Optimal Dynamic Majority-Quorum Rules: Theory and Evidence

Bhargav Nagaraja Bhatt[1] Jonas Gehrlein[2] Kremena Valkanova[2]

[1]Morpho [2]Parity Technologies

Swiss Theory Day
September 10, 2026

Collective decisions trade off accuracy against time

	◀ **act too early**	**act too late** ▶
	too little evidence	too much delay
DAO governance	enact an opposed upgrade	prolong a vulnerability window
Crowd annotation	mislabel the item	burn budget, slow the pipeline
Election audits	certify the wrong winner	delay the transfer of power

The value of a decision depends on *what* is chosen and *when*.

Collective decisions trade off accuracy against time

	◀ **act too early**	**act too late** ▶
	too little evidence	too much delay
DAO governance	enact an opposed upgrade	prolong a vulnerability window
Crowd annotation	mislabel the item	burn budget, slow the pipeline
Election audits	certify the wrong winner	delay the transfer of power

The value of a decision depends on *what* is chosen and *when*.

Collective decisions trade off accuracy against time

	◀ **act too early**	**act too late** ▶
	too little evidence	too much delay
DAO governance	enact an opposed upgrade	prolong a vulnerability window
Crowd annotation	mislabel the item	burn budget, slow the pipeline
Election audits	certify the wrong winner	delay the transfer of power

The value of a decision depends on *what* is chosen and *when*.

Collective decisions trade off accuracy against time

	◀ **act too early**	**act too late** ▶
	too little evidence	too much delay
DAO governance	enact an opposed upgrade	prolong a vulnerability window
Crowd annotation	mislabel the item	burn budget, slow the pipeline
Election audits	certify the wrong winner	delay the transfer of power

The value of a decision depends on *what* is chosen and *when*.

The design opportunity

- In many modern settings the path of votes is observable, not just the final tally
- Rules could adapt to the evolving informational state — instead of fixed voting windows, participation thresholds, or predetermined sample sizes

How should collective decision rules exploit the vote path
to balance accuracy and timeliness?

The design opportunity

- In many modern settings the path of votes is observable, not just the final tally
- Rules could adapt to the evolving informational state — instead of fixed voting windows, participation thresholds, or predetermined sample sizes

How should collective decision rules exploit the vote path
to balance accuracy and timeliness?

The design opportunity

- In many modern settings the path of votes is observable, not just the final tally
- Rules could adapt to the evolving informational state — instead of fixed voting windows, participation thresholds, or predetermined sample sizes

How should collective decision rules exploit the vote path
to balance accuracy and timeliness?

The design opportunity

- In many modern settings the path of votes is observable, not just the final tally
- Rules could adapt to the evolving informational state — instead of fixed voting windows, participation thresholds, or predetermined sample sizes

How should collective decision rules exploit the vote path
to balance accuracy and timeliness?

This paper

We study optimal collective decision rules: rules that identify the majority preference by acting on the observed vote stream.

- Three structural features define the problem

finite, exhaustible pool	vs. Wald-style tests: inexhaustible signal source
exogenous arrivals	vs. dynamic voting design: planner solicits votes at will
hard deadline	vs. finite-population audits: no binding horizon

This paper

We study optimal collective decision rules: rules that identify the majority preference by acting on the observed vote stream.

- Three structural features define the problem

finite, exhaustible pool	vs. Wald-style tests: inexhaustible signal source
exogenous arrivals	vs. dynamic voting design: planner solicits votes at will
hard deadline	vs. finite-population audits: no binding horizon

This paper

We study optimal collective decision rules: rules that identify the majority preference by acting on the observed vote stream.

- Three structural features define the problem

finite, exhaustible pool	vs. Wald-style tests: inexhaustible signal source
exogenous arrivals	vs. dynamic voting design: planner solicits votes at will
hard deadline	vs. finite-population audits: no binding horizon

This paper

We study optimal collective decision rules: rules that identify the majority preference by acting on the observed vote stream.

- Three structural features define the problem

finite, exhaustible pool	vs. Wald-style tests: inexhaustible signal source
exogenous arrivals	vs. dynamic voting design: planner solicits votes at will
hard deadline	vs. finite-population audits: no binding horizon

This paper

We study optimal collective decision rules: rules that identify the majority preference by acting on the observed vote stream.

- Three structural features define the problem

finite, exhaustible pool	vs. Wald-style tests: inexhaustible signal source
exogenous arrivals	vs. dynamic voting design: planner solicits votes at will
hard deadline	vs. finite-population audits: no binding horizon

Related literature

- Statistical social choice: votes as noisy signals of a latent quantity
  $\rightarrow$ we keep the inference view, but the rule is the choice variable
  Condorcet 1785; Young '88; Nitzan–Paroush '84; Austen-Smith–Banks '96; Feddersen–Pesendorfer '96; Myerson '98
- Optimal design of (dynamic) voting rules: planner gathers costly information
  $\rightarrow$ here: the sample is given, not built.
  Nitzan–Paroush '82; Persico '04; Gershkov–Szentes '09; Bognar–Börgers–Meyer-ter-Vehn '15; Gershkov–Moldovanu–Shi '17
- Sequential statistical decisions & optimal stopping: optimal stopping with a cost, collapsing boundaries
  $\rightarrow$ counterpart in collective choice under finite population and fixed deadline
  Wald '48; Chernoff '61; Lai '79; Siegmund '85; Drugowitsch et al. '12; Fudenberg–Strack–Strzalecki '18; Stark '08

Related literature

- Statistical social choice: votes as noisy signals of a latent quantity
  $\rightarrow$ we keep the inference view, but the rule is the choice variable
  Condorcet 1785; Young '88; Nitzan–Paroush '84; Austen-Smith–Banks '96; Feddersen–Pesendorfer '96; Myerson '98
- Optimal design of (dynamic) voting rules: planner gathers costly information
  $\rightarrow$ here: the sample is given, not built.
  Nitzan–Paroush '82; Persico '04; Gershkov–Szentes '09; Bognar–Börgers–Meyer-ter-Vehn '15; Gershkov–Moldovanu–Shi '17
- Sequential statistical decisions & optimal stopping: optimal stopping with a cost, collapsing boundaries
  $\rightarrow$ counterpart in collective choice under finite population and fixed deadline
  Wald '48; Chernoff '61; Lai '79; Siegmund '85; Drugowitsch et al. '12; Fudenberg–Strack–Strzalecki '18; Stark '08

Related literature

- Statistical social choice: votes as noisy signals of a latent quantity
  $\rightarrow$ we keep the inference view, but the rule is the choice variable
  Condorcet 1785; Young '88; Nitzan–Paroush '84; Austen-Smith–Banks '96; Feddersen–Pesendorfer '96; Myerson '98
- Optimal design of (dynamic) voting rules: planner gathers costly information
  $\rightarrow$ here: the sample is given, not built.
  Nitzan–Paroush '82; Persico '04; Gershkov–Szentes '09; Bognar–Börgers–Meyer-ter-Vehn '15; Gershkov–Moldovanu–Shi '17
- Sequential statistical decisions & optimal stopping: optimal stopping with a cost, collapsing boundaries
  $\rightarrow$ counterpart in collective choice under finite population and fixed deadline
  Wald '48; Chernoff '61; Lai '79; Siegmund '85; Drugowitsch et al. '12; Fudenberg–Strack–Strzalecki '18; Stark '08

Environment

- Society of N voters (N odd), each with a fixed binary preference: $+$ or $-$
- Unknown number of supporters $M^+ \in \{0, \dots, N\}$, $\quad M^- = N - M^+$
- Benevolent planner: accept iff $M^+ > N/2$
- Uniform prior $M^+ \sim \mathrm{Unif}\{0, \dots, N\}$

Environment

- Society of N voters (N odd), each with a fixed binary preference: $+$ or $-$
- Unknown number of supporters $M^+ \in \{0, \ldots, N\}$, $\quad M^- = N - M^+$
- Benevolent planner: accept iff $M^+ > N/2$
- Uniform prior $M^+ \sim \mathrm{Unif}\{0, \ldots, N\}$

Environment

- Society of N voters (N odd), each with a fixed binary preference: $+$ or $-$
- Unknown number of supporters $M^+ \in \{0, \ldots, N\}$, $\quad M^- = N - M^+$
- Benevolent planner: accept iff $M^+ > N/2$
- Uniform prior $M^+ \sim \mathrm{Unif}\{0, \ldots, N\}$

Sequential vote arrival

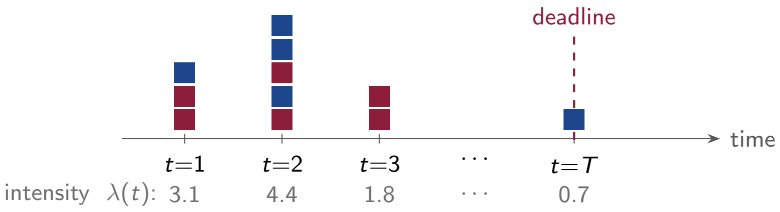

- Discrete dates $t = 1, \ldots, T$; batch of a_t^+ favorable, a_t^- unfavorable votes at t
- Cumulative counts s_t^+, s_t^-, turnout $s_t = s_t^+ + s_t^-$, lead $d_t = |s_t^+ - s_t^-|$
- Batch sizes: $A_t \sim \mathrm{Poisson}(\lambda(t))$ independent across t; $\Lambda(t) = \sum_{\tau \leq t} \lambda(\tau)$
- Each voter votes at most once $\Rightarrow$ censoring $a_t = \min\{A_t, N - s_{t-1}\}$
- Sampling without replacement, $a^+ \mid a, M^+ \sim \mathrm{Hypergeom}(N - s,\ M^+ - s^+,\ a)$

