Online Appendix to
Search in the AI Age: Free Summaries and Costly Verification

Dheer Avashia

August 26, 2026

Abstract

This online appendix collects the detailed algebra, posterior derivations, aggregation-rule extensions, scalar proofs, multidimensional derivations, quantitative-illustration formulas, and notation tables supporting the main manuscript. Cross-document references link directly to the corresponding equations, propositions, theorems, and sections in the main paper.

Contents

How to read this appendix. The appendix is organized by use not by the order in which the results were developed by me or in the paper. Section A collects finite-state Bayesian and dynamic-programming facts used throughout. Section B records maintained assumptions and reference objects. Section C derives the scalar benchmark, aggregation-rule extensions, and scalar proofs. Section D gives the multidimensional posterior, cutoff, and finite-slot proofs, including the formulas behind the quantitative illustration. Section E collects notation. Numbered cross-references to the main manuscript are clickable in the PDF.

Result-to-proof audit map. The shortest route through the mathematics is the following. Each row lists the body result, the appendix block containing its full proof, and the main identities worth checking first.

Main-paper result	Proof or derivation here	Audit starting point
Shared-pool attenuation and depletion	Section C.6	Posterior odds, weighted-tail likelihood ratios, and remaining summary content.
Finite-pool verification and precision regions	Sections C.3 and C.6	Likelihood-normalized transition, posterior martingale, and the two-branch Jensen gap.
Large-pool scalar disappearance	Section C.6	Parity-robust majority tails at every fixed post-summary history.
Self-concealing omission and search suppression	Section D.1	Bayes likelihood-ratio ordering of coverage and the topic-arrival decomposition.
Slot scarcity and multidimensional persistence	Section D.1	Integrated sparse posterior, terminal-support geometry, and the omitted-dimension Jensen gap.
Prevalence bound and welfare nonordering	Section D.1	Additive relevance objective, approximation ratio, and the two explicit open-set constructions.

A Technical Preliminaries

This appendix collects the standard finite-state facts used repeatedly in the proofs below. Stating them once keeps the later derivations short while making clear where each Bayesian, dynamic-programming, and approximation step enters.

Lemma 1 (Finite integrated likelihoods and posterior martingales). Let Ω be a finite latent state space with prior $p_0(\omega) > 0$. Let z be an observed history whose likelihood under latent state ω is $L_z(\omega) \geq 0$, and suppose

$$\mathcal{L}(z) = \sum_{\omega\in\Omega} p_0(\omega)L_z(\omega) > 0.$$

Then the posterior is

$$p(\omega \mid z) = \frac{p_0(\omega)L_z(\omega)}{\sum_{\omega' \in \Omega} p_0(\omega')L_z(\omega')}.$$

If x is a feasible next observation with joint likelihood $L_{z,x}(\omega)$, then

$$\mathbb{P}(x \mid z) = \frac{\sum_{\omega \in \Omega} p_0(\omega)L_{z,x}(\omega)}{\sum_{\omega \in \Omega} p_0(\omega)L_z(\omega)}.$$

For any bounded latent payoff component $g : \Omega \to \mathbb{R}$, the posterior mean

$$\mu_z(g) = \sum_{\omega \in \Omega} g(\omega)p(\omega \mid z)$$

satisfies the martingale identity

$$\mathbb{E}[\mu_{z,X}(g) \mid z] = \mu_z(g),$$

where the expectation is over the posterior-predictive distribution of the next observation X.

Proof. The first two displays are Bayes' rule and the law of total probability on a finite latent state space. For the martingale identity, write

$$\mathbb{E}[\mu_{z,X}(g) \mid z] = \sum_x \mathbb{P}(x \mid z) \sum_{\omega \in \Omega} g(\omega)p(\omega \mid z, x).$$

Substituting Bayes' rule gives

$$\sum_x \sum_{\omega \in \Omega} g(\omega)\frac{p_0(\omega)L_{z,x}(\omega)}{\mathcal{L}(z)} = \sum_{\omega \in \Omega} g(\omega)\frac{p_0(\omega)\sum_x L_{z,x}(\omega)}{\mathcal{L}(z)}.$$

Because the next observation partitions the continuation histories, $\sum_x L_{z,x}(\omega) = L_z(\omega)$. The expression therefore equals $\sum_\omega g(\omega)p(\omega \mid z) = \mu_z(g)$. □

Lemma 2 (Finite-state Bellman cutoff and one-step lower bound). Consider a finite-horizon state s with stopping payoff $S(s)$, marginal inspection cost $c(s)$, and posterior-predictive distribution over successor states s'. If the value function satisfies

$$V(s) = \max\{S(s),\ -c(s) + \mathbb{E}[V(s') \mid s]\},$$

then continuation is optimal at s if and only if

$$c(s) \le \bar{c}(s) \equiv \mathbb{E}[V(s') \mid s] - S(s).$$

Moreover, for any feasible continuation policy σ that pays the next inspection cost and then follows a continuation payoff $V^\sigma(s') \le V(s')$,

$$\mathbb{E}[V(s') \mid s] - S(s) \ge \mathbb{E}[V^\sigma(s') \mid s] - S(s).$$

In particular, the one-step-and-stop policy gives a lower bound on the exact dynamic continuation gain.

