Online Appendix to "Endogenous Model Uncertainty and Curiosity under Cognitive Limits"

Jimmy Lee*

November 2026

Click here for the latest version

This appendix collects the formal statements referred to in the text but not essential to the argument, grouped by the section of the text they support; each is followed by its proof. Proofs of the main-text results are in Appendix A of the paper. Throughout, ϕ_ν denotes the $\mathcal{N}(0,\nu^2)$ density, $h_k := b_k'\Sigma^{-1}b_k = \nu_k^{-2}$ the precision with which model k's direction can be read, and "the two-model comparison" the evidence constraint (5) restricted to two nodes, under which a node's evidence cap exceeds one half if and only if its log Bayes factor against the other is positive.

The evidence constraint (Section 2.2.3). Let a single alternative k be held, let $\pi = (\pi_0, \pi_k)$ be the prior the automatic assignment carries, and let deliberation choose how much of the alternative's prior plausibility to engage, $\tau \in [0,1]$ ($\tau = 0$: not entertained), and what fraction $\alpha \in [0,1]$ of the evidence to process. The weight it can hold is the corresponding posterior $p_k(\alpha,\tau) := \tau\pi_k\Lambda_k(x)^\alpha/(\pi_0 + \tau\pi_k\Lambda_k(x)^\alpha)$.

Lemma OA.1 (Partial processing)**.** *Under the prior that defines* (5)*, the weights reachable by deliberation are* $[0,\overline{W}_k(x)]$ *if* $\Lambda_k(x) > 1$ *and* $[0,\pi_k] \supset [0,\overline{W}_k(x)]$ *if* $\Lambda_k(x) < 1$*. With favorable evidence the cap follows from partial processing; with unfavorable evidence it is a restriction, since neglect would lift the weight back toward the prior.*

Proof of Lemma OA.1. Write $\Lambda := \Lambda_k(x)$. The posterior p_k is continuous on $[0,1]^2$, zero at $\tau = 0$, increasing in τ, and increasing (decreasing) in α if $\Lambda > 1$ ($\Lambda < 1$). Its maximum is $p_k(1,1) = \pi_k\Lambda/(\pi_0 + \pi_k\Lambda)$, which is (5) under the prior that defines it, if $\Lambda > 1$, and $p_k(0,1) = \pi_k$ if $\Lambda < 1$; every value in between is reached. □

Next replace the constraint by a penalty. Let the DM choose any $w_k \in [0,1]$ and pay $c_e\,(w_k - \overline{W}_k)^+$, $c_e \geq 0$, for weight in excess of the cap, so that her objective is $V_e(w_k,n) := V(w_k,n) - c_e(w_k - \overline{W}_k)^+$ with V as in (12). Define the largest average return to exceeding the cap,

$$\bar{M}_k \;:=\; \sup_{n\geq 0}\;\sup_{w\in(\overline{W}_k,1]}\;\frac{V(w,n) - V(\overline{W}_k,n)}{w - \overline{W}_k},$$

which is finite because $\psi_k(n) \leq 1/P_k(0)$.

Proposition OA.1 (A soft cap)**.** *(i) If* $c_e > \bar{M}_k$*, every optimum of* V_e *has* $w_k \leq \overline{W}_k$*, and the solution coincides with that of the capped problem; in particular the information threshold* $\underline{w}_k$ *of Corollary 1 and the curiosity failures of Section 3.2 are unchanged.*

(ii) In the graded regime, $\gamma > 2D_k$*, with* $c_e < M_k(0) := \Delta U_k - D_k(1 - 2\overline{W}_k) - \gamma\overline{W}_k$ *and* $n = 0$*, the optimal weight exceeds the cap by* $\min\{(M_k(0) - c_e)/(\gamma - 2D_k),\; 1 - \overline{W}_k\}$*, continuous and nonincreasing in* c_e*; as* $c_e \to 0$ *the weight is set by value alone, akin to unconstrained motivated beliefs.*

*Global Poverty Research Lab, Northwestern University. Email: jimmylee@u.northwestern.edu.

(iii) The information threshold $\underline{w}_k$ does not depend on c_e.

Proof of Proposition OA.1. (i) For any n and $w > \overline{W}_k$, $V_e(w,n) - V_e(\overline{W}_k,n) = [V(w,n) - V(\overline{W}_k,n)] - c_e(w - \overline{W}_k) \leq (\bar{M}_k - c_e)(w - \overline{W}_k) < 0$, so w is strictly dominated by $\overline{W}_k$ at the same n, and the maximization of V_e over $[0,1] \times \mathbb{Z}_{\geq 0}$ has the same value and solutions as the maximization of V over $[0,\overline{W}_k] \times \mathbb{Z}_{\geq 0}$. (ii) For $w > \overline{W}_k$ and $n = 0$, $\partial_w V_e = \Delta U_k - D_k(1-2w) - \gamma w - c_e$, which is strictly decreasing in w when $\gamma > 2D_k$ and equals $M_k(0) - c_e > 0$ at $w = \overline{W}_k$. The first-order condition gives $w = (\Delta U_k - D_k - c_e)/(\gamma - 2D_k)$, and projecting onto $[\overline{W}_k, 1]$ and subtracting $\overline{W}_k$ yields the stated excess, since $M_k(0) = \Delta U_k - D_k - (\gamma - 2D_k)\overline{W}_k$. As $c_e \to 0$ the condition becomes $\partial_w V = 0$, the unconstrained optimum. (iii) Whether the first signal is bought turns on $w_k \psi_k(1)/(2\lambda) \geq c$, which does not involve c_e (Corollary 1). □

Updating (Section 2.2.4).

