymbol{\lambda}'}$ denote the value functions under arrival rates $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}'$ respectively. We prove $V_t^{\boldsymbol{\lambda}'}(s^+,s^-) \leqslant V_t^{\boldsymbol{\lambda}}(s^+,s^-)$ for all $(s^+,s^-) \in \mathcal{S}$ and all t by backward induction on t.

Terminal state. $V_T^{\boldsymbol{\lambda}'} = U = V_T^{\boldsymbol{\lambda}}$, since the terminal value is independent of $\boldsymbol{\lambda}$.

Inductive step. Suppose $V_{t+1}^{\boldsymbol{\lambda}'}(s^+,s^-) \geqslant V_{t+1}^{\boldsymbol{\lambda}}(s^+,s^-)$ for all (s^+,s^-). It suffices to show $\mathcal{C}_t^{\boldsymbol{\lambda}'}(s^+,s^-) \geqslant \mathcal{C}_t^{\boldsymbol{\lambda}}(s^+,s^-)$ for all (s^+,s^-), since then $V_t^{\boldsymbol{\lambda}'} = \max\{U,\mathcal{C}_t^{\boldsymbol{\lambda}'}\} \geqslant \max\{U,\mathcal{C}_t^{\boldsymbol{\lambda}}\} = V_t^{\boldsymbol{\lambda}}$.

The κ-terms cancel, giving

$$\mathcal{C}_t^{\boldsymbol{\lambda}'} - \mathcal{C}_t^{\boldsymbol{\lambda}} = \mathbb{E}_{\lambda'(t+1)}\Big[V_{t+1}^{\boldsymbol{\lambda}'}\Big] - \mathbb{E}_{\lambda'(t+1)}\big[V_{t+1}^{\boldsymbol{\lambda}}\big] + \mathbb{E}_{\lambda'(t+1)}\big[V_{t+1}^{\boldsymbol{\lambda}}\big] - \mathbb{E}_{\lambda(t+1)}\big[V_{t+1}^{\boldsymbol{\lambda}}\big]\,.$$

By the inductive hypothesis, $V_{t+1}^{\boldsymbol{\lambda}'} \geqslant V_{t+1}^{\boldsymbol{\lambda}}$ everywhere, so averaging over any kernel preserves the inequality, thus $\mathbb{E}_{\lambda'(t+1)}\big[V_{t+1}^{\boldsymbol{\lambda}'}\big] - \mathbb{E}_{\lambda'(t+1)}\big[V_{t+1}^{\boldsymbol{\lambda}}\big] \geqslant 0$.

By assumption $\lambda'(t+1) \geqslant \lambda(t+1)$. By monotonicity of truncation, the censored Poisson batch size $\min\{A',M\}$ with $A' \sim \text{Poisson}(\lambda'(t+1))$ first-order stochastically dominates $\min\{A,M\}$ with $A \sim \text{Poisson}(\lambda(t+1))$. Since $\mathcal{I}_{t+1}^{\boldsymbol{\lambda}}(a) := \mathbb{E}[V_{t+1}^{\boldsymbol{\lambda}} \mid a, s^+, s^-]$ is non-decreasing in a by Lemma 4, $\mathbb{E}_{\lambda'(t+1)}\big[V_{t+1}^{\boldsymbol{\lambda}}\big] - \mathbb{E}_{\lambda(t+1)}\big[V_{t+1}^{\boldsymbol{\lambda}}\big] \geqslant 0$.

Hence $\mathcal{C}_t^{\boldsymbol{\lambda}'} \geqslant \mathcal{C}_t^{\boldsymbol{\lambda}}$ and $V_t^{\boldsymbol{\lambda}'} \geqslant V_t^{\boldsymbol{\lambda}}$ everywhere, completing the induction.

Threshold comparison. From $\mathcal{C}_t^{\boldsymbol{\lambda}'}(\mu,s) \geqslant \mathcal{C}_t^{\boldsymbol{\lambda}}(\mu,s)$ pointwise, it follows that $G_t^{\boldsymbol{\lambda}'} := U - \mathcal{C}_t^{\boldsymbol{\lambda}'} \leqslant U - \mathcal{C}_t^{\boldsymbol{\lambda}} =: G_t^{\boldsymbol{\lambda}}$ pointwise. By the proof of Theorem 1, both $G_t^{\boldsymbol{\lambda}'}(\cdot,s)$ and $G_t^{\boldsymbol{\lambda}}(\cdot,s)$ are non-decreasing in μ on $[\frac{1}{2},1]$. In the non-degenerate case, at $\mu = \mu_t^{*,\boldsymbol{\lambda}}(s)$ we have $G_t^{\boldsymbol{\lambda}}(\mu_t^{*,\boldsymbol{\lambda}},s) = 0$, hence $G_t^{\boldsymbol{\lambda}'}(\mu_t^{*,\boldsymbol{\lambda}},s) \leqslant 0$; since $G_t^{\boldsymbol{\lambda}'}$ is non-decreasing with zero-crossing at $\mu_t^{*,\boldsymbol{\lambda}'}(s)$, this gives $\mu_t^{*,\boldsymbol{\lambda}'}(s) \geqslant \mu_t^{*,\boldsymbol{\lambda}}(s)$. In the degenerate case ($G_t^{\boldsymbol{\lambda}}(\frac{1}{2},s) \geqslant 0$, so $d_t^{*,\boldsymbol{\lambda}}(s) = 0$), the inequality $d_t^{*,\boldsymbol{\lambda}'}(s) \geqslant 0 = d_t^{*,\boldsymbol{\lambda}}(s)$ holds trivially. By the strict monotonicity of μ in $d = |s^+ - s^-|$ for fixed s, $\mu_t^{*,\boldsymbol{\lambda}'}(s) \geqslant \mu_t^{*,\boldsymbol{\lambda}}(s)$ translates to $d_t^*(s;\boldsymbol{\lambda}') \geqslant d_t^*(s;\boldsymbol{\lambda})$.

Part (iii). Assume $\lambda(\tau) \geqslant \lambda(\tau+1)$ for all $\tau \in \{1,\ldots,T-1\}$. We prove $V_t(s^+,s^-) \geqslant$

$V_{t+1}(s^+, s^-)$ for all $(s^+, s^-) \in \mathcal{S}$ and $t < T$ by backward induction on $T - t$.

Terminal state. $V_{T-1} = \max\{U, \mathcal{C}_{T-1}\} \geqslant U = V_T$.

