Online Appendix

B Temporary Monopoly: Additional Results

Section 4.1 discussed how temporary monopoly protection can alleviate the crowding-out effect. This appendix provides some formal results about the optimal length τ^* of temporary monopoly protection.

A key simplification is that all equilibrium objects that relate to τ^* depend on δ and τ only through $z \equiv \delta^\tau$. To emphasize this dependence, we will abuse notation by denoting the incumbent's cutoff as $q^*(z) \equiv q^*(n)$ and welfare as $W(z) \equiv W(n)$. The continuation values from Proposition 4 become

$$\begin{aligned} V_1(n,z) &= (1-z)\, E[q_1 \mid q_1 > q^*(z)] + z\, E[q^{\text{win}}(n) \mid q_1 > q^*(z)], \\ V_0(n,z) &= (1-z)\, E[q_1 \mid q_1 \le q^*(z)] + z\, E[q^{\text{win}}(n) \mid q_1 \le q^*(z)], \end{aligned}$$

with region thresholds $\underline{n}(z) \equiv \min\{n \ge 1 : V_1(n,z) - \bar{h} > q^{\text{SQ}}\}$ and $\overline{n}(z) \equiv \min\{n \ge 1 : V_0(n,z) - \bar{h} \ge q^{\text{SQ}}\}$. Welfare in region (2) decomposes as

$$(1-\delta)\, W(z) \;=\; \underbrace{\int_0^\infty \max\{q, q^{\text{SQ}}\}\, f(q)\, dq \;-\; \bar{h}\big[1 - F(q^{\text{SQ}})\big]}_{\text{baseline: perpetual monopoly}} \;-\; \underbrace{\int_{q^{\text{SQ}}}^{q^*(z)} \big[(q - q^{\text{SQ}}) - \bar{h}\big]\, f(q)\, dq}_{\text{extensive margin: suppressed innovations}}$$

$$+ \underbrace{\int_{q^*(z)}^\infty z\big(E[q^{\text{win}}(n) \mid q_1 = q] - q\big)\, f(q)\, dq}_{\text{intensive margin: post-monopoly disruption}}\,. \qquad (7)$$

The first term is welfare under perpetual monopoly protection, where the incumbent reveals truthfully and the regulator always follows its recommendation. The remaining two terms formalize the trade-off discussed in Section 4.1: on the extensive margin, increasing τ (decreasing z) contracts the interval $(q^{\text{SQ}}, q^*(z))$ and fewer innovations are suppressed; on the intensive margin, a longer monopoly delays competitive entry, discounting the gains from disruption by z.

Proposition 5. *Let $\tau^* \in \{0, 1, 2, \dots\} \cup \{\infty\}$ be the shortest monopoly length that maximizes welfare $W(\delta^\tau)$.*

1. *If $\tau^* > 0$, then $\underline{n}(\tau^*) \le n < \overline{n}(\tau^*)$.*

2. *Infinite monopoly protection is never optimal: $0 \le \tau^* < \infty$.*

Since the purpose of temporary monopoly protection is to elicit the incumbent's information, it is only worth providing if the regulator responds to that information—hence part (1). Part (2) follows because even when region (1) is empty, the marginal delay cost on the intensive margin eventually swamps the marginal information benefit on the extensive margin.

To state the comparative static on τ^*, we formalize the regulator's optimization problem. For the discrete problem the feasible set is $\mathcal{Z} \equiv \{1, \delta, \delta^2, \ldots\} \cup \{0\}$, where $z = 0$ corresponds to $\tau = \infty$. For the continuous relaxation, we extend the same formulas to all $z \in [0, 1]$. Define the normalized welfare function $J : [0, 1] \to \mathbb{R}$ by $J(z) \equiv (1-\delta)\, W(z)$ for a fixed market size n:

$$J(z) = \begin{cases} q^{\mathrm{SQ}}, & n < \underline{n}(z), \\ \text{RHS of (7)}, & \underline{n}(z) \leq n < \overline{n}(z), \\ (1 - z)\big(E[q_1] - \bar{h}\big) + z\big(E[q^{\mathrm{win}}(n)] - \bar{h}\big), & n \geq \overline{n}(z). \end{cases} \tag{8}$$

Since $\underline{n}(z)$ and $\overline{n}(z)$ are weakly decreasing in z (Proposition 4), smaller values of z raise both thresholds, moving a fixed n weakly from region (3) to (2) to (1). At the endogenous boundaries the adjacent expressions coincide, so J is continuous on $[0, 1]$. Each branch depends on (δ, τ) only through $z = \delta^\tau$.

Proposition 6. *In the continuous relaxation $\tau \in [0, \infty)$, the regulator solves $\max_{z \in (0,1]} J(z)$, where J is defined in (8). Since J depends on (δ, τ) only through $z = \delta^\tau$, the maximizer z^* is independent of δ. If $z^* \in (0, 1)$, the optimal protection length $\tau^* = \ln z^* / \ln \delta$ is strictly increasing in the discount factor δ.*

The mapping $\tau^* = \ln z^* / \ln \delta$ fully characterizes the δ-dependence of τ^*. The two forces identified in Section 4.1—a more patient incumbent requiring longer protection to reveal information, and a more patient regulator willing to tolerate longer delays—are both consequences of this mechanical relationship.

