Supplementary Online Appendix to
"Reputation and Disclosure in Dynamic Networks"

I. Sebastian Buhai*
SOFI at Stockholm University
Instituto de Economia at UC Chile
NIPE at Minho University

This version: May 24, 2026. Latest version.

Guide

This supplementary online appendix contains three separately labelled appendices: OA for worked benchmarks, OB for supplementary extensions and primitive formulas, and OC for institutional fit. Proofs of the main text theorems and core technical material appear in the manuscript appendix.

OA Supplementary censoring benchmarks

Example OA.1 (One-intermediary Ornstein-Uhlenbeck Poisson terminal threshold when $h = 0$)**.** *Consider the terminal path* $e \to i \to 0$ *with no continuous public learning* ($h = 0$) *and public review dates arriving as a Poisson process with intensity* $\bar{\lambda}_i$*. Fix a feasible public slice* ξ *on which the entire continuation and reputational term is independent of* d*, so* $\Psi_i(d, \xi) = \bar{\Psi}_i(\xi)$*. This is the constant-intercept restriction used in Proposition OB.5. By Proposition OB.5,*

$$\Delta_i^{\text{term}}(d, \xi) = \alpha_i\Big(\frac{d^2}{\rho + 2\kappa} + \frac{2b_i d}{\rho + \kappa}\Big) + \bar{\Psi}_i(\xi).$$

---

[*] Contact email: `sebastian.buhai@sofi.su.se`. Website: sebastianbuhai.com

Completing the square gives

$$\Delta_i^{\mathrm{term}}(d,\xi) = \alpha_i \mathfrak{a}_i \big(d - p_i^{\mathrm{PV}}\big)^2 - H_i(\xi),$$

where

$$\mathfrak{a}_i = (\rho + 2\kappa)^{-1}, \qquad p_i^{\mathrm{PV}} = -b_i \frac{\rho + 2\kappa}{\rho + \kappa}, \qquad H_i(\xi) := \alpha_i \mathfrak{a}_i \big(p_i^{\mathrm{PV}}\big)^2 - \bar{\Psi}_i(\xi).$$

Hence, immediate disclosure is optimal for every feasible d when $H_i(\xi) \le 0$. If $H_i(\xi) > 0$, the waiting region on that slice is the explicit interval

$$I_i^\star(\xi) = K(\xi) \cap \big(p_i^{\mathrm{PV}} - r_i(\xi),\, p_i^{\mathrm{PV}} + r_i(\xi)\big), \qquad r_i(\xi) := \sqrt{\frac{H_i(\xi)}{\alpha_i \mathfrak{a}_i}}.$$

If the d-independent term has the form $\bar{\Psi}_i(\xi) = \bar{\Psi}_{0,i}(\xi) + \beta_i \delta_i(\xi)$ with $\delta_i(\xi) > 0$ representing the direct review date reputational gain from disclosure, then $H_i(\xi)$ and therefore the waiting radius $r_i(\xi)$ are decreasing in β_i. Holding $\bar{\Psi}_i(\xi)$ fixed, the pivot magnitude $|p_i^{\mathrm{PV}}|$ is linear in $|b_i|$ and the slack term $\alpha_i \mathfrak{a}_i (p_i^{\mathrm{PV}})^2$ is quadratic in $|b_i|$, so the waiting interval shifts away from zero and weakly widens as $|b_i|$ increase. This is the fixed slice benchmark underlying the monotone interval-policy comparison recorded in Online Appendix OB.

Example OA.2 (Positive public learning, age envelope, and waiting intervals)**.** *Keep the terminal path $e \to i \to 0$ and allow $h > 0$ in the public signal process. On the Kalman-Bucy steady-state reference slice,*

$$\bar{v}(h) = \frac{\sigma_Y^2}{h^2}\left(\sqrt{\kappa^2 + \frac{h^2 \sigma_X^2}{\sigma_Y^2}} - \kappa\right), \qquad \lambda_{\mathrm{eff}}(h) := \kappa + h^2 \sigma_Y^{-2} \bar{v}(h) = \sqrt{\kappa^2 + \frac{h^2 \sigma_X^2}{\sigma_Y^2}}.$$

Let d denote the current centered gap. If the file is disclosed at age a, the expected correction at horizon u is

$$\mathbb{E}_t[\Delta_u] = e^{-\lambda_{\mathrm{eff}}(h) u}(d - \Delta m_t^{\mathrm{ret}}),$$

with martingale innovation terms having zero conditional mean. The age-envelope subclass writes the current disclosure residual as $d - \Delta m_t^{\mathrm{ret}} = e^{-\lambda_{\mathrm{eff}}(h) a} r$ and assumes $|r| \le D$ on unresolved slices. Equivalently, the feasible current residual support is contained in $e^{-\lambda_{\mathrm{eff}}(h) a}[-D, D]$. Hence,

$$G_i^{\mathrm{mat}}(d,\xi;\varrho) \le C_i e^{-\lambda_{\mathrm{eff}}(h) a}, \qquad C_i := \alpha_i \left(\frac{D^2}{\rho + 2\lambda_{\mathrm{eff}}(h)} + \frac{2|b_i| D}{\rho + \lambda_{\mathrm{eff}}(h)}\right),$$

where the quadratic term carries the stronger factor $e^{-2\lambda_{\text{eff}}(h)a}$*. If the audit discipline margin is* $\beta_i\underline{\delta}_i$*, then*

$$a_i^\dagger = \left[\frac{\log(C_i/(\beta_i\underline{\delta}_i))}{\lambda_{\text{eff}}(h)}\right]_+$$

verifies Assumption 6. A transparent calibrated subclass takes the review comparison to be quadratic around a pivot p_i*,*

$$\Delta_i^{\text{term}}(d,a) = \eta_i(d-p_i)^2 + \beta_i\underline{\delta}_i - C_i e^{-\lambda_{\text{eff}}(h)a},$$

so waiting is optimal exactly on

$$I_i^*(a) = K \cap [p_i - r_i(a), p_i + r_i(a)], \qquad r_i(a) = \left[\frac{C_i e^{-\lambda_{\text{eff}}(h)a} - \beta_i\underline{\delta}_i}{\eta_i}\right]_+^{1/2}.$$

For $(\kappa, h, \sigma_X, \sigma_Y, C_i, \beta_i\underline{\delta}_i, \eta_i, p_i) = (0.30, 0.80, 1, 1, 0.80, 0.20, 0.60, -0.70)$*,* $\lambda_{\text{eff}}(h) = 0.854$ *and* $a_i^\dagger = 1.624$*. With* $K = [-1.5, 1.5]$*,*

$$I_i^*(0) = [-1.50, 0.30], \quad I_i^*(0.5) = [-1.43, 0.03], \quad I_i^*(1.0) = [-1.18, -0.22], \quad I_i^*(1.7) = \varnothing.$$

Thus positive public learning pins down both the effective age rate and the finite age at which the waiting interval vanishes.

Example OA.3 (Consistent downstream support-hole update on $e \to i \to j \to k \to 0$)**.** *Suppose the file has already been relayed to nonterminal intermediary* j*, whose successor is* k *rather than the decision maker. After affine transport between public events, the current feasible support at* j *is the single interval*

$$K^{(1)} = (a, b).$$

Let j*'s waiting interval at her next public review be* $I_j = (\ell_j, u_j)$ *with* $a < \ell_j < u_j < b$*. If* j *retains, the post-review support is the single translated interval*

$$K^{(2),\text{ret}} = \big((a,b) \cap I_j\big) - \Delta m_j^{\text{ret}} = (\ell_j - \Delta m_j^{\text{ret}},\ u_j - \Delta m_j^{\text{ret}}).$$

If instead j *relays to* k*, then*

$$K^{(2),\text{relay}} = \big((a,b) \cap I_j^c\big) - \Delta m_j^{\text{relay}} = (a - \Delta m_j^{\text{relay}},\ \ell_j - \Delta m_j^{\text{relay}}) \cup (u_j - \Delta m_j^{\text{relay}},\ b - \Delta m_j^{\text{relay}}),$$

a genuine two-component interval union. If j *were the last intermediary before node* 0*, the* *same forward action would be terminal disclosure and the episode would instead collapse to the resolved support* $\{0\}$. *Thus the displayed union is precisely the nonterminal relay update.*

Example OA.4 (Exact censoring versus a Gaussian shortcut: details)**.** *Example 1 gives the main calculation. The variance formula follows from Lemma 7: for* $K = (-.6, -.2] \cup [.2, .6)$ *and* $d \sim \mathcal{N}(0, 1)$,

$$\mathrm{Var}(d \mid d \in K) = \frac{2\int_{.2}^{.6} x^2\phi(x)\,dx}{2\int_{.2}^{.6} \phi(x)\,dx} \approx .169.$$

A moment-matched Gaussian preserves the first two moments and loses the support hole, assigning positive mass to states that the exact public history rules out.