Inductive step. Suppose $V_{t+1}(s^+, s^-) \geqslant V_{t+2}(s^+, s^-)$ for all (s^+, s^-). It suffices to show $\mathcal{C}_t(s^+, s^-) \geqslant \mathcal{C}_{t+1}(s^+, s^-)$ for all (s^+, s^-), since then $V_t = \max\{U, \mathcal{C}_t\} \geqslant \max\{U, \mathcal{C}_{t+1}\} = V_{t+1}$.

Note that $\mathcal{C}_t = -\kappa + \mathbb{E}_{\lambda(t+1)}[V_{t+1}]$ and $\mathcal{C}_{t+1} = -\kappa + \mathbb{E}_{\lambda(t+2)}[V_{t+2}]$, and

$$\mathcal{C}_t - \mathcal{C}_{t+1} = \mathbb{E}_{\lambda(t+1)}[V_{t+1} - V_{t+2}] + \mathbb{E}_{\lambda(t+1)}[V_{t+2}] - \mathbb{E}_{\lambda(t+2)}[V_{t+2}]$$

By the inductive hypothesis, $V_{t+1} \geqslant V_{t+2}$ everywhere, so averaging under any kernel preserves the inequality, thus $\mathbb{E}_{\lambda(t+1)}[V_{t+1} - V_{t+2}] \geqslant 0$.

Since $\lambda(t+1) \geqslant \lambda(t+2)$ by assumption, the censored Poisson batch size $\min\{A_{t+1}, M\}$ with $A_{t+1} \sim \text{Poisson}(\lambda(t+1))$ first-order stochastically dominates $\min\{A_{t+2}, M\}$ with $A_{t+2} \sim \text{Poisson}(\lambda(t+2))$, by the monotonicity of truncation and the standard FOSD ordering of Poisson distributions in their rate parameter. Since $\mathcal{I}_{t+2}(a) := \mathbb{E}[V_{t+2} \mid a, s^+, s^-]$ is non-decreasing in a by Lemma 4, it follows that $\mathbb{E}_{\lambda(t+1)}[V_{t+2}] - \mathbb{E}_{\lambda(t+2)}[V_{t+2}] \geqslant 0$. Hence, $\mathcal{C}_t \geqslant \mathcal{C}_{t+1}$ and $V_t \geqslant V_{t+1}$ everywhere, completing the induction.

Threshold comparison. From $\mathcal{C}_t(\mu, s) \geqslant \mathcal{C}_{t+1}(\mu, s)$ pointwise, $G_t := U - \mathcal{C}_t \leqslant U - \mathcal{C}_{t+1} =: G_{t+1}$ pointwise. By the proof of Theorem 1, both $G_t(\cdot, s)$ and $G_{t+1}(\cdot, s)$ are non-decreasing in μ on $[\frac{1}{2}, 1]$.

In the non-degenerate case, at $\mu = \mu^*_{t+1}(s)$ we have $G_{t+1}(\mu^*_{t+1}(s), s) = 0$, hence $G_t(\mu^*_{t+1}(s), s) \leqslant 0$. Since G_t is non-decreasing with zero-crossing at $\mu^*_t(s)$, this gives $\mu^*_t(s) \geqslant \mu^*_{t+1}(s)$.

In the degenerate case, where $G_{t+1}(\frac{1}{2}, s) \geqslant 0$ (so $d^*_{t+1}(s) = 0$), the inequality $d^*_t(s) \geqslant 0 = d^*_{t+1}(s)$ holds trivially. Note that if $G_t(\frac{1}{2}, s) \geqslant 0$ (so $d^*_t(s) = 0$), then since $G_t \leqslant G_{t+1}$ we also have $G_{t+1}(\frac{1}{2}, s) \geqslant 0$, so both thresholds are zero and equal. The inequality is satisfied with equality. By the strict monotonicity of μ in d for fixed s, $\mu^*_t(s) \geqslant \mu^*_{t+1}(s)$ implies $d^*_{t+1}(s) \leqslant d^*_t(s)$. □

A.4 Proof of Theorem 2

The proof shows that the closed-form rule $\hat{\phi}$ coincides with the optimal rule ϕ^* to the stated order. It is divided in five steps. Step 1 reduces the planner's two-dimensional tally to a single sufficient statistic: her confidence in the current leader depends on the counts only through the turnout-scaled lead, and does so in a Gaussian way. Step 2 identifies the time scale on which the statistic evolves and derives, from the exact posterior, the planner's predictive law for the terminal statistic. Step 3 records the Gaussian belief dynamics this law implies. Step 4 derives two benchmark strategies whose indifference points bracket ϕ^*. Step 5 evaluates those two points and finds that both equal $2\log\rho - 2\log\log\rho$ up to an additive constant, thus proving the theorem. By relabeling the alternatives we take $d \geqslant 0$ throughout w.l.o.g.

Step 1: Local quadratic approximation of the evidence statistic. We define a tractable scalar statistic for the posterior in order to evaluate the optimal stopping rule where the decision is non-trivial, i.e. for $s \asymp N$ and vote shares near the boundary $\hat{p} = \frac{s^+}{s} \approx \frac{1}{2}$. By Lemma 1 the posterior distribution of unseen supporters is BetaBinomial$(N-s,\, s^+ + 1,\, s^- + 1)$, which is log-concave. Let $l(m) := \log(\mathrm{Prob}(M^+ = m \mid s^+, s^-))$ and define the evidence statistic $\ell(s^+, s^-)$ as the difference in log-posterior mass between the mode $\hat{m}$ and the tie state $m_0 = N/2$:

$$\ell(s^+, s^-) := l(\hat{m}) - l(m_0).$$

Although m takes integer values, the mass function is a ratio of Gamma functions and so extends to a smooth (real-analytic) function of a real argument. Taylor expanding $l(m)$ around the mode $\hat{m}$ yields:

$$l(m_0) = l(\hat{m}) + l'(\hat{m})(m_0 - \hat{m}) + \tfrac{1}{2} l''(\hat{m})(m_0 - \hat{m})^2 + O\big((m_0 - \hat{m})^3 \cdot l'''(\hat{m})\big). \tag{13}$$