When τ is restricted to integers, the regulator optimizes over $\mathcal{Z}$. If J is strictly quasiconcave in τ—which is the case for common distributions and parametrizations—then the optimal protection length τ^* remains weakly increasing in δ.[31] Without quasiconcavity, however, the monotone comparative static may fail: the optimal grid point can jump non-monotonically across local maxima as δ shifts the grid.

Proofs

Proof of Proposition 5. Step 1: Reparametrization. As established above, all equilibrium objects—$q^*(z)$, $V_1(n, z)$, $V_0(n, z)$, $\underline{n}(z)$, $\overline{n}(z)$—depend on (δ, τ) only through $z = \delta^\tau$, and the

[31]With quasiconcave J, the discrete optimum must be one of the two grid points flanking z^*; that is, $\tau^* \in \{\lfloor t(\delta) \rfloor, \lceil t(\delta) \rceil\}$ where $t(\delta) = \ln z^* / \ln \delta$. Since $t(\delta)$ is strictly increasing in δ, both its floor and ceiling are weakly increasing, so the discrete τ^* is weakly increasing in δ.

continuous extension $J : [0,1] \to \mathbb{R}$ defined in (8) is continuous and independent of δ. The regulator's discrete problem is $\max_{z\in\mathcal{Z}} J(z)$.

Step 2: $q^*(z)$ is strictly increasing in z. Write the cutoff equation as $g(q,z) \equiv q[1 - z(1 - F(q)^{n-1})] - q^{\mathrm{SQ}} = 0$. Then $\partial g/\partial z = -q(1 - F(q)^{n-1}) < 0$ and $\partial g/\partial q = 1 - z(1 - F(q)^{n-1}) + z(n-1)qF(q)^{n-2}f(q) > 0$. By the implicit function theorem, $dq^*/dz = -(\partial g/\partial z)/(\partial g/\partial q) > 0$.

Step 3: If $\tau^* > 0$ then $\underline{n}(\tau^*) \le n < \overline{n}(\tau^*)$. Let $z^* = \delta^{\tau^*}$ with $\tau^* > 0$, so $z^* < 1$. Since τ^* is the smallest maximizer, z^* is the largest maximizer in $\mathcal{Z}$, so $J(z^*) > J(z)$ for every feasible $z > z^*$; in particular $J(z^*) > J(1)$. We rule out regions (1) and (3).

Region (1) is ruled out. If $n < \underline{n}(z^*)$, then $J(z^*) = q^{\mathrm{SQ}}$. But $J(z) \ge q^{\mathrm{SQ}}$ for every $z \in [0,1]$, since in regions (2) and (3) the regulator approves innovations with strictly positive net benefit, and in region (1) $J = q^{\mathrm{SQ}}$. In particular $J(1) \ge q^{\mathrm{SQ}} = J(z^*)$, contradicting $J(z^*) > J(1)$.

Region (3) is ruled out. If $n \ge \overline{n}(z^*)$, the regulator ignores the incumbent's recommendation and always approves, so monopoly protection serves no informational purpose and only delays competitive entry. Formally, $J(z^*) = (1 - z^*)(E[q_1] - \bar{h}) + z^*(E[q^{\mathrm{win}}(n)] - \bar{h})$. Because $z^* < 1$ and $E[q^{\mathrm{win}}(n)] > E[q_1]$ for $n \ge 2$, we have $J(z^*) < E[q^{\mathrm{win}}(n)] - \bar{h}$. At $z = 1$, unconditional approval is feasible after every message and yields $E[q^{\mathrm{win}}(n)] - \bar{h}$, so $J(1) \ge E[q^{\mathrm{win}}(n)] - \bar{h} > J(z^*)$—contradicting $J(z^*) > J(1)$.

Since regions (1) and (3) are ruled out, n must lie in region (2): $\underline{n}(z^*) \le n < \overline{n}(z^*)$, i.e., $\underline{n}(\tau^*) \le n < \overline{n}(\tau^*)$.

Step 4: Existence and $\tau^* < \infty$. $\mathcal{Z} \subset [0,1]$ is compact (a convergent sequence together with its limit), and J is continuous on $[0,1]$, so J attains a maximum on $\mathcal{Z}$. We show $z = 0$ is not optimal. Differentiating the region-(2) expression in (7) from the right at $z = 0$ (where $q^*(0) = q^{\mathrm{SQ}}$):

$$J'(0^+) = f(q^{\mathrm{SQ}})\,\frac{dq^*}{dz}\bigg|_{z=0} \cdot \bar{h} \;+\; \int_{q^{\mathrm{SQ}}}^{\infty} \big(E[q^{\mathrm{win}}(n) \mid q_1 = q] - q\big)\, f(q)\, dq.$$

From Step 2, $dq^*/dz\big|_{z=0} = q^{\mathrm{SQ}}(1 - F(q^{\mathrm{SQ}})^{n-1}) > 0$ for $n \ge 2$. The integral is strictly positive since $q^{\mathrm{win}}(n) = \max\{q_1, M_{n-1}\} \ge q_1$ with strict inequality on a set of positive measure. Hence $J'(0^+) > 0$, so $J(z) > J(0)$ for all sufficiently small $z > 0$. Since $\delta^\tau \to 0$, there exists a feasible $z = \delta^\tau \in \mathcal{Z}$ in this neighborhood, so $z = 0$ is not optimal and $\tau^* < \infty$. ■

Proof of Proposition 6. In the piecewise definition (8), the three case expressions and the region thresholds $\underline{n}(z)$, $\overline{n}(z)$ that determine which case applies all depend on (δ, τ) only through $z = \delta^\tau$. Hence J is a function of z alone, and the maximizer z^* is independent of δ.