Lemma OA.2 (Updating preserves the discipline)**.** *(i) If $w_k \leq \overline{W}_k(x)$ for every k, then the odds of every alternative against the default never exceed its Bayes factor, before or after the signals: $w_{k,n}/w_{0,n} \leq \Lambda_k(x)\Lambda_k^y(y^n)$ for every realization y^n. In particular, when a single alternative k is held against the default, $w_{k,n} \leq \Lambda_k(x)\Lambda_k^y(y^n)/\big(1 + \Lambda_k(x)\Lambda_k^y(y^n)\big)$, the updated bound of that two-model comparison: a weight that respects the evidence constraint before the signals respects it after them.*

(ii) Under the predictive distribution implied by the mixture $\mathbf{w}$, $\mathbb{E}_{\mathbf{w}}[\mathbf{w}_n] = \mathbf{w}$: the weights are a martingale, so acquiring signals cannot be expected to move attention weights in either direction, only to resolve it.

Proof of Lemma OA.2. (i) By (6), $w_{k,n}/w_{0,n} = (w_k/w_0)\Lambda_k^y(y^n)$. If $w_j \leq \overline{W}_j(x)$ for every j, then $w_0 \geq 1 - \sum_j \overline{W}_j(x) = L_0/(L_0 + \sum_j L_j)$, so $w_k/w_0 \leq \Lambda_k(x)$ and the odds bound follows. When a single alternative is held, $w_{k,n} = g(w_k)$ with $g(z) := z\Lambda_k^y/(1 - z + z\Lambda_k^y)$ strictly increasing on $[0,1]$; evaluating g at $w_k \leq \overline{W}_k(x) \leq \Lambda_k/(1+\Lambda_k)$ gives the bound. (ii) $w_{k,n}$ is the posterior probability of model k under the prior $\mathbf{w}$, so its expectation under the corresponding predictive distribution is w_k by the tower property. □

Beliefs and welfare (Section 2.3).

Remark OA.1 (The welfare loss from neglect). Under (8) with φ linear, suppose alternative k is true with posterior mean m_k and the DM acts on the default, $a_0 = (m_0 - \chi)/\lambda$. Her expected payoff is $a_0(m_k - \chi) - \frac{\lambda}{2}a_0^2$, against $(m_k - \chi)^2/(2\lambda)$ from the act the alternative recommends, and the difference is $(m_k - m_0)^2/(2\lambda) = D_k(x)$. The per-stage welfare loss from a curiosity failure when the alternative is right is therefore the disagreement D_k, to which is added the value $\psi_k(n)/(2\lambda)$ of the signals she would have bought under the alternative; when the default is right the loss is zero and the DM has saved the cognitive cost of holding the alternative. Under the Bayesian posterior the expected per-stage loss is $\overline{W}_k(x)D_k(x)$ plus the foregone information value weighted by $\overline{W}_k(x)$, which is what a planner weighs against the cost of lifting the cap through the initial signal. Because the constraint caps weight at the posterior, an intervention that supplies evidence can reduce neglect but cannot induce weight the evidence does not justify, so by this measure it weakly improves on the DM's choice whenever its cost is below that expected loss.

Alternative activation (Section 3.1).

Remark OA.2 (What the regimes depend on). The classification needs neither the quadratic payoff, nor the Gaussian structure, nor a linear φ. Write $M(w_k) := \max_a\{(1 - w_k)\varphi(\mathbb{E}_0 u) + w_k\varphi(\mathbb{E}_k u)\}$ for the value of the single act, so that V is $M(w_k)$ plus a term affine in w_k, less $\frac{\gamma}{2}w_k^2 + cn$. As a maximum of functions affine in w_k, M is convex, strictly so wherever the models' optimal acts differ, and its curvature measures how far apart those acts are. The objective is thus a difference of convex functions: where γ exceeds M's curvature the weight moves continuously, and where M's curvature exceeds γ it sits at 0 or at the bound. The quadratic case makes that curvature constant at $2D_k(x)$, which is what delivers the closed forms; the dichotomy needs only that one act be chosen before the model is known.

Corollary OA.1 (Flukes and the fragility of activation)**.** *In the all-or-nothing regime, no observation produces a small durable weight: an isolated favorable observation that raises ΔU_k or $\overline{W}_k$ without crossing $T_k(\overline{W}_k)$ leaves $w_k^* = 0$ exactly. In the graded regime, under Assumption 1, a perturbation raising ΔU_k just above D_k produces a weight below the information threshold of Corollary 1, so no signal is acquired and the weight does not grow.*

Proof of Corollary OA.1. In the all-or-nothing regime the optimum lies in $\{0, \overline{W}_k\}$, so no perturbation of $(\Delta U_k, \overline{W}_k)$ that keeps $\Delta U_k \le T_k(\overline{W}_k)$ moves it from zero. In the graded regime, $w_k^* = (\Delta U_k - D_k)/(\gamma - 2D_k) < \underline{w}_k$ whenever $\Delta U_k - D_k < (\gamma - 2D_k)\,\underline{w}_k$. Then $n^* = 0$ by Corollary 1, so nothing refines the alternative or moves its cap, and the next stage repeats the same problem. □

Complementarity and curiosity failures (Section 3.2). Let a single alternative k be held under Assumption 1 with φ linear, and let the stage repeat once. A *forward-looking* DM chooses (w_k, n) at stage 1 to maximize

$$V(w_k, n) \;+\; \beta\, \mathbb{E}_{w_k}\big[\phi(\overline{W}_{k,n}; n)\big], \qquad \phi(\overline{W}; n) := \max_{0 \le w' \le \overline{W}} V(w', n) + cn, \qquad \beta \in [0,1],$$

where V is (12), $\overline{W}_{k,n} = \Lambda_k \Lambda_k^y(y^n)/(1 + \Lambda_k \Lambda_k^y(y^n))$ is the stage-2 cap of (6), ϕ is the stage-2 value when the DM re-chooses her weight under the updated cap without further signals (the stage-1 signals are paid for once), and the expectation is under the stage-1 mixture. Model-contingent values enter stage 2 at their stage-1 expectations, so signals matter there only through the cap, the channel through which looking ahead could overturn a curiosity failure. Write $\underline{w}_k = 2\lambda \nu_y^2 c\, P_k(0) P_k(1)$ for the information threshold of Corollary 1.