Proof. The Bellman maximum chooses continuation exactly when the continuation branch is weakly larger than the stopping branch:

$$-c(s) + \mathbb{E}[V(s') \mid s] \ge S(s).$$

Rearranging gives $c(s) \le \mathbb{E}[V(s') \mid s] - S(s)$. The lower-bound statement follows because V is the supremum over feasible continuation payoffs, so $V(s') \ge V^\sigma(s')$ state by state and hence also in conditional expectation. □

Lemma 3 (Ordered coefficient ratios). Let

$$A(t) = \sum_{j \in J} a_j t^j, \qquad B(t) = \sum_{j \in J} b_j t^j,$$

where J is a finite ordered subset of nonnegative integers, $b_j > 0$, $a_j \ge 0$, and at least one $a_j > 0$. If the coefficient ratios a_j/b_j are weakly increasing in j, then $A(t)/B(t)$ is weakly increasing on $t > 0$. If the coefficient ratios are strictly increasing on two indices that receive positive weight, then $A(t)/B(t)$ is strictly increasing on $t > 0$.

Proof. Let $r_j = a_j/b_j$. Differentiating and collecting unordered pairs gives

$$B(t)^2 \frac{d}{dt}\left(\frac{A(t)}{B(t)}\right) = \sum_{i<j} b_i b_j (r_j - r_i)(j - i) t^{i+j-1}.$$

Each summand is nonnegative when $r_j \ge r_i$. If at least one strict coefficient-ratio inequality receives positive weight, the derivative is strictly positive for all $t > 0$. □

Lemma 4 (Majority tails, neutral ties, and fixed post-summary evidence). Fix $\rho \in (1/2, 1)$. Let $C_N \sim \mathrm{Bin}(N, \rho)$ and let

$$\zeta_N^\tau(\rho) = \mathbb{P}(C_N > N/2) + \tau\, \mathbb{P}(C_N = N/2)$$

with the tie term present only when N is even. Then

$$\zeta_N^\tau(\rho) \to 1 \quad \text{for every fixed } \tau \in [0, 1],$$

and the low-state counterpart

$$\mathbb{P}(\mathrm{Bin}(N, 1-\rho) > N/2) \; + \; \tau\, \mathbb{P}(\mathrm{Bin}(N, 1-\rho) = N/2)$$

converges to zero. The even-N tie mass is $O(N^{-1/2})$. Under neutral tie-breaking, $\tau = 1/2$, the even-size and preceding odd-size majority accuracies coincide:

$$\zeta_{2k}^{1/2}(\rho) = \zeta_{2k-1}(\rho).$$

Finally, under this neutral-majority experiment, for every fixed $\bar{m} < \infty$, every posterior reached after at most $\bar{m}$ post-summary inspections converges to the certainty limit selected by the summary, uniformly over the finitely many such histories.

Proof. Since $\rho > 1/2$, Hoeffding's inequality gives

$$\mathbb{P}(C_N \le N/2) \le \exp\{-2N(\rho - 1/2)^2\},$$

so the strict-majority tail converges to one. Replacing ρ by $1 - \rho < 1/2$ gives convergence of the low-state upper tail to zero. The tie probability is a single binomial mass at the center. Stirling's formula gives $\mathbb{P}(C_N = N/2) = O(N^{-1/2})$ for even N, so the weighted tie term vanishes as $N \to \infty$.

For the neutral-tie identity, imagine first drawing $2k$ signals and then, if there is a tie, resolving it by an independent fair coin. This rule is equivalent to deleting one of the $2k$ signals uniformly at random and taking the strict majority of the remaining $2k - 1$ signals: if the original profile is not tied, deleting one signal cannot change the majority sign; if it is tied, the deleted signal is positive with probability $1/2$ and negative with probability $1/2$. The retained $2k - 1$ signals are still i.i.d. with success probability ρ, which proves $\zeta_{2k}^{1/2}(\rho) = \zeta_{2k-1}(\rho)$.

For the last claim, compression consistency must be retained. Fix $m \le \bar{m}$ and a history with $y \le m$ positive signals. Under a favorable summary, posterior odds are

$$\frac{\mu_0}{1-\mu_0}\left(\frac{\rho}{1-\rho}\right)^{2y-m} \frac{H_1^{1/2,N}(m, y; \rho)}{H_1^{1/2,N}(m, y; 1-\rho)}.$$

For fixed (m, y), the majority threshold differs from $N/2$ by only a constant. Hence the numerator consistency probability converges to one and the denominator consistency probability converges to zero by the same tail argument used at entry. The raw-history factor is finite and positive, so the odds diverge to $+\infty$. Under an unfavorable summary the corresponding $H_0^{1/2,N}$ ratio converges to zero, so posterior odds converge to zero. There are only finitely many histories with $m \le \bar{m}$, which makes both convergences uniform over that collection. □

Lemma 5 (Bounded convergence for finite-state posterior objects). Let $X_N \in [0, 1]$ under a sequence of probability laws P_N. If $X_N \to x$ in P_N-probability for some constant x, then $\mathbb{E}_{P_N}[X_N] \to x$.