For $z^* \in (0,1)$, the continuous protection length is $\tau^* = \ln z^* / \ln \delta$. Then

$$\frac{d\tau^*}{d\delta} = -\frac{\ln z^*}{\delta(\ln \delta)^2}.$$

Since $z^* \in (0,1)$ implies $\ln z^* < 0$, and $\delta(\ln \delta)^2 > 0$, we have $d\tau^*/d\delta > 0$. ■

C Enriching the Policy Space: Additional Results

Section 4.2 enriched the baseline policy space with $k \geq 2$ innovative policies. Example 2 constructed a specific regular equilibrium which highlighted that equilibrium behavior becomes much richer in this setting. This appendix generalizes and formalizes the findings from Section 4.2.

We start by repeating an observation about entrants' regular strategies. The key difference from the baseline ($k = 1$) setting is that entrants' regular messages become informative: each entrant's recommendation reveals which policy yields its highest quality draw.

Lemma 7 (Entrant regular strategy)**.** *With $k \geq 2$ innovative policies, an entrant playing a regular strategy recommends the innovative policy under which its quality draw is highest:*

$$m_j = \operatorname*{argmax}_{p \in \{1,\ldots,k\}} q_j^p \qquad \textit{for each } j \geq 2.$$

Proof. If innovative policy p is adopted, entrant j's expected profit is $q_j^p \, F(q_j^p)^{n-1}$, which is strictly increasing in q_j^p. Under the status quo the entrant's payoff is zero. Hence the regular strategy—recommending the policy that maximizes the entrant's payoff if implemented—selects the innovative policy with the highest quality draw. Since draws are i.i.d. from an atomless distribution, ties occur with probability zero. ■

One feature of Example 2's equilibrium is that the incumbent's message rule is not regular—for instance, an incumbent's regular message rule would not exhibit the informational spillovers present in Example 2. In general, in a regular equilibrium, the incumbent does not play a regular strategy. To make progress, we will impose some additional structure on the incumbent's recommendation rules.

Definition 1 (Well-behaved equilibrium)**.** *A regular equilibrium is* well-behaved *if the incumbent's strategy satisfies the following conditions:*

(i) *The incumbent sends $m_1 \in \{0, p_1^{\max}\}$—either the status quo or its best innovative policy. The recommendation rule is label-symmetric: the event $\{m_1 = 0\}$ depends on the quality vector $(q_1^1, \ldots, q_1^k)$ only through $(q_1^{\max}, q_1^{-\max})$, where $q_1^{-\max} = (q_1^{(2)}, \ldots, q_1^{(k)})$ denotes the incumbent's non-maximal qualities in decreasing order.*

(ii) *If* $\Pr(m_1{=}0) > 0$, *let* $\tilde{F}^p$ *denote the CDF of the incumbent's policy-p quality* q_1^p *conditional on the incumbent recommending the status-quo* ($m_1 = 0$). *Then* $\tilde{F}^p(t) \geq F(t)$ *for all* t.

Condition (i) specifies that the incumbent's recommendation is quasi-binary: it lobbies either for the status quo or its best innovative policy. Importantly, although the incumbent never recommends a dominated innovative policy, the message rule can depend on the incumbent's full quality vector, generating the informational spillover discussed in Section 4.2—while requiring that the recommendation rule treat policy labels symmetrically conditional on the order statistics $(q_1^{\max}, q_1^{-\max})$. This label-symmetry ensures that, conditional on $m_1 = 0$, each individual quality q_1^p shares a common posterior distribution. Condition (ii) rules out perverse posteriors by requiring that a status-quo recommendation is *unfavorable* news about the incumbent's quality for each innovative policy.[32]

With multiple innovative policies, the regulator uses messages to learn not only whether innovation is attractive, but also which innovation is most promising. Entrant support for policy p is favorable news about that policy's quality, as is incumbent endorsement.

Proposition 7. *In any well-behaved equilibrium, the regulator's decision rule takes the following form.*

There exist policy-independent cutoffs $\bar{c}_e, \bar{c}_i, \bar{c}_0 \in \{0, 1, \ldots, n{-}1\} \cup \{\infty\}$ *such that the regulator approves innovation if and only if: for at least one innovative policy* $p \geq 1$, *the number of entrant recommendations* c_p *of* p *meets a threshold that depends on the incumbent's message:*

1. $c_p \geq \bar{c}_e$ *and the incumbent recommends* p *(endorsement);*

2. $c_p \geq \bar{c}_i$ *and the incumbent recommends another innovative policy* $p' \neq p$; *or*

3. $c_p \geq \bar{c}_0$ *and the incumbent recommends the status quo.*

If the incumbent does not babble ($\Pr(m_1 = 0) \in (0, 1)$*), then its endorsement of innovation weakly lowers the threshold:* $\bar{c}_e \leq \min\{\bar{c}_i, \bar{c}_0\}$.