Proposition OA.2 (Curiosity failures with general payoffs)**.** *Let φ be linear, let the payoff u be bounded, let a single alternative k be held at weight w, and let the subsequent signals be uninformative under the default, so that $\mathbb{E}_{0,n} u(a, \theta)$ does not depend on n (Assumption 1). Write $U_k(n) := \mathbb{E}_y^k \max_a \mathbb{E}_{k,n} u(a, \theta)$ and $B_k := \sup_n U_k(n) - U_k(0) < \infty$.*

(i) The value of the n signals at weight w, the change in (4) *gross of their cost, is $w[U_k(n) - U_k(0)] \le wB_k$.*

(ii) Hence $n^* = 0$ whenever $w \le c/B_k$. If $\overline{W}_k(x) \le c/B_k$, the alternative is never investigated whatever its value: a curiosity failure.

Proof of Proposition OA.2. (i) The single-act term of (4) does not depend on n, and under Assumption 1 neither does U_0, so the value of the signals is $w[U_k(n) - U_k(0)] \le wB_k$; B_k is finite because u is bounded. (ii) The gain from $n \ge 1$ signals is at most $wB_k - c \le \overline{W}_k(x) B_k - c$. □

Proposition OA.3 (Curiosity failures with a forward-looking DM)**.** *(i)* (The cap is a martingale.) *For every n, $\mathbb{E}_{f_0}[\overline{W}_{k,n}] \le \overline{W}_k$ under the default's density and $\mathbb{E}_k[\overline{W}_{k,n}] \ge \overline{W}_k$ under the alternative's; under the stage-1 mixture, $\mathbb{E}_{w_k}[\overline{W}_{k,n}]$ is increasing in w_k and equals $\overline{W}_k$ at $w_k = \overline{W}_k$.*

(ii) (Graded regime.) *If $\gamma > 2D_k$, then for every $w_k \le \overline{W}_k$ and every n, $\mathbb{E}_{w_k}[\phi(\overline{W}_{k,n}; 0)] \le \phi(\overline{W}_k; 0)$: relaxing the cap has no option value. The forward-looking DM acquires no signal whenever $\overline{W}_k \le \underline{w}_k/(1+\beta)$.*

(iii) (All-or-nothing regime.) *Let $\gamma < 2D_k$ and $\Delta U_k^\infty \le T_k(\overline{W}_k)$, where $\Delta U_k^\infty := \Delta U_k + 1/(2\lambda P_k(0))$ is the value gap under full refinement, so that the myopic DM holds the default alone whatever the number of signals. Then, at any stage-1 weight, $n \ge 1$ signals gain at most $\beta\, G_k(\Lambda_k/\Lambda_k^a + \overline{W}_k) - cn$ over the default alone, where $\Lambda_k^a \ge \Lambda_k$ is the odds at which the cap reaches $T_k^{-1}(\Delta U_k^\infty)$ if $\Delta U_k^\infty > \gamma/2$, and $\Lambda_k^a = \infty$ if $\Delta U_k^\infty \le \gamma/2$, when no cap makes the alternative worth activating; and $G_k := \sup_{\overline{W}, n} \phi(\overline{W}; n) - V(0,0)$. The forward-looking DM therefore stays with the default alone whenever $c \ge \beta G_k(\Lambda_k/\Lambda_k^a + \overline{W}_k)$.*

Proof of Proposition OA.3. Write $\Lambda^y := \Lambda_k^y(y^n)$ and $g(\Lambda^y) := \Lambda_k \Lambda^y/(1 + \Lambda_k \Lambda^y)$, so that $\overline{W}_{k,n} = g(\Lambda^y)$ with g increasing and concave and $g(1) = \overline{W}_k$.

(i) Λ^y is a likelihood ratio, so $\mathbb{E}_{f_0}[\Lambda^y]=1$ and $\mathbb{E}_k[1/\Lambda^y]=1$. Jensen's inequality gives $\mathbb{E}_{f_0}[g(\Lambda^y)]\le g(1)$ and, since $x\mapsto 1/(1+x/\Lambda_k)$ is convex, $\mathbb{E}_k[g(\Lambda^y)]\ge g(1)$. Under the mixture, $\mathbb{E}_{w_k}=(1-w_k)\mathbb{E}_{f_0}+w_k\mathbb{E}_k$ is increasing in w_k because $\mathbb{E}_k[g]\ge\mathbb{E}_{f_0}[g]$. At $w_k=\overline{W}_k=\Lambda_k/(1+\Lambda_k)$, (6) gives $w_{k,n}=g(\Lambda^y)=\overline{W}_{k,n}$, and $\mathbb{E}_{w_k}[w_{k,n}]=w_k$ by Lemma OA.2(ii).