Observe that $l'(\hat{m}) = 0$ because $\hat{m}$ is the maximum a posteriori estimate, and any $O(1)$ integer-rounding difference between $\hat{m}$ and $N\hat{p}$ is absorbed into the remainder. Note that

$$-l''(\hat{m})\,(m_0 - \hat{m})^2 = \frac{(m_0 - \hat{m})^2}{\sigma^2}\,(1 + o(1)), \tag{14}$$

following by the local Gaussian/Laplace approximation to the Beta–Binomial posterior, which implies that $-l''(\hat{m})$ is asymptotically the inverse posterior variance (Van der Vaart, 2000, Ch. 10). We evaluate the two posterior moments in (14): the variance σ^2 and the offset $m_0 - \hat{m}$. Evaluating the variance of BetaBinomial$(N - s,\, s^+ + 1,\, s^- + 1)$ and noting that $\hat{p} - \frac{1}{2} = O(\sqrt{\log \rho / s}) = o(1)$ on the region of interest,

$$\sigma^2 = (N - s) \cdot \frac{(s^+ + 1)(s^- + 1)}{(s + 2)^2} \cdot \frac{N + 2}{s + 3} = \frac{N(N - s)}{4s}\big(1 + o(1)\big). \tag{15}$$

The posterior mean evaluates to $\mathbb{E}[M^+ \mid s^+, s^-] = s^+ + (N - s)\frac{s^+ + 1}{s+2} = N\frac{s^+}{s} + O(1)$. For the near-symmetric posterior the mode agrees with the mean to $O(1)$, thus

$$m_0 - \hat{m} = -\frac{Nd}{2s} + O(1). \tag{16}$$

Substituting these terms into (14) yields $-l''(\hat{m})\,(m_0 - \hat{m})^2 = \frac{d^2}{s} \cdot \frac{N}{N-s}(1 + o(1))$.

Next, we show that the third-order Taylor remainder in (13) is negligible relative to the leading quadratic term on the region of interest. By standard properties of log-density expansions around the mode (Van der Vaart, 2000, Ch. 10), the third derivative $l'''(\hat{m})$ can be expressed as $l'''(\hat{m}) = \frac{\gamma}{\sigma^3}\big(1 + o(1)\big)$, where γ is the posterior skewness. Combining with $-l''(\hat{m}) = \sigma^{-2}\big(1 + o(1)\big)$, the relative ratio of the cubic remainder to the quadratic term in (13) simplifies to:

$$\frac{|l'''(\hat{m})|\,|m_0 - \hat{m}|^3}{|l''(\hat{m})|\,|m_0 - \hat{m}|^2} = \frac{1}{3}|\gamma| Z\big(1 + o(1)\big), \tag{17}$$

where $Z := \frac{\hat{m}-m_0}{\sigma}$ is the standardized score. Substituting the offset from (16) and the variance from (15):

$$Z = \frac{Nd/(2s)}{\sqrt{N(N-s)}/(2\sqrt{s})}\bigl(1+o(1)\bigr) = \frac{d\sqrt{N}}{\sqrt{s(N-s)}}\bigl(1+o(1)\bigr). \tag{18}$$

The skewness γ of the BetaBinomial$(N-s, s^+ + 1, s^- + 1)$ posterior, evaluated near the boundary ($\hat{p} \approx 1/2$, so $a \approx s/2$ and $b \approx s/2$) is:

$$\gamma = -d\frac{2N-s+2}{s+4} \cdot \frac{2}{s}\sqrt{\frac{s+3}{(N-s)(N+2)}} = -\frac{2d(2N-s)}{s^{3/2}\sqrt{N(N-s)}}\bigl(1+o(1)\bigr). \tag{19}$$

Therefore,

$$\frac{1}{3}|\gamma|Z\bigl(1+o(1)\bigr) = \frac{1}{3}\frac{2d^2N(2N-s)}{s^2N(N-s)}\bigl(1+o(1)\bigr) = \frac{1}{3}\frac{2Z^2(2N-s)}{Ns}\bigl(1+o(1)\bigr). \tag{20}$$

Since $2N-s < 2N$, the scaling factor satisfies $(2N-s)/(Ns) < 2/s$. On $\mathcal{S}_\varepsilon$, turnout satisfies $s \geqslant \varepsilon N$ and $Z^2 \leqslant C_\varepsilon^2 \log\rho$, so the cubic-to-quadratic ratio (17) is bounded above by a constant multiple of $\log\rho/N$. Since $\log\rho \leqslant \varepsilon N^{1/3}$, this bound vanishes as $N \to \infty$. Hence $2\ell(s^+, s^-) = \frac{d^2}{s}\frac{N}{N-s}\,(1+o(1))$ uniformly on $\mathcal{S}_\varepsilon$.

For the posterior mass function, however, the exponent must be controlled in absolute rather than relative terms. Multiplying the preceding relative error by the quadratic term $Z^2 \leqslant C_\varepsilon^2 \log\rho$ yields an absolute exponent error bounded by a constant multiple of $(\log\rho)^2/N$, which also vanishes under the growth condition. Therefore, uniformly for $|m-\hat{m}| \leqslant C_\varepsilon \sigma\sqrt{\log\rho}$,

$$\Pr(M^+ = m \mid s^+, s^-) = e^{l(\hat{m})}\exp\Bigl\{-\frac{(m-\hat{m})^2}{2\sigma^2}\Bigr\}\bigl(1+o(1)\bigr). \tag{21}$$

By log-concavity, the posterior mass outside this window is at most $\rho^{-C_\varepsilon^2/3}$.

Summing (21) over the window gives $e^{l(\hat{m})} = (\sigma\sqrt{2\pi})^{-1}(1+o(1))$. Summing between the tie state $m_0 = N/2$ and $\hat{m}$, and using that the lattice sum differs from the corresponding Gaussian integral by at most its largest term, yields

$$\mu - \tfrac{1}{2} = \big(\Phi(Z) - \tfrac{1}{2}\big)\big(1 + o(1)\big) + O(N^{-1/2}), \qquad Z := \operatorname{sgn}(d)\sqrt{2\ell}, \tag{22}$$

uniformly on $\mathcal{S}_\varepsilon$ with $Z^2 \leqslant C_\varepsilon^2 \log \rho$. The local approximation (21) will also be used in Step 2 to obtain relative tail asymptotics at distances of order $\sqrt{\log \rho}$ from the mode.

Step 2: The information clock and the predictive law. We now derive the dynamics of the evidence scale 2ℓ. Parameterize the unknown state as $M^+ = N/2 + \delta$ and set $\theta := 2\delta/\sqrt{N}$. This is the local-alternative scaling: a majority with $\delta \gg \sqrt{N}$ is identifiable from any substantial sample, while one with $\delta \ll \sqrt{N}$ is undetectable even at full turnout, so $\theta = O(1)$ is the non-trivial range. Let $v_1, v_2, \dots \in \{0, 1\}$ be the votes in arrival order, $\mathcal{F}_s$ the information after s of them, and $\tilde{d}_s = \sum_{k \leqslant s}(2v_k - 1)$ the signed lead.