Proof. For any $\varepsilon > 0$,

$$\mathbb{E}_{P_N}[|X_N - x|] \leq \varepsilon + \mathbb{P}_{P_N}(|X_N - x| > \varepsilon),$$

because $|X_N - x| \leq 1$. The probability term converges to zero by assumption. Letting $\varepsilon \downarrow 0$ gives L^1 convergence and therefore convergence of expectations. □

Lemma 6 (Second-order expansion for martingale belief movements). Let Λ' be a scalar posterior index satisfying $\mathbb{E}[\Lambda' \mid s] = \Lambda$, and define $\Delta = \Lambda' - \Lambda$. If S_ε is three times continuously differentiable with $\overline{S}_\varepsilon^{(3)} = \sup_\lambda |S_\varepsilon^{(3)}(\lambda)| < \infty$, then

$$\mathbb{E}[S_\varepsilon(\Lambda') - S_\varepsilon(\Lambda) \mid s] = \frac{1}{2} S_\varepsilon''(\Lambda) \operatorname{Var}(\Lambda' \mid s) + R_\varepsilon(s),$$

where

$$|R_\varepsilon(s)| \leq \frac{\overline{S}_\varepsilon^{(3)}}{6} \mathbb{E}[|\Delta|^3 \mid s].$$

Proof. Taylor's theorem around Λ gives

$$S_\varepsilon(\Lambda') = S_\varepsilon(\Lambda) + S_\varepsilon'(\Lambda)\Delta + \frac{1}{2} S_\varepsilon''(\Lambda)\Delta^2 + r, \qquad |r| \leq \frac{\overline{S}_\varepsilon^{(3)}}{6} |\Delta|^3.$$

Taking conditional expectations and using $\mathbb{E}[\Delta \mid s] = 0$ proves the claim. □

B Reference Tables and Maintained Assumptions

This section collects compact reference tables that were removed from the submission PDF to keep the main text focused on the theorem spine. No assumption or notation is changed.

Table 1: Maintained Assumptions in the Within-Product Theory

Assumption	Role in the theory	Deferred margin
Finite review pool N known	Makes the summary and later inspections jointly consistent objects from the same evidence set	Unknown review-pool size or ambiguity about what the summary covers
Authentic and representative pool; honest compression	Interprets raw signals as genuine reports representative of the modeled conditional experience distribution, and the summary as truthful coarsening	Endogenous authenticity, review manipulation, and selection into reviewing
Oriented binary signals with $\rho > 1/2$	Gives a transparent majority benchmark and interpretable high/low quality evidence	Heterogeneous signal precisions, continuous sentiment, or unknown signal orientation
Bayesian expected-utility maximization and known primitives	The consumer knows her tastes, priors, precisions, costs, coverage prior, and summary rule and can compute the induced posterior	Misspecified beliefs, rule learning, or strategic platform design
Interior quality priors and baseline $\mathbb{P}(\theta, M) = \mathbb{P}(\theta)G(M)$	Keeps likelihood ratios nondegenerate and separates latent-quality uncertainty from coverage uncertainty in the core model	Boundary priors or correlated quality and coverage
Passive inspection without replacement	Makes each click an unopened raw signal drawn from the finite pool	Targeted review search, ranking effects, or endogenous inspection order
Exit normalized to zero and exogenous alternative-product value $\bar{M}$	Isolates within-product verification from across-product search	Endogenous across-option continuation value in the companion numerical paper
Affine expected utility in beliefs	Delivers the scalar stopping margin and the multidimensional decision index	Nonlinear utility in posterior beliefs or richer risk/taste interactions

Table 2: Scalar Benchmark Objects and Where They Enter

Object	Economic meaning	Main use
$H_a^\tau(m, y; p)$	Probability that the uninspected scalar signal pool can still generate summary cell a	Compression-consistent posterior and transition probabilities
$\mathcal{I}^a(\bar{M})$	Gross one-step value of unpacking one more raw signal after summary a	Primitive verification threshold
$b(\bar{M})$	Stopping margin induced by exit/current product/outside opportunity	Jensen-gap and decision-margin results
$[\mu^{\min}, \mu^{\max}]$	Terminal belief interval still reachable from the remaining finite scalar pool	Dynamic reachability and sandwich

Table 3: Core Multidimensional Objects and Where They Enter

Object	Economic meaning	Main use
$\eta_{d,t}$	Posterior probability that the next unopened sparse review discusses dimension d	Residual topic-arrival channel
$\mathrm{Cove}(\gamma; K)$	Share of the consumer's taste weight placed on displayed dimensions	Endpoint regimes and taste-region interpretation
RL_t	Preference-weighted topic-arrival probability on omitted dimensions	Coverage-residual ordering
$\Omega_{d,t}$	One-step gross payoff gain per unit of taste weight from learning about dimension d	Taste hyperplanes and laboratory cutoff
$\Lambda_t = \sum_d \gamma_d \mu_{d,t}$	Taste-weighted multidimensional decision index	Dynamic reachability and absorption
$\mathcal{R}_{d,t}$	Taste-weighted reduction in residual verification pressure from covering dimension d	Consumer-relevance ranking of summary slots
$\mathcal{W}(\phi; c), \mathrm{Read}(\phi; c)$	Ex ante consumer value and review-opening probability under truthful kernel Q_ϕ	Welfare–click nonordering