(ii) When $\gamma>2D_k$, $V(\cdot,0)$ is strictly concave in w', so $\phi(\overline{W};0)=V(\min\{\overline{W},w^u\},0)$ with $w^u:=\max\{0,(\Delta U_k-D_k)/(\gamma-2D_k)\}$, and $\phi(\cdot;0)$ is nondecreasing and concave in $\overline{W}$. By (i), $\mathbb{E}_{w_k}[\overline{W}_{k,n}]\le\overline{W}_k$ for $w_k\le\overline{W}_k$, so Jensen's inequality and monotonicity give $\mathbb{E}_{w_k}[\phi(\overline{W}_{k,n};0)]\le\phi(\mathbb{E}_{w_k}[\overline{W}_{k,n}];0)\le\phi(\overline{W}_k;0)$. For the threshold, decompose the gain from n signals. Within stage 1, $V(w_k,n)-V(w_k,0)=w_k\psi_k(n)/(2\lambda)-cn\le w_k n\,\psi_k(1)/(2\lambda)-cn$, since ψ_k is concave with $\psi_k(0)=0$. In stage 2, $V(w',n)+cn=V(w',0)+w'\psi_k(n)/(2\lambda)$, so $\phi(\overline{W};n)\le\phi(\overline{W};0)+\overline{W}n\,\psi_k(1)/(2\lambda)$, and taking expectations with $\mathbb{E}_{w_k}[\overline{W}_{k,n}]\le\overline{W}_k$ and the first part of (ii), $\mathbb{E}_{w_k}[\phi(\overline{W}_{k,n};n)]-\phi(\overline{W}_k;0)\le\overline{W}_k n\,\psi_k(1)/(2\lambda)$. The total gain is therefore at most $n[(w_k+\beta\overline{W}_k)\,\psi_k(1)/(2\lambda)-c]\le n[(1+\beta)\overline{W}_k\,\psi_k(1)/(2\lambda)-c]$, which is nonpositive for every n when $\overline{W}_k\le\underline{w}_k/(1+\beta)$, since $\psi_k(1)/(2\lambda)=[2\lambda\nu_y^2P_k(0)P_k(1)]^{-1}=c/\underline{w}_k$.

(iii) Write $\Delta U_k(n):=\Delta U_k+\psi_k(n)/(2\lambda)$. Since $\psi_k(n)<1/P_k(0)$, $\Delta U_k(n)<\Delta U_k^\infty\le T_k(\overline{W}_k)$ for every n, so by Proposition 1(ii)(b) the alternative is not worth activating at the stage-1 cap whatever the signals: $\phi(\overline{W}_k;n)=V(0,n)+cn=V(0,0)$ for every n. Fix any $w_k\le\overline{W}_k$ and n. The stage-1 term satisfies $V(w_k,n)-V(0,0)\le\phi(\overline{W}_k;n)-cn-V(0,0)=-cn$. In stage 2, $\phi(\overline{W};n)=V(0,0)$ unless the cap activates the alternative, which requires $\overline{W}>T_k^{-1}(\Delta U_k(n))$; if $\Delta U_k^\infty\le\gamma/2=T_k(1)$, no cap activates it; otherwise, T_k^{-1} being decreasing on $(\gamma/2,D_k]$, activation requires $\overline{W}_{k,n}\ge\overline{W}^a:=T_k^{-1}(\Delta U_k^\infty)$, that is, $\Lambda^y\ge\Lambda_k^a/\Lambda_k$ with $\Lambda_k^a:=\overline{W}^a/(1-\overline{W}^a)$, and $\overline{W}^a\ge\overline{W}_k$ gives $\Lambda_k^a\ge\Lambda_k$. Also $\phi-V(0,0)\le G_k$. Under the stage-1 mixture, whose density is $f_0(1-w_k+w_k\Lambda^y)$, Markov's inequality and $\mathbb{E}_{f_0}[\Lambda^y]=1$ give $\Pr_{w_k}(\Lambda^y\ge\Lambda_k^a/\Lambda_k)\le(1-w_k)\Lambda_k/\Lambda_k^a+w_k\le\Lambda_k/\Lambda_k^a+\overline{W}_k$. The total gain is therefore at most $\beta G_k(\Lambda_k/\Lambda_k^a+\overline{W}_k)-cn$; with $n=0$ the cap does not move and the gain is $V(w_k,0)-V(0,0)\le 0$. □

Marginal likelihood (Section 3.3).

Lemma OA.3 (Marginal likelihood)**.** *Under* (1)*, $x\sim\mathcal{N}\left(\mu_k b_k,\ \Sigma+\sigma_k^2 b_k b_k'\right)$ under mental model k, so $L_k(x)$ is the density of that distribution evaluated at x.*

Proof of Lemma OA.3. $x=\theta b_k+\varepsilon$ is Gaussian with mean $\mu_k b_k$ and covariance $\Sigma+\sigma_k^2 b_k b_k'$. □

Endogenous exposure (Section 3.6).

Proposition OA.4 (Ambiguity aversion tilts search toward discriminating signals)**.** *Under the criterion of Proposition 6, let alternative k be held and let $N:=\pi-\mathbb{E}_{\mathbf{w}}[\pi_n]$ be the non-instrumental component of a signal's value.*

(i) A signal that moves only the weight, discriminating between the models without refining either, has $N\ge 0$ for every strictly concave φ.

(ii) A signal that refines both models without moving the weight, making $(U_{0,n},U_{k,n})$ a mean-preserving spread of (U_0,U_k), has, to second order, $N=\frac{\xi}{2}w_k(1-w_k)\left[(\Delta U_k)^2-\mathbb{E}(\Delta U_{k,n})^2\right]\le 0$.

Proof of Proposition OA.4. (i) With predictions unchanged, $\pi_n=\pi(w_{k,n})$ for the function π of Proposition 6(ii), which is concave, and $\mathbb{E}_{\mathbf{w}}[w_{k,n}]=w_k$ (Lemma OA.2), so Jensen's inequality gives $\mathbb{E}_{\mathbf{w}}[\pi(w_{k,n})]\le\pi(w_k)$. (ii) Apply the second-order form of the premium at the fixed weight; a mean-preserving spread raises $\mathbb{E}(\Delta U_{k,n})^2$. □

Robustness (Section 3.7). Fix an alternative j and a coordinate e of the initial signal that only j reads: $\alpha_{k,e}=0$ for every $k\ne j$, including the default, and $\alpha_{j,e}\ne 0$, with α_k the attention mapping of Lemma 1.