Moreover, when $k \geq 2$, *there exists a well-behaved equilibrium.*

Unlike the baseline, the regulator now conditions on both the incumbent's message and the entrant count c_p. Sufficiently strong entrant support can lead the regulator to override the incumbent's recommendation (case 2). The incumbent's support remains

[32]Without condition (ii), a status-quo recommendation can perversely be *good* news about a non-endorsed quality component: the event $m_1 = 0$ may select quality vectors in a way that makes some component unusually high, even though labels are treated symmetrically. Condition (ii) rules out exactly these favorable-selection effects by requiring the conditional distribution of each component under $m_1 = 0$ to be weakly worse than the prior.

relevant—endorsement lowers the required entrant threshold among on-path messages, and the incumbent-endorsed policy is preferred when entrant counts are tied—but entrant messages are no longer irrelevant.

For existence, the proof constructs a well-behaved equilibrium in which the incumbent always recommends its favorite innovative policy, so the status-quo message is off path. Even in this case, the regulator's decision remains sensitive to entrant support across policies, with different count thresholds depending on whether a policy is incumbent-endorsed.

Although Definition 1(i) does not impose a threshold rule on the incumbent's recommendation, Proposition 7 implies that such a threshold structure arises endogenously in any equilibrium where the status-quo message is used on path.

Corollary 5. *In any well-behaved equilibrium where the incumbent recommends both the status quo and innovation with positive probability (*$\Pr(m_1{=}0) \in (0,1)$*) and its endorsement strictly affects the adoption probability of the endorsed policy, there exists a function* $\kappa\colon \mathbf{R}_+^{k-1} \to \mathbf{R}_+$ *such that the incumbent recommends its best innovative policy if and only if* $q_1^{\max} > \kappa(q_1^{-\max})$*.*

Proofs

Notation. Write $r = p_1^{\max}$ for the incumbent's best innovative policy. Let $c_p = \sum_{j\geq 2} \mathbf{1}[m_j = p]$ denote the number of entrants recommending policy p, with $\sum_{p=1}^{k} c_p = n-1$. By Lemma 7, if entrants play regular strategies, then each entrant recommends $\operatorname{argmax}_p q_j^p$, and since F is atomless, each entrant independently recommends each policy with probability $\frac{1}{k}$, so $\boldsymbol{c} = (c_1, \dots, c_k)$ follows a Multinomial$(n-1, \frac{1}{k}, \dots, \frac{1}{k})$ distribution.

By Bayes' rule, the posterior CDF of entrant j's quality for policy p, conditional on recommending p, is $G(x) = F(x)^k$; conditional on *not* recommending p (regardless of which other policy was recommended), it is $H(x) = \frac{k\,F(x) - F(x)^k}{k-1}$. For $F(x) \in (0,1)$, $G(x) < F(x) < H(x)$, so recommendation is good news and non-recommendation is bad news in the FOSD sense. Note that H does not depend on which of the other $k-1$ policies the entrant recommended; this follows from the symmetry of the remaining policies under i.i.d. draws.

The regulator's expected peak quality for policy p, conditional on the incumbent's message m_1 and entrant count c_p, is

$$V_p(m_1, c_p) = \int_0^\infty \bigl[1 - K(t)\, G(t)^{c_p}\, H(t)^{n-1-c_p}\bigr]\, dt, \tag{9}$$

where K is the CDF of q_1^p conditional on m_1. This expression has two key properties: *count monotonicity*—V_p is strictly increasing in c_p (since $G(t) < H(t)$ for all t with $F(t) \in (0,1)$, so each replacement of an H-draw by a G-draw strictly lowers the integrand on a set of positive measure)—and *FOSD monotonicity*—if $K_1(t) \leq K_2(t)$ for all t, then V_p under K_1 weakly exceeds V_p under K_2.

When the incumbent recommends r, write $\tilde{G}_1$ for the CDF of $q_1^{\max}$ conditional on recommendation, and $\tilde{H}_1^*$ for the CDF of any single non-endorsed quality q_1^s ($s \neq r$) conditional on recommendation. By symmetry of the i.i.d. draws, $\tilde{H}_1^*$ is the same for every non-endorsed policy. Write $V_r^{c_r} \equiv V_r(m_1{=}r,\, c_r)$ and $V_s^{c_s} \equiv V_s(m_1{=}r,\, c_s)$ for any non-endorsed policy s. When the incumbent recommends the status quo, if the recommendation event has positive probability, the label-symmetry requirement in Definition 1(i) ensures all $\tilde{F}^p$ coincide; write $\tilde{F}$ for their common CDF and $V_0(c) \equiv V_p(m_1{=}0,\, c)$.

Lemma 8 (Monotonicity and endorsement)**.** *In any well-behaved equilibrium at any $n \geq 2$:*

(a) Entrant switch. *For any policy p and any fixed incumbent message, if one additional entrant switches to recommending p (from some other policy $p' \neq p$), the regulator's expected utility from adopting p strictly increases.*

(b) Incumbent endorsement. *Suppose the incumbent doesn't babble:* $\Pr(m_1 = 0) \in (0, 1)$. *For any policy p and any fixed entrant messages, the regulator's expected utility from adopting p is strictly higher if the incumbent recommended p than if the incumbent recommended any other policy.*

Proof of Lemma 8. Part (a). Raising c_p by one replaces an H-draw with a G-draw for policy p, strictly raising V_p by count monotonicity.