To eliminate the mean-reverting pull-back effect caused by sampling without replacement, we first establish that the following centered, rescaled state process is a martingale.

Lemma 5. *Conditional on M^+, the process $M_s := (\tilde{d}_s - 2\delta)/(N - s)$, $s = 0, \dots, N-1$, is a martingale with respect to $(\mathcal{F}_s)$. Its scaled increments are*

$$\sqrt{N}\,(M_{k+1} - M_k) = \frac{\sqrt{N}\,(2v_{k+1} - 1)}{N - k - 1} + \frac{\sqrt{N}\,(\tilde{d}_k - 2\delta)}{(N - k - 1)(N - k)}, \tag{23}$$

with conditional variance

$$\operatorname{Var}\big(\sqrt{N}(M_{k+1} - M_k) \mid \mathcal{F}_k\big) = \frac{N \cdot 4p_k(1 - p_k)}{(N - k - 1)^2}, \qquad p_k = \frac{M^+ - s_k^+}{N - k}. \tag{24}$$

Proof. Given $\mathcal{F}_k$, $M^+ - s_k^+$ supporters remain among $N - k$ unseen voters. The conditional probability of a positive vote is $p_k = \frac{M^+ - s_k^+}{N-k}$. Using $2M^+ = N + 2\delta$ and $2s_k^+ = k + \tilde{d}_k$:

$$\mathbb{E}[\tilde{d}_{k+1} - \tilde{d}_k \mid \mathcal{F}_k] = 2p_k - 1 = \frac{2\delta - \tilde{d}_k}{N-k}.$$

Taking the conditional expectation of M_{k+1}:

$$\mathbb{E}[M_{k+1} \mid \mathcal{F}_k] = \frac{\tilde{d}_k + \mathbb{E}[\tilde{d}_{k+1} - \tilde{d}_k \mid \mathcal{F}_k] - 2\delta}{N-k-1} = \frac{\tilde{d}_k - 2\delta}{N-k} = M_k.$$

Thus, $\mathbb{E}[M_{k+1} \mid \mathcal{F}_k] = M_k$ holds for all N.

Note that the increment $2v_{k+1} - 1$ takes values ± 1 with probabilities p_k and $1 - p_k$, so $\mathrm{Var}(2v_{k+1} - 1 \mid \mathcal{F}_k) = 4p_k(1-p_k)$. Using $\tilde{d}_{k+1} = \tilde{d}_k + (2v_{k+1} - 1)$, we express the scaled martingale increment:

$$\sqrt{N}(M_{k+1} - M_k) = \frac{\sqrt{N}(2v_{k+1} - 1)}{N-k-1} + \frac{\sqrt{N}(\tilde{d}_k - 2\delta)}{(N-k-1)(N-k)}.$$

Thus, the conditional variance is:

$$\mathrm{Var}\left(\sqrt{N}(M_{k+1} - M_k) \,\middle|\, \mathcal{F}_k\right) = \mathrm{Var}\left(\frac{\sqrt{N}(2v_{k+1} - 1)}{N-k-1} \,\middle|\, \mathcal{F}_k\right) = N\frac{4p_k(1-p_k)}{(N-k-1)^2}.$$

□

Summing the conditional variance (24), the predictable quadratic variation of $\sqrt{N}(M_s - M_0)$ is $\langle\sqrt{N}M\rangle_s = N\sum_{k=0}^{s-1} 4p_k(1-p_k)/(N-k-1)^2$. In the local regime where $\theta = O(1)$, the proportion $p_k = \frac{1}{2} + o(1)$ uniformly in probability, so $4p_k(1-p_k) = 1 + o(1)$. Re-indexing $j = N - k - 1$ and applying integral comparison to the tail sum yields:

$$\begin{aligned}\langle\sqrt{N}M\rangle_s &= N\sum_{j=N-s}^{N-1} \frac{1+o(1)}{j^2} = N\left(\int_{N-s}^{N} \frac{dx}{x^2} + O\left(\frac{1}{(N-s)^2}\right)\right)(1 + o(1)) \\ &= N\left(\frac{1}{N-s} - \frac{1}{N}\right) + o(1) = \frac{s}{N-s} + o(1).\end{aligned}$$

We call $\tau(s) = s/(N-s)$ the information clock, as it accumulates at unit rate. We use the following notation

$$\tau := \tau(s_t) = \frac{s_t}{N - s_t}, \qquad \tau_T := \tau\big(\hat{s}_T\big) = \frac{\hat{s}_T}{N - \hat{s}_T}, \qquad \Delta\tau(t) := \tau_T - \tau. \tag{25}$$

On $\mathcal{S}_\varepsilon$ all three quantities are bounded away from 0 and ∞. The realized terminal turnout s_T differs from $\hat{s}_T$ by the Poisson fluctuation of the arrivals, which is $O(\sqrt{N})$. Let $\widehat{X}_s := \sqrt{N}\,\tilde{d}_s/(N-s)$ denote the current lead rescaled by the residual pool. Since $\tau(\cdot)$ is a bijection we also index $\widehat{X}$ by the clock, writing $\widehat{X}_\tau$ for $\widehat{X}_s$ at $\tau = \tau(s)$. Since $\tilde{d}_s = (N-s)M_s + 2\delta$ and $\sqrt{N}M_0 = -\theta$, we have $\widehat{X}_s = \theta\,\tau(s) + \sqrt{N}\,(M_s - M_0)$, so by Lemma 5 and the quadratic-variation computation above,

$$\mathbb{E}[\widehat{X}_s \mid \theta] = \theta\,\tau(s), \qquad \operatorname{Var}(\widehat{X}_s \mid \theta) = \mathbb{E}\langle\sqrt{N}M\rangle_s = \tau(s)\big(1 + o(1)\big), \tag{26}$$

the mean exactly. The lead $\tilde{d}_s$ given M^+ is a shifted hypergeometric variable, whose mass function is log-concave and a ratio of Gamma functions, so the Laplace expansion of Step 1 applies verbatim and $\widehat{X}_s \mid \theta$ is Gaussian to the accuracy established there. We write $\widehat{X}_\tau \mid \theta \sim N(\theta\tau, \tau)$.