Sequential vote arrival

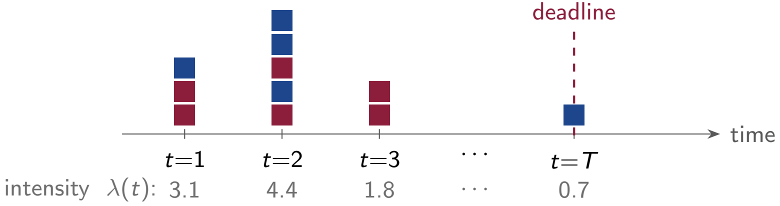

- Discrete dates $t = 1, \ldots, T$; batch of a_t^+ favorable, a_t^- unfavorable votes at t
- Cumulative counts s_t^+, s_t^-, turnout $s_t = s_t^+ + s_t^-$, lead $d_t = |s_t^+ - s_t^-|$
- Batch sizes: $A_t \sim \mathrm{Poisson}(\lambda(t))$ independent across t; $\Lambda(t) = \sum_{\tau \leq t} \lambda(\tau)$
- Each voter votes at most once $\Rightarrow$ censoring $a_t = \min\{A_t,\, N - s_{t-1}\}$
- Sampling without replacement, $a^+ \mid a, M^+ \sim \mathrm{Hypergeom}(N - s,\ M^+ - s^+,\ a)$

Sequential vote arrival

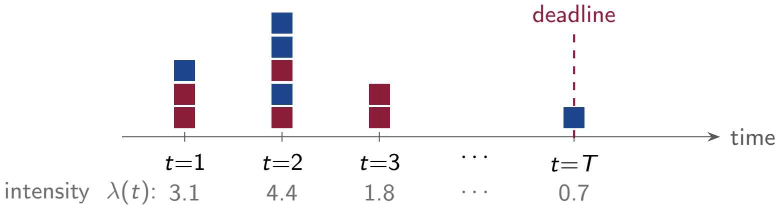

- Discrete dates $t = 1, \ldots, T$; batch of a_t^+ favorable, a_t^- unfavorable votes at t
- Cumulative counts s_t^+, s_t^-, turnout $s_t = s_t^+ + s_t^-$, lead $d_t = |s_t^+ - s_t^-|$
- Batch sizes: $A_t \sim \mathrm{Poisson}(\lambda(t))$ independent across t; $\Lambda(t) = \sum_{\tau \leq t} \lambda(\tau)$
- Each voter votes at most once $\Rightarrow$ censoring $a_t = \min\{A_t, N - s_{t-1}\}$
- Sampling without replacement, $a^+ \mid a, M^+ \sim \mathrm{Hypergeom}(N - s,\ M^+ - s^+,\ a)$

Learning

- Voters exchangeable $\Rightarrow (s_t^+, s_t^-)$ is a sufficient statistic

Lemma (Predictive distribution)

For a batch of size $a \leq N - s$,

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

- Beta–Binomial(a, $s^+{+}1$, $s^-{+}1$): independent of N, which enters only via censoring and the majority cutoff
- Posterior belief $\quad \mu(s^+, s^-) = P\left(M^+ > \frac{N}{2} \,\middle|\, s^+, s^-\right) = \displaystyle\sum_{m^+ > N/2} \frac{\binom{m^+}{s^+}\binom{N - m^+}{s^-}}{\binom{N+1}{s+1}}$
- Equivalent states: $(s^+, s^-) \leftrightarrow (\mu, s) \leftrightarrow (d, s)$

Learning

- Voters exchangeable $\Rightarrow$ (s_t^+, s_t^-) is a sufficient statistic

Lemma (Predictive distribution)

For a batch of size $a \leq N - s$,

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

- Beta–Binomial(a, s^++1, s^-+1): independent of N, which enters only via censoring and the majority cutoff
- Posterior belief $\mu(s^+, s^-) = P\left(M^+ > \frac{N}{2} \,\middle|\, s^+, s^-\right) = \sum_{m^+ > N/2} \frac{\binom{m^+}{s^+}\binom{N - m^+}{s^-}}{\binom{N+1}{s+1}}$
- Equivalent states: $(s^+, s^-) \leftrightarrow (\mu, s) \leftrightarrow (d, s)$

Learning

- Voters exchangeable $\Rightarrow$ (s_t^+, s_t^-) is a sufficient statistic

Lemma (Predictive distribution)

For a batch of size $a \leq N - s$,

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

- Beta–Binomial(a, $s^+\!+1$, $s^-\!+1$): independent of N, which enters only via censoring and the majority cutoff
- Posterior belief $\quad \mu(s^+, s^-) = P\left(M^+ > \frac{N}{2} \,\middle|\, s^+, s^-\right) = \displaystyle\sum_{m^+ > N/2} \frac{\binom{m^+}{s^+}\binom{N-m^+}{s^-}}{\binom{N+1}{s+1}}$
- Equivalent states: $(s^+, s^-) \leftrightarrow (\mu, s) \leftrightarrow (d, s)$

Learning

- Voters exchangeable $\Rightarrow$ (s_t^+, s_t^-) is a sufficient statistic

Lemma (Predictive distribution)

For a batch of size $a \le N - s$,

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

- Beta–Binomial(a, $s^+ + 1$, $s^- + 1$): independent of N, which enters only via censoring and the majority cutoff
- Posterior belief $\quad \mu(s^+, s^-) = P\left(M^+ > \frac{N}{2} \,\middle|\, s^+, s^-\right) = \sum_{m^+ > N/2} \frac{\binom{m^+}{s^+}\binom{N-m^+}{s^-}}{\binom{N+1}{s+1}}$
- Equivalent states: $(s^+, s^-) \leftrightarrow (\mu, s) \leftrightarrow (d, s)$

Planner's problem

At each t, after observing (s_t^+, s_t^-): stop or continue at flow cost $\kappa > 0$.

- Stopping payoff (implement the leading alternative):

$$U(s^+, s^-) = R \cdot \max\{\mu(s^+, s^-), 1 - \mu(s^+, s^-)\}, \qquad R > 0$$

- Bellman equation, terminal condition $V_T = U$:

$$V_t(s^+, s^-) = \max\Big\{ U(s^+, s^-), \underbrace{-\kappa + \mathbb{E}\big[V_{t+1}(\tilde{s}_{t+1}^+, \tilde{s}_{t+1}^-) \mid s^+, s^-\big]}_{\mathcal{C}_t(s^+, s^-)} \Big\}$$

expectation over censored-Poisson batch size and Beta–Binomial composition

- Dynamic voting rule $\phi : (s^+, s^-, t) \mapsto \{\text{stop}, \text{continue}\}$; ϕ^* optimal

Planner's problem

At each t, after observing (s_t^+, s_t^-): stop or continue at flow cost $\kappa > 0$.

- Stopping payoff (implement the leading alternative):

$$U(s^+, s^-) = R \cdot \max\{\mu(s^+, s^-), 1 - \mu(s^+, s^-)\}, \qquad R > 0$$

- Bellman equation, terminal condition $V_T = U$:

$$V_t(s^+, s^-) = \max\Big\{ U(s^+, s^-), \underbrace{-\kappa + \mathbb{E}\big[V_{t+1}(\tilde{s}^+_{t+1}, \tilde{s}^-_{t+1}) \mid s^+, s^-\big]}_{\mathcal{C}_t(s^+, s^-)}\Big\}$$

expectation over censored-Poisson batch size and Beta–Binomial composition

- Dynamic voting rule $\phi : (s^+, s^-, t) \mapsto \{\text{stop}, \text{continue}\}$; $\quad \phi^*$ optimal

Planner's problem

At each t, after observing (s_t^+, s_t^-): stop or continue at flow cost $\kappa > 0$.

- Stopping payoff (implement the leading alternative):

$$U(s^+, s^-) = R \cdot \max\{\mu(s^+, s^-), 1 - \mu(s^+, s^-)\}, \qquad R > 0$$

- Bellman equation, terminal condition $V_T = U$:

$$V_t(s^+, s^-) = \max\Big\{ U(s^+, s^-), \underbrace{-\kappa + \mathbb{E}\big[V_{t+1}(\tilde{s}^+_{t+1}, \tilde{s}^-_{t+1}) \mid s^+, s^-\big]}_{\mathcal{C}_t(s^+, s^-)}\Big\}$$

  expectation over censored-Poisson batch size and Beta–Binomial composition

- Dynamic voting rule $\phi : (s^+, s^-, t) \mapsto \{\text{stop}, \text{continue}\}$; ϕ^* optimal

Planner's problem

At each t, after observing (s_t^+, s_t^-): stop or continue at flow cost $\kappa > 0$.