Part (b). Fix policy p and entrant count c_p. In (9) the entrant terms are held fixed, so V_p depends on the incumbent's message only through the posterior CDF K of q_1^p. By FOSD monotonicity, it suffices to compare K across messages.

If the incumbent recommended $m_1 = s \in \{1, \ldots, k\} \setminus \{p\}$: then $K = \tilde{H}_1^*$. Conditional on the same recommendation event $\mathcal{E} = \{m_1 = r\}$, the endorsed quality satisfies $q_1^{\max} > q_1^s$ a.s. (atomless F), so $\tilde{G}_1(t) \leq \tilde{H}_1^*(t)$ for all t, with strict inequality on a set of positive measure. FOSD monotonicity gives $V_p(m_1{=}p,\, c_p) > V_p(m_1{=}s,\, c_p)$.

If the incumbent recommended $m_1 = 0$: then $K = \tilde{F}$. For any label p,

$$F(t) = \tfrac{\Pr(m_1 \neq 0)}{k} \left[\tilde{G}_1(t) + (k-1)\, \tilde{H}_1^*(t)\right] + \Pr(m_1{=}0)\, \tilde{F}(t),$$

since conditional on endorsement each label is equally likely to be endorsed. Since $\tilde{G}_1 \leq \tilde{H}_1^*$ (previous case) and $\tilde{F} \geq F$ (condition (ii)), the decomposition gives $\tilde{G}_1(t) \leq F(t) \leq \tilde{F}(t)$. Hence $V_p(m_1{=}p,\, c_p) \geq V_p(m_1{=}0,\, c_p)$. Moreover, the inequality is strict: since $\tilde{G}_1(t) < \tilde{H}_1^*(t)$ on a set of positive measure, the decomposition gives $\tilde{G}_1(t) < F(t) \leq \tilde{F}(t)$ on that set, so FOSD monotonicity yields $V_p(m_1{=}p,\, c_p) > V_p(m_1{=}0,\, c_p)$. ■

Proof of Proposition 7. For each policy p, V_p depends on entrant messages only through c_p and on the incumbent's message only through its type relative to p. Count monotonicity (Lemma 8(a)) gives a cutoff on c_p in each case:

- $m_1 = p$: cutoff $\bar{c}_e$;
- $m_1 = p' \neq p, 0$: cutoff $\bar{c}_i$;
- $m_1 = 0$: cutoff $\bar{c}_0$.

By label symmetry, these cutoffs are policy-independent. By Lemma 8(b), V_p under endorsement weakly exceeds V_p under any other incumbent message at the same c_p, so $\bar{c}_e \leq \min(\bar{c}_i,\ \bar{c}_0)$.

Among eligible policies, the regulator selects the highest V_p. At equal entrant count c, for any non-endorsed policy $s \neq r$, Lemma 8(b) applied to policy r gives $V_r(m_1{=}r,\ c) > V_r(m_1{=}s,\ c)$. By exchangeability of policy labels, both $V_s(m_1{=}r,\ c)$ and $V_r(m_1{=}s,\ c)$ are computed from the same incumbent posterior $\tilde{H}_1^*$ and the same entrant posteriors G, H, so $V_s(m_1{=}r,\ c) = V_r(m_1{=}s,\ c)$. Hence $V_r^c > V_s^c$: the endorsed policy is strictly preferred. Symmetry of $\tilde{H}_1^*$ across non-endorsed policies gives $V_s^c = V_{s'}^c$ for $s, s' \neq r$. ■

Proof of Proposition 7 (existence). By Proposition 7, the regulator's strategy takes the count-cutoff form. It remains to verify entrant optimality and construct an equilibrium.

The equilibrium takes the following form. The incumbent always lobbies for its best innovative policy $p_1^{\max}$. Entrants lobby for the policy that maximizes their quality. The regulator plays the count-cutoff strategy from Proposition 7.

Entrant optimality. Lemma 7 identifies the form of a regular entrant strategy; here we verify that this strategy is optimal in the constructed equilibrium. Suppose $q_j^p > q_j^\ell$; we show recommending p is weakly preferred to ℓ. Let x denote the messages of the other $n-1$ firms, let ρ swap labels p and ℓ, and let $\alpha_a(b,x)$ denote the probability the regulator adopts policy a given $m_j = b$ and x. By label symmetry, $\Pr(x) = \Pr(\rho(x))$ and $\alpha_{\rho(a)}(\rho(b), \rho(x)) = \alpha_a(b,x)$.