Lemma OA.4 (Varying plausibility and recommendation separately)**.** *Varying x_e leaves $\Lambda_k(x)$ unchanged for every $k \neq j$, and μ_j enters no likelihood other than L_j and no posterior mean other than m_j. The map $(x_e, \mu_j) \mapsto (\log \Lambda_j(x), m_j)$ has Jacobian determinant*

$$\frac{\alpha_{j,e}\,\tau_j\,\mu_j}{1+\tau_j\nu_j^2}.$$

Hence, whenever $\mu_j \neq 0$, joint local adjustments of x_e and μ_j move the plausibility of model j while holding its recommended act $(m_j-\chi)/\lambda$ fixed, and move its recommended act while holding its plausibility fixed, leaving the default's act and every other Bayes factor unchanged.

Proof of Lemma OA.4. For $k \neq j$, $s_k = \alpha_k' x$ does not involve x_e. By (13), $\log\Lambda_k = \frac{1}{2}[(R_0 - R_k) + (q_0 - q_k) + (c_0 - c_k)]$, where $R_0 - R_k = s_k^2/\nu_k^2 - s_0^2/\nu_0^2$, the q's depend on x only through the s's, and the c's are constants; so $\log\Lambda_k$ is free of x_e. Likewise, $\log\Lambda_j$ depends on (x_e, μ_j) only through $\frac{1}{2}[s_j^2/\nu_j^2 - (s_j-\mu_j)^2/(\sigma_j^2+\nu_j^2)]$, with $\partial s_j/\partial x_e = \alpha_{j,e}$. Writing $d := (s_j-\mu_j)/(\sigma_j^2+\nu_j^2)$ and $P := P_j(0)$,

$$\frac{\partial \log\Lambda_j}{\partial x_e} = \alpha_{j,e}\Big(\frac{s_j}{\nu_j^2} - d\Big), \quad \frac{\partial \log\Lambda_j}{\partial \mu_j} = d, \quad \frac{\partial m_j}{\partial x_e} = \frac{\alpha_{j,e}}{\nu_j^2 P}, \quad \frac{\partial m_j}{\partial \mu_j} = \frac{\tau_j}{P}.$$

The determinant is $(\alpha_{j,e}/P)\,[\tau_j s_j/\nu_j^2 - dP]$. Since $P/(\sigma_j^2+\nu_j^2) = \tau_j/\nu_j^2$, $dP = \tau_j(s_j-\mu_j)/\nu_j^2$, so the determinant equals $\alpha_{j,e}\tau_j\mu_j/(\nu_j^2 P) = \alpha_{j,e}\tau_j\mu_j/(1+\tau_j\nu_j^2)$. When it is nonzero, the inverse function theorem gives the local adjustments. □

Interference effects (Section 4.2).

Corollary OA.2 (Curiosity and the number of rivals)**.** *Under Assumption 1, fix x and let the menu contain K alternatives that are symmetric given x, each with likelihood L, value gap ΔU, and disagreement $D > 0$, so that each bound is $\overline{W}(K) = L/(L_0 + KL)$, strictly decreasing in K and tending to zero. Curiosity A and information acquired n^* are zero at $K = 0$ and positive at $K = 1$ whenever the alternative is activated at a weight above the level in Corollary 1.*

(i) If $\sup_n \Delta U < \min\{D, \eta/2\}$, then beyond a finite K no alternative is held, so $A = 0$ and $n^* = 0$.

(ii) If $\Delta U > \min\{D, \eta/2\}$, rivals that make identical predictions are held at a total weight bounded away from zero. Curiosity then stays bounded away from zero if and only if $D < \eta/2$ and $\Delta U < \eta - D$; otherwise the total weight tends to one and curiosity to zero, as the pooled rivals displace the default.

(iii) Part (i) does not need symmetry. For any menu, let $\bar{D} := \max_k D_k$ and $a := \sup_n \max_k(\Delta U_k - D_k)$. If $a < 0$ and $a + \bar{D} < \eta/2$, then no alternative is held once $\max_k \overline{W}_k$ is small enough, as when the rivals' total likelihood $\sum_k L_k$ grows without bound.

Proof of Corollary OA.2. With K symmetric alternatives, (5) gives $\overline{W}(K) = L/(L_0 + KL)$, strictly decreasing in K with limit zero. At $K = 0$, $A = n^* = 0$; at $K = 1$, an activated alternative has $0 < w^* < 1$, so $A > 0$, and $n^* > 0$ once w^* exceeds the level in Corollary 1.

For general K, fix n and write $W := \sum_{k\geq 1} w_k$ and $d_k := m_k - m_0$, so that $d_k^2 = 2\lambda D$ for every k. Since $\sum_k w_k d_k^2 = 2\lambda D W$ and $(\sum_k w_k d_k)^2 \leq (\sum_k w_k |d_k|)^2 = 2\lambda D W^2$, the hedging cost satisfies $A(\mathbf{w}) \geq W(1-W)D$, with equality when the rivals make identical predictions. By the identity in the proof of Lemma 2 and $\sum_k w_k^2 \leq W \max_k w_k \leq W\,\overline{W}(K)$, $C(\mathbf{w}) \geq \frac{\eta}{2}W^2 - \frac{\eta-\gamma}{2}W\,\overline{W}(K)$. Hence, by (11),

$$V(\mathbf{w}, n) - V(0, n) \;\leq\; W\big[\Delta U - D + \tfrac{\eta-\gamma}{2}\overline{W}(K)\big] + W^2\big(D - \tfrac{\eta}{2}\big).$$

The right side is a quadratic in W that vanishes at zero. Its slope at zero is $\Delta U - D + \frac{\eta-\gamma}{2}\overline{W}(K)$ and its value at $W = 1$ is $\Delta U - \frac{\eta}{2} + \frac{\eta-\gamma}{2}\overline{W}(K)$. If $\Delta U^\infty < \min\{D, \eta/2\}$, both are negative for every n once $\overline{W}(K) <$

$2(\min\{D,\eta/2\}-\Delta U^{\infty})/(\eta-\gamma)$, which holds beyond a finite K, and a quadratic that vanishes at zero, has negative slope there, and is negative at one is negative on $(0,1]$. Then $\mathbf{w}^*=0$ at every n, so $A=0$, and $n^*=0$ under Assumption 1 (Proposition 2).