- Stopping payoff (implement the leading alternative):

$$U(s^+, s^-) = R \cdot \max\{\mu(s^+, s^-),\, 1 - \mu(s^+, s^-)\}, \qquad R > 0$$

- Bellman equation, terminal condition $V_T = U$:

$$V_t(s^+, s^-) = \max\Big\{ U(s^+, s^-),\, \underbrace{-\kappa + \mathbb{E}\big[V_{t+1}(\tilde{s}^+_{t+1}, \tilde{s}^-_{t+1}) \mid s^+, s^-\big]}_{\mathcal{C}_t(s^+,s^-)}\Big\}$$

  expectation over censored-Poisson batch size and Beta–Binomial composition

- Dynamic voting rule $\phi :\ (s^+, s^-, t) \mapsto \{\text{stop}, \text{continue}\}$; $\quad \phi^*$ optimal

The optimal rule is a dynamic majority-quorum rule

Theorem (Threshold structure)

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 $d_t^*(s_t) \geq 0$ such that ϕ^* prescribes continuation if and only if $|s_t^+ - s_t^-| < d_t^*(s_t)$. When $s_t = N$, $d_t^*(N) = 0$.

- The threshold depends on the lead, turnout, *and* time; a given lead may be decisive at high turnout but not at low
- Proof idea: The proof establishes that a larger current lead makes future beliefs stochastically more favorable. Consequently, the net gain from stopping is monotone in the current lead. ▸ details

The optimal rule is a dynamic majority-quorum rule

Theorem (Threshold structure)

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 $d_t^*(s_t) \geq 0$ such that ϕ^* prescribes continuation if and only if $|s_t^+ - s_t^-| < d_t^*(s_t)$. When $s_t = N$, $d_t^*(N) = 0$.

- The threshold depends on the lead, turnout, *and* time; a given lead may be decisive at high turnout but not at low
- Proof idea: The proof establishes that a larger current lead makes future beliefs stochastically more favorable. Consequently, the net gain from stopping is monotone in the current lead. details

The optimal rule is a dynamic majority-quorum rule

Theorem (Threshold structure)

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 $d_t^*(s_t) \geq 0$ such that ϕ^* prescribes continuation if and only if $|s_t^+ - s_t^-| < d_t^*(s_t)$. When $s_t = N$, $d_t^*(N) = 0$.

- The threshold depends on the lead, turnout, *and* time; a given lead may be decisive at high turnout but not at low
- Proof idea: The proof establishes that a larger current lead makes future beliefs stochastically more favorable. Consequently, the net gain from stopping is monotone in the current lead. ▸ details

The exact boundary: a collapsing lens

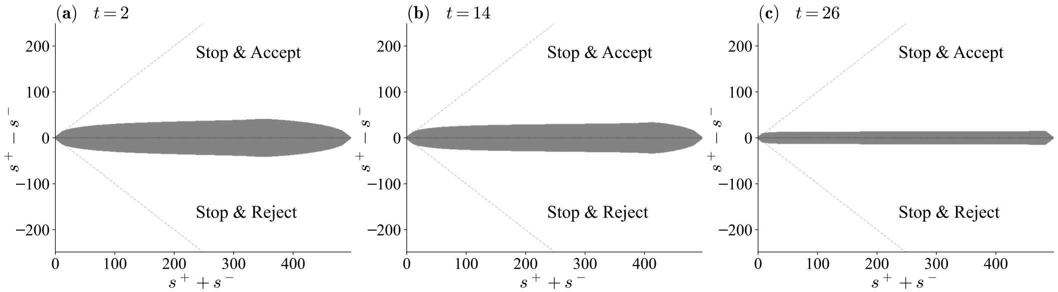

Exact rule ϕ^* in lead–turnout space at several dates; shaded: waiting is optimal. $N = 497$, $T = 28$, $\kappa = 0.001$, $R = 100$ ($\rho = 10^5$), uniform arrivals with $\Lambda(T) = 159$.

- Two forces shape the lens: statistical precision (small leads informative only at sufficient turnout $\Rightarrow$ widening) vs. finite-pool depletion (surviving leads become decisive as $s_t \to N \Rightarrow$ collapse), while the whole boundary sinks as $t \to T$

Comparative statics

Proposition

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

1. $d_t^*(s; \rho', \boldsymbol{\lambda}) \leq d_t^*(s; \rho, \boldsymbol{\lambda})$ *for all* $\rho' \leq \rho$, t, s;
2. $d_t^*(s; \rho, \boldsymbol{\lambda}') \leq d_t^*(s; \rho, \boldsymbol{\lambda})$ *whenever* $\lambda'(t) \leq \lambda(t)$ *for all* $t \in \{1, \ldots, T\}$;
3. $d_{t+1}^*(s; \rho, \boldsymbol{\lambda}) \leq d_t^*(s; \rho, \boldsymbol{\lambda})$ *for all* $t < T$, s, *whenever* $\lambda(\tau) \geq \lambda(\tau + 1)$ *for all* $\tau \in \{t + 1, \ldots, T - 1\}$.

- One common force: the option value of continued sampling
- The conditions are tight: cumulative dominance does *not* suffice in (ii), and back-loaded arrivals can reverse monotonicity in (iii) counterexamples

Comparative statics

Proposition

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

1. $d_t^*(s;\rho',\boldsymbol{\lambda}) \leq d_t^*(s;\rho,\boldsymbol{\lambda})$ *for all* $\rho' \leq \rho$, t, s;
2. $d_t^*(s;\rho,\boldsymbol{\lambda}') \leq d_t^*(s;\rho,\boldsymbol{\lambda})$ *whenever* $\lambda'(t) \leq \lambda(t)$ *for all* $t \in \{1,\ldots,T\}$;
3. $d_{t+1}^*(s;\rho,\boldsymbol{\lambda}) \leq d_t^*(s;\rho,\boldsymbol{\lambda})$ *for all* $t < T$, s, *whenever* $\lambda(\tau) \geq \lambda(\tau+1)$ *for all* $\tau \in \{t+1,\ldots,T-1\}$.

- One common force: the option value of continued sampling
- The conditions are tight: cumulative dominance does *not* suffice in (ii), and back-loaded arrivals can reverse monotonicity in (iii) ▸ counterexamples

Comparative statics

Proposition

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

1. $d_t^*(s;\rho',\boldsymbol{\lambda}) \leq d_t^*(s;\rho,\boldsymbol{\lambda})$ *for all* $\rho' \leq \rho$, t, s;
2. $d_t^*(s;\rho,\boldsymbol{\lambda}') \leq d_t^*(s;\rho,\boldsymbol{\lambda})$ *whenever* $\lambda'(t) \leq \lambda(t)$ *for all* $t \in \{1,\ldots,T\}$;
3. $d_{t+1}^*(s;\rho,\boldsymbol{\lambda}) \leq d_t^*(s;\rho,\boldsymbol{\lambda})$ *for all* $t < T$, s, *whenever* $\lambda(\tau) \geq \lambda(\tau+1)$ *for all* $\tau \in \{t+1,\ldots,T-1\}$.

- One common force: the option value of continued sampling
- The conditions are tight: cumulative dominance does *not* suffice in (ii), and back-loaded arrivals can reverse monotonicity in (iii) counterexamples

Comparative statics

Proposition

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

1. $d_t^*(s; \rho', \boldsymbol{\lambda}) \leq d_t^*(s; \rho, \boldsymbol{\lambda})$ *for all* $\rho' \leq \rho$, t, s;
2. $d_t^*(s; \rho, \boldsymbol{\lambda}') \leq d_t^*(s; \rho, \boldsymbol{\lambda})$ *whenever* $\lambda'(t) \leq \lambda(t)$ *for all* $t \in \{1, \ldots, T\}$;
3. $d_{t+1}^*(s; \rho, \boldsymbol{\lambda}) \leq d_t^*(s; \rho, \boldsymbol{\lambda})$ *for all* $t < T$, s, *whenever* $\lambda(\tau) \geq \lambda(\tau + 1)$ *for all* $\tau \in \{t + 1, \ldots, T - 1\}$.