Example OA.5 (Concave proxy for downstream recentering)**.** *This illustration uses a concave downstream continuation value, matching the curvature of expected quadratic loss. Let*

$$U_j(x; b_j) = -\omega_j(x - p_j)^2, \qquad \omega_j > 0,$$

where p_j *is the downstream continuation pivot. Take affine review maps*

$$T_i^{\mathrm{relay}}(d) = d - r, \qquad T_i^{\mathrm{ret}}(d) = d - k, \qquad r > k > 0.$$

The relay event creates the larger public mean correction, so the post-update downstream gap is smaller after relay than after retention. The downstream contribution to intermediary i*'s relay gain is*

$$U_j(T_i^{\mathrm{relay}}(d); b_j) - U_j(T_i^{\mathrm{ret}}(d); b_j).$$

Its derivative is

$$U_j'(d - r; b_j) - U_j'(d - k; b_j) = 2\omega_j(r - k) > 0.$$

Thus the downstream term is order preserving in d*. More generally, the same conclusion follows for any concave* U_j *whenever the two affine maps are common-slope shifts with* $T_i^{\mathrm{relay}}(d) \leq T_i^{\mathrm{ret}}(d)$*. This example illustrates the mechanism used in Lemma OB.1: the physicaly correct map ordering and the concavity of the downstream continuation value work together to preserve the upstream interval comparison.*

OA.1 Poisson race timing illustration

The structural Poisson/ two-route subclass of Corollary OB.16 also yields a compact timing picture. Let us fix $\bar{\lambda}_i = 1.10$, $\rho = 0.10$, $\kappa = 0.40$, public age $a = 0.5$, $\bar{M} = 1$, and first disclosure premium $\varpi_i = 0.05$. Now let us consider a common public slice on which the equilibrium policies in both topologies prescribe disclosure at the first active review opportunity on the relevant branch. On that slice,

$$\mathbb{E}[\tau^{\mathrm{SL}}] = \frac{1}{\bar\lambda_i} + \frac{1}{\bar\lambda_j}, \qquad \mathbb{E}[\tau^{\mathrm{TR}}] = \frac{1}{\bar\lambda_i + \bar\lambda_j}.$$

The corollary's sufficient direct route gain is

$$G(\bar\lambda_j) := \varpi_i + \bar M e^{-\kappa a} \frac{\rho + \kappa}{\rho + \kappa + \bar\lambda_j}.$$

If one compares the corollary's reduced form direct route gain $G(\bar\lambda_j)$ to a benchmark combined cost $c^{\Delta} = 0.23$, the corresponding private reduced form cutoff is

$$\bar\lambda_j < \bar\lambda_j^{\mathrm{dir}} := (\rho + \kappa)\left(\frac{\bar M e^{-\kappa a}}{c^{\Delta} - \varpi_i} - 1\right) \approx 1.77.$$

This is a private reduced form direct route cutoff. It contains the direct branch's first-credit term. The corresponding planner comparison is ΔW^P in the metrics below, which excludes private route credit rents.

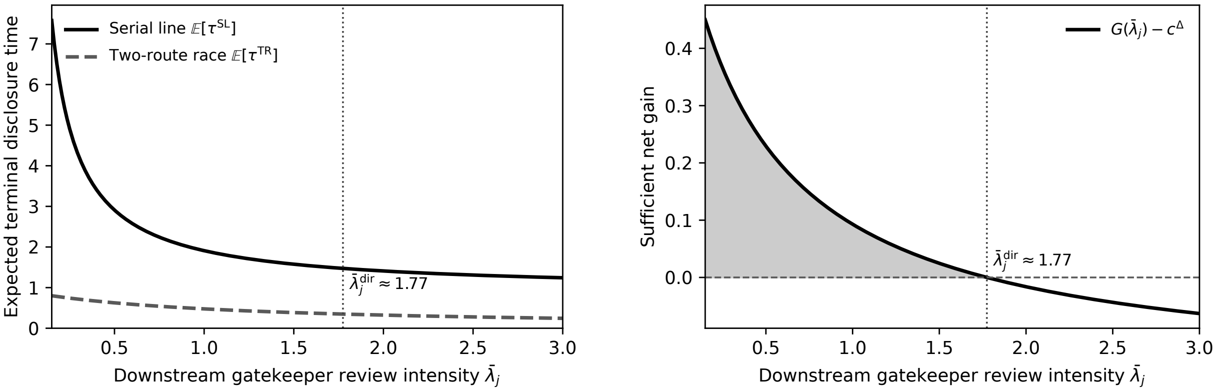

Figure 1: Poisson/ two-route numerical illustration. Left: expected terminal disclosure time in the serial line and the two-route race as the downstream gatekeeper's review intensity $\bar\lambda_j$ varies, holding $\bar\lambda_i = 1.10$ fixed. Right: the reduced form direct route gain $G(\bar\lambda_j) - c^{\Delta}$ with benchmark combined cost $c^{\Delta} = 0.23$; the vertical marker is the direct route cutoff $\bar\lambda_j^{\mathrm{dir}} \approx 1.77$. The code used to generate the two figure panels can be downloaded from the author's website.

OA.2 Formation metrics for the structural Poisson race

The formation proposition is stated at the level of wedges. The structural Poisson/two-route subclass puts those wedges on one primitive footing with an operative incumbent route.

Example OA.6 (Timing and formation metrics from common primitives)**.** *Take the feasible link set*

$$L^{\max} = \{(e,i),(i,0),(e,j),(j,0)\}$$

and let the incumbent network be $G^0 = \{(e,i),(i,0)\}$*. The two links* (e,j) *and* $(j,0)$ *are complementary: neither link, on its own, creates a second route from the expert to the decision maker, but together they create the two-route race* $G^1 = G^0 \cup \{(e,j),(j,0)\}$*. In the structural Poisson subclass,*

$$\Delta V_0^j = \bar{M} e^{-\kappa a} \left(\frac{\bar{\lambda}_i + \bar{\lambda}_j}{\rho + \kappa + \bar{\lambda}_i + \bar{\lambda}_j} - \frac{\bar{\lambda}_i}{\rho + \kappa + \bar{\lambda}_i} \right) > 0,$$

and convenient productive assignment increments are

$$\Delta A_j = \eta_j \frac{\bar{\lambda}_j}{\rho + \bar{\lambda}_i + \bar{\lambda}_j}, \qquad \Delta A_i = \eta_i \left(\frac{\bar{\lambda}_i}{\rho + \bar{\lambda}_i + \bar{\lambda}_j} - \frac{\bar{\lambda}_i}{\rho + \bar{\lambda}_i} \right) < 0.$$

Thus the two-link expansion raises planner value whenever

$$\Delta V_0^j + \Delta A_j + \Delta A_i > c_{ej} + c_{j0} + r_{j0}^{\text{rec}}.$$

Nevertheless G^0 *can be pairwise stable because adding* (e,j) *alone is useless to* e *once* $(j,0)$ *is absent, and adding* $(j,0)$ *alone is useless to the decision maker once* j *has no access to the record. This is the complementary-link coordination mechanism behind insufficient connection.*

For excess connection, consider starting from the completed network G^1*. The planner prefers deletion of branch* j *whenever the reverse strict inequality with* ΔA_i *included holds. But the completed branch can remain pairwise stable if private route credit increments* ΔC_e *and* ΔC_j *satisfy*

$$\alpha_e \Delta V_0^j + \Delta C_e > c_{ej}, \qquad \alpha_j \Delta V_0^j + \phi_j \Delta A_j + \Delta C_j > c_{j0}, \qquad \Delta V_0^j > r_{j0}^{\text{rec}}.$$

These inequalities are jointly feasible because route credit rents enter private formation values while planner-centric welfare counts productive surplus.

OA.3 Static benchmark proof details

Proof of Proposition 3. Let D be the set disclosed with positive probability and suppose $F(\Theta \setminus D) > 0$. Following silence, quadratic loss gives $y(\varnothing) = \mathbb{E}[\theta \mid \theta \in \Theta \setminus D]$; atomlessness makes equilibrium mixing relevant only on the finite indifference set. For a withheld state, the gain from truthful disclosure is

$$h(\theta) = \alpha_e[(y(\varnothing) - \theta - b_e)^2 - b_e^2].$$

Averaging over the pooling set gives

$$\mathbb{E}[h(\theta) \mid \theta \in \Theta \setminus D] = \alpha_e \operatorname{Var}(\theta \mid \theta \in \Theta \setminus D) > 0.$$

Thus some withheld states profitably disclose, contradicting equilibrium. Hence, the pooling set has zero measure. □

Proof of Proposition 4. For part (i), aligned biases let the standard unraveling argument iterate down the path: if any intermediary pools a positive-measure set, a marginal state strictly benefits from separating and being passed on. Backward induction gives a full-disclosure PBE.

For part (ii), suppose a full-disclosure PBE exists on a bias-reversing path and let $y(\varnothing)$ be the decision maker's action after silence. For a sender with bias b, blocking a truthful message is strictly preferred whenever

$$|y(\varnothing) - \theta - b| < |b|.$$

Take the first nonzero positive bias on the path and the first later negative bias; the reverse case is symmetric. Their blocking intervals are

$$(y(\varnothing) - 2b_+, y(\varnothing)) \quad \text{and} \quad (y(\varnothing), y(\varnothing) - 2b_-), \qquad b_+ > 0 > b_-.$$

Since Θ is an interval with full support and $y(\varnothing)$ lies in its convex hull, at least one of these open intervals has positive prior mass inside Θ. On that set the corresponding sender strictly blocks, contradicting full disclosure. Thus every PBE leaves positive mass undisclosed. □

OB Supplementary benchmarks and extensions

The main text studies pure Markov perfect Bayesian equilibria of the institutional model in the case of the single file: public docketing, one operative pending file, public review calendars, public review dates, timestamped hard records, and path or two-route state variables. This section develops bounded-memory variants for institutions that also track repeated arrivals or queues.

Repeated arrivals and raw evidence. Experts may receive raw evidence over time. A sender may therefore hold more than one undisclosed record, and institutional rules may allow queueing, supersession, or archival. A pure strategy for each non-Decision-Maker sender specifies, for every private history, a disclosure rule at each review opportunity and, when the extension allows endogenous public-state measurable calendars, such a calendar rule as well. The DM's strategy specifies a public-history measurable action y_t at each date. A belief system specifies, after every public history, a posterior over latent states and latent types. A perfect Bayesian equilibrium is defined in the standard way.