Conversely, if the rivals make identical predictions, the hedging cost is $W(1-W)D$ however W is split and the cost of attention weights is within $\frac{\eta-\gamma}{2}W\overline{W}(K)$ of $\frac{\eta}{2}W^2$, so the best value with total W is within $\frac{\eta-\gamma}{2}\overline{W}(K)$ of $V(0,n)+g(W)$, $g(W):=W(\Delta U-D)+W^2(D-\frac{\eta}{2})$, and the feasible totals increase to $[0,1)$. Since $\Delta U>\min\{D,\eta/2\}$, $\max g>0$, so W_K^* is bounded away from zero and its limit points maximize g. If $D<\eta/2$ and $D<\Delta U<\eta-D$, g has the unique interior maximizer $W^\circ=(\Delta U-D)/(\eta-2D)$, so $A\to W^\circ(1-W^\circ)D>0$; otherwise g is uniquely maximized at $W=1$ (convex with $g(1)=\Delta U-\eta/2>0$ if $D\geq\eta/2$, nondecreasing on $[0,1]$ if not), so $W_K^*\to 1$ and $A\to 0$.

For part (iii), write $\varepsilon:=\max_k\overline{W}_k$. The Cauchy–Schwarz inequality $(\sum_k w_k|d_k|)^2\leq W\sum_k w_k d_k^2$ gives $A(\mathbf{w})\geq(1-W)\sum_k w_k D_k$ for any menu, and the cost of attention weights is at least $\frac{\eta}{2}W^2-\frac{\eta-\gamma}{2}W\varepsilon$ as above, so $V(\mathbf{w},n)-V(0,n)\leq W[a+\frac{\eta-\gamma}{2}\varepsilon]+W^2(\bar{D}-\frac{\eta}{2})$. For small ε this quadratic has negative slope at zero and is negative at $W=1$, hence negative on $(0,1]$; so $\mathbf{w}^*=0$ at every n, and $n^*=0$ under Assumption 1. By (5), $\varepsilon\to 0$ as $\sum_j L_j\to\infty$ with each L_k bounded; in the symmetric case $a+\bar{D}=\Delta U$, which recovers part (i). □

Model an informational treatment as a budget $B\geq 0$ allocated across alternatives, $\beta=(\beta_1,\ldots,\beta_K)$ with $\beta_k\geq 0$ and $\sum_k\beta_k=B$, which multiplies L_k by e^{β_k}. This normalization holds fixed the total log-likelihood content delivered and varies only its concentration.

Proposition OA.5 (The ceiling on untargeted information)**.** *Let $\overline{W}_k(\beta)$ denote the post-treatment bound* (5)*.*

(i) (Targeting.) *Order the alternatives so that $L_1=\max_k L_k$. For fixed B, $\max_k\overline{W}_k(\beta)$ is maximized at the vertex $\beta_1=B$. Along that allocation $\overline{W}_1$ is strictly increasing in B and $\overline{W}_1(B)\to 1$ as $B\to\infty$.*

(ii) (Breadth.) *Under the uniform allocation $\beta_k=B/K$, $\max_k\overline{W}_k$ is strictly increasing in B but*

$$\sup_{B\geq 0}\max_k\overline{W}_k(B)\;=\;\frac{\max_k L_k}{\sum_{j=1}^K L_j}\;<\;1.$$

(iii) (Non-comparability.) *Consequently, for any activation threshold $T\in\big(\max_k L_k/\sum_j L_j,\ 1\big)$, a targeted treatment activates some alternative for B large enough, while a uniform treatment fails to activate any alternative at every budget, however large.*

Proof of Proposition OA.5. Write $S:=\sum_{j\geq 1}e^{\beta_j}L_j$, so that $\overline{W}_k=e^{\beta_k}L_k/(L_0+S)$. (i) $\overline{W}_1$ is strictly increasing in β_1 and strictly decreasing in each β_j, $j\neq 1$, so subject to $\sum_j\beta_j=B$ it is maximized at $\beta_1=B$, where $\overline{W}_1(B)=e^B L_1/(L_0+e^B L_1+\sum_{j\geq 2}L_j)\to 1$. (ii) Under $\beta_k=B/K$, $\overline{W}_k(B)=L_k/(L_0e^{-B/K}+\sum_j L_j)$ is strictly increasing in B with supremum $L_k/\sum_j L_j$, not attained and strictly below one when $K\geq 2$ and every $L_j>0$. (iii) Immediate from (i) and (ii). □

Proposition OA.6 (Stacking premium)**.** *Holding two alternatives at once at weights w_1,w_2, rather than in two separate episodes, costs a premium of $w_1w_2[\eta-(m_1-m_0)(m_2-m_0)/\lambda]$, the first term from attention weights and the second from hedging, which rises when the two mental models depart from the default in opposite directions and falls when they depart in the same direction. Absent discounting, sequencing therefore weakly dominates stacking whenever $\eta\geq(m_1-m_0)(m_2-m_0)/\lambda$, in particular whenever the departures have opposite signs.*

Proof of Proposition OA.6. Two separate episodes yield $f_1(w_1)+f_2(w_2)$ in the notation of the proof of Proposition 9(iii); holding both at once adds the pairwise term $w_1w_2[d_1d_2/\lambda-\eta]$, whose negative is the premium. □