- One common force: the option value of continued sampling
- The conditions are tight: cumulative dominance does *not* suffice in (ii), and back-loaded arrivals can reverse monotonicity in (iii) ▸ counterexamples

Comparative statics

Proposition

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

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

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

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

- One common force: the option value of continued sampling
- The conditions are tight: cumulative dominance does *not* suffice in (ii), and back-loaded arrivals can reverse monotonicity in (iii) ▸ counterexamples

Illustration: how the boundary moves with ρ and $\boldsymbol{\lambda}$

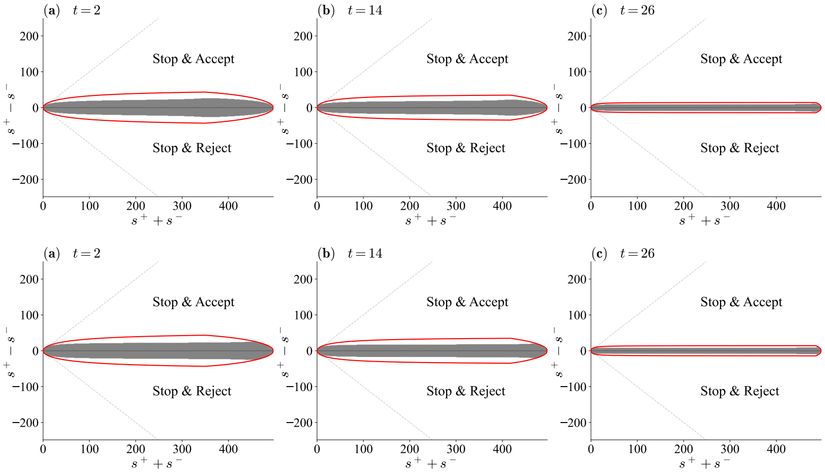

Red curve: benchmark parameterization.

Lower stakes (top) and thinner arrivals (bottom) each compress the continuation lens, as in parts (i)–(ii).

From backward induction to a mechanism

- No closed form for $d_t^*(s)$: the value of waiting is defined recursively through later stopping decisions, on a two-dimensional state, over a finite horizon, with a transition law that shifts as the pool depletes
- Backward induction is exact but is too opaque to be used as a governance mechanism
- Strategy: derive an approximate closed-form rule

From backward induction to a mechanism

- No closed form for $d_t^*(s)$: the value of waiting is defined recursively through later stopping decisions, on a two-dimensional state, over a finite horizon, with a transition law that shifts as the pool depletes
- Backward induction is exact but is too opaque to be used as a governance mechanism
- Strategy: derive an approximate closed-form rule

From backward induction to a mechanism

- No closed form for $d_t^*(s)$: the value of waiting is defined recursively through later stopping decisions, on a two-dimensional state, over a finite horizon, with a transition law that shifts as the pool depletes
- Backward induction is exact but is too opaque to be used as a governance mechanism
- Strategy: derive an approximate closed-form rule

A closed-form rule with a constant threshold

Censored forecast of terminal turnout: $\hat{s}_T := \min\{s_t + (\Lambda(T) - \Lambda(t)),\ N\}$.

Definition (Adaptive rule)

Let $\hat{\phi}$ be the majority-quorum rule stopping at (d_t, s_t, t) iff $t = T$, or $s_t = N$, or $t < T$ and

$$\frac{(d_t/s_t)^2}{1/s_t - 1/\hat{s}_T} \ \geq\ \left[2\log\rho - 2\log\log\rho\right].$$

- LHS: squared signal-to-noise ratio — strength of the sample lead over the reducible sampling variance: current noise $1/s_t$ *net of* the noise $1/\hat{s}_T$ that survives to the deadline
- RHS: a constant — the Wald–Chernoff stakes level, with a log log correction

A closed-form rule with a constant threshold

Censored forecast of terminal turnout: $\hat{s}_T := \min\{s_t + (\Lambda(T) - \Lambda(t)),\ N\}$.

Definition (Adaptive rule)

Let $\hat{\phi}$ be the majority-quorum rule stopping at (d_t, s_t, t) iff $t = T$, or $s_t = N$, or $t < T$ and

$$\frac{(d_t/s_t)^2}{1/s_t - 1/\hat{s}_T} \geq [2 \log \rho - 2 \log \log \rho].$$

- LHS: squared signal-to-noise ratio — strength of the sample lead over the reducible sampling variance: current noise $1/s_t$ *net of* the noise $1/\hat{s}_T$ that survives to the deadline
- RHS: a constant — the Wald–Chernoff stakes level, with a log log correction

A closed-form rule with a constant threshold

Censored forecast of terminal turnout: $\hat{s}_T := \min\{s_t + (\Lambda(T) - \Lambda(t)),\ N\}$.

Definition (Adaptive rule)

Let $\hat{\phi}$ be the majority-quorum rule stopping at (d_t, s_t, t) iff $t = T$, or $s_t = N$, or $t < T$ and

$$\frac{(d_t/s_t)^2}{1/s_t - 1/\hat{s}_T} \ \geq\ \left[2 \log \rho - 2 \log\log \rho\right].$$

- LHS: squared signal-to-noise ratio — strength of the sample lead over the reducible sampling variance: current noise $1/s_t$ *net of* the noise $1/\hat{s}_T$ that survives to the deadline
- RHS: a constant — the Wald–Chernoff stakes level, with a $\log\log$ correction

The closed form tracks the exact boundary

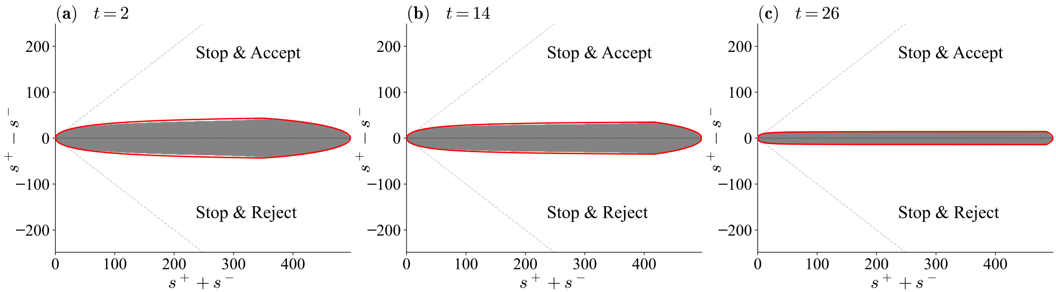

Exact rule ϕ^* in lead–turnout space at several dates and $\hat{\phi}$ (red); $N = 497$, $T = 28$, $\kappa = 0.001$, $R = 100$ ($\rho = 10^5$), uniform arrivals with $\Lambda(T) = 159$.

The adaptive rule replicates the optimal policy

Theorem (Approximation)

Fix T. For every $\varepsilon \in (0, \frac{1}{2})$ there exist $N_\varepsilon, \rho_\varepsilon$ such that, for all $N \geq N_\varepsilon$ and $\rho \geq \rho_\varepsilon$ with $\log \rho \leq \varepsilon N^{1/3}$, 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 $\mathcal{S}_\varepsilon := \left\{(d_t, s_t, t) : \ t < T, \ \varepsilon N \leq s_t, \ s_t + \varepsilon N \leq \hat{s}_T \leq (1-\varepsilon)N\right\}$.

- Disagreement possible only on a band around the stopping boundary (where the planner is closest to indifferent between acting and waiting) whose relative width $\to 0$ as $\rho \to \infty$ ▸ proof sketch

The adaptive rule replicates the optimal policy

Theorem (Approximation)

Fix T. For every $\varepsilon \in (0, \frac{1}{2})$ there exist $N_\varepsilon, \rho_\varepsilon$ such that, for all $N \geq N_\varepsilon$ and $\rho \geq \rho_\varepsilon$ with $\log \rho \leq \varepsilon N^{1/3}$, 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 $\mathcal{S}_\varepsilon := \left\{(d_t, s_t, t) : \ t < T, \ \varepsilon N \leq s_t, \ s_t + \varepsilon N \leq \hat{s}_T \leq (1-\varepsilon)N\right\}$.

- Disagreement possible only on a band around the stopping boundary (where the planner is closest to indifferent between acting and waiting) whose relative width $\to 0$ as $\rho \to \infty$ ▸ proof sketch

A postable form of the rule

$\hat{\phi}$ cannot be represented as a dated schedule; the factor $N/(N-s_t)$ cancels in ratio form; carrying it through and evaluating the threshold on the mean path $s_t \to \Lambda(t)$ fixes it ex ante:

Definition (Posted rule)

Let $\tilde{\phi}$ be the majority-quorum rule stopping at (d_t, s_t, t) iff $t = T$, or $s_t = N$, or $t < T$ and

$$\frac{d_t^2}{s_t} \cdot \frac{N}{N-s_t} \geq \left[2\log\rho - 2\log\log\rho\right] \frac{N\big(\Lambda(T)-\Lambda(t)\big)}{\big(N-\Lambda(t)\big)\Lambda(T)}.$$

- Evidence: LLR statistic d_t^2/s_t for a Bernoulli parameter under sampling without replacement, multiplied by $N/(N-s_t)$
- Threshold: Wald–Chernoff level $\times$ residual informativeness
- Informational time: only $\Lambda(t), \Lambda(T)$ matter — monotone collapse under *any* arrival process

A postable form of the rule

$\hat{\phi}$ cannot be represented as a dated schedule; the factor $N/(N - s_t)$ cancels in ratio form; carrying it through and evaluating the threshold on the mean path $s_t \to \Lambda(t)$ fixes it ex ante:

Definition (Posted rule)

Let $\tilde{\phi}$ be the majority-quorum rule stopping at (d_t, s_t, t) iff $t = T$, or $s_t = N$, or $t < T$ and

$$\frac{d_t^2}{s_t} \cdot \frac{N}{N - s_t} \geq \left[2 \log \rho - 2 \log \log \rho\right] \frac{N(\Lambda(T) - \Lambda(t))}{(N - \Lambda(t))\Lambda(T)}.$$