Bounded-memory benchmark MPBE. A finite recursive class handles bounded queues. A *bounded-memory benchmark MPBE* is a PBE in which:

- there exists a finite-dimensional public state $\widetilde{S}_t$ that summarizes the payoff-relevant public history;
- along the operative path or network under study, each non-decision-maker sender has at most $K_i < \infty$ simultaneously tracked undisclosed records inside the recursive description;
- whenever a raw arrival or relay receipt would create more than K_i tracked records, a fixed measurable projection operator Π_i maps the raw post-arrival inventory into the admissible tracked state before continuation values are evaluated. Such projections are institutional queue-management rules. Most recent record and priority rules are examples of bounded-memory institutional summaries;
- each non-decision-maker sender's continuation value is Markov in $(\widetilde{S}_t, \widetilde{Z}_t^i)$, where $\widetilde{Z}_t^i$ is the bounded-memory private summary induced by (K_i, Π_i).

The bounded-memory benchmark extends the core theorem environment to institutions that track a bounded queue of operative files. It preserves the timing and network forces while replacing file-specific terminal disclosure with resolution of the bounded-memory operative status.

Supersession and archival. Mandatory supersession and automatic archival are queue-management rules. They naturally lead to eventual resolution of operative status. The queue variant preserves the timing logic of the single-file episode.

OB.1 Further extensions

The following extensions preserve the paper's core public-censoring logic while changing the state space.

Multidimensional fundamentals. Let the fundamental instead be a d-dimensional diffusion $(X_t)_{t\geq 0}$ on $\mathbb{R}^d$. Linear-Gaussian state dynamics provide a finite-dimensional Gaussian public filter. Under public review date censoring, the natural exact family is multivariate truncated normal. A scalar ladder survives when a one-dimensional sufficient index summarizes the current holder's private information relative to public beliefs. With several payoff-relevant indices, the continuation region becomes a set in $\mathbb{R}^d$ and the ladder becomes a collection of threshold surfaces.

Noisy but verifiable records. If disclosed records are noisy but verifiable, public uncertainty after disclosure remains positive even at the disclosure time. In the standard linear-Gaussian version of that extension, the one-dimensional scalar coordinate remains the holder's private minus public gap in posterior means. Under quadratic loss the residual posterior variance cancels from the disclosure-versus-retention material comparison, so the instantaneous material gain remains

$$\alpha_i d_t^i (d_t^i + 2b_i).$$

The exact public posterior in this extension is censored by review date nondisclosure. The underlying uncensored reference kernel remains linear-Gaussian. Since the holder perfectly observes the noisy disclosed record, her private posterior can then be written against that reference kernel, and every continuation-relevant private statistic remains an affine function of the same one-dimensional gap. Hence, the noisy-record Gaussian benchmark preserves the exact one-dimensional threshold logic; the coefficients of the Gaussian recursion change.

Multiple audiences. The baseline model has one single decision maker. However, with several downstream audiences, each audience may maintain its own public posterior and its own reputation state for the intermediary. The problem remains option-theoretic and the public state becomes higher dimensional. For any fixed public state, however, the sealed record gap is still scalar, so the review date comparison remains a one-dimensional interval comparison. The policy manifold indexed by audiences and public states becomes high-dimensional, while each realized retention rule is still an interval rule in the private gap.

Collective disclosure institutions. Committees, panels, and disclosure teams generate collective release decisions. The group-disclosure paper of Onuchic and Ramos [2025] shows how voting or deliberation rules shape verifiable release. Embedding those procedures inside a dynamic network would be a natural extension of the present framework.

OB.2 Interval-policy verification details

In this subsection, we record the order and affine sufficient conditions behind the compact benchmark in the main text.

Lemma OB.1 (Concave shift comparison on ordered slices)**.** *Fix connected one-sided slices* D_i *and* D_j *with one-sided orders* $\preceq_i$ *and* $\preceq_j$*. Let* $\varphi_j : D_j \to \mathbb{R}$ *be an order coordinate, so that* $x \preceq_j y$ *if and only if* $\varphi_j(x) \le \varphi_j(y)$*. Let* $T^{\mathrm{relay}}, T^{\mathrm{ret}} : D_i \to D_j$ *be order preserving from* $\preceq_i$ *into* $\preceq_j$ *and suppose the retained proxy weakly dominates the relay image:*

$$T^{\mathrm{relay}}(d) \preceq_j T^{\mathrm{ret}}(d) \qquad \text{for every } d \in D_i.$$

Let $g(d) := \varphi_j(T^{\mathrm{ret}}(d)) - \varphi_j(T^{\mathrm{relay}}(d)) \ge 0$*. Let us now fix a sign* $\sigma_j \in \{1, -1\}$*, where* $\sigma_j = 1$ *is the increasing-continuation case and* $\sigma_j = -1$ *is the decreasing-continuation case. Assume the signed retained-minus-relay gap condition*

$$d_1 \preceq_i d_2 \quad \Longrightarrow \quad \sigma_j\big(g(d_2) - g(d_1)\big) \le 0.$$

Now let $U_j : D_j \to \mathbb{R}$ *be the downstream continuation value. We say that* U_j *has signed concave ordered interval increments with sign* σ_j *on* D_j *if, whenever* $x_1 \preceq_j x_2$*,* $y_1 \preceq_j y_2$*,* $x_\nu \preceq_j y_\nu$ *for* $\nu = 1, 2$*, and the interval lengths* $\ell_\nu = \varphi_j(y_\nu) - \varphi_j(x_\nu)$ *satisfy* $\sigma_j(\ell_2 - \ell_1) \le 0$*, one has*

$$U_j(y_2) - U_j(x_2) \le U_j(y_1) - U_j(x_1).$$

If U_j *has signed concave ordered interval increments with sign* σ_j*, then*

$$d \mapsto U_j\big(T^{\mathrm{relay}}(d)\big) - U_j\big(T^{\mathrm{ret}}(d)\big)$$

is order preserving on D_i *with respect to* $\preceq_i$*. Equivalently, a proof may impose this final composite order preservation directly; the signed concave shift and signed gap hypotheses are a primitive suficient route.*

Proof. Take $d_1 \preceq_i d_2$. Order preservation then yields

$$T^{\mathrm{relay}}(d_1) \preceq_j T^{\mathrm{relay}}(d_2), \qquad T^{\mathrm{ret}}(d_1) \preceq_j T^{\mathrm{ret}}(d_2),$$

and pointwise ordering gives $T^{\mathrm{relay}}(d_\nu) \preceq_j T^{\mathrm{ret}}(d_\nu)$ for $\nu = 1, 2$. The signed gap assumption gives $\sigma_j(g(d_2) - g(d_1)) \le 0$, exactly the signed length premise for the two retained-relay intervals below. Apply signed concave ordered interval increments to

$$x_\nu = T^{\mathrm{relay}}(d_\nu), \qquad y_\nu = T^{\mathrm{ret}}(d_\nu), \qquad \nu = 1, 2.$$

This yields

$$U_j\bigl(T^{\mathrm{ret}}(d_2)\bigr) - U_j\bigl(T^{\mathrm{relay}}(d_2)\bigr) \le U_j\bigl(T^{\mathrm{ret}}(d_1)\bigr) - U_j\bigl(T^{\mathrm{relay}}(d_1)\bigr).$$

Multiplying by -1 finally gives

$$U_j\bigl(T^{\mathrm{relay}}(d_2)\bigr) - U_j\bigl(T^{\mathrm{ret}}(d_2)\bigr) \ge U_j\bigl(T^{\mathrm{relay}}(d_1)\bigr) - U_j\bigl(T^{\mathrm{ret}}(d_1)\bigr),$$

so the displayed relay-minus-retention difference is order preserving in $\preceq_i$. □

Lemma OB.2 (Pivot stitching on interval union supports)**.** *Let $K = \bigcup_{n=1}^{N} K_n \subseteq \mathbb{R}$ be a finite ordered union of pairwise disjoint connected components, with $K_1 < \cdots < K_N$, and let $p \in \mathbb{R}$. Suppose $\Delta : K \to \mathbb{R}$ is continuous, weakly decreasing on every connected component of $K \cap (-\infty, p]$, and weakly increasing on every connected component of $K \cap [p, \infty)$. Assume in addition the following cross component no-skipping condition: whenever $x < z$ are points in K with $\Delta(x) < 0$ and $\Delta(z) < 0$, one has $\Delta(y) < 0$ for every $y \in K \cap [x, z]$. Then there exists an interval $I \subseteq \mathbb{R}$, possibly empty, such that*

$$\{d \in K : \Delta(d) < 0\} = K \cap I.$$

Proof. Let

$$S := \{d \in K : \Delta(d) < 0\}.$$

If $S = \varnothing$, take $I = \varnothing$ and stop. For each connected component K_n, continuity and one-sided monotonicity imply that $S \cap K_n$ is an interval, possibly empty. The cross component no-skipping condition then makes S order convex relative to K: whenever $x, z \in S$ and $y \in K$ satisfies $x \le y \le z$, one has $y \in S$.