When policy a is adopted, entrant j's expected payoff is $\psi(q_j^a, a, x) \equiv q_j^a \cdot \Pr(\max_{i \neq j} q_i^a < q_j^a \mid x)$, strictly increasing in q_j^a, with $\psi(q, \rho(a), \rho(x)) = \psi(q, a, x)$. Rewriting $U_j(\ell)$ via the substitution $x \mapsto \rho(x)$, $a \mapsto \rho(a)$ gives

$$U_j(p) - U_j(\ell) = \sum_x \Pr(x) \sum_a \alpha_a(p,x) \big[\psi(q_j^a, a, x) - \psi(q_j^{\rho(a)}, a, x)\big].$$

For $a \notin \{p, \ell\}$ the bracket vanishes. Applying the symmetry identities $\psi(q, \ell, x) = \psi(q, p, \rho(x))$ and $\alpha_\ell(p, \rho(x)) = \alpha_p(\ell, x)$ to combine the $a = p$ and $a = \ell$ terms yields

$$U_j(p) - U_j(\ell) = \sum_x \Pr(x)\ \underbrace{\big[\psi(q_j^p, p, x) - \psi(q_j^\ell, p, x)\big]}_{>0}\ \underbrace{\big[\alpha_p(p,x) - \alpha_p(\ell,x)\big]}_{\geq 0},$$

where the second factor is non-negative because changing j's message from ℓ to p raises p's recommendation count by one; by Lemma 8(a) this strictly raises the regulator's continuation value from adopting p while leaving all other policies' values unchanged, so under the regulator's argmax rule the set of profiles at which p is chosen weakly expands. Hence recommending $\operatorname{argmax}_p q_j^p$ is optimal.

Existence by construction. Let the incumbent always recommend $p_1^{\max}$, so $m_1 = 0$ is off path. This strategy trivially satisfies Definition 1(i), and condition (ii) is vacuous since $\Pr(m_1{=}0) = 0$. Assign the off-path belief and continuation strategy so that the regulator responds to $m_1 = 0$ exactly as to $m_1 = 1$ (admissible in PBE).

On-path posteriors: $\tilde{G}_1(t) = F(t)^k$ and $\tilde{H}_1^*(t) = [kF(t) - F(t)^k]/(k-1)$, with $\Pr(m_1 = r) = 1/k$. Regulator optimality and the count-cutoff form (cases 1–2 of Proposition 7) follow from the on-path posteriors; case 3 is vacuous since $m_1 = 0$ is off path. Entrant optimality follows from the argument above.

Incumbent optimality. Under this strategy, recommendation occurs a.s., so $\tilde{G}_1 = G$ and $\tilde{H}_1^* = H$. For brevity, write

$$w(c) \equiv \int_0^\infty \big[1 - G(t)^c\, H(t)^{n-c}\big]\, dt, \qquad c = 0, 1, \ldots, n,$$

for the regulator's expected peak quality of a policy with c recommenders out of n firms. Since replacing an H-draw by a G-draw strictly raises expected peak quality, w is strictly increasing. Hence $V_r^{c_r} = w(c_r + 1)$ and $V_s^{c_s} = w(c_s)$ for $s \neq r$.

(i) Deviation to a different innovative policy. Suppose $q_1^r > q_1^s$ and the incumbent deviates from $m_1 = r$ to $m_1 = s$. Let ρ swap labels r and s. By exchangeability, $\Pr(\boldsymbol{c}) = \Pr(\rho(\boldsymbol{c}))$, and by label symmetry of the regulator's rule,

$$\text{choice}(m_1{=}s,\ \rho(\boldsymbol{c})) = \rho(\text{choice}(m_1{=}r,\ \boldsymbol{c})).$$

If policy p is adopted at entrant count c_p, the incumbent's payoff is $\tilde{\phi}(q_1^p, c_p) \equiv q_1^p\, G(q_1^p)^{c_p}\, H(q_1^p)^{n-1-c_p}$, strictly increasing in q_1^p. Substituting $\boldsymbol{c} \mapsto \rho(\boldsymbol{c})$ in $U(s)$ and applying the symmetry relation: if $a \equiv \text{choice}(m_1{=}r,\ \boldsymbol{c})$, then $\rho(a)$ is adopted under $(m_1{=}s,\ \rho(\boldsymbol{c}))$ with the same entrant count $c_{\rho(a)}(\rho(\boldsymbol{c})) = c_a(\boldsymbol{c})$, giving

$$U(r) - U(s) = \sum_{\boldsymbol{c}} \Pr(\boldsymbol{c})\, \big[\tilde{\phi}(q_1^a, c_a) - \tilde{\phi}(q_1^{\rho(a)}, c_a)\big].$$

Each summand is positive when $a = r$ (since $q_1^r > q_1^s$), negative when $a = s$, and zero when $a \notin \{r, s\}$.

To sign the sum, we show: *if s is adopted at $\boldsymbol{c}$, then r is adopted at $\rho(\boldsymbol{c})$.* If s is adopted at $\boldsymbol{c}$ under $m_1 = r$, then $w(c_s) \geq w(c_r+1)$ (since s beats r) and $w(c_s) \geq w(c_t)$ for every $t \notin \{r,s\}$ (since s beats all rivals), so $c_s > c_r$ and $c_s \geq c_t$. Since s is adopted over the status quo, $w(c_s) - \bar{h} \geq q^{\mathrm{SQ}}$. At $\rho(\boldsymbol{c})$, policy r has c_s entrant recommenders and value $w(c_s+1)$, policy s has value $w(c_r)$, and every $t \notin \{r,s\}$ has value $w(c_t)$. Therefore

$$w(c_s+1) > w(c_s) \geq \max\big\{w(c_r+1),\, w(c_t) : t \notin \{r,s\}\big\},$$

and $w(c_s+1) - \bar{h} > w(c_s) - \bar{h} \geq q^{\mathrm{SQ}}$. Hence r is the unique maximizer of the regulator's objective at $\rho(\boldsymbol{c})$ and is adopted.