- Evidence: LLR statistic d_t^2/s_t for a Bernoulli parameter under sampling without replacement, multiplied by $N/(N - s_t)$
- Threshold: Wald–Chernoff level $\times$ residual informativeness
- Informational time: only $\Lambda(t), \Lambda(T)$ matter — monotone collapse under *any* arrival process

A postable form of the rule

$\hat{\phi}$ cannot be represented as a dated schedule; the factor $N/(N - s_t)$ cancels in ratio form; carrying it through and evaluating the threshold on the mean path $s_t \to \Lambda(t)$ fixes it ex ante:

Definition (Posted rule)

Let $\tilde{\phi}$ be the majority-quorum rule stopping at (d_t, s_t, t) iff $t = T$, or $s_t = N$, or $t < T$ and

$$\frac{d_t^2}{s_t} \cdot \frac{N}{N - s_t} \geq \left[2 \log \rho - 2 \log \log \rho\right] \frac{N(\Lambda(T) - \Lambda(t))}{(N - \Lambda(t))\Lambda(T)}.$$

- Evidence: LLR statistic d_t^2/s_t for a Bernoulli parameter under sampling without replacement, multiplied by $N/(N - s_t)$
- Threshold: Wald–Chernoff level $\times$ residual informativeness
- Informational time: only $\Lambda(t), \Lambda(T)$ matter — monotone collapse under *any* arrival process

A postable form of the rule

$\hat{\phi}$ cannot be represented as a dated schedule; the factor $N/(N - s_t)$ cancels in ratio form; carrying it through and evaluating the threshold on the mean path $s_t \to \Lambda(t)$ fixes it ex ante:

Definition (Posted rule)

Let $\tilde{\phi}$ be the majority-quorum rule stopping at (d_t, s_t, t) iff $t = T$, or $s_t = N$, or $t < T$ and

$$\frac{d_t^2}{s_t} \cdot \frac{N}{N - s_t} \geq \left[2 \log \rho - 2 \log \log \rho\right] \frac{N(\Lambda(T) - \Lambda(t))}{(N - \Lambda(t))\Lambda(T)}.$$

- Evidence: LLR statistic d_t^2/s_t for a Bernoulli parameter under sampling without replacement, multiplied by $N/(N - s_t)$
- Threshold: Wald–Chernoff level $\times$ residual informativeness
- Informational time: only $\Lambda(t), \Lambda(T)$ matter — monotone collapse under *any* arrival process

The two rules, and the shape of the collapse

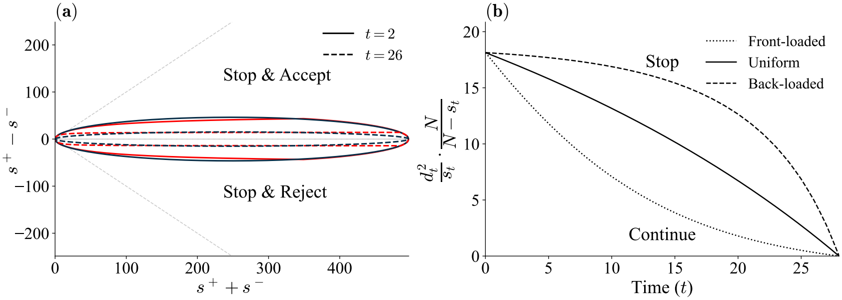

- (a) $\tilde{\phi}$'s lens peaks at $s_t = N/2$; $\hat{\phi}$'s peak shifts right as updates to $\hat{s}_T$ delay pool saturation — yet the two boundaries stay quantitatively close
- (b) Front-loading concentrates information early: rapid initial decline, flat near the deadline; back-loading reverses the pattern; totals held fixed

The two rules, and the shape of the collapse

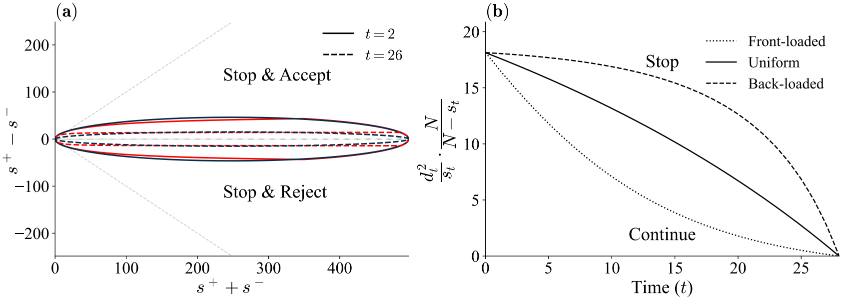

- (a) $\tilde{\phi}$'s lens peaks at $s_t = N/2$; $\hat{\phi}$'s peak shifts right as updates to $\hat{s}_T$ delay pool saturation — yet the two boundaries stay quantitatively close
- (b) Front-loading concentrates information early: rapid initial decline, flat near the deadline; back-loading reverses the pattern; totals held fixed

The two rules, and the shape of the collapse

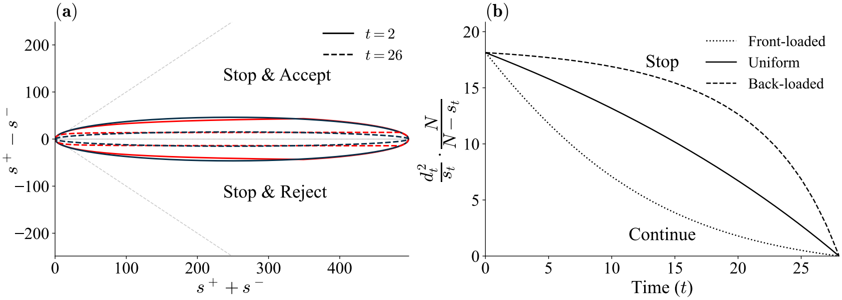

- (a) $\tilde{\phi}$'s lens peaks at $s_t = N/2$; $\hat{\phi}$'s peak shifts right as updates to $\hat{s}_T$ delay pool saturation — yet the two boundaries stay quantitatively close
- (b) Front-loading concentrates information early: rapid initial decline, flat near the deadline; back-loading reverses the pattern; totals held fixed

The rule as a sequential significance test

$\rho = R/\kappa$ is hard to reason about in practice; restate the same rule via an error bound:

Corollary (Calibration)

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

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

where $\alpha_{\bar{t}} \in (0, \frac{1}{2})$ is the error tolerance at $\bar{t}$: the posterior probability of implementing the minority preference is at most $\alpha_{\bar{t}}$ for any decision reached at or before $\bar{t}$, given by $\alpha_{\bar{t}} = 1 - \Phi\left(\sqrt{\left[2 \log \rho - 2 \log \log \rho\right] \frac{N(\Lambda(T) - \Lambda(\bar{t}))}{(N - \Lambda(\bar{t}))\Lambda(T)}}\right)$.

- Stop as soon as P(leading alternative is the minority preference) $\leq \alpha_{\bar{t}}$

The rule as a sequential significance test

$\rho = R/\kappa$ is hard to reason about in practice; restate the same rule via an error bound:

Corollary (Calibration)

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

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

where $\alpha_{\bar{t}} \in (0, \frac{1}{2})$ is the error tolerance at $\bar{t}$: the posterior probability of implementing the minority preference is at most $\alpha_{\bar{t}}$ for any decision reached at or before $\bar{t}$, given by $\alpha_{\bar{t}} = 1 - \Phi\left(\sqrt{\left[2\log\rho - 2\log\log\rho\right] \frac{N(\Lambda(T)-\Lambda(\bar{t}))}{(N-\Lambda(\bar{t}))\Lambda(T)}}\right)$.