Let $a := \inf S$ and $b := \sup S$ in the extended real line. If $y \in K \cap (a, b)$, then there exist $x, z \in S$ with $x < y < z$, so order convexity gives $y \in S$. Conversely, $S \subseteq K \cap [a, b]$ by definition of a and b. Because S is relatively open in K, the endpoint issue is then whether a and b themselves belong to S. Choose I to be (a, b), $[a, b)$, $(a, b]$, or $[a, b]$ according to those endpoint inclusions, with the corresponding half-line convention if $a = -\infty$ or $b = \infty$. Then $S = K \cap I$, as claimed. □

Definition OB.3 (Affine invariant buffered benchmark family)**.** *Fix a common-sign finite* *path $P = (i_0 = e, i_1, \ldots, i_k, 0)$, a common pivot p, a decreasing buffer sequence $\eta_0 > \eta_1 > \cdots > \eta_k > 0$, a compact unresolved reputation box $\mathcal{R}^U$, and an age bound $\bar{A} < \infty$. Let $\widetilde{\mathcal{V}}_{\bar{A}}$ and $\mathcal{Q}_{\bar{A}}$ denote the compact ranges of the reference variance $\tilde{v}$ and smoother variance q generated by the exact single-file Ornstein-Uhlenbeck public signal filter on the age band $[0, \bar{A}]$. For each stage $m \in \{0, \ldots, k\}$, let*

$$D_m = \begin{cases} [p + 2\eta_m, \infty), & \text{if all path biases are weakly positive,} \\ (-\infty, p - 2\eta_m], & \text{if all path biases are weakly negative.} \end{cases}$$

On these one-sided slices, $\preceq_{i_m}$ denotes the order away from the common pivot: in the positive case it is the usual order, and in the negative case it is the reverse usual order. Write $\overline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, \infty\}$. The explicit buffered affine family

$$\mathfrak{S}^{\mathrm{BA}}_{k+1}(\bar{A}, \mathcal{R}^U, \eta)$$

consists of public slices $\xi = (\tilde{v}, R, c, a, \bar{d}, q, K)$ such that $c = i_m$ for some m, $a \in [0, \bar{A}]$, $R \in \mathcal{R}^U$, $\tilde{v} \in \widetilde{\mathcal{V}}_{\bar{A}}$, $q \in \mathcal{Q}_{\bar{A}}$, and K is a union of at most $k+1$ ordered closed intervals with endpoints in $\overline{\mathbb{R}}$. Exterior rays are therefore admitted through $\pm\infty$; in particular, $K_\tau = \mathbb{R}$ belongs to the family. If $m < k$, we additionally require that on $K \cap D_m$ the relay gap and retained gap proxy maps into stage $m+1$ have the affine form

$$T^{\mathrm{relay}}_m(d) = a_m + \gamma_m d, \qquad T^{\mathrm{ret}}_m(d) = \tilde{a}_m + \tilde{\gamma}_m d, \qquad 0 < \gamma_m, \tilde{\gamma}_m \leq 1,$$

that both images lie in D_{m+1} and that the retained proxy weakly dominates the relay image in the downstream order,

$$T^{\mathrm{relay}}_m(d) \preceq_{i_{m+1}} T^{\mathrm{ret}}_m(d)$$

pointwise on $K \cap D_m$. In the positive-bias coordinate, this is the economic case in which relay creates the larger public mean correction and hence the smaller post-update gap.

Lemma OB.4 (Primitive invariant-domain check)**.** *Fix the buffered affine family in Definition OB.3. Now suppose that, at every active interface and on every feasible component of $K \cap D_m$, the affine relay and retained-proxy maps satisfy the endpoint-coherence inequalities*

$$T^{\mathrm{relay}}_m(K \cap D_m) \cup T^{\mathrm{ret}}_m(K \cap D_m) \subseteq D_{m+1},$$

with the displayed inclusion holding after the public mean translations generated by the review update. Suppose also that recertification maps have weights $\chi_i \in (0, 1)$, so audit posteriors *remain in a compact interior interval, and that the Ornstein-Uhlenbeck signal coordinates remain in their compact age band ranges. Then between-review transport, retention, and relay keep unresolved public slices inside the buffered affine family whenever the review policy intervals are admissible.*

Proof of Lemma OB.4. Between public review events, the support coordinate is transported by an affine map and interval unions with endpoints in $\overline{\mathbb{R}}$ are preserved. At a review date, retention intersects the current support with an admissible policy interval and relay intersects it with the complement of that interval; both operations preserve finite interval unions, and the subsequent public mean correction is a translation. The endpoint-coherence inclusions in the statement put the relay and retained proxy images inside the next buffered one-sided slice, so the interface coordinates remain in the declared stage domains. The Ornstein-Uhlenbeck age band keeps the public signal and smoother variances in the compact ranges used in Definition OB.3. Finally, $\chi_i \in (0,1)$ implies $M_i([0,1]) = [\chi_i \pi_i^0, 1 - \chi_i(1-\pi_i^0)] \subset (0,1)$, so recertification returns audit beliefs to a compact interior interval. Each coordinate of the successor slice therefore satisfies the definition of the buffered affine family. ☐

Proposition OB.5 (Terminal and local comparison objects)**.** *Fix a feasible public slice* $\xi = (\tilde{v}, R, c = i, a, \bar{d}, q, K)$ *and a connected subinterval* $D \subseteq K$ *for current custodian* i.

(i) In the one-intermediary Ornstein-Uhlenbeck path $e \to i \to 0$, *let* $\Delta m_i^{\mathrm{ret}}(\xi)$ *be the public mean jump that would follow retention on the slice, and set* $\widetilde{d}_i(d,\xi) := d - \Delta m_i^{\mathrm{ret}}(\xi)$. *The immediate-terminal-disclosure gain has the decomposition*

$$\Delta_i^{\mathrm{term}}(d,\xi) = G_i^{\mathrm{PV}}(d,\xi) + \Psi_i(d,\xi), \qquad G_i^{\mathrm{PV}}(d,\xi) = \alpha_i\big(\mathfrak{a}_i(\xi)\widetilde{d}_i(d,\xi)^2 + 2b_i\mathfrak{b}_i(\xi)\widetilde{d}_i(d,\xi)\big),$$

where $\mathfrak{a}_i(\xi), \mathfrak{b}_i(\xi) > 0$ *are exact in the no-future-public-learning benchmark and in affine public-learning subclasses; with* $h > 0$, $\Psi_i(d,\xi)$ *also absorbs the nonlinear truncated-filter residual generated by future public signals. Hence the terminal waiting set on* D *is interval valued whenever the full gain* $\Delta_i^{\mathrm{term}}(\cdot,\xi)$ *is continuous and one-sided monotone away from some pivot* $p_i^{\star}(\xi)$ *on* D. *If the slice is written in the post-retention centered coordinate,* $\Delta m_i^{\mathrm{ret}}(\xi) = 0$ *and the displayed formula reduces to the simpler quadratic in* d.

(ii) In the two intermediary path $e \to i \to j \to 0$, *write the relay versus retention gain at* i *as*

$$\Delta_i^{\mathrm{relay}}(d,\xi) = \widehat{\Delta}_i^{\mathrm{relay}}(d,\xi) + \Omega_i(d,\xi),$$

where $\widehat{\Delta}_i^{\mathrm{relay}}$ *is the downstream continuation term from handing the file to* j, *and* Ω_i *collects* i*'s own retention side terms. If intermediary* j *uses interval retention on every* *continuation slice, the relay gap map and retained gap proxy send D into a common downstream one-sided slice D_j, are order preserving from $(D, \preceq_i)$ into $(D_j, \preceq_j)$, the retained proxy weakly dominates the relay image in the downstream order, and the downstream continuation comparison satisfies Lemma OB.1 on D_j, then $\widehat{\Delta}_i^{\mathrm{relay}}(\cdot, \xi)$ is order preserving on D in $\preceq_i$. If $\Omega_i(\cdot, \xi)$ is also order preserving on D in the same direction, then so is the full gain $\Delta_i^{\mathrm{relay}}(\cdot, \xi)$.*

Proof of Proposition OB.5. For part (i), fix the current public slice ξ and realized gap d. Immediate terminal disclosure reveals the dated file X_τ once and for all. Relative to retention, let $\Delta m_i^{\mathrm{ret}}(\xi)$ be the public mean jump that would follow retention on the slice and set $\widetilde{d}_i(d, \xi) := d - \Delta m_i^{\mathrm{ret}}(\xi)$. Let $\delta_{i,\xi}(u; d)$ denote the induced correction in the decision maker's conditional mean of X_{t+u} at horizon increment $u \geq 0$. When $h = 0$, public learning after the review date disappears and the correction is exactly

$$\delta_{i,\xi}(u; d) = e^{-\kappa u}\widetilde{d}_i(d, \xi).$$

With $h > 0$, retention leaves a truncated Gaussian prior for future public learning, so the later posterior correction need not be affine in $\widetilde{d}_i(d, \xi)$. The decomposition below is therefore used as an affine-component representation: in subclasses where

$$\delta_{i,\xi}(u; d) = a_i(u, \xi)\widetilde{d}_i(d, \xi) + \varepsilon_i(u, \xi)$$

with d-independent second moment of ε_i, the quadratic part is exact; otherwise all remaining d-dependent public-learning terms are assigned to $\Psi_i(d, \xi)$. Under quadratic loss and discounting this gives

$$\Delta_i^{\mathrm{term}}(d, \xi) = \alpha_i \Big(\mathfrak{a}_i(\xi)\widetilde{d}_i(d, \xi)^2 + 2b_i\mathfrak{b}_i(\xi)\widetilde{d}_i(d, \xi)\Big) + \Psi_i(d, \xi),$$

where, in the affine subclass,

$$\mathfrak{a}_i(\xi) := \int_0^\infty e^{-\rho u} a_i(u, \xi)^2 du, \qquad \mathfrak{b}_i(\xi) := \int_0^\infty e^{-\rho u} a_i(u, \xi) du.$$

Thus the verification argument uses the one-sided monotonicity of the full gain, not an exact affine formula for every $h > 0$ filter. The material turning point in the centered correction coordinate is

$$\widetilde{p}_i^{\mathrm{PV}}(\xi) = -b_i \frac{\mathfrak{b}_i(\xi)}{\mathfrak{a}_i(\xi)}.$$

In the original gap coordinate the pivot is shifted by the retained public mean jump, $p_i^{\mathrm{PV}}(\xi) =$

$\Delta m_i^{\mathrm{ret}}(\xi) + \widetilde{p}_i^{\mathrm{PV}}(\xi)$. In the boundary case $h = 0$, the loadings reduce to $a_i(u, \xi) = e^{-\kappa u}$, so $\mathfrak{a}_i(\xi) = (\rho + 2\kappa)^{-1}$ and $\mathfrak{b}_i(\xi) = (\rho + \kappa)^{-1}$. If the full gain is continuous and one-sided monotone away from a slice-specific pivot on the connected slice under study, then its strict sublevel set is interval valued on that slice. This is exactly the terminal benchmark stated in the proposition.