Following Lemma 1, conditional on a future arrivals $a^+ \sim \mathrm{BetaBinomial}(a, s^+ + 1, s^- + 1)$, the same family as the Step 1 posterior with $N - s$ replaced by a. The Laplace expansion (21) and the tail bound $\rho^{-C_\varepsilon^2/3}$ therefore hold verbatim for a^+. Writing $\tau' = \tau(s+a)$, the terminal statistic $\widehat{X}_{s+a}$ has, by the martingale identity (26) applied from date t,

$$\mathbb{E}[\widehat{X}_{s+a} \mid \mathcal{F}_t] = \widehat{X}_s\,\frac{\tau'}{\tau}\,(1 + O(1/N)), \qquad \operatorname{Var}(\widehat{X}_{s+a} \mid \mathcal{F}_t) = (\tau' - \tau)\frac{\tau'}{\tau}\,(1 + O(1/N)). \tag{27}$$

A direct computation links this back to Step 1: since $\tilde{d}_s = d$,

$$\frac{\widehat{X}_s^2}{\tau(s)} = \frac{N\tilde{d}_s^2/(N-s)^2}{s/(N-s)} = \frac{d^2}{s}\cdot\frac{N}{N-s} = 2\ell, \tag{28}$$

so $Z_\tau = \widehat{X}_\tau/\sqrt{\tau}$ is exactly the z-score of Step 1.

Step 3: Predictive belief dynamics. In Step 2 we derived the planner's predictive law for the terminal statistic. We now specify the belief dynamics it implies, in the Gaussian form of Steps 1 and 2.

In the local scaling of Step 2 the uniform prior on M^+ is a flat prior $p(\theta)\propto 1$ on θ, and the Gaussian predictive of Step 2 reads $\widehat{X}_\tau \mid \theta \sim N(\theta\tau, \tau)$ at the current date. Updating the flat prior via Bayes' rule and using that the path $\mathcal{F}_\tau$ informs θ only through its endpoint $\widehat{X}_\tau$ gives the posterior density:

$$p(\theta \mid \mathcal{F}_\tau) \propto \exp\left(-\frac{(\widehat{X}_\tau - \theta\tau)^2}{2\tau}\right) = \exp\left(-\frac{(\theta - \widehat{X}_\tau/\tau)^2}{2(1/\tau)}\right).$$

This kernel implies the Normal posterior $\theta \mid \mathcal{F}_\tau \sim N\left(\frac{\widehat{X}_\tau}{\tau}, \frac{1}{\tau}\right)$, so τ is the posterior precision. Thus, the planner choosing the leading alternative expects to be correct with probability $\mathrm{Prob}(\theta > 0 \mid \mathcal{F}_\tau) = \Phi(Z_\tau)$, leaving a residual error risk of $\Phi(-|Z_\tau|)$ upon stopping.

The conditional mean and variance of $\widehat{X}_{\tau'}$ given $\mathcal{F}_\tau$ were computed in (27): they are $\widehat{X}_\tau\tau'/\tau$ and $(\tau' - \tau)\tau'/\tau$, up to $1 + O(1/N)$. (The decomposition into drift uncertainty $(\tau'-\tau)^2/\tau$ and fresh noise $\tau'-\tau$ is the Gaussian reading of the same two moments.) Dividing by $\sqrt{\tau'}$,

$$Z_{\tau'} \mid \mathcal{F}_\tau \sim N\left(Z_\tau\sqrt{\frac{\tau'}{\tau}}, \frac{\tau' - \tau}{\tau}\right). \tag{29}$$

This has the following implications:

(i) The belief is a martingale. For any normal variable with mean m and variance s^2, the expectation of its standard normal CDF is $\Phi\left(m/\sqrt{1+s^2}\right)$. Applying this to (29) yields:

$$\mathbb{E}\left[\Phi(Z_{\tau'}) \mid \mathcal{F}_\tau\right] = \Phi\left(\frac{Z_\tau\sqrt{\tau'/\tau}}{\sqrt{\tau'/\tau}}\right) = \Phi(Z_\tau), \qquad \text{i.e.,} \qquad \mathbb{E}[\mu_{\tau'} \mid \mathcal{F}_\tau] = \mu_\tau. \tag{30}$$

(ii) Capped variance of belief updates. Recall that the posterior mean of θ is $m_\tau := \widehat{X}_\tau/\tau$. The variance it accumulates between the present τ and the deadline τ_T is

$$\operatorname{Var}\big(m_{\tau_T} - m_\tau \,\big|\, \mathcal{F}_\tau\big) \;=\; \frac{1}{\tau} - \frac{1}{\tau_T} \;=\; \frac{\Delta\tau}{\tau\,\tau_T}, \tag{31}$$

a quantity known at τ and vanishing as $\tau \to \tau_T$. The leader is overturned only if m changes sign, so the relevant state variable compares the estimate's distance from zero with the remaining movement available to it:

$$\beta \;:=\; \frac{|m_\tau|}{\sqrt{1/\tau - 1/\tau_T}} \;=\; |Z_\tau|\sqrt{\frac{\tau_T}{\Delta\tau}}, \tag{32}$$

measuring the current lead in standard deviations of remaining learning. Since $m_{\tau_T} \mid \mathcal{F}_\tau \sim N\left(m_\tau,\, \frac{\Delta\tau}{\tau\tau_T}\right)$, the probability that the deadline verdict reverses the current leader is exactly $\Phi(-\beta)$. By (28), $2\ell = Z_\tau^2 = \beta^2\,\Delta\tau/\tau_T$, and by (25) a direct computation gives the statistic in (6):

$$\beta^2 = \frac{(d_t/s_t)^2}{1/s_t - 1/\hat{s}_T}. \tag{33}$$

Step 4: Bracketing the optimal boundary. We now derive upper and lower bounds on the continuation region of ϕ^*, by comparing the payoff from stopping with that of candidate continuation strategies. Evaluating such strategies requires the belief martingale property of Lemma 2 at random stopping times rather than fixed dates.

Because the belief process $\mu_u \in [0,1]$, $u = t, \dots, T$, is a bounded martingale under the planner's filtration (Lemma 2) and every stopping time satisfies $\sigma \leqslant T$, Doob's optional sampling theorem applies (e.g., Karatzas and Shreve, 1991, Thm. 1.3.22). Combined with the tower property, this yields, for any stopping time σ with values in $\{t, \dots, T\}$, writing μ_{τ_T}

for the deadline belief,

$$\mathbb{E}\big[\mu_{\tau_T} \mid \mathcal{F}_\sigma\big] = \mu_\sigma, \qquad \text{and} \qquad \mathbb{E}\big[\mu_\sigma \mid \mathcal{F}_\tau\big] = \mu_\tau. \tag{34}$$

To evaluate candidate continuation strategies, we isolate the maximal option value of future learning over $[\tau, \tau_T]$. Let $\mathcal{G}_N := \mathbb{E}\left[\max\{0, 1 - 2\mu_{\tau_T}\} \mid \mathcal{F}_\tau\right]$ denote the normalized expected gain from terminal leadership reversals. Thus, the maximum economic option value of waiting until the deadline is $R\mathcal{G}_N$.