- Stop as soon as $P(\text{leading alternative is the minority preference}) \leq \alpha_{\bar{t}}$

Evidence: Polkadot OpenGov

- Among the largest DAOs by market cap; all protocol changes and treasury spending via binding on-chain referenda, 28-day decision periods
- Full vote path observed at twice-daily resolution $\rightarrow$ maps directly to (s_t^+, s_t^-)
- Incumbent rule is itself dynamic: the Approval–Support Mechanism (ASM)
- $\Rightarrow$ strategic timing is possible and likely
- Data: 1,824 referenda (06/2023–10/2025) $\rightarrow$ 952 cleaned $\rightarrow$ 933 in the stopping evaluation
- Participation in accounts, not tokens ▸ cleaning & wedge

Evidence: Polkadot OpenGov

- Among the largest DAOs by market cap; all protocol changes and treasury spending via binding on-chain referenda, 28-day decision periods
- Full vote path observed at twice-daily resolution → maps directly to (s_t^+, s_t^-)
- Incumbent rule is itself dynamic: the Approval–Support Mechanism (ASM)
- ⇒ strategic timing is possible and likely
- Data: 1,824 referenda (06/2023–10/2025) → 952 cleaned → 933 in the stopping evaluation
- Participation in accounts, not tokens ▸ cleaning & wedge

Evidence: Polkadot OpenGov

- Among the largest DAOs by market cap; all protocol changes and treasury spending via binding on-chain referenda, 28-day decision periods
- Full vote path observed at twice-daily resolution $\to$ maps directly to (s_t^+, s_t^-)
- Incumbent rule is itself dynamic: the Approval–Support Mechanism (ASM)
- $\Rightarrow$ strategic timing is possible and likely
- Data: 1,824 referenda (06/2023–10/2025) $\to$ 952 cleaned $\to$ 933 in the stopping evaluation
- Participation in accounts, not tokens ‣ cleaning & wedge

Evidence: Polkadot OpenGov

- Among the largest DAOs by market cap; all protocol changes and treasury spending via binding on-chain referenda, 28-day decision periods
- Full vote path observed at twice-daily resolution $\rightarrow$ maps directly to (s_t^+, s_t^-)
- Incumbent rule is itself dynamic: the Approval–Support Mechanism (ASM)
- $\Rightarrow$ strategic timing is possible and likely
- Data: 1,824 referenda (06/2023–10/2025) $\rightarrow$ 952 cleaned $\rightarrow$ 933 in the stopping evaluation
- Participation in accounts, not tokens ▸ cleaning & wedge

The incumbent: the Approval–Support Mechanism

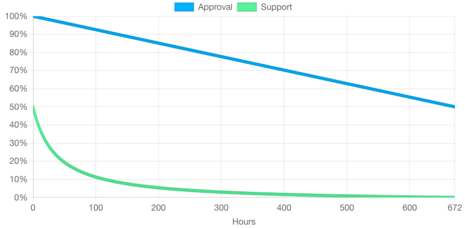

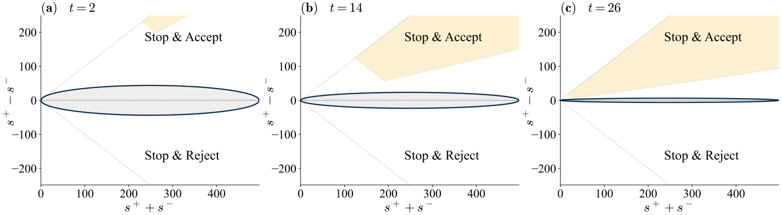

Both rules relax over time — the ASM by construction, $\tilde{\phi}$ because residual informativeness vanishes.

But the ASM permits early approval only; $\tilde{\phi}$ conditions on $|d_t|$ and can stop in either direction.

Is the arrival model tenable?

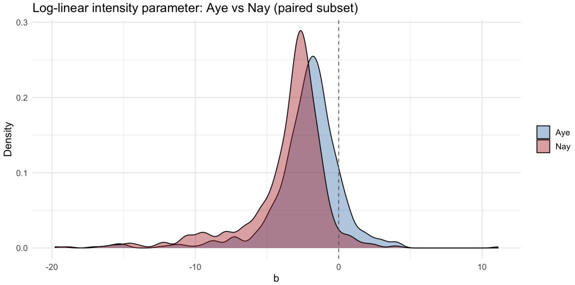

- Non-homogeneous Poisson, log-linear intensity in normalized time:

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

- Fits are tight across referenda; participation strongly front-loaded: pooled median $b = -3.05$
- Rule inputs: $\Lambda(t) = \Lambda(T)\, F(t/T; b)$, with $\Lambda(T) = 159$, $N = 496$ from the training half

Empirical strategy: strictly out of sample

- Chronological split at the median referendum: train (calibrate everything) / test (evaluate, frozen parameters)
- Planner states a tolerance α; the training data pin down the earliest reference date $\bar{t}^*(\alpha)$ whose threshold keeps realized training error within α
- Errors = divergent decisions: the rule's early call differs from the deadline account majority
- Replay each test referendum's realized path; record when $\tilde{\phi}$ with frozen $(\alpha, \bar{t}^*)$ first fires; time saved measured against realized ASM duration

Empirical strategy: strictly out of sample

- Chronological split at the median referendum: train (calibrate everything) / test (evaluate, frozen parameters)
- Planner states a tolerance α; the training data pin down the earliest reference date $\bar{t}^*(\alpha)$ whose threshold keeps realized training error within α
- Errors = divergent decisions: the rule's early call differs from the deadline account majority
- Replay each test referendum's realized path; record when $\tilde{\phi}$ with frozen $(\alpha, \bar{t}^*)$ first fires; time saved measured against realized ASM duration

Empirical strategy: strictly out of sample

- Chronological split at the median referendum: train (calibrate everything) / test (evaluate, frozen parameters)
- Planner states a tolerance α; the training data pin down the earliest reference date $\bar{t}^*(\alpha)$ whose threshold keeps realized training error within α
- Errors = divergent decisions: the rule's early call differs from the deadline account majority
- Replay each test referendum's realized path; record when $\tilde{\phi}$ with frozen $(\alpha, \bar{t}^*)$ first fires; time saved measured against realized ASM duration

Empirical strategy: strictly out of sample

- Chronological split at the median referendum: train (calibrate everything) / test (evaluate, frozen parameters)
- Planner states a tolerance α; the training data pin down the earliest reference date $\bar{t}^*(\alpha)$ whose threshold keeps realized training error within α
- Errors $=$ divergent decisions: the rule's early call differs from the deadline account majority
- Replay each test referendum's realized path; record when $\tilde{\phi}$ with frozen $(\alpha, \bar{t}^*)$ first fires; time saved measured against realized ASM duration

Out-of-sample performance

Tolerance α (%)	$\bar{t}^*$ (day) (train)	Divergent decisions (%) train	Divergent decisions (%) test	Median time saved (days) train	Median time saved (days) 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

- At $\alpha = 2\%$: median referendum decided on day 10.8 of 28, with 60% of final turnout cast (IQR 50–77%), 15 days faster than the ASM, diverging in 1.7% of test referenda
- $\bar{t}^*$ stable across an order of magnitude in α (days 20.3–24.2): a compact summary of how unrepresentative early participation currently is

Out-of-sample performance

Tolerance α (%)	$\bar{t}^*$ (day) (train)	Divergent decisions (%) train	test	Median time saved (days) 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

- At $\alpha = 2\%$: median referendum decided on day 10.8 of 28, with 60% of final turnout cast (IQR 50–77%), 15 days faster than the ASM, diverging in 1.7% of test referenda
- $\bar{t}^*$ stable across an order of magnitude in α (days 20.3–24.2): a compact summary of how unrepresentative early participation currently is

Out-of-sample performance

Tolerance α (%)	$\bar{t}^*$ (day) (train)	Divergent decisions (%) train	test	Median time saved (days) 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

- At $\alpha = 2\%$: median referendum decided on day 10.8 of 28, with 60% of final turnout cast (IQR 50–77%), 15 days faster than the ASM, diverging in 1.7% of test referenda
- $\bar{t}^*$ stable across an order of magnitude in α (days 20.3–24.2): a compact summary of how unrepresentative early participation currently is

The speed–accuracy frontier

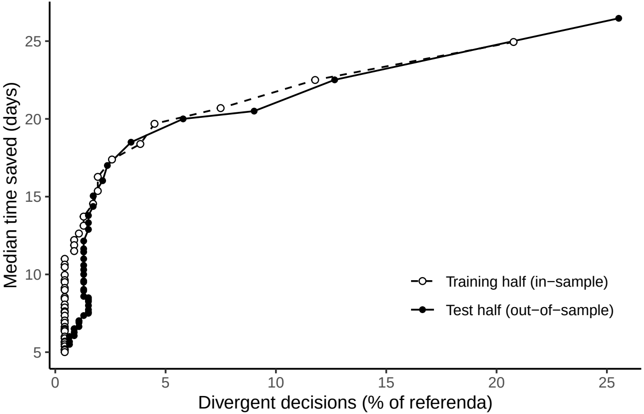

- Train and test curves nearly coincide
- Steep at the accurate end: the first basis points of error buy nearly all the delay savings; exact zero divergence forfeits the mechanism ($\bar{t} \to T \Rightarrow$ never stop early)

Takeaways

- Optimal aggregation of asynchronous votes under a finite pool and a hard deadline: a dynamic majority-quorum rule with a lens-shaped continuation region, collapsing toward the deadline
- Comparative statics governed by the option value of learning
- Large-population limit: one sufficient statistic; a sequential significance test on informational time, run by a single error tolerance
- Polkadot, out of sample: half the decision time at a 2% error tolerance

Thank you! bhargav@morpho.org {jonas, kremena}@parity.io

Takeaways

- Optimal aggregation of asynchronous votes under a finite pool and a hard deadline: a dynamic majority-quorum rule with a lens-shaped continuation region, collapsing toward the deadline
- Comparative statics governed by the option value of learning
- Large-population limit: one sufficient statistic; a sequential significance test on informational time, run by a single error tolerance
- Polkadot, out of sample: half the decision time at a 2% error tolerance

Thank you!