Since ρ is an involution and $c_s > c_r$ whenever s is adopted at $\boldsymbol{c}$, such profiles satisfy $\boldsymbol{c} \neq \rho(\boldsymbol{c})$. Therefore every negative summand (s adopted at $\boldsymbol{c}$) is paired with the summand at $\rho(\boldsymbol{c})$ (r adopted with the same entrant count c_s), and they cancel:

$$\big[\tilde{\phi}(q_1^s, c_s) - \tilde{\phi}(q_1^r, c_s)\big] + \big[\tilde{\phi}(q_1^r, c_s) - \tilde{\phi}(q_1^s, c_s)\big] = 0.$$

All remaining summands have $a = r$ or $a \notin \{r,s\}$, contributing non-negative terms. Hence $U(r) \geq U(s)$.

(ii) Deviation to the status quo. Under the off-path belief, $m_1 = 0$ is treated as $m_1 = 1$. If $p_1^{\max} = 1$, the deviation is payoff-equivalent to on-path play. If $p_1^{\max} = r \neq 1$, it is equivalent to deviating from r to 1, which is weakly unprofitable by (i) since $q_1^r \geq q_1^1$. ■

Proof of Corollary 5. Fix a well-behaved equilibrium with $\Pr(m_1 = 0) > 0$. By Proposition 7, the regulator uses a count-cutoff strategy. Since entrants' strategies depend only on their own quality vectors, the distribution of entrant recommendations is independent of the incumbent's realized qualities. Hence, conditional on the incumbent's message, the adoption probability of each policy depends only on that message.

Fix $q_1^{-\max}$ and let $r = p_1^{\max}$. By label symmetry, define

$$\begin{aligned} \pi_e &\equiv \Pr(\text{adopt } r \mid m_1 = r), \\ \pi_i &\equiv \Pr(\text{adopt } s \mid m_1 = r) \qquad \text{for any } s \neq r, \\ \pi_0 &\equiv \Pr(\text{adopt } p \mid m_1 = 0) \qquad \text{for any } p \geq 1. \end{aligned}$$

These probabilities do not depend on the incumbent's realized quality vector. By assumption, the incumbent's endorsement strictly affects the adoption probability of the endorsed policy, i.e., $\pi_e \neq \pi_0$. This implies $\pi_e > \pi_0$ because $\pi_e \geq \pi_0$ always holds weakly. To see the weak inequality, we show that if r is adopted under $m_1 = 0$ at some count vector $\boldsymbol{c}$, then r is also adopted under $m_1 = r$ at the same $\boldsymbol{c}$. Indeed, if r is adopted under $m_1 = 0$, then $c_r \geq \bar{c}_0 \geq \bar{c}_e$, so r is eligible under endorsement. Moreover, c_r must weakly exceed the count of every other eligible policy under $m_1 = 0$. Any non-endorsed policy s that is eligible under $m_1 = r$ (with $c_s \geq \bar{c}_i$) either was also eligible under $m_1 = 0$ (so $c_s \leq c_r$) or became newly eligible (so $c_s < \bar{c}_0 \leq c_r$). In either case $c_s \leq c_r$, so by count monotonicity and the endorsement advantage (Proposition 7), $V_r^{c_r} \geq V_r^{c_s} > V_s^{c_s}$. Hence r is adopted under $m_1 = r$. Summing over $\boldsymbol{c}$ gives $\pi_e \geq \pi_0$.

Now consider the incumbent's payoff difference Δ between recommending r and recommending the status quo. For each $\boldsymbol{c}$, the regulator's adoption decision depends on $\boldsymbol{c}$ and the incumbent's message but not on the incumbent's realized quality vector. For fixed $q_1^{-\max}$, only the incumbent's payoff from the endorsed policy r being adopted varies with $q_1^{\max}$, and this payoff is strictly increasing in $q_1^{\max}$. Since $\pi_e > \pi_0$, there exist count vectors at which r is adopted under $m_1 = r$ but not under $m_1 = 0$. At such count vectors the incumbent's gain from endorsement is strictly increasing in $q_1^{\max}$, so Δ is strictly increasing in $q_1^{\max}$ for fixed $q_1^{-\max}$. The set of $q_1^{\max}$ for which the incumbent prefers to recommend innovation is therefore an upper interval, defining the threshold function κ. ■

Example 2: Equilibrium Verification

Proposition 8. *The strategy profile described in Example 2 constitutes an equilibrium for* $\bar{h} \in (0.751, 0.762)$.

Proof. Set $k = 2$, $n = 3$, $F = \text{Uniform}[0,1]$, $q^{\text{SQ}} = 0.05$. The entrant and incumbent constraints admit closed-form verification; the regulator's posterior values are checked numerically.

Entrant IC. By symmetry across policies, each entrant's payoff from recommending p depends on p only through q_j^p and is increasing in that quality. Hence recommending $\text{argmax}\{q_j^1, q_j^2\}$ is optimal.