Lemma 6. *Suppose $\mu_\tau \geqslant \frac{1}{2}$. For every stopping time $\tau' \in [\tau, \tau_T]$,*

$$\mathbb{E}\big[\max\{\mu_{\tau'}, 1 - \mu_{\tau'}\} \mid \mathcal{F}_\tau\big] - \mu_\tau \;\leqslant\; \mathcal{G}_N, \tag{35}$$

with equality holding for $\tau' \equiv \tau_T$.

Proof. Using the payoff identity $\max\{x, 1 - x\} = x + \max\{0, 1 - 2x\}$ for $x \in [0, 1]$, the expected payoff under policy τ' expands to:

$$\begin{aligned} \mathbb{E}\big[\max\{\mu_{\tau'}, 1 - \mu_{\tau'}\} \mid \mathcal{F}_\tau\big] &= \mathbb{E}[\mu_{\tau'} \mid \mathcal{F}_\tau] + \mathbb{E}\left[\max\{0, 1 - 2\mu_{\tau'}\} \mid \mathcal{F}_\tau\right] \\ &= \mu_\tau + \mathbb{E}\left[\max\{0, 1 - 2\mu_{\tau'}\} \mid \mathcal{F}_\tau\right], \end{aligned}$$

where the second equality follows from (34). Because $x \mapsto \max\{0, 1-2x\}$ is convex, applying conditional Jensen's inequality to (34) yields:

$$\max\{0, 1 - 2\mu_{\tau'}\} = \max\{0, 1 - 2\mathbb{E}[\mu_{\tau_T} \mid \mathcal{F}_{\tau'}]\} \;\leqslant\; \mathbb{E}\left[\max\{0, 1 - 2\mu_{\tau_T}\} \mid \mathcal{F}_{\tau'}\right]. \tag{36}$$

Taking conditional expectations $\mathbb{E}[\,\cdot \mid \mathcal{F}_\tau]$ on both sides of (36) and applying the tower property yields $\mathbb{E}\left[\max\{0, 1 - 2\mu_{\tau'}\} \mid \mathcal{F}_\tau\right] \leqslant \mathcal{G}_N$. Subtracting μ_τ establishes (35). Equality for $\tau' \equiv \tau_T$ holds by definition of $\mathcal{G}_N$. □

Note that any continuation strategy costs at least κ, while by Lemma 6 its expected gain over stopping now is at most $R\,\mathcal{G}_N$. Hence, if $R\,\mathcal{G}_N < \kappa$, stopping is optimal. If, however, the maximal cost of waiting is lower than the gain of waiting until the end, so if $R\,\mathcal{G}_N > \kappa(T-t)$, continuing is optimal. Both implications hold exactly for ϕ^*.

We express $\mathcal{G}_N$ through the current evidence score β. Let $\mathcal{G}(\beta, r)$ denote the Gaussian option value, in which $\mu_{\tau_T} = \Phi(Z_{\tau_T})$ with $Z_{\tau_T} \sim N(\beta\sqrt{r},\, r)$ and $r = \Delta\tau/\tau$.

Lemma 7. *Uniformly on* $\{(d, s, t) \in \mathcal{S}_\varepsilon : \ \beta^2 \leqslant 4 \log \rho + 2\}$,

$$\mathcal{G}_N = \mathcal{G}\big(\beta, \Delta\tau/\tau\big)\big(1 + o(1)\big). \tag{37}$$

Proof. Fix $C = C_\varepsilon \geqslant 3$ in the windows of Steps 1 and 2 large enough that $C_\varepsilon^2 \geqslant 9 \sup_{\mathcal{S}_\varepsilon} \Delta\tau/\tau$. The cutoff $Z_{\tau_T} < -C\sqrt{\log \rho}$ below then lies at least $3\sqrt{\log \rho}$ standard deviations of Z_{τ_T} below its mean, and the window of Step 2, taken as $C(1 + \sup_{\mathcal{S}_\varepsilon}(\tau/\Delta\tau)^{1/2})\sqrt{\log \rho}$ standard deviations, contains the range $[-C\sqrt{\log \rho}, 0]$ of Z_{τ_T}. Condition first on the number a of arriving votes and let $\tau_a := \tau(s+a)$, $\beta_a := |Z_\tau|\sqrt{\tau_a/(\tau_a - \tau)}$. The integrand $\max\{0, 1-2\mu_{\tau_T}\}$ vanishes when $Z_{\tau_T} \geqslant 0$. On $Z_{\tau_T} < -C\sqrt{\log \rho}$ it is at most one, and the event has probability at most ρ^{-4} by Step 2 and the choice of C_ε. On the remaining range $-C\sqrt{\log \rho} \leqslant Z_{\tau_T} < 0$ it equals $2(\Phi(-Z_{\tau_T}) - \frac{1}{2})(1 + o(1)) + O(N^{-1/2})$ by (22). Writing $h(z) := \max\{2(\Phi(-z) - \frac{1}{2}), 0\}$, integration by parts gives $\mathbb{E}[h(Z_{\tau_T}) \mid a] = \int_{-\infty}^{0} 2\phi(z)\mathrm{Prob}(Z_{\tau_T} \leqslant z \mid a)\, dz$, and inserting the relative-error statement of Step 2 shows that this equals its Gaussian counterpart $\mathcal{G}(\beta_a, (\tau_a - \tau)/\tau)$ up to a relative $o(1)$ plus an additive $O(N^{-1/2}\phi(\beta_a))$. Since $\mathcal{G} \asymp \sqrt{\Delta\tau/\tau}\,\phi(\beta)/(1 + \beta^2)$ by (42), and $\phi(\beta_a) \geqslant \rho^{-3}$ on the stated range, all additive errors are $o(\mathcal{G})$ because $\rho^{-3} \gg \rho^{-4}$ and, since $\log \rho = o(N^{1/2})$, $\rho^{-3} \gg e^{-cN^{1/2}}$. We now average over a. The arrivals after t are independent Poisson variables with total mean $\bar{a} := \hat{s}_T - s \asymp N$, so $\mathrm{Prob}(|a - \bar{a}| > N^{3/4}) \leqslant e^{-cN^{1/2}}$, which is negligible by the same comparison, and on $|a - \bar{a}| \leqslant N^{3/4}$ no censoring occurs. On that event $\tau_a = \tau_T(1 + O(N^{-1/4}))$, so the prefactor of $\mathcal{G}$ moves by a relative