C Scalar Benchmark

C.1 Tie-Breaking Under Even N

The main text uses a randomized-majority kernel H_a^τ that covers odd and even N in one object. When N is even, ties occur with positive probability and the aggregation rule must specify how tied evidence states map into the binary overview. A convenient parameterization is

$$\mathbb{P}(a = 1 \mid \text{tie}) = \tau, \qquad \tau \in [0, 1], \tag{C.1}$$

where a tie means that exactly $N/2$ of the underlying human signals are positive. The neutral rule is $\tau = \frac{1}{2}$, while $\tau > \frac{1}{2}$ breaks ties toward a positive high-quality-indicating overview and $\tau < \frac{1}{2}$ breaks ties toward a negative low-quality-indicating overview.

Under this parameterization, the high-state overview accuracy becomes

$$\begin{aligned}\mathbb{P}(a = 1 \mid \theta = 1) = &\sum_{m=\lfloor N/2 \rfloor + 1}^{N} \binom{N}{m} \rho^m (1-\rho)^{N-m} \\ &+ \begin{cases} \tau \binom{N}{N/2} \rho^{N/2} (1-\rho)^{N/2}, & N \text{ even}, \\ 0, & N \text{ odd}. \end{cases}\end{aligned} \tag{C.2}$$

The first term is the probability that a strict majority of the underlying signals indicate high quality. The second term is the tie probability multiplied by the probability that the designer resolves a tied state in favor of a positive overview. By symmetry,

$$\mathbb{P}(a = 0 \mid \theta = 0) = \sum_{m=\lfloor N/2 \rfloor + 1}^{N} \binom{N}{m} \rho^m (1-\rho)^{N-m} + \begin{cases} (1-\tau) \binom{N}{N/2} \rho^{N/2} (1-\rho)^{N/2}, & N \text{ even}, \\ 0, & N \text{ odd}, \end{cases} \tag{C.3}$$

so the neutral rule $\tau = \frac{1}{2}$ restores state symmetry.

This tie-breaking margin is small but economically meaningful. Increasing τ shifts probability mass toward positive overviews precisely in the most ambiguous states, which can expand consideration and alter verification by making tied evidence appear more favorable. The welfare sign is not monotone in τ: it depends on where the stopping margin lies relative to the tie-cell posteriors and on whether verification would otherwise have occurred. For that reason, the main text fixes neutral tie-breaking, $\tau = 1/2$, and treats strategic tie-breaking as a simple illustration of the broader endogenous-ϕ agenda.

For the high-precision elimination results in Corollaries 5 and 7, even-N tie-breaking does not introduce a qualitatively new force: favorable overviews still become nearly sufficient statistics as $\rho \to 1$. The monotone derivative result in Corollary 6 uses the parity-robust simple-zero property in Lemma 8. Figure 1 illustrates the neutral even-N case numerically by comparing the one-step verification geometry for $N = 3$, $N = 4$ with neutral tie-breaking $\tau = 1/2$, and $N = 5$. As in Figure 3, the left panel reports the cost-free value object, while the right panel translates it into a policy threshold for an illustrative inspection cost.