bhargav@morpho.org {jonas, kremena}@parity.io

Takeaways

- Optimal aggregation of asynchronous votes under a finite pool and a hard deadline: a dynamic majority-quorum rule with a lens-shaped continuation region, collapsing toward the deadline
- Comparative statics governed by the option value of learning
- Large-population limit: one sufficient statistic; a sequential significance test on informational time, run by a single error tolerance
- Polkadot, out of sample: half the decision time at a 2% error tolerance

Thank you!

bhargav@morpho.org {jonas, kremena}@parity.io

Takeaways

- Optimal aggregation of asynchronous votes under a finite pool and a hard deadline: a dynamic majority-quorum rule with a lens-shaped continuation region, collapsing toward the deadline
- Comparative statics governed by the option value of learning
- Large-population limit: one sufficient statistic; a sequential significance test on informational time, run by a single error tolerance
- Polkadot, out of sample: half the decision time at a 2% error tolerance

Thank you!

bhargav@morpho.org {jonas, kremena}@parity.io

Takeaways

- Optimal aggregation of asynchronous votes under a finite pool and a hard deadline: a dynamic majority-quorum rule with a lens-shaped continuation region, collapsing toward the deadline
- Comparative statics governed by the option value of learning
- Large-population limit: one sufficient statistic; a sequential significance test on informational time, run by a single error tolerance
- Polkadot, out of sample: half the decision time at a 2% error tolerance

Thank you! bhargav@morpho.org {jonas, kremena}@parity.io

Takeaways

- Optimal aggregation of asynchronous votes under a finite pool and a hard deadline: a dynamic majority-quorum rule with a lens-shaped continuation region, collapsing toward the deadline
- Comparative statics governed by the option value of learning
- Large-population limit: one sufficient statistic; a sequential significance test on informational time, run by a single error tolerance
- Polkadot, out of sample: half the decision time at a 2% error tolerance

Thank you! bhargav@morpho.org {jonas, kremena}@parity.io

Appendix

Results at a glance

1. Optimal policy = dynamic majority-quorum rule: unique threshold $d_t^*(s_t)$, collapsing at the deadline (Theorem 1)
2. Monotone comparative statics in stakes, arrival intensity, and time, driven by the option value of sampling (Proposition 1)
3. Large-population, high-stakes limit: closed-form rule on a single statistic; equivalently a deadline-indexed sequential significance test (Theorem 2, Corollary 1)
4. Polkadot out of sample: at $\alpha = 2\%$, decision times halved, $< 2\%$ divergence from deadline outcomes

Proof of Theorem 1: architecture

Induction statement: for each s_t, the net stopping gain $G_t(\mu, s) = U(\mu, s) - \mathcal{C}_t(\mu, s)$ is non-decreasing in μ on $[\frac{1}{2}, 1]$.

- Martingale lemma: $\mathbb{E}[\mu(s^+ + a^+, s^- + a^-) \mid s^+, s^-] = \mu(s^+, s^-)$ for any random batch

  (iterated expectations on $\mu = \mathbb{E}[1\{M^+ > N/2\} \mid \cdot])$
- FOSD lemma: for $\mu_{1,t} > \mu_{2,t} \geq \frac{1}{2}$, $P(\mu_{t+1} \leq x \mid \mu_{1,t}, s_t) \leq P(\mu_{t+1} \leq x \mid \mu_{2,t}, s_t)$ for all x
  - likelihood ratio of the predictive for consecutive a^+: $\frac{k+1+s^+}{k+1} \cdot \frac{a-k}{a-k+s^-}$ increasing in s^+, decreasing in s^- at fixed s $\Rightarrow$ MLRP $\Rightarrow$ FOSD, conditional on a
  - averaging over a preserves FOSD because the censored-Poisson batch-size law is independent of μ_t
- Base case $(t = T-1)$: $G_{T-1} = \kappa - R\,\mathbb{E}[\max\{0, 1 - 2\mu_T\} \mid \mu_{T-1}, s_{T-1}]$, non-decreasing by FOSD; inductive step decomposes G_t into terms inheriting monotonicity

back

Counterexample: cumulative dominance is not enough (ii)

$N = 5$, $T = 2$, $R = 10$, $\kappa = \frac{1}{5}$; $\boldsymbol{\lambda} = (3, 2)$ dominates $\boldsymbol{\lambda}' = (1, 3)$ cumulatively; state $(s_1^+, s_1^-) = (2, 0)$.

- Continuation at $t = 1$ hinges on the *future* rate only: $\lambda(2) = 2 < 3 = \lambda'(2)$
- Evaluating the thresholds:

$$\begin{aligned} \lambda(2) = 2: \ & p_2(3) = 1 - 5e^{-2} \approx 0.323 \leq \tfrac{2}{5} \ \Rightarrow \ \text{stop at } t = 1 \\ \lambda'(2) = 3: \ & p_3(3) = 1 - \tfrac{17}{2}e^{-3} \approx 0.577 > \tfrac{2}{5} \ \Rightarrow \ \text{continue} \end{aligned}$$

  so $d_1^*(2; \boldsymbol{\lambda}) < d_1^*(2; \boldsymbol{\lambda}')$ despite cumulative dominance of $\boldsymbol{\lambda}$
- Volume vs. front-loading: more total arrivals, but fewer *remaining*, lowers the option value now $\Rightarrow$ part (ii) requires dominance over every future sub-interval

back

Counterexample: back-loading breaks monotone collapse (iii)

$N = 7,\ T = 3,\ \boldsymbol{\lambda} = (1, 1, 5),\ R = 10,\ \kappa = \frac{1}{2}$; state $(s^+, s^-) = (3, 1)$.

$$U(3,1) = R\,\mu(3,1) = 10 \cdot \tfrac{13}{14} = \tfrac{65}{7} \approx 9.286$$

$$\begin{aligned} \mathcal{C}_1(3,1) &\approx 9.197 < U(3,1) \ \Rightarrow\ \text{stop at } t = 1 \\ \mathcal{C}_2(3,1) &\approx 9.411 > U(3,1) \ \Rightarrow\ \text{continue at } t = 2 \end{aligned}$$

- Hence $d_1^*(4) = 2 < d_2^*(4) = 4$: the threshold *rises* in anticipation of the informative last period
- Front-loading in part (iii) rules exactly this out: each period is then the most informative remaining

back

Theorem 2, sketch I: the evidence statistic and the information clock

Step 1: Laplace expansion of the posterior. Posterior of M^+ is log-concave $\mathrm{BetaBinomial}(N-s,\, s^+{+}1,\, s^-{+}1)$; define the evidence statistic as the log-posterior gap between mode $\hat{m}$ and tie state $m_0 = N/2$: $\ell := l(\hat{m}) - l(m_0)$.

- Taylor at the mode ($l'(\hat{m}) = 0$): with $\sigma^2 = \frac{N(N-s)}{4s}(1+o(1))$ and $m_0 - \hat{m} = -\frac{Nd}{2s} + O(1)$, $2\ell = \frac{d^2}{s} \cdot \frac{N}{N-s}(1+o(1)) =: Z^2$
- Cubic-to-quadratic remainder ratio $\asymp |\gamma| Z \lesssim \log \rho / N \to 0$ under $\log \rho \le \varepsilon N^{1/3}$; absolute exponent error $\lesssim (\log \rho)^2/N \to 0$
- Summing the Gaussian window: $\mu - \frac{1}{2} = (\Phi(Z) - \frac{1}{2})(1+o(1)) + O(N^{-1/2})$, $\quad Z = \mathrm{sgn}(d)\sqrt{2\ell}$; log-concavity bounds the tails by $\rho^{-c^2/3}$

Step 2: the information clock. Local alternatives $M^+ = N/2 + \delta$, $\theta := 2\delta/\sqrt{N}$ ($\theta = O(1)$ is the non-trivial range).

- $M_s := (\tilde{d}_s - 2\delta)/(N-s)$ is an exact martingale given M^+: kills the mean-reverting pull-back of sampling without replacement
- Its quadratic variation accumulates at unit rate in $\tau(s) := \frac{s}{N-s}$, the information clock; rescaled lead $\widehat{X}_\tau := \sqrt{N}\, \tilde{d}_s/(N-s)$ satisfies $\widehat{X}_\tau \mid \theta \sim N(\theta\tau, \tau)$, and $Z_\tau^2 = \widehat{X}_\tau^2/\tau = 2\ell$

back

Theorem 2, sketch II: belief dynamics and bracketing

Step 3: belief dynamics. Flat prior on $\theta \Rightarrow \theta \mid \mathcal{F}_\tau \sim N\big(\widehat{X}_\tau/\tau,\ 1/\tau\big)$: the clock is the posterior precision; correctness probability upon stopping is $\Phi(|Z_\tau|)$.

- Remaining learning budget: $\mathrm{Var}(m_{\tau_T} - m_\tau \mid \mathcal{F}_\tau) = \frac{1}{\tau} - \frac{1}{\tau_T}$, known at τ, vanishing at the deadline
- Lead in standard deviations of remaining learning: $\beta := \frac{|m_\tau|}{\sqrt{1/\tau - 1/\tau_T}}$, so the deadline verdict reverses the leader with probability exactly $\Phi(-\beta)$, and $\beta^2 = \dfrac{(d_t/s_t)^2}{1/s_t - 1/\hat{s}_T}$ — the statistic of $\hat{\phi}$

Step 4: bracketing ϕ^*. Deadline option value $\mathcal{G}_N := \mathbb{E}\big[\max\{0,\ 1 - 2\mu_{\tau_T}\} \mid \mathcal{F}_\tau\big]$.