For part (ii), fix the two intermediary path $e \to i \to j \to 0$ and hold fixed a downstream interval policy for j. By Proposition 5, relay by i maps the current gap d into a downstream continuation slice through a relay gap transition $T^{\mathrm{relay}}(d)$. Likewise, retention by i induces a retained gap proxy $T^{\mathrm{ret}}(d)$ summarizing the downstream continuation component after the public censoring update. Under the maintained hypotheses, both maps land in a common downstream one-sided slice, are order preserving, and satisfy $T^{\mathrm{relay}}(d) \preceq_j T^{\mathrm{ret}}(d)$ there. Lemma OB.1 therefore implies that the downstream continuation component

$$\widehat{\Delta}_i^{\mathrm{relay}}(d, \xi)$$

is order preserving in $\preceq_i$ on the slice. Any additional current custody term $\Omega_i(d, \xi)$ must be handled separately; when it is order preserving on the same slice and in the same direction, the full gain $\Delta_i^{\mathrm{relay}}(d, \xi) = \widehat{\Delta}_i^{\mathrm{relay}}(d, \xi) + \Omega_i(d, \xi)$ is order preserving as well. Because a relay is informative, the decision maker's action can jump at the review date through the public mean update generated by the relay event. There is nevertheless no additional terminal material term to add: the episode continues after relay, so the material consequences of that jump are already embedded in the downstream continuation value difference $\widehat{\Delta}_i^{\mathrm{relay}}$. □

Theorem OB.6 (Verified interval-policy MPBE on the buffered affine family)**.** *Fix the single-file review benchmark of Assumption 2 and a common-sign finite path whose feasible slices belong to the explicit buffered affine family* $\mathfrak{S}_{k+1}^{\mathrm{BA}}(\bar{A}, \mathcal{R}^U, \eta)$ *from Definition OB.3. Restrict attention to admissible interval-policy assessments in Definition 1. If the terminal stage satisfies the one-sided monotonicity condition in Proposition OB.5(i) and its componentwise strict sublevel sets satisfy the no-skipping condition in Lemma OB.2, and if at each upstream stage the relevant downstream continuation comparison satisfies Lemma OB.1, the additional upstream own term is order preserving on the current buffered slice in the same direction, and the componentwise strict sublevel sets satisfy the same no-skipping condition, then the constructed assessment is a pure interval-policy MPBE on* $\mathfrak{S}_{k+1}^{\mathrm{BA}}(\bar{A}, \mathcal{R}^U, \eta)$*. For every sender* i_m *and every feasible public slice* ξ *with current custodian* $c = i_m$*, the equilibrium review date retention region is* $K(\xi) \cap I_{i_m}^*(\xi)$ *for a slice-specific interval* $I_{i_m}^*(\xi)$*.*

Corollary OB.7 (Primitive affine verification)**.** *In the common-sign affine benchmark, suppose each downstream one-sided slice* D_{m+1} *admits an order coordinate* φ_{m+1}*, the maps*

T_m^{relay} and T_m^{ret} have the affine form in Definition OB.3, their images stay inside D_{m+1}, and the retained proxy weakly dominates the relay image in the downstream order. Suppose the retained-minus-relay coordinate gap

$$\varphi_{m+1}\big(T_m^{\text{ret}}(d)\big) - \varphi_{m+1}\big(T_m^{\text{relay}}(d)\big)$$

satisfies the signed gap condition in Lemma OB.1. If the downstream continuation value U_{m+1} satisfies that lemma's signed concave shift comparison on D_{m+1}, then the continuation comparison hypothesis in Theorem OB.6 holds at stage m. Hence primitive common-sign affine review environments populate the buffered affine class.

Proof of Theorem 1. The Ornstein-Uhlenbeck and linear Gaussian assumptions give the smoother representation and the Gaussian likelihood kernel in Lemma 3. Public review outcomes multiply that likelihood by the event determined by the interval policy and translate the public mean, so Proposition 5 and Corollary 1 give the exact finite interval union state.

The common-sign affine discipline verifies the continuation comparison hypothesis of Theorem OB.6 through Corollary OB.7. The assumptions on no skipping and endpoint coherence make the strict waiting sets order convex on the finite interval-union support. Theorem OB.6 and Proposition 2 then verify the constructed assessment as a pure MPBE on the compact benchmark class.

Public audit recertification and Bayes-plausible information dispersion verify interior review discipline by Corollary 3; call the resulting discipline constant $\underline{\delta}_i$. The exponential age envelope sets $\overline{G}_i(a) = C_i e^{-\lambda_i^{\text{age}} a}$ in Assumption 6, so $\overline{G}_i(a) \le \beta_i \underline{\delta}_i$ for all $a \ge [\log(C_i/(\beta_i \underline{\delta}_i))/\lambda_i^{\text{age}}]_+$. The old-record resolution conclusion is therefore Theorem 2 applied to this primitive subclass. $\square$

OB.3 Supplementary invariant-domain benchmark

Example OB.8 (A buffered-pivot benchmark class with $k \ge 3$)**.** *Fix any path length $k \ge 3$ and a common-sign path $P = (e, i_1, \ldots, i_k, 0)$. Choose a common pivot p and a buffer $\eta > 0$. Suppose there is a decreasing sequence of buffers $(\eta_m)_{m=0}^k$ with $\eta_m > 0$. Define $D_m = [p + 2\eta_m, \infty)$ and take the active-side feasible slices to be $K \cap D_m$. Suppose supports are admitted as unions of at most $k+1$ closed intervals in $\overline{\mathbb{R}}$ and the relay gap and retained gap proxy maps take the affine forms*

$$T_m^{\text{relay}}(d) = a_m + \gamma_m d, \qquad T_m^{\text{ret}}(d) = \tilde{a}_m + \tilde{\gamma}_m d, \qquad 0 < \gamma_m, \tilde{\gamma}_m \le 1,$$

with coefficients chosen so that $T_m^{\mathrm{relay}}(D_m) \cup T_m^{\mathrm{ret}}(D_m) \subseteq [p + 2\eta_{m+1}, \infty) = D_{m+1}$ *and* $T_m^{\mathrm{relay}}(d) \leq T_m^{\mathrm{ret}}(d)$ *on* D_m*. The equality* $D_{m+1} = [p + 2\eta_{m+1}, \infty)$ *is deliberate: it keeps the target slice inside the same two-buffer convention that defines the stage-*$m+1$ *domain. Assume the downstream continuation value and affine retained-minus-relay gap satisfy the signed concave-shift condition of Lemma OB.1, every upstream own term is order preserving on* D_m *in the same direction, and the componentwise strict sublevel sets are coherent in the sense of Lemma OB.2. Then the common-sign affine comparison conditions hold stage by stage, and the maintained componentwise no-skipping condition supplies the interval stitching step. This gives an explicit nonempty benchmark class for arbitrary* $k \geq 3$.

OB.4 Supplementary audit and formation primitives

Example OB.9 (Two-signal information dispersion audit experiment)**.** *Fix a compact interior posterior interval* $[\underline{\pi}_i, \overline{\pi}_i] \subset (0,1)$ *and choose* $\varepsilon_P > \varepsilon_N > 0$ *small enough that* $\pi \pm \varepsilon_P$ *remains in* $(0,1)$ *on the interval. Let the skipped review audit experiment produce posterior values* $\pi + \varepsilon_N$ *and* $\pi - \varepsilon_N$ *with probability one half each, and let the prompt audit experiment produce* $\pi + \varepsilon_P$ *and* $\pi - \varepsilon_P$ *with probability one half each. Both experiments are Bayes plausible. If* $G_i(\pi) = \pi^2$*, then*

$$\mathbb{E}[G_i(\Pi_i^P) - G_i(\Pi_i^N)] = \varepsilon_P^2 - \varepsilon_N^2 > 0.$$

For a general payoff with $G_i'' \geq \gamma_i$*, the same calculation gives the lower bound* $\frac{1}{2}\gamma_i(\varepsilon_P^2 - \varepsilon_N^2)$*. The example illustrates the equal-mean information route: prompt resolution makes the audit record more informative, and strict convexity of reputation value prices that information.*

Example OB.10 (Boundary drift and audit refresh)**.** *Continue Example OB.9, but this time allow the prompt experiment's dispersion radius to shrink near the posterior boundary. If* $\varepsilon_P(\pi) - \varepsilon_N(\pi) \to 0$ *as* $\pi \downarrow 0$ *or as* $\pi \uparrow 1$*, the one-review Jensen gain can vanish at the boundary even though prompt audit remains more informative at every interior posterior. Audit refresh, equilibrium disclosure before long strings of skipped reviews, or an auxiliary public audit signal keeps the discipline margin uniformly positive on the region used by the old-record resolution theorem.*

Remark OB.11 (A route to verifying the unresolved interior region). The information-dispersion benchmark shows why a uniform review date margin is naturally an interior statement. Corollary 3 gives one primitive stabilizing force: public recertification mixes beliefs back toward an interior benchmark before audit outcomes are evaluated. Other mechanisms, such as equilibrium disclosure before too many consecutive skipped reviews accumulate or an auxiliary public type signal, can play the same role. The main text uses invariant interiority as a maintained region or verifies it through audit refresh.