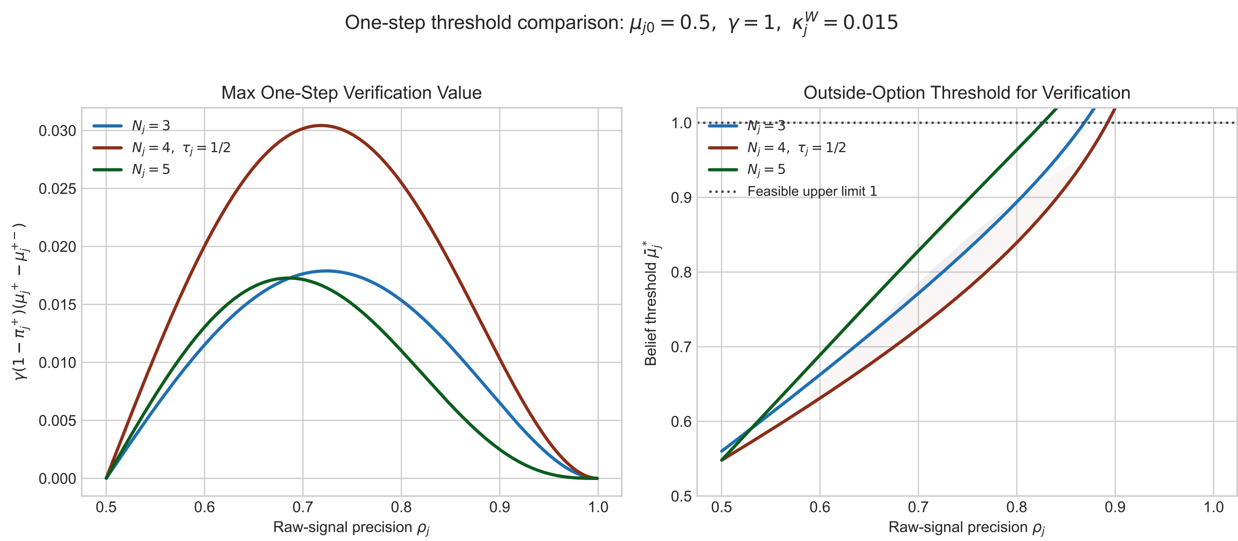

Figure 1: One-step verification thresholds with odd and even signal counts. The left panel plots the maximal one-step verification value after a favorable overview for $N = 3$, $N = 4$ with neutral tie-breaking, and $N = 5$ under $\mu_0 = 0.5$ and $\gamma = 1$. The right panel translates those value objects into outside-option thresholds $\bar{\mu}^*$ for illustrative inspection cost $\kappa^W = 0.015$. Verification is possible only where $\bar{\mu}^* \le \mu^+$, so the shaded bands mark the outside-option beliefs for which one more inspection is optimal. The even-N case is qualitatively similar: the verification region shrinks and eventually exits the feasible belief space as ρ approaches one.

C.2 Scalar Derivations

C.2.1 Parity-robust majority consistency

This subsection collects in one place the identities that make every scalar posterior and transition formula valid for odd and even pool sizes. Fix a history with m opened signals, y of them positive, and let

$$n = N - m, \qquad s(m,y) = \left\lfloor \frac{N}{2} \right\rfloor + 1 - y, \qquad \omega_N = \mathbf{1}\{N \text{ is even}\}\tau.$$

Let $B \sim \mathrm{Bin}(n,p)$. The favorable-summary consistency probability is

$$\begin{aligned} H_1^\tau(m,y;p) &= \mathbb{P}(B \ge s(m,y)) + \omega_N\, \mathbb{P}(B = s(m,y)-1) && \text{(C.4)} \\ &= (1-\omega_N)\, \mathbb{P}(B \ge s(m,y)) + \omega_N\, \mathbb{P}(B \ge s(m,y)-1). && \text{(C.5)} \end{aligned}$$

The first term in (C.4) is the strict-majority tail. When N is even, $B = s(m,y) - 1 = N/2 - y$ is exactly the tie event and receives weight τ; when N is odd, $\omega_N = 0$ and the point-mass term vanishes. Equation (C.5) follows from

$$\mathbb{P}(B = s - 1) = \mathbb{P}(B \geq s - 1) - \mathbb{P}(B \geq s).$$

Thus the parity-robust kernel is a convex combination of two adjacent binomial tails. We use the conventions $\mathbb{P}(B \geq q) = 1$ for $q \leq 0$ and $\mathbb{P}(B \geq q) = 0$ for $q > n$, so the formula also covers histories that already force a summary cell.

Because the unfavorable cell is the complementary output of the same randomized binary kernel,

$$H_0^\tau(m, y; p) = 1 - H_1^\tau(m, y; p). \tag{C.6}$$

Define the binomial mass

$$\mathrm{b}_{n,p}(q) = \begin{cases} \binom{n}{q} p^q (1-p)^{n-q}, & q \in \{0, \ldots, n\}, \\ 0, & \text{otherwise.} \end{cases}$$

Equivalently, expanding the binomial probabilities gives

$$\begin{aligned} H_1^\tau(m, y; p) = &\sum_{b=\max\{0, s(m,y)\}}^{n} \binom{n}{b} p^b (1-p)^{n-b} \\ &+ \omega_N \mathrm{b}_{n,p}(s(m, y) - 1), \end{aligned} \tag{C.7}$$

where a sum whose lower endpoint exceeds n is empty.

The transition recursion follows by designating the next unopened signal $X \sim \text{Bernoulli}(p)$ and conditioning on its realization. If $X = 1$, the successor history is $(m + 1, y + 1)$; if $X = 0$, it is $(m + 1, y)$. The law of total probability therefore gives, for $a \in \{0, 1\}$,

$$H_a^\tau(m, y; p) = pH_a^\tau(m + 1, y + 1; p) + (1 - p)H_a^\tau(m + 1, y; p). \tag{C.8}$$

This argument conditions on the same terminal randomized-majority event in both branches, so it remains valid at even N; the tie point mass needs no separate correction after it has been included in H^τ.

At a terminal history $m = N$, no signs remain to be integrated out and

$$H_1^\tau(N, y; p) = \mathbf{1}\{y > N/2\} + \mathbf{1}\{N \text{ even}\}\tau\mathbf{1}\{y = N/2\}, \qquad H_0^\tau(N, y; p) = 1 - H_1^\tau(N, y; p). \tag{C.9}$$

At entry, the same object is overview precision:

$$H_1^\tau(0, 0; \rho) = \zeta_N^\tau(\rho). \tag{C.10}$$

Under the maintained neutral rule $\tau = 1/2$, state symmetry additionally gives

$$H_1^{1/2}(0, 0; 1 - \rho) = H_0^{1/2}(0, 0; \rho) = 1 - \zeta_N^{1/2}(\rho), \qquad H_0^{1/2}(0, 0; 1 - \rho) = \zeta_N^{1/2}(\rho). \tag{C.11}$$

Equations (C.6), (C.8), and (C.10) are the internal probability checks used by all subsequent posterior, transition, martingale, and reachability calculations.