- Doob optional sampling (bounded belief martingale) + Jensen: *every* stopping time's expected gain over stopping now is $\le \mathcal{G}_N \quad \Rightarrow \quad$ stop if $R\,\mathcal{G}_N < \kappa$; continue if $R\,\mathcal{G}_N > \kappa(T-t)$ — both exact for ϕ^*
- Gaussian evaluation $\mathcal{G}_N = \mathcal{G}(\beta,\ \Delta\tau/\tau)(1 + o(1))$ uniformly on the relevant window (Poisson averaging over arrivals contributes $1 + o(1)$)
- MLRP of the predictive $\Rightarrow \mathcal{G}_N$ non-increasing in $\beta \Rightarrow$ the stopping set is an up-set: ϕ^* is a threshold rule in β, with cutoff bracketed by the roots $\beta_{T-t}(t) \le \beta^*(t) \le \beta_1(t)$ of $\mathcal{G} = c/\rho$, $c \in \{T - t,\ 1\}$

back

Theorem 2, sketch III: evaluating the bracket

Step 5: Laplace evaluation at the reversal edge. Re-centering the reversal integral at $\xi = -\beta$:

$$\mathcal{G}(\beta, \tfrac{\Delta\tau}{\tau}) = \frac{1}{\pi}\sqrt{\frac{\Delta\tau}{\tau}}\,\frac{e^{-\beta^2/2}}{\beta^2}\,(1 + o(1))$$

- The prefactor β^{-2} vs. β^{-1} in $\Phi(-\beta)$: the option value is an order β *smaller* than the reversal probability — only marginal reversals pay — the source of the $-2\log\log\rho$ correction
- Inverting $\mathcal{G} = c/\rho$ by iterated logarithms:

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

- Bracket width: $\beta_1^2 - \beta_{T-t}^2 = 2\log(T-t) + o(1) = O(1)$ (T fixed); both ends equal $L_\rho + O(1)$ with $L_\rho := 2\log\rho - 2\log\log\rho$
- The $O(1)$ band K_ε satisfies $K_\varepsilon \le L_\rho\,\frac{\log\log\rho}{\log\rho}$ for large ρ: ϕ^* and $\hat{\phi}$ agree wherever $\left|\beta^2/L_\rho - 1\right| > \log\log\rho/\log\rho$ — the band of the theorem
- Posted form: $N/(N-s_t)$ cancels inside β^2; carrying it through and evaluating the threshold on the mean path $s_t \to \Lambda(t)$ yields $\tilde{\phi}$

back

Data construction and the conviction wedge

- From 1,824 raw referenda: keep substantive dispositions on six focus tracks; drop proposer-instructed rejections (faulty metadata) and degenerate tally paths (too few snapshots, single jumps, large non-monotonicities) $\rightarrow$ 952
- Exclude `Root` track from the stopping evaluation (single-slot queue distorts timing) $\rightarrow$ 933 (467 train / 466 test)
- Delegated votes register in bulk when the delegate votes; arrivals measured by directly voting accounts
- Conviction-weighted ASM outcome vs. final account majority: disagreement in $\approx$ 17% of referenda; hence errors are benchmarked against the *deadline account majority*, which any account-based rule could at best replicate

back

Caveat: would behavior change under $\tilde{\phi}$?

- The counterfactual replays historical paths; adopting $\tilde{\phi}$ changes timing incentives (Lucas critique)
- The rule's inference needs only that arrival timing is uninformative about preferences: a *symmetric* shift (everyone moves earlier) leaves the tally representative; recalibrate $\Lambda(\cdot)$ and the rule survives
- Structural reason to expect symmetry: $\tilde{\phi}$ conditions on $|d_t| \Rightarrow$ early acceptance *and* early rejection, neutral between alternatives; incentives attach to being ahead or behind, not to the alternative favored
- The incumbent ASM allows early *approval only* $\Rightarrow$ adopting $\tilde{\phi}$ removes a built-in asymmetry rather than introducing one
- Still a hypothesis, not a theorem: validation requires deliberately generated variation (randomized stopping rules, staged low-stakes deployment with the deadline as backstop)

Caveat: would behavior change under $\tilde{\phi}$?

- The counterfactual replays historical paths; adopting $\tilde{\phi}$ changes timing incentives (Lucas critique)
- The rule's inference needs only that arrival timing is uninformative about preferences: a *symmetric* shift (everyone moves earlier) leaves the tally representative; recalibrate $\Lambda(\cdot)$ and the rule survives
- Structural reason to expect symmetry: $\tilde{\phi}$ conditions on $|d_t| \Rightarrow$ early acceptance *and* early rejection, neutral between alternatives; incentives attach to being ahead or behind, not to the alternative favored
- The incumbent ASM allows early *approval only* $\Rightarrow$ adopting $\tilde{\phi}$ removes a built-in asymmetry rather than introducing one
- Still a hypothesis, not a theorem: validation requires deliberately generated variation (randomized stopping rules, staged low-stakes deployment with the deadline as backstop)

Caveat: would behavior change under $\tilde{\phi}$?

- The counterfactual replays historical paths; adopting $\tilde{\phi}$ changes timing incentives (Lucas critique)
- The rule's inference needs only that arrival timing is uninformative about preferences: a *symmetric* shift (everyone moves earlier) leaves the tally representative; recalibrate $\Lambda(\cdot)$ and the rule survives
- Structural reason to expect symmetry: $\tilde{\phi}$ conditions on $|d_t|$ $\Rightarrow$ early acceptance *and* early rejection, neutral between alternatives; incentives attach to being ahead or behind, not to the alternative favored
- The incumbent ASM allows early *approval only* $\Rightarrow$ adopting $\tilde{\phi}$ removes a built-in asymmetry rather than introducing one
- Still a hypothesis, not a theorem: validation requires deliberately generated variation (randomized stopping rules, staged low-stakes deployment with the deadline as backstop)

Caveat: would behavior change under $\tilde{\phi}$?

- The counterfactual replays historical paths; adopting $\tilde{\phi}$ changes timing incentives (Lucas critique)
- The rule's inference needs only that arrival timing is uninformative about preferences: a *symmetric* shift (everyone moves earlier) leaves the tally representative; recalibrate $\Lambda(\cdot)$ and the rule survives
- Structural reason to expect symmetry: $\tilde{\phi}$ conditions on $|d_t| \Rightarrow$ early acceptance *and* early rejection, neutral between alternatives; incentives attach to being ahead or behind, not to the alternative favored
- The incumbent ASM allows early *approval only* $\Rightarrow$ adopting $\tilde{\phi}$ removes a built-in asymmetry rather than introducing one
- Still a hypothesis, not a theorem: validation requires deliberately generated variation (randomized stopping rules, staged low-stakes deployment with the deadline as backstop)

Caveat: would behavior change under $\tilde{\phi}$?

- The counterfactual replays historical paths; adopting $\tilde{\phi}$ changes timing incentives (Lucas critique)
- The rule's inference needs only that arrival timing is uninformative about preferences: a *symmetric* shift (everyone moves earlier) leaves the tally representative; recalibrate $\Lambda(\cdot)$ and the rule survives
- Structural reason to expect symmetry: $\tilde{\phi}$ conditions on $|d_t|$ $\Rightarrow$ early acceptance *and* early rejection, neutral between alternatives; incentives attach to being ahead or behind, not to the alternative favored
- The incumbent ASM allows early *approval only* $\Rightarrow$ adopting $\tilde{\phi}$ removes a built-in asymmetry rather than introducing one
- Still a hypothesis, not a theorem: validation requires deliberately generated variation (randomized stopping rules, staged low-stakes deployment with the deadline as backstop)

References I

[1] Wald, A., Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*.

[2] Chernoff, H. (1961). Sequential tests for the mean of a normal distribution. *Proc. Fourth Berkeley Symposium*.

[3] Lai, T.L. (1979). Sequential tests for hypergeometric distributions.

[4] Gershkov, A., Szentes, B. (2009). Optimal voting schemes with costly information acquisition. *JET*.

[5] Bognar, K., Börgers, T., Meyer-ter-Vehn, M. (2015). An optimal voting procedure when voting is costly. *JET*.

[6] Gershkov, A., Moldovanu, B., Shi, X. (2017). Optimal voting rules. *REStud*.

[7] Drugowitsch, J., et al. (2012). The cost of accumulating evidence in perceptual decision making. *J. Neuroscience*.

[8] Fudenberg, D., Strack, P., Strzalecki, T. (2018). Speed, accuracy, and the optimal timing of choices. *AER*.

References II

[9] Stark, P.B. (2008). Conservative statistical post-election audits. *Annals of Applied Statistics*.

[10] Myerson, R. (1998). Population uncertainty and Poisson games. *IJGT*.

[11] Wood, G. (2022). Polkadot: vision for a heterogeneous multi-chain framework.