Example OB.12 (Complementary-link formation wedge on the same primitive class)**.** *Take the feasible link set*

$$L^{\max} = \{(e,i),(i,0),(e,j),(j,0)\}$$

and the structural Poisson/two-route subclass above. The incumbent direct route $e \to i \to 0$ *is always present. Let*

$$G^0 = \{(e,i),(i,0)\}, \qquad G^1 = G^0 \cup \{(e,j),(j,0)\},$$

so branch j *is activated when the two complementary links form together. Let us now write*

$$\begin{aligned} \Delta V_0^j &= V_0(G^1) - V_0(G^0), & \Delta A_j &= A_j(G^1) - A_j(G^0), \\ \Delta A_i &= A_i(G^1) - A_i(G^0) < 0, & \Delta C_e, \Delta C_j &\text{ are private route-credit increments.} \end{aligned}$$

These are the decision-maker, productive-assignment, incumbent-cannibalization, and private route-credit increments created by completing branch j*.*

For insufficient connection*, the planner wants the two-link expansion whenever*

$$\Delta V_0^j + \Delta A_j + \Delta A_i > c_{ej} + c_{j0} + r_{j0}^{\text{rec}}.$$

Yet, do remark that pairwise stability can keep G^0 *in place: adding* (e,j) *alone leaves* j *without a path to the decision maker and costs* e *the positive amount* c_{ej}*; adding* $(j,0)$ *alone leaves* j *without access to the record and imposes the reception cost* r_{j0}^{rec} *on the decision maker. Hence, at least one endpoint rejects each one-link deviation even though the two links together would raise planner value.*

For excess connection*, start from* G^1 *and choose parameters so that*

$$\Delta V_0^j + \Delta A_j + \Delta A_i < c_{ej} + c_{j0} + r_{j0}^{\text{rec}}, \qquad \Delta V_0^j > r_{j0}^{\text{rec}},$$

while the total private retention surplus satisfies

$$\alpha_e \Delta V_0^j + \Delta C_e > c_{ej}, \qquad \alpha_j \Delta V_0^j + \phi_j \Delta A_j + \Delta C_j > c_{j0}.$$

Thus material value, assignment value, and route-credit rents jointly make both private senders keep the redundant branch, while the decision maker also keeps the incoming link. The planner deletes it because route credit rents are redistributive and the incumbent's lost *assignment value is counted by $\Delta A_i < 0$. The same feasible topology can therefore exhibit insufficient or excess connection under different primitive parameter values.*

Example OB.13 (A route-credit overconnection variant)**.** *Take the completed two-branch network and suppose the planner counts decision-maker value, productive assignment value, and link costs, while senders also receive private route-credit rents from being publicly salient on the winning route. Choose route-credit increments $\Delta C_e, \Delta C_j$ so that the senders' total material, assignment, and route-credit retention surplus exceeds their private link costs, while the planner gain from that branch net of productive assignment and reception costs is negative. Then both sender endpoints keep their links under pairwise stability, while the planner deletes the branch. This supplementary example isolates the rent channel. The main formation proposition therefore highlights underconnection created by common-record censoring value.*

OB.5 Supplementary reputational and formation formulas

In this section, we record the primitive audit and formation objects used in the main text. On a path, $R_t = (R_t^j)_{j \in P \setminus \{0\}}$ is the public reputation vector, C_t is the current custodian, A_t is the public file age, and

$$S_t = (m_t, \tilde{v}_t, \mathcal{H}_t^{\tau}, R_t, C_t, A_t, \bar{d}_t, q_t, K_t)$$

is the public state when the holder's gap law is a one-dimensional censored Gaussian with interval-union support.

Audit experiments and separated reputation updating. Fix intermediary i and let the public track audit quality $\theta_i \in \{H, L\}$ through posterior $\pi_t^i = \mathbb{P}(\theta_i = H \mid \mathcal{F}_t^p)$ and reputation $R_t^i = \rho_i(\pi_t^i)$. At an active review, the public observes the custody outcome and an audit mark s whose likelihood $\ell_\theta(s)$ is conditionally independent of the sealed file gap given the public slice and custody outcome. Bayes' rule then gives

$$\pi_t^i = \frac{\pi_{t-}^i \ell_H(s)}{\pi_{t-}^i \ell_H(s) + (1 - \pi_{t-}^i)\ell_L(s)}.$$

Thus the custody record censors the file gap and the audit record prices review quality. The single-kernel benchmark uses one custody region at each public slice; type-indexed custody regions are the weighted-support case in Lemma OB.14.

Lemma OB.14 (Type-indexed custody and weighted support)**.** *If audit types H and L use type-indexed retention intervals $I_H(\xi)$ and $I_L(\xi)$ at the same public review slice, an observed* *retention gives*

$$p(d \mid \text{ret}, \xi) \propto \phi(d; \bar{d}, q)\big[\pi \mathbf{1}\{d \in K \cap I_H(\xi)\} + (1-\pi)\mathbf{1}\{d \in K \cap I_L(\xi)\}\big].$$

If the bracketed weight is constant on its positive reference-mass support, the posterior is one censored Gaussian on that support. Otherwise the exact state is a finite weighted support, equivalently a type-indexed mixture. The support-censoring economics remains: retention removes states and a rival branch inherits the resulting feasible set in a common-record race.

Proof. Bayes' rule multiplies the reference density by the probability that the realized type retains at gap d. A single censored Gaussian has density proportional to $\phi(d; \bar{d}, q)\mathbf{1}\{d \in K'\}$, so the weighted posterior is single-kernel exactly when its positive cells all carry the same weight. Otherwise the public statistic must retain the finite collection of weighted cells or, equivalently, the type mixture. □

Assignment value and route credit. Future projects arrive to intermediary i at rate ν_i. Conditional on an opportunity, i receives the mandate with public probability $p_i(R_t^i)$, and a correct assignment to high audit quality creates surplus S_i^{proj}. At docket date τ,

$$A_i(G) = \mathbb{E}\Big[\int_\tau^\infty e^{-\rho(t-\tau)} \nu_i p_i(R_t^i) \pi_t^i S_i^{\text{proj}}\, dt \;\Big|\; \mathcal{F}_\tau^p\Big].$$

The private component is $\phi_i A_i(G)$, $\phi_i \in [0,1]$. Visible participation on the route that first resolves the episode also creates private route credit. If $\tau^*(G)$ is the first terminal resolution time and $\Gamma^*(G)$ is the winning route, write

$$C_i(G) = \mathbb{E}\Big[e^{-\rho(\tau^*(G)-\tau)} \varpi_i^{\text{cred}}(\tau^*(G), \Gamma^*(G))\mathbf{1}\{i \in \Gamma^*(G)\} \;\Big|\; \mathcal{F}_\tau^p\Big].$$

Planner welfare counts productive surplus; private formation values add route-credit rents.

Bayes-plausible information discipline. Let $G_i(\pi)$ be the unweighed discounted continuation contribution generated by the audit-reputation state after posterior π, and also assume $G_i''(\pi) \geq \gamma_i > 0$ on the compact unresolved audit region. At a review date, let $\Pi_i^P(\pi)$ and $\Pi_i^N(\pi)$ be the posterior random variables generated by the audit experiments attached to prompt resolution and skipped active review, with support in that compact region. If both experiments are Bayes plausible,

$$\mathbb{E}[\Pi_i^P(\pi) \mid \pi] = \mathbb{E}[\Pi_i^N(\pi) \mid \pi] = \pi,$$

and the prompt experiment is uniformly more informative,

$$\Pi_i^P(\pi) \text{ is a mean-preserving spread of } \Pi_i^N(\pi), \qquad \mathrm{Var}(\Pi_i^P(\pi)) - \mathrm{Var}(\Pi_i^N(\pi)) \geq \eta_i > 0,$$

then

$$\mathbb{E}[G_i(\Pi_i^P(\pi)) - G_i(\Pi_i^N(\pi)) \mid \pi] \geq \frac{1}{2}\gamma_i\eta_i.$$

This is the unweighted discipline margin in Assumption 5; the payoff-weighted margin is $\beta_i\underline{\delta}_i$.

Verification of Corollary 3. With public recertification $M_i(\pi) = (1-\chi_i)\pi + \chi_i\pi_i^0$, $\chi_i \in (0,1)$ and $\pi_i^0 \in (0,1)$, we have

$$M_i([0,1]) = [\chi_i\pi_i^0, 1-\chi_i(1-\pi_i^0)] \subset (0,1).$$

Thus, all realized audit posteriors lie in a compact interior interval. Bayes plausibility, the variance gap, and the second-order integral form of convexity gives the displayed continuation-value lower bound, which equals $\underline{\delta}_i$ in Assumption 5. Multiplication by β_i only forms the payoff-weighted penalty used in the resolution theorem. □

Age envelope. In the dated-record Ornstein-Uhlenbeck subclass, revealing a record of public age a changes future posterior means by at most $C_i e^{-\lambda_i^{\mathrm{eff}} a}$ in conditional L^2 on compact unresolved slices. Quadratic material losses and bounded conditional second moments imply

$$G_i^{\mathrm{mat}}(d,\xi;\varrho) \leq \overline{G}_i(a) := C_i' e^{-\lambda_i^{\mathrm{eff}} a},$$

so that $\overline{G}_i(a) \downarrow 0$. Theorem 2 applies once this material envelope falls below the review-discipline margin.