$O(N^{-1/4})$, while the exponent needs exact expansion: from $\beta_a^2 = d^2/s + d^2/a$,

$$e^{-\beta_a^2/2} = e^{-\beta^2/2}\, e^{\zeta(a-\bar{a})}\, e^{-d^2(a-\bar{a})^2/(2\bar{a}^2 a)}, \qquad \zeta := \frac{d^2}{2\bar{a}^2} = O\Big(\frac{\log\rho}{N}\Big),$$

where the last factor is $1+O(\log\rho\, N^{-1/2}) = 1+o(1)$ on the event, and $\mathbb{E}[e^{\zeta(a-\bar{a})}] = \exp\{\bar{a}(e^{\zeta} - 1 - \zeta)\} = \exp\{O((\log\rho)^2/N)\} = 1 + o(1)$ by the Poisson moment generating function. $\square$

Next, we express $\mathcal{G}$ as a function of the current evidence score β. Evaluating (29) at $\tau' = \tau_T$ using (32),

$$Z_{\tau_T} \mid \mathcal{F}_\tau \;\sim\; N\bigg(\beta\sqrt{\tfrac{\Delta\tau}{\tau}},\; \tfrac{\Delta\tau}{\tau}\bigg), \tag{38}$$

and since $\mu_{\tau_T} = \Phi(Z_{\tau_T})$, we can write $\mathcal{G} = \mathcal{G}(\beta, \Delta\tau/\tau)$. Moreover, $\mathcal{G}(\beta, .)$ is continuous and strictly decreasing in β, with $\mathcal{G} \leqslant \Phi(-\beta) \to 0$ as $\beta \to \infty$ and $\mathcal{G}(0,\cdot) > 0$.[15]. For every $c \in [1, T]$ and all ρ such that $c/\rho < \mathcal{G}(0, \Delta\tau/\tau)$, the equation

$$\mathcal{G}\big(\beta_c(t),\, \Delta\tau/\tau\big) = \frac{c}{\rho} \tag{39}$$

has a unique root $\beta_c(t)$. In this notation, the two benchmarks read: stopping is optimal whenever $\beta \geqslant \beta_1(t)$, and continuing whenever $\beta < \beta_{T-t}(t)$. Since $\beta_c(t)$ decreases in c and $T - t \geqslant 1$, the cutoffs are ordered $\beta_{T-t}(t) \leqslant \beta_1(t)$, and

$$\big\{\beta < \beta_{T-t}(t)\big\} \;\subseteq\; \text{continuation region} \;\subseteq\; \big\{\beta < \beta_1(t)\big\}, \tag{40}$$

Moreover, $\mathcal{G}_N$ is non-increasing in d at fixed (s,t): conditional on the arrivals, raising s^+ by one and lowering s^- by one moves the predictive law $\mathrm{BetaBinomial}(a, s^+ + 1, s^- + 1)$ of Lemma 1 up in the likelihood-ratio order (the ratio of mass functions is $(k+s^+ +1)/(a-k+s^-)$, increasing in k), hence μ_{τ_T} increases stochastically by the monotone likelihood ratio property

[15] Immediate from the integral representation (41).

(Karlin and Rubin, 1956), and $\max\{0, 1-2\mu_{\tau_T}\}$ is non-increasing in μ_{τ_T}. Since β increases with d at fixed (s,t) by (33), the map $\beta \mapsto R\mathcal{G}_N$ is non-increasing, so the stopping region of ϕ^* is an up-set $\{\beta \geqslant \beta^*(s,t)\}$: the optimal policy is a threshold rule in β, with cutoff placed by the exact bracket (40) in $[\beta_{T-t}(t), \beta_1(t)]$, i.e. $\beta_{T-t}(t) \leqslant \beta^*(t) \leqslant \beta_1(t)$.[16]

This monotonicity also extends the inclusion (40) from the window to all of $\mathcal{S}_\varepsilon$. By (33) a step of two in d changes β^2 by at most $(4\beta+4)/(\varepsilon\sqrt{N}) + 4/(\varepsilon^2 N) \leqslant 1$ on the window, so any state with $\beta^2 > 4\log\rho + 2$ has $\mathcal{G}_N$ no larger than at a state with $\beta^2 \in (4\log\rho+1, 4\log\rho+2]$ at the same (s,t), where stopping is optimal. Hence the right inclusion in (40) holds on all of $\mathcal{S}_\varepsilon$, and the left one trivially so.

Step 5: Evaluating the bounds. Both ends of (40) are roots of (39), at $c = 1$ and $c = T-t$, respectively. Therefore, we solve (39) asymptotically as $\rho \to \infty$, for fixed $c \in [1, T]$, by first deriving the leading-order size of $\mathcal{G}$ for large β, and then inverting.

To evaluate $\mathcal{G} = \mathbb{E}\left[\max\{0, 1-2\mu_{\tau_T}\} \mid \mathcal{F}_\tau\right]$, note that $1-2\mu_{\tau_T} = 2\big(\frac{1}{2} - \Phi(Z_{\tau_T})\big)$ is non-zero only on reversal paths where $Z_{\tau_T} < 0$. Following from (38), we can write $Z_{\tau_T} = \sqrt{\Delta\tau/\tau}(\beta+\xi)$ with $\xi \sim N(0,1)$, such that $Z_{\tau_T} < 0$ corresponds to $\xi < -\beta$. Integrating over these paths yields the explicit representation:

$$\mathcal{G}\left(\beta, \frac{\Delta\tau}{\tau}\right) = 2\int_{-\infty}^{-\beta} \left[\frac{1}{2} - \Phi\big(\sqrt{\Delta\tau/\tau}(\beta+\xi)\big)\right] \phi(\xi)\, d\xi. \tag{41}$$

We re-center (41) at the edge of the reversal region. Substituting $\xi = -\beta - u$ with $u > 0$, the density at ξ factorizes as $\phi(\beta+u) = \phi(\beta)e^{-\beta u - u^2/2}$ and $Z_{\tau_T} = -\sqrt{\Delta\tau/\tau}\, u$, so

$$\mathcal{G} = 2\int_0^\infty \left[\Phi\big(\sqrt{\Delta\tau/\tau}\, u\big) - \tfrac{1}{2}\right] \phi(\beta)\, e^{-\beta u - u^2/2}\, du.$$

Split at $u_0 = \beta^{-1/2}$. On $u > u_0$ the integrand is at most $\frac{1}{2}\phi(\beta)e^{-\beta u}$, contributing $o(\phi(\beta)\beta^{-2})$ uniformly. On $u \leqslant u_0$, $e^{-u^2/2} = 1 + O(\beta^{-1})$ and, by a second-order Taylor bound with

[16] The $(1+o(1))$ shifts each root by $o(1)$ in β^2 (Step 5), which the band K_ε absorbs.