Corollary OA.3 (Eliminating rivals one at a time)**.** *If multiple alternatives are plausible after x and one clears its threshold, it is investigated first; if the evidence goes against it, its likelihood falls and by Proposition 9(i) every other alternative's bound rises, so a second alternative may clear its threshold in turn. If no alternative is worth holding after x, alone or jointly, as under the conditions of Proposition 9(iii), no investigation starts, no likelihood changes, and no bound moves: the joint failure of Proposition 9(iii) is absorbing until a new exogenous observation arrives.*

Proof of Corollary OA.3. Under Assumption 1 no information is acquired about an alternative without weight (Corollary 1). If the signals go against the one investigated, $\Lambda_k^y \to 0$ by Proposition 4(b)(ii), so L_k falls in the denominator of (5) and every $\overline{W}_j$, $j \neq k$, rises (Proposition 9(i)). If $\mathbf{w}^* = 0$ after x, then $n^* = 0$ and the bounds stay as they were. □

Creative thinking (Sections 5.2 and 5.3).

Corollary OA.4 (What cannot move decisions can justify invention)**.** *Let ζ be a component of the initial signal that node S does not read, $B_S'\Sigma^{-1}\zeta = 0$. Then:*

(i) The posterior over the parameters of S, and so every decision made under S, is the same at x and at $x+\zeta$.

(ii) For any template t, ζ shifts the standardized surprise along t by $b_t'\Sigma^{-1}\zeta/\sqrt{h_t}$, which is nonzero whenever t reads ζ.

Proof of Corollary OA.4. (i) Given the parameters ϑ of S, the log likelihood of x is $\vartheta' B_S'\Sigma^{-1}x - \frac{1}{2}\vartheta' B_S'\Sigma^{-1}B_S\vartheta$ up to a term free of ϑ, so the posterior under S depends on x only through $B_S'\Sigma^{-1}x$, which ζ leaves unchanged. (ii) ζ shifts the residual r by ζ and hence z_t by $b_t'V_S^{-1}\zeta/\sqrt{h_t}$. By the Woodbury identity, $V_S^{-1} = \Sigma^{-1} - \Sigma^{-1}B_S(\Sigma_S^{-1} + B_S'\Sigma^{-1}B_S)^{-1}B_S'\Sigma^{-1}$, whose second term vanishes on ζ because $B_S'\Sigma^{-1}\zeta = 0$. The shift is therefore $b_t'\Sigma^{-1}\zeta/\sqrt{h_t}$, nonzero exactly when t reads ζ. □

Proposition OA.7 (Horizontal step, general ρ)**.** *With reading noise v, let $a := v + (1-\rho)\sigma^2$ and $g := v + \sigma^2 + (N-1)\rho\sigma^2$. The Bayes factor of the general node against the local node is*

$$\log \Lambda_{\text{gen:loc}} = \underbrace{\tfrac{1}{2}\Big[N\log(v+\sigma^2) - (N-1)\log a - \log g\Big]}_{\textit{Occam gain: a shared parameter for N situations}} - \tfrac{1}{2}SS\Big[\frac{1}{a} - \frac{1}{v+\sigma^2}\Big] + \tfrac{N}{2}\bar{d}^2\Big[\frac{1}{v+\sigma^2} - \frac{1}{g}\Big]. \tag{OA.1}$$

It is zero when $N = 1$. For $N \geq 2$ and $\rho > 0$, it is increasing in the shared departure $\bar{d}^2$ and decreasing in the dispersion SS, and the general node is admitted if and only if SS falls below a threshold that is increasing in $\bar{d}^2$. At $SS = \bar{d} = 0$ it is increasing in ρ. At $\rho = 1$ the Occam gain is increasing in N. At $\rho = 1$ and $v = 1$ it reduces to (19)*.*

Proof of Proposition OA.7. Under the local node the departures d_ℓ are independent $\mathcal{N}(0, v+\sigma^2)$. Under the general node, each ϑ_ℓ has variance σ^2 and any two have covariance $\rho\sigma^2$, so $d = (d_1, \ldots, d_N)' \sim \mathcal{N}(0,\ aI + \rho\sigma^2\mathbf{1}\mathbf{1}')$ with $a := v + (1-\rho)\sigma^2$. Writing $g := a + N\rho\sigma^2$, this covariance has determinant $a^{N-1}g$ and inverse $a^{-1}[I - \rho\sigma^2\mathbf{1}\mathbf{1}'/g]$, so the quadratic form is $SS/a + N\bar{d}^2/g$, using $\sum_\ell d_\ell^2 = SS + N\bar{d}^2$. Subtracting the local log likelihood, $-\frac{1}{2}[N\log(v+\sigma^2) + (SS + N\bar{d}^2)/(v+\sigma^2)]$, gives (OA.1), which at $\rho = v = 1$ is (19). *Signs.* When $N = 1$, $SS = 0$ and $g = v + \sigma^2$, so every term vanishes. For $N \geq 2$ and $\rho > 0$, $g - (v+\sigma^2) = (N-1)\rho\sigma^2 > 0$ makes the coefficient on $\bar{d}^2$ positive, and $a < v + \sigma^2$ makes the coefficient on SS negative. The expression is affine and decreasing in SS, so admission in the two-node comparison holds if and only if SS is below a threshold that is affine and increasing in $\bar{d}^2$. *The Occam term.* At $SS = \bar{d} = 0$ the expression reduces to an

Occam term $O_N(\rho)$ with $O'_N(\rho) = \frac{1}{2}(N-1)\sigma^2[1/a - 1/g] > 0$ for $\rho > 0$, since $g - a = N\rho\sigma^2$. At $\rho = 1$, $a = v$ and $g = v + N\sigma^2$, and $O_{N+1}(1) - O_N(1) = \frac{1}{2}[\log(1+\sigma^2/v) - \log(1+\sigma^2/(v+N\sigma^2))] > 0$. □

Critical thinking (Section 6). Suppose the default was adopted after searching a class $\mathscr{S}_0$ and an alternative t^* after searching a class $\mathscr{S}$. A critical thinker discounts each model's prior weight in the automatic assignment $p_k(x)$ of Section 2.2.3 by its search: $\pi_0 \propto |\mathscr{S}_0|^{-\kappa}$ and $\pi_{t^*} \propto |\mathscr{S}|^{-\kappa}$.