OB.6 Supplementary direct-access benchmark

Along a fixed route in the maintained old-record resolution regime, the relevant comparison is disclosure speed. Proposition 10 gives common gap censoring: observed retention on one branch reshapes the rival branch's feasible set before the rival next moves. The comparison below gives conditions under which a highly reputation-sensitive intermediary discloses on a direct route to the decision maker before a slower downstream gatekeeper moves. Compare the serial line $e \to i \to j \to 0$ with the race formed by direct branches $e \to i \to 0$ and $e \to j \to 0$.

Theorem OB.15 (Direct access and waiting set inclusion)**.** *Let us suppose there is one expert e, one decision maker 0, and two intermediaries i and j with identical biases $b_i = b_j$ but different reputational intensities $\beta_i > \beta_j$. Consider the serial line G^{SL} with $e \to i \to j \to 0$ and the two-route race G^{TR} with routes $e \to i \to 0$ and $e \to j \to 0$. Assume the two topologies are compared under the same docket date law of the common public slice and hidden gap, that all primitives other than topology are identical across them, and that the public review protocol of Section 2, the common gap recursion of Lemma 2, and the unresolved announced state recursion of Proposition 8 hold.*

For a common public slice ξ, let $I_i^{\mathrm{SL}}(\xi)$ be intermediary i's waiting region on the serial line and $I_i^{\mathrm{TR}}(\xi)$ her waiting region on the direct branch in the race. Let $\Delta_i^{\mathrm{SL}}(d,\xi)$ and $\Delta_i^{\mathrm{TR}}(d,\xi)$ denote the corresponding full review date forward gains: in the serial line this is the gain from relay to j, while on the direct branch it is the gain from terminal disclosure to the decision maker. Suppose now there exist publicly measurable nonnegative functions $\Gamma_i(\xi)$, $\Pi_i(\xi)$, and $\Lambda_i(\xi)$ such that, for every feasible slice and every feasible d,

$$\Delta_i^{\mathrm{TR}}(d,\xi) - \Delta_i^{\mathrm{SL}}(d,\xi) \geq \Gamma_i(\xi) + \Pi_i(\xi) - \Lambda_i(\xi), \tag{1}$$

where $\Gamma_i(\xi)$ is the gain from removing the downstream gatekeeper's review delay, $\Pi_i(\xi)$ is the marginal first disclosure credit gain from direct terminal access in the race, and $\Lambda_i(\xi)$ is an upper bound on any increase in the continuation value of retaining on the direct branch relative to the serial line. If

$$\Gamma_i(\xi) + \Pi_i(\xi) - \Lambda_i(\xi) \geq 0$$

on every feasible slice, then

$$I_i^{\mathrm{TR}}(\xi) \subseteq I_i^{\mathrm{SL}}(\xi)$$

for every feasible slice. If, in addition, there is a positive probability set of feasible slices on which $I_i^{\mathrm{SL}}(\xi)$ contains a feasible subset $B(\xi)$ of positive conditional mass such that

$$\Delta_i^{\mathrm{TR}}(d,\xi) \geq 0 > \Delta_i^{\mathrm{SL}}(d,\xi) \qquad \textit{for every } d \in B(\xi),$$

then $I_i^{\mathrm{TR}}(\xi) \subsetneq I_i^{\mathrm{SL}}(\xi)$ on that set. If, on some slice, the inequality in (1) is strong enough that $\Delta_i^{\mathrm{TR}}(d,\xi) \geq 0$ for every feasible d, then $I_i^{\mathrm{TR}}(\xi) = \varnothing$ on that slice.

Theorem OB.15 is a sufficient slice by slice comparison: it delivers waiting set inclusion on common public slices, even though later retention events generally move the line and the race to different public states.

Corollary OB.16 (Structural Poisson race benchmark)**.** *In the two-route benchmark of*

Theorem OB.15, suppose public review dates on branch $k \in \{i, j\}$ *arrive as Poisson events with intensity* $\bar{\lambda}_k$*, the remaining material value of terminal disclosure at public age* $a + s$ *is* $M(a+s) = \bar{M}e^{-\kappa(a+s)}$*, branch* i *earns a first disclosure premium* $\varpi_i \geq 0$*, and the direct branch has* $\Lambda_i(\xi) = 0$*. Then* (1) *holds with*

$$\Gamma_i(\xi) = \bar{M}e^{-\kappa a}\frac{\rho+\kappa}{\rho+\kappa+\bar{\lambda}_j}, \qquad \Pi_i(\xi) = \varpi_i, \qquad \Lambda_i(\xi) = 0.$$

Hence

$$\Delta_i^{\mathrm{TR}}(d,\xi) - \Delta_i^{\mathrm{SL}}(d,\xi) \geq \varpi_i + \bar{M}e^{-\kappa a}\frac{\rho+\kappa}{\rho+\kappa+\bar{\lambda}_j} > 0$$

on every feasible slice. The direct branch waiting region is weakly smaller on every slice, and it is empty whenever this lower bound dominates the serial-line retention hurdle uniformly on that slice.

Proof of Theorem OB.15. Fix a feasible public slice ξ. The pointwise full-gain comparison in (1) directly implies

$$\Delta_i^{\mathrm{TR}}(d,\xi) \geq \Delta_i^{\mathrm{SL}}(d,\xi)$$

for every feasible d. Hence every d for which retention is optimal on the direct branch is also a d for which retention is optimal on the serial line, so

$$I_i^{\mathrm{TR}}(\xi) \subseteq I_i^{\mathrm{SL}}(\xi).$$

Under the theorem's additional nondegeneracy condition, there is a positive probability set of slices on which $I_i^{\mathrm{SL}}(\xi)$ contains a feasible subset $B(\xi)$ of positive conditional mass satisfying

$$\Delta_i^{\mathrm{TR}}(d,\xi) \geq 0 > \Delta_i^{\mathrm{SL}}(d,\xi) \qquad \text{for every } d \in B(\xi).$$

Every $d \in B(\xi)$ is therefore retained on the serial line and forwarded on the direct branch, so $I_i^{\mathrm{TR}}(\xi) \subsetneq I_i^{\mathrm{SL}}(\xi)$ on that set. If the direct branch gain is nonnegative for every feasible d on some slice, the direct branch waiting set is empty on that slice.

The theorem proves the exact slice-by-slice waiting-set inclusion above. After the first non-disclosure event, the serial line and race generally jump to different public states, so later review comparisons are evaluated on their realized slices. The gain Δ_i^{TR} is the full direct-branch gain, including its equilibrium continuation and retention operator. □

Proof of Corollary OB.16. Let $T_j \sim \mathrm{Exp}(\bar{\lambda}_j)$ be the downstream review delay removed by direct access. In the structural subclass, the serial line loses the discounted age-decay term during that delay, direct terminal access adds ϖ_i, and $\Lambda_i(\xi) = 0$. Hence

$$\Gamma_i(\xi) = \mathbb{E}[M(a) - e^{-\rho T_j} M(a + T_j)] = \bar{M} e^{-\kappa a} \left(1 - \frac{\bar{\lambda}_j}{\rho + \kappa + \bar{\lambda}_j}\right),$$

which give the displayed bound. The remaining conclusions are Theorem OB.15 applied to this positive lower bound. □

Structural Poisson/two-route subclass. For the direct-access comparison, take the serial line $e \to i \to j \to 0$ and the two-route race $e \to i \to 0$, $e \to j \to 0$. Let branch-j review delay be $T_j \sim \mathrm{Exp}(\bar{\lambda}_j)$, direct terminal access create first-disclosure premium $\varpi_i \geq 0$, remaining material value at age a be $M(a) = \bar{M} e^{-\kappa a}$, and the direct branch have $\Lambda_i(\xi) = 0$. Then

$$\Delta_i^{\mathrm{TR}}(d, \xi) - \Delta_i^{\mathrm{SL}}(d, \xi) \geq \Gamma_i(\xi) + \Pi_i(\xi) - \Lambda_i(\xi),$$

where

$$\Gamma_i(\xi) = \bar{M} e^{-\kappa a} \frac{\rho + \kappa}{\rho + \kappa + \bar{\lambda}_j}, \qquad \Pi_i(\xi) = \varpi_i, \qquad \Lambda_i(\xi) = 0.$$

This primitive subclass verifies Corollary OB.16 and supplies the metric terms used in the formation construction.

OB.7 Hazard-weighted review and mixture boundaries

If review occurrence itself depends on the hidden gap through a hazard multiplier $\lambda(d)$, review arrival adds a likelihood factor before custody censoring:

$$p(d \mid \text{review}) \propto \lambda(d) \phi(d; \bar{d}, q) \mathbf{1}\{d \in K\},$$

$$p(d \mid \text{review, retain}) \propto \lambda(d) \phi(d; \bar{d}, q) \mathbf{1}\{d \in K \cap I_i\}.$$

The public-retention censoring logic is unchanged; the single Gaussian endpoint recursion is replaced by a hazard-weighted kernel unless λ lies in a closed exponential-quadratic family.

With two audit types, a shared gap reference density, and type-contingent custody intervals I_H and I_L, retention gives

$$p(d \mid R) \propto \phi(d; \bar{d}, q)\{\omega_H \mathbf{1}\{d \in K \cap I_H\} + \omega_L \mathbf{1}\{d \in K \cap I_L\}\},$$

where ω_H, ω_L are the public type weights after the audit mark. Type-specific gap kernels would be a broader mixture extension; the weighted-support object used here keeps the single reference density. The main text's common custody-region assumption is the one-kernel specialization $I_H = I_L$.