C.2.2 Majority-rule overview precision

This section derives equation (3.3). Conditional on the latent quality state θ, define the indicator that signal k is correct:

$$X_k = 1\{r_k = \theta\}. \tag{C.12}$$

By assumption,

$$\mathbb{P}(X_k = 1 \mid \theta) = \rho, \tag{C.13}$$

and conditional independence of the underlying human signals implies that the total number of correct signals,

$$C = \sum_{k=1}^{N} X_k, \tag{C.14}$$

is binomial:

$$C \mid \theta \sim \text{Binomial}(N, \rho). \tag{C.15}$$

Under randomized majority aggregation, the AI overview is correct if a strict majority of the underlying signals are correct, and an exact tie is reported correctly with probability τ. Therefore

$$\zeta(\rho) = \mathbb{P}(a = \theta \mid \theta) = \mathbb{P}(C > N/2 \mid \theta) + \begin{cases} \tau\, \mathbb{P}(C = N/2 \mid \theta), & N \text{ even}, \\ 0, & N \text{ odd}. \end{cases} \tag{C.16}$$

Using the binomial probability mass function,

$$\mathbb{P}(C = m \mid \theta) = \binom{N}{m} \rho^m (1-\rho)^{N-m}, \tag{C.17}$$

and summing over all values of m for which the majority is strictly correct, with the tie point added when N is even, yields

$$\zeta(\rho) = \sum_{m=\lfloor N/2 \rfloor + 1}^{N} \binom{N}{m} \rho^m (1-\rho)^{N-m} + \begin{cases} \tau \binom{N}{N/2} \rho^{N/2} (1-\rho)^{N/2}, & N \text{ even}, \\ 0, & N \text{ odd}, \end{cases} \tag{C.18}$$

which is equation (3.3) in the main text.

Theorem 1 (Condorcet improvement under neutral majority aggregation). Suppose the underlying signals are conditionally independent, with each signal correct with probability

$p \in (1/2, 1)$. Let

$$\zeta_N^{1/2}(p) = \mathbb{P}(\mathrm{Bin}(N, p) > N/2) + \mathbf{1}\{N \text{ is even}\}\frac{1}{2}\,\mathbb{P}(\mathrm{Bin}(N, p) = N/2)$$

denote neutral-majority accuracy for a pool of size N. If $N = 2$, $\zeta_2^{1/2}(p) = p$. If $N \geq 3$, then

$$\zeta_N^{1/2}(p) > p. \tag{C.19}$$

Moreover, for every $s \geq 1$,

$$\zeta_{2s}^{1/2}(p) = \zeta_{2s-1}^{1/2}(p),$$

so the even neutral-tie case is exactly the adjacent odd-pool case.

This theorem is the information-aggregation benchmark behind the summary. If signals are correctly oriented and individually better than noise, majority aggregation is more accurate than any one signal. The result is not the main novelty of the paper; its role is to justify the benchmark summary as a truthful compression that can be highly informative before the consumer reads any raw reviews.

Proof. First consider $N = 2s + 1$. Let

$$h_{2s+1}(p) = \mathbb{P}(\mathrm{Bin}(2s + 1, p) \geq s + 1)$$

be strict-majority accuracy. Let S be the number of correct signals among the other $2s$ signals after singling out one signal. Majority accuracy exceeds single-signal accuracy by

$$h_{2s+1}(p) - p = (1 - p)\,\mathbb{P}(S \geq s + 1) - p\,\mathbb{P}(S \leq s - 1).$$

Pair each term $k \in \{s + 1, \ldots, 2s\}$ in the first probability with $2s - k \in \{0, \ldots, s - 1\}$ in the second. Since the binomial coefficients are equal and $p > 1 - p$ when $p > 1/2$, every paired difference is strictly positive. Hence $h_{2s+1}(p) > p$ for $p > 1/2$, with equality only at the one-signal case $s = 0$. At $p = 1/2$, symmetry gives $h_{2s+1}(1/2) = 1/2$.