Proposition OA.8 (Questioning the default)**.** *In a comparison of the default with an alternative t^*, the bound in the alternative's evidence constraint under critical thinking κ is*

$$\overline{W}_{t^*} \;=\; \frac{\Lambda\,(|\mathscr{S}_0|/|\mathscr{S}|)^{\kappa}}{1+\Lambda\,(|\mathscr{S}_0|/|\mathscr{S}|)^{\kappa}},$$

where Λ is the alternative's Bayes factor against the default. It is increasing in κ if and only if $|\mathscr{S}_0| > |\mathscr{S}|$. In particular, against an alternative held before the initial signal ($|\mathscr{S}| = 1$), critical thinking never makes the default stickier, and it loosens any default that was itself chosen after a search.

Proof of Proposition OA.8. With prior weights $\pi_0 \propto |\mathscr{S}_0|^{-\kappa}$ and $\pi_{t^*} \propto |\mathscr{S}|^{-\kappa}$ in the automatic assignment of Section 2.2.3, the alternative's cap in the two-model comparison is $\Lambda\rho/(1+\Lambda\rho)$ with $\rho := (|\mathscr{S}_0|/|\mathscr{S}|)^{\kappa}$. This is strictly increasing in ρ, and $\log\rho = \kappa\log(|\mathscr{S}_0|/|\mathscr{S}|)$ is strictly increasing in κ if and only if $|\mathscr{S}_0| > |\mathscr{S}|$. If $|\mathscr{S}| = 1$, $\rho = |\mathscr{S}_0|^{\kappa}$ is nondecreasing in κ, strictly so when $|\mathscr{S}_0| > 1$. □

Critical thinking is thus not a bias toward the familiar: whether the default becomes stickier or looser depends on which mental model owes more to after-the-fact search.

Proposition OA.9 (Fresh evidence)**.** *Let t^* be chosen on x, and let $x'_1, \ldots, x'_m$ be later observations independent of x given the process that generates them, such as observations from new plots. Then the critical thinker's log odds for $S \cup \{t^*\}$ against S are*

$$\log\Lambda_{S\cup t^*:S}(x) \;-\; \kappa\log|\mathscr{S}| \;+\; \sum_{i=1}^{m}\log\Lambda_{S\cup t^*:S}(x'_i),$$

so they differ from the naive DM's by the constant $\kappa\log|\mathscr{S}|$. Each fresh observation changes the log odds in expectation by the amount by which $S \cup \{t^*\}$ fits it better than S, in Kullback–Leibler divergence from the process that generates it, as in Corollary 3. As $m \to \infty$, the bounds in both DMs' evidence constraints for $S \cup \{t^*\}$ converge to one if t^* improves the fit and to zero if it worsens it.

Proof of Proposition OA.9. By Definition 3 the discount is charged only on the initial signal x. Since $x'_1, \ldots, x'_m$ are independent of x and of one another under each node, the log odds of $S \cup \{t^*\}$ against S are $\log\Lambda_{S\cup t^*:S}(x) - \kappa\log|\mathscr{S}| + \sum_i \log\Lambda_{S\cup t^*:S}(x'_i)$, and the naive DM's differ only by the constant $\kappa\log|\mathscr{S}|$. If each x'_i is drawn from the generating process g, then $\mathbb{E}_g[\log\Lambda_{S\cup t^*:S}(x'_i)] = \mathrm{KL}(g \,\|\, p_S) - \mathrm{KL}(g \,\|\, p_{S\cup t^*}) =: \Delta$ as in Corollary 3, so by the strong law the sum tends to $+\infty$ if $\Delta > 0$ and to $-\infty$ if $\Delta < 0$. Both DMs' log odds therefore diverge in the same direction, and the cap $e^{\ell}/(1+e^{\ell})$ of the two-node comparison tends to one or to zero accordingly. □

Corollary OA.5 (Untested insights)**.** *Let t be a centered template with precision υ, chosen by searching a class $\mathscr{S}$ after x, under the conditions of Proposition 13. Suppose that observations bearing on t are acquired only by a DM who places positive weight on $S \cup \{t\}$, as under Assumption 1. If $c_0(\upsilon, |\mathscr{S}|) < z_t^2 \le c_\kappa(\upsilon, |\mathscr{S}|)$, then the naive DM admits $S \cup \{t\}$ and the critical thinker does not. The critical thinker then acquires no observation bearing on t, and her log odds for $S \cup \{t\}$ against S remain at $\log\Lambda_{S\cup t:S}(x) - \kappa\log|\mathscr{S}| \le 0$ in every later stage, whether or not t is correct.*

Proof of Corollary OA.5. By Proposition 13, a DM with discount κ admits $S \cup \{t\}$ after x if and only if $z_t^2 > c_\kappa(\upsilon, |\mathscr{S}|)$. Since $c_0 < z_t^2 \leq c_\kappa$, the naive DM admits it and the critical thinker does not, her discounted log odds being nonpositive. A node not admitted carries no weight, so by hypothesis any observation she acquires has the same distribution under $S \cup \{t\}$ and S; its Bayes factor is one and her log odds do not move. The discount is charged once (Definition 3), so the log odds stay where x left them, the node is not admitted at the next stage, and the argument repeats, whether or not t is correct. □