Proposition OB.17 (Bivariate Gaussian robustness of common-record censoring)**.** *Let* (g_i, g_j) *be jointly Gaussian on an unresolved public slice, with means* (μ_i, μ_j)*, variances* (σ_i^2, σ_j^2)*, and correlation* $\rho \in [0, 1]$*. If branch* j *retains on interval* I_j*, then*

$$g_i \mid g_j = x \sim \mathcal{N}\left(\mu_i + \rho\frac{\sigma_i}{\sigma_j}(x - \mu_j), (1 - \rho^2)\sigma_i^2\right), \qquad x \in I_j,$$

with x *distributed according to the marginal Gaussian law of* g_j *truncated to* I_j*. For any continuation functional* H_i *with quadratic growth, define now the imperfect-copy censoring value by integrating* H_i *under this conditional law and subtracting the independent-record value. Then, this value is continuous in* ρ*, equals zero at* $\rho = 0$*, and converges to the certified-copy value as* $\rho \uparrow 1$ *when the two branch records have the same marginal law.*

If $H_i(z) = a_i + b_i z + c_i z^2$*, the retained-event value is closed form. Writing*

$$\bar{x}_j = \mathbb{E}[g_j \mid g_j \in I_j], \qquad s_j^2 = \mathrm{Var}(g_j \mid g_j \in I_j),$$

one has

$$\mathbb{E}[H_i(g_i) \mid g_j \in I_j] = a_i + b_i\bar{m}_i(\rho) + c_i\{\bar{m}_i(\rho)^2 + (1 - \rho^2)\sigma_i^2 + \rho^2\sigma_i^2\sigma_j^{-2}s_j^2\},$$

where $\bar{m}_i(\rho) = \mu_i + \rho\sigma_i\sigma_j^{-1}(\bar{x}_j - \mu_j)$*. Thus the common-record wedge is the continuous high-correlation limit of a bivariate Gaussian update, and it vanishes as branch gaps become conditionally independent.*

Proof. The conditional Gaussian formula gives the displayed law. Continuity in ρ follows from dominated convergence, using the quadratic-growth bound and bounded conditional second moments of the truncated Gaussian family. At $\rho = 0$, g_i is independent of branch j's retention event, so the retained-event law of g_i equals its independent-record law. As $\rho \uparrow 1$ with common marginal law, the conditional variance of g_i given g_j vanishes and the bivariate update collapses to the scalar certified-copy restriction $g_i = g_j \in I_j$. The quadratic formula follows by applying the law of iterated expectations and the first two truncated-normal moments of g_j. □

OB.8 Supplementary old-record and formation variants

Corollary OB.18 (Boundary-compatible old-record resolution by rate dominance)**.** *Consider the same single-file path game and review recurrence as in Theorem 2. Replace the* *uniform interior penalty by an age-dependent lower bound: along unresolved histories at public age a, a skipped active review lowers the reputational continuation contribution by at least $\delta_i(a)$ relative to forward movement. Suppose that for some constants $C_i, c_i > 0$, $a_i^0 < \infty$, and effective age rates $\lambda_i^{\mathrm{age}} > \nu_i^{\mathrm{age}} \geq 0$,*

$$G_i^{\mathrm{mat}}(d, \xi; \varrho) \leq C_i e^{-\lambda_i^{\mathrm{age}} a} \quad \textit{and} \quad \beta_i \delta_i(a) \geq \beta_i c_i e^{-\nu_i^{\mathrm{age}} a}$$

for all feasible unresolved states with $a \geq a_i^0$. Then, there is a finite age

$$a_i^{\mathrm{RD}} := \max\left\{a_i^0, 0, \frac{\log(C_i/(\beta_i c_i))}{\lambda_i^{\mathrm{age}} - \nu_i^{\mathrm{age}}}\right\}$$

after which forward movement is optimal at any active review of sender i. Hence, on the maintained unresolved-slice event, review recurrence and finite downstream custody again imply terminal disclosure in finite calendar time almost surely. This corollary allows Bayesian reputation penalties to drift toward the boundary, provided that dated record's material value decays faster than the review-discipline penalty.

Proof of Corollary OB.18. Fix sender i and an unresolved public history at file age a. By hypothesis, the material gain from one more retention is at most $C_i e^{-\lambda_i^{\mathrm{age}} a}$ for all $a \geq a_i^0$, while a skipped active review costs at least $\beta_i c_i e^{-\nu_i^{\mathrm{age}} a}$ in reputational continuation value. Since $\lambda_i^{\mathrm{age}} > \nu_i^{\mathrm{age}}$, there is a finite age

$$a_i^{\mathrm{RD}} = \max\left\{a_i^0, 0, \frac{\log(C_i/(\beta_i c_i))}{\lambda_i^{\mathrm{age}} - \nu_i^{\mathrm{age}}}\right\}$$

(with the last term omitted when $C_i \leq \beta_i c_i$ and the displayed maximum interpreted with the corresponding convention) such that

$$C_i e^{-\lambda_i^{\mathrm{age}} a} \leq \beta_i c_i e^{-\nu_i^{\mathrm{age}} a} \qquad \text{for all } a \geq a_i^{\mathrm{RD}}.$$

Thus, after age a_i^{RD}, retention cannot strictly dominate forward movement at any active review. Under the maintained tie-breaking convention, the sender relays or terminally discloses at the first active review after that age.

The review-recurrence argument from the proof of Theorem 2 applies verbatim with a_i^{RD} in place of $a_i^{\dagger}$: if the file remains outstanding at i, the first active review after public age a_i^{RD} occurs in finite calendar time almost surely. On the maintained unresolved-slice event, custody moves only downstream on a finite path, so iterating over custodians yields terminal disclosure in finite time almost surely. □

Corollary OB.19 (Route-credit overconnection)**.** *In the formation benchmark, suppose completing a redundant branch generates private first-disclosure route credit for the involved senders but creates planner value below its link and receiver costs once productive assignment effects are counted. Then, for an open set of positive link costs and route-credit rents, the completed network is pairwise stable while the planner deletes the redundant branch. This secondary wedge is driven by private credit; the main-text wedge is the public-information value of a second certified route.*

Definition OB.20 (Branchwise stability)**.** *A network is* branchwise stable without transfers *if it is pairwise stable and, in addition, no missing branch* $e \to j \to 0$ *can be added as a coordinated two link deviation such that all directly involved parties* $e, j, 0$ *weakly gain and at least one strictly gains, using the private payoffs in Definition 4.*

Corollary OB.21 (Censoring wedge under branchwise stability)**.** *In the same formation benchmark, allow coordinated addition of a missing branch as in Definition OB.20. If the planner gain from completing branch* j *is driven by* $\Delta\Theta_0^j$ *while the contracting environment cannot price the counterfactual continuation value created by branch-*j *retention, then the branch can fail to form without transfers even though it raises* W^P*. Pairwise stability adds a two-link activation problem; branchwise stability removes that coordination problem but still leaves an incidence problem under the limited-contracting environment.*

Proof of Corollary OB.21. Branchwise stability removes one-link coordination failure but preserves consent by directly involved parties under the same limited-contracting benchmark. If $\Delta\Theta_0^j$ generates the planner gain and the required parties cannot contract on that counterfactual continuation value, choose costs so the planner inequality is strict while the consent constraint fails. The coordinated branch is rejected although W^P rises. ☐

OC Additional docket-compatible institutions

The core environment is designed for institutions with docketing, sealed records, and public or common-knowledge review processes. It clarifies the relevant "public" and the single-file review protocol in each case.

OC.1 Audit, compliance, and case management dockets

Many audit and compliance systems create a chain of custody once a certified issue is opened. A case receives a docket date, a current handler, and a review schedule, while the underlying evidentiary content remains sealed to outsiders until formal release. Here the relevant public is typically a regulator, a board, an internal monitor, or any set of observers who see the docket and review trail before they see the substantive file. The single-file episode fits here naturally: the age of the pending file is known, custody transfers are logged, and the central question is when the current custodian releases the certified record from an already known docket. Protected escalation rights inside firms or public agencies are a natural interpretation of the direct access comparison.

OC.2 Protected technical reporting in AI and safety review

A similar structure appears in model release, safety, and incident-reporting pipelines. A technical finding can be logged, timestamped, and placed under review while its substantive content remains restricted to the review chain. Here the relevant public may be an internal safety board, an external regulator, a designated incident-response committee, or a precommitted set of governance monitors. This interpretation is complementary to Buhai [2025], which studies release cadence, observable quiet windows, and release ladders in a single-sender real-options model of AI deployment. The direct access result here speaks to protected reporting rights for high-accountability auditors or red teams that would otherwise sit behind a slower internal release gate.

OC.3 Rating surveillance and advisory bottlenecks

Rating surveillance, watch-list procedures, and advisory review chains also fit the model when review activity is recorded in a regulated or internally auditable trail. A review can be opened and timestamped, the fact of review can be observable to the relevant audience, and the operative evidence can remain sealed until a formal action is taken. Here, the relevant public might be the market, a regulator, a supervisory board, or internal governance staff. Proposition 11 isolates underconnection from common-record censoring value; the surrounding formation discussion and Online Appendix OB separate this force from links driven mainly by private route-credit rents.

References

I. S. Buhai. Real Option AI: Reversibility, Silence, and the Release Ladder. arXiv preprint, 2025. arXiv:2511.16958 [econ.TH].

P. Onuchic and J. Ramos. Disclosure by groups. CEPR Discussion Paper No. DP20057, 2025. CEPR Press, 16 March 2025.