$\sup_z |z\phi(z)| \leqslant 1$, $|\Phi(\sqrt{\Delta\tau/\tau}u) - \frac{1}{2} - \phi(0)\sqrt{\Delta\tau/\tau}u| \leqslant \frac{1}{2}(\Delta\tau/\tau)\,u^2 \leqslant \frac{1}{2}C_\varepsilon u^2$ with $C_\varepsilon :=$ $\sup_{\mathcal{S}_\varepsilon} \Delta\tau/\tau$, so the linearization error contributes $O(\phi(\beta)\beta^{-3})$, again uniformly. Both corrections are of relative order $O(\beta^{-1}) = O((\log\rho)^{-1/2})$. Using $\int_0^\infty u\,e^{-\beta u}\,du = \beta^{-2}$ and $\phi(0) = 1/\sqrt{2\pi}$,

$$\mathcal{G} = 2\phi(\beta)\,\phi(0)\sqrt{\tfrac{\Delta\tau}{\tau}} \cdot \frac{1}{\beta^2}\big(1+o(1)\big) = \sqrt{\frac{2\Delta\tau}{\pi\,\tau}}\,\frac{\phi(\beta)}{\beta^2}\big(1+o(1)\big). \tag{42}$$

The prefactor β^{-2} refines Step 3(ii): the reversal probability $\Phi(-\beta)$ carries only β^{-1}, so the option value is smaller by a factor of order β.

With $\sqrt{2/\pi}\cdot(2\pi)^{-1/2} = 1/\pi$, $\mathcal{G} = \pi^{-1}\sqrt{\Delta\tau/\tau}\,e^{-\beta^2/2}\beta^{-2}(1+o(1))$. Setting $\mathcal{G} = c/\rho$ and taking logarithms,

$$-\log\pi + \tfrac{1}{2}\log\tfrac{\Delta\tau}{\tau} - \tfrac{\beta^2}{2} - 2\log\beta + o(1) = \log c - \log\rho,$$

that is, multiplying by -2,

$$\beta^2 + 2\log\beta^2 = 2\log\rho - 2\log c + \log\tfrac{\Delta\tau}{\tau} - 2\log\pi + o(1). \tag{43}$$

The logarithm is of smaller order than β^2, so a first pass gives $\beta^2 = 2\log\rho\,(1+o(1))$, hence $\log\beta^2 = \log 2 + \log\log\rho + o(1)$. Returning this to (43),

$$\beta_c(t)^2 = 2\log\rho - 2\log\log\rho - 2\log c + \log\tfrac{\Delta\tau}{\tau} - 2\log\pi - 2\log 2 + o(1). \tag{44}$$

Substituting (44) back into the logarithm changes $\log\beta^2$ by $O(\log\log\rho/\log\rho) = o(1)$, so a further iteration alters only the $o(1)$ term and the displayed terms are exact. The expansion is uniform over $c \in [1,T]$ and over $\mathcal{S}_\varepsilon$, since each root diverges like $\sqrt{2\log\rho}$ and $\Delta\tau/\tau$ is bounded away from 0 and ∞ by (25). Evaluating (44) at $c = 1$ and $c = T - t$, we obtain

$$\big[\beta_1(t)\big]^2 - \big[\beta_{T-t}(t)\big]^2 = 2\log(T-t) + o(1) \;\leqslant\; 2\log T + o(1), \tag{45}$$

which is $O(1)$ because T is fixed. The remaining terms in (44) are likewise bounded uniformly over $\mathcal{S}_\varepsilon$, so both ends of the bracket equal $L_\rho + O(1)$, where $L_\rho := 2\log\rho - 2\log\log\rho$. Write K_ε for a bound on this $O(1)$, depending only on ε and T. By (40), ϕ^* stops at every state of $\mathcal{S}_\varepsilon$ with $\beta^2 > L_\rho + K_\varepsilon$ and continues at every state with $\beta^2 < L_\rho - K_\varepsilon$, while $\hat{\phi}$ stops exactly when $\beta^2 \geqslant L_\rho$. Since $K_\varepsilon \leqslant L_\rho \frac{\log\log\rho}{\log\rho}$ for $\rho \geqslant \rho_\varepsilon$, the two rules agree wherever $|\beta^2/L_\rho - 1| > \log\log\rho/\log\rho$, which by (33) gives the band of the theorem.

A.5 Proof of Corollary 1

Write $\Delta\tau(t)/\tau_T$ for the factor in (7). Fix $\bar{t}$ and let α be given by (9), which implies

$$\big[\Phi^{-1}(1-\alpha)\big]^2 = \big[2\log\rho - 2\log\log\rho\big]\,\frac{\Delta\tau(\bar{t})}{\tau_T}. \tag{46}$$

Taking the ratio of the shape factor at t and at $\bar{t}$, the common $N/\Lambda(T)$ cancels and

$$\frac{\Delta\tau(t)/\tau_T}{\Delta\tau(\bar{t})/\tau_T} = \frac{\Lambda(T)-\Lambda(t)}{N-\Lambda(t)}\cdot\frac{N-\Lambda(\bar{t})}{\Lambda(T)-\Lambda(\bar{t})}, \tag{47}$$

the factor appearing in (8). Multiplying (46) by (47) gives $\big[2\log\rho - 2\log\log\rho\big]\Delta\tau(t)/\tau_T$, the right-hand side of (7). Hence, (8) and (7) define the same stopping set for every $\bar{t}$ and the reference date reparameterizes the rule without changing it.

By (22) the posterior probability that the leading alternative is the minority preference is $1-\mu = \Phi(-|Z|)$, and by (28) with (7) a decision at date t requires $|Z| \geqslant \big\{[2\log\rho - 2\log\log\rho]\Delta\tau(t)/\tau_T\big\}^{1/2}$. At $t = \bar{t}$ this lower bound equals $\Phi^{-1}(1-\alpha)$ by (46), so a decision taken at $\bar{t}$ carries posterior error probability at most $\Phi\big(-\Phi^{-1}(1-\alpha)\big) = \alpha$. Because Λ is nondecreasing, $\Delta\tau(t)/\tau_T$ is nonincreasing, so at any earlier date the required evidence is weakly larger and the implied error weakly smaller. The bound holds for any decision reached at or before $\bar{t}$.