For even $N = 2s$ with neutral tie-breaking, write $C = X + Z$, where $X \sim \mathrm{Bin}(2s - 1, p)$ and $Z \sim \mathrm{Bernoulli}(p)$ are independent. Expanding the strict-majority and tie events gives

$$\begin{aligned}\mathbb{P}(C > s) + \frac{1}{2}\,\mathbb{P}(C = s) &= \mathbb{P}(X \geq s + 1) + p\,\mathbb{P}(X = s) \\ &\quad + \frac{1}{2}\big[(1 - p)\,\mathbb{P}(X = s) + p\,\mathbb{P}(X = s - 1)\big].\end{aligned} \tag{C.20}$$

The adjacent central binomial coefficients are equal: $\binom{2s-1}{s-1} = \binom{2s-1}{s}$. Consequently,

$$p\,\mathbb{P}(X = s - 1) = p\binom{2s - 1}{s - 1}p^{s-1}(1 - p)^s = (1 - p)\,\mathbb{P}(X = s).$$

Substituting this identity into (C.20) yields

$$\mathbb{P}(C > s) + \frac{1}{2}\mathbb{P}(C = s) = \mathbb{P}(X \geq s+1) + [p + (1-p)]\,\mathbb{P}(X = s) \tag{C.21}$$

$$= \mathbb{P}(X \geq s) = \zeta^{1/2}_{2s-1}(p). \tag{C.22}$$

Thus $N = 2$ has accuracy $\zeta^{1/2}_1(p) = p$, while every even $N \geq 4$ inherits the strict improvement from the adjacent odd case $2s - 1 \geq 3$. In the notation of this paper, $p = \rho$ and $\zeta^{1/2}_N(\rho) = \zeta(\rho)$. □

C.3 Scalar Posterior, Transition, and Benchmark Derivations

C.3.1 No-AI posterior after m signals

This subsection derives the no-summary posterior in (C.32). In the no-AI environment, the consumer learns only from costly raw-signal inspection. Fix the product. Let $r_1 \in \{0, 1\}$ denote the first inspected raw signal, where a positive realization $r_1 = 1$ is high-quality-indicating and a negative realization $r_1 = 0$ is low-quality-indicating.

After one positive signal, Bayes' rule gives

$$\mathbb{P}(\theta = 1 \mid r_1 = 1) = \frac{\mathbb{P}(r_1 = 1 \mid \theta = 1)\,\mathbb{P}(\theta = 1)}{\mathbb{P}(r_1 = 1 \mid \theta = 1)\,\mathbb{P}(\theta = 1) + \mathbb{P}(r_1 = 1 \mid \theta = 0)\,\mathbb{P}(\theta = 0)} \tag{C.23}$$

$$= \frac{\mu_0\rho}{\mu_0\rho + (1-\mu_0)(1-\rho)}. \tag{C.24}$$

After one negative signal,

$$\mathbb{P}(\theta = 1 \mid r_1 = 0) = \frac{\mathbb{P}(r_1 = 0 \mid \theta = 1)\,\mathbb{P}(\theta = 1)}{\mathbb{P}(r_1 = 0 \mid \theta = 1)\,\mathbb{P}(\theta = 1) + \mathbb{P}(r_1 = 0 \mid \theta = 0)\,\mathbb{P}(\theta = 0)} \tag{C.25}$$

$$= \frac{\mu_0(1-\rho)}{\mu_0(1-\rho) + (1-\mu_0)\rho}. \tag{C.26}$$

Now suppose the consumer has inspected m raw signals and observed y positives. Then $m - y$ inspected signals are negative. Conditional on $\theta = 1$, the probability of any realized history with exactly y positives and $m - y$ negatives is

$$\rho^y(1-\rho)^{m-y}, \tag{C.27}$$

while conditional on $\theta = 0$ it is

$$(1-\rho)^y\rho^{m-y}. \tag{C.28}$$

If one writes the likelihood of the count y explicitly, the binomial coefficient appears on both sides:

$$\mathbb{P}(Y = y \mid \theta = 1, m) = \binom{m}{y} \rho^y (1-\rho)^{m-y}, \tag{C.29}$$

$$\mathbb{P}(Y = y \mid \theta = 0, m) = \binom{m}{y} (1-\rho)^y \rho^{m-y}. \tag{C.30}$$

Applying Bayes' rule,

$$\mathbb{P}(\theta = 1 \mid m, y) = \frac{\mu_0 \binom{m}{y} \rho^y (1-\rho)^{m-y}}{\mu_0 \binom{m}{y} \rho^y (1-\rho)^{m-y} + (1-\mu_0)\binom{m}{y}(1-\rho)^y \rho^{m-y}} \tag{C.31}$$

$$= \frac{\mu_0 \rho^y (1-\rho)^{m-y}}{\mu_0 \rho^y (1-\rho)^{m-y} + (1-\mu_0)(1-\rho)^y \rho^{m-y}}. \tag{C.32}$$

The only part of the count likelihood that drops out is the common combinatorial term $\binom{m}{y}$, so the posterior depends on the sufficient statistic (m, y).

C.3.2 AI posterior after an overview and partial inspection

This subsection derives equation (3.7). In the AI environment, the consumer first observes the overview realization $a \in \{0, 1\}$ and then, if desired, inspects m underlying raw signals of which y are positive. The key difference from the no-AI environment is that the overview and the inspected raw signals must be jointly consistent.