Incumbent IC. Fix type $(q_1^1, q_1^2) = (a, b)$ with $a \geq b$, and suppose the incumbent endorses policy 1. Under $F = \text{Uniform}[0,1]$, the entrant posteriors are $G(x) = x^2$ (recommended policy) and $H(x) = 2x - x^2$ (other policy). Policy 1 is adopted unless both entrants recommend 2. Let $A_j = \{q_j^1 \geq q_j^2\}$ and $B_j = \{q_j^1 \leq a\}$. The incumbent wins under policy 1 iff $(A_2 \cup A_3) \cap B_2 \cap B_3$, giving

$$\Pr(\text{win under 1}) = a^2 - (a - \tfrac{a^2}{2})^2 = a^3 - \tfrac{a^4}{4}.$$

Similarly, $\Pr(\text{win under 2}) = b^4/4$. Hence endorsing the better policy yields

$$\Pi(a, b) = a^4 - \tfrac{a^5}{4} + \tfrac{b^5}{4}. \tag{10}$$

Policy choice. Since $\Pi(a,b) - \Pi(b,a) = \frac{1}{2}[g(a) - g(b)]$ with $g(x) = x^4(2-x)$ strictly increasing on $[0,1]$, the incumbent prefers to endorse its higher-quality policy.

Innovation versus status quo. Define $\kappa(b)$ by $\Pi(\kappa, b) = q^{\mathrm{SQ}}$:

$$\kappa^4 - \tfrac{\kappa^5}{4} + \tfrac{b^5}{4} = q^{\mathrm{SQ}}. \tag{11}$$

Since Π is strictly increasing in a, the incumbent prefers innovation iff $a > \kappa(b)$. Implicit differentiation gives $\kappa'(b) = -5b^4/[\kappa^3(16-5\kappa)] < 0$, so the threshold is strictly decreasing in the non-maximal quality. The incumbent therefore follows the proposed rule: $m_1 = p$ iff $q_1^p \geq \max\{q_1^{-p}, \kappa(q_1^{-p})\}$.

Regulator IC. Given the proposed strategies, six on-path message configurations arise. For each, the regulator's expected peak quality $E[\max_j q_j^p \mid \boldsymbol{m}]$ is computed using G, H, and the incumbent posteriors implied by κ. The resulting continuation values, verified numerically, are:

m_1	(m_2, m_3)	$E[\max q^p]$ (best p)	Eq. action
0	both for p	0.801	SQ
0	split	0.705	SQ
0	both against p	0.494	SQ
1	$(1,1)$	0.874 (p=1)	$p = 1$
1	split	0.841 (p=1)	$p = 1$
1	$(2,2)$	0.812 (p=2)	$p = 2$

In each case, the proposed action selects the policy with the highest continuation value. After $m_1 = 1$, cross-policy comparisons confirm this: $0.874 > 0.563$, $0.841 > 0.731$, and $0.812 > 0.790$.

The binding constraints are: (i) after $m_1 = 0$, the regulator rejects iff $\bar{h} \geq 0.801 - q^{\mathrm{SQ}} = 0.751$; and (ii) in the override case (m_1=1, m_2=m_3=2), the regulator approves iff $\bar{h} < 0.812 - q^{\mathrm{SQ}} = 0.762$. Both hold for $\bar{h} \in (0.751, 0.762)$.

Off-path beliefs. The only off-path event is an entrant sending $m_j = 0$. Let p^* denote the policy entrant j would recommend on path, so $q_j^{p^*} \geq q_j^{-p^*}$. Assign the off-path belief $(q_j^1, q_j^2) \sim F^2$, so the message is uninformative about entrant j's quality: its posterior for each policy is F, replacing the on-path posteriors G (recommended policy) and H (other policy). Since G FOSD-dominates F, which FOSD-dominates H, and $E[\max_i q_i^p \mid \boldsymbol{m}]$ is increasing in each firm's quality distribution under FOSD, the deviation weakly lowers the continuation value of policy p^* and weakly raises that of $-p^*$.

If $m_1 \neq 0$, innovation is adopted on path. Any best response of the regulator after the deviation selects an action weakly less favorable to entrant j: the chosen policy can shift from p^* to $-p^*$ or to SQ, but never toward p^*. Since the entrant's expected profit conditional on adoption is increasing in its own quality and $q_j^{p^*} \geq q_j^{-p^*}$, entrant j weakly prefers p^* to $-p^*$, and prefers either to SQ (payoff 0). Hence the deviation is weakly unprofitable.

If $m_1 = 0$, SQ is adopted on path. Although the deviation raises the continuation value of $-p^*$, it cannot induce the regulator to approve innovation. For any policy p, the deviating entrant's posterior F and the other entrant's posterior $K \in \{G, H\}$ are each FOSD-dominated by G, so replacing either with G can only raise $E[\max_i q_i^p \mid \boldsymbol{m}]$. The upper bound—both entrants having posterior G for policy p—is exactly the "both for p" row in the table, with value 0.801. Since $0.801 - q^{\text{SQ}} = 0.751 < \bar{h}$, the regulator still prefers SQ after any deviation with $m_1 = 0$. Hence deviating to $m_j = 0$ is weakly unprofitable in all cases. ■