Fix the product. Let $\iota_N = \mathbf{1}\{N \text{ is even}\}$, let $\tau \in [0, 1]$ be the tie-breaking probability, and let $B \sim \text{Bin}(N - m, p)$. Given an overview outcome $a \in \{0, 1\}$ and an inspected history (m, y), define

$$H_a^\tau(m, y; p) = \begin{cases} \mathbb{P}(B > N/2 - y) + \iota_N \tau\, \mathbb{P}(B = N/2 - y), & a = 1, \\ \mathbb{P}(B < N/2 - y) + \iota_N (1-\tau)\, \mathbb{P}(B = N/2 - y), & a = 0. \end{cases}$$

For the benchmark, $\tau = 1/2$. In the rest of this derivation, write H_a for this tie-adjusted consistency probability. This is the probability that the remaining uninspected signals complete the finite signal set in a way that is consistent with overview outcome a when each remaining signal is positive with probability p. In the actual model, p is not primitive: it is implied by the latent-quality state and the correctness parameter ρ. Under $\theta = 1$, $p = \rho$; under $\theta = 0$, $p = 1 - \rho$. Thus the H-function compactly writes the probability that the unseen remainder of the finite signal pool can still produce the observed overview under each latent-quality state.

Consider first the overview-only case $(a, 0, 0)$. Under $\theta = 1$, a favorable overview occurs with probability

$$L_1(1, 0, 0) = H_1(0, 0; \rho) = \zeta, \tag{C.33}$$

while under $\theta = 0$ it occurs with probability

$$L_0(1,0,0) = H_1(0,0;1-\rho) = 1-\zeta. \tag{C.34}$$

Hence

$$\widetilde{\mu}^{AI}(1,0,0) = \frac{\mu_0\zeta}{\mu_0\zeta + (1-\mu_0)(1-\zeta)}, \tag{C.35}$$

Under neutral tie-breaking, state symmetry in (C.11) gives

$$L_1(0,0,0) = H_0(0,0;\rho) = 1-\zeta, \qquad L_0(0,0,0) = H_0(0,0;1-\rho) = \zeta.$$

Therefore the unfavorable-overview posterior is explicitly

$$\widetilde{\mu}^{AI}(0,0,0) = \frac{\mu_0(1-\zeta)}{\mu_0(1-\zeta) + (1-\mu_0)\zeta}. \tag{C.36}$$

Now let the consumer inspect m raw signals after seeing overview outcome a. Conditional on $\theta = 1$, the inspected portion of the history contributes

$$\binom{m}{y}\rho^y(1-\rho)^{m-y}. \tag{C.37}$$

The remaining $N-m$ uninspected signals must still be consistent with the overview, which contributes the factor

$$H_a(m,y;\rho). \tag{C.38}$$

Therefore the full likelihood under $\theta = 1$ is

$$L_1(a,m,y) = \binom{m}{y}\rho^y(1-\rho)^{m-y}H_a(m,y;\rho). \tag{C.39}$$

By the same logic, under $\theta = 0$ the inspected signals contribute

$$\binom{m}{y}(1-\rho)^y\rho^{m-y}, \tag{C.40}$$

and the consistency of the remaining uninspected signals contributes

$$H_a(m,y;1-\rho), \tag{C.41}$$

so

$$L_0(a,m,y) = \binom{m}{y}(1-\rho)^y\rho^{m-y}H_a(m,y;1-\rho). \tag{C.42}$$

Applying Bayes' rule to the joint history (a,m,y) yields

$$\widetilde{\mu}^{AI}(a,m,y) = \mathbb{P}(\theta = 1 \mid a,m,y) \tag{C.43}$$

$$= \frac{\mu_0 L_1(a,m,y)}{\mu_0 L_1(a,m,y) + (1-\mu_0)L_0(a,m,y)}. \tag{C.44}$$

This is equation (3.7). Relative to the no-AI posterior, the new object is the H-function: it enforces joint consistency between the inspected signals and the overview.

C.3.3 Next-signal probabilities in the no-AI and AI environments

This subsection derives equations (C.47) and (3.9).

In the no-AI environment, after inspecting m raw signals and observing y positives, the posterior is $\widetilde{\mu}(m,y)$. The next raw signal is positive with probability ρ if $\theta = 1$ and with probability $1-\rho$ if $\theta = 0$. Therefore the law of total probability gives

$$\begin{aligned}
\widetilde{\pi}^{+}(m,y) &= \mathbb{P}(r_{m+1} = 1 \mid m,y) && \text{(C.45)}\\
&= \mathbb{P}(r_{m+1} = 1 \mid \theta = 1, m, y)\,\mathbb{P}(\theta = 1 \mid m,y)\\
&\qquad + \mathbb{P}(r_{m+1} = 1 \mid \theta = 0, m, y)\,\mathbb{P}(\theta = 0 \mid m,y) && \text{(C.46)}\\
&= \widetilde{\mu}(m,y)\rho + \big(1-\widetilde{\mu}(m,y)\big)(1-\rho), && \text{(C.47)}
\end{aligned}$$

which is the no-summary predictive probability used as the comparison object.

In the AI environment, fix a reachable history (a,m,y). Conditional on latent precision p, the probability that the next inspected signal is positive must account for the fact that the remaining uninspected signals still have to be consistent with the observed overview. By conditional exchangeability, it is enough to designate one of the remaining $N-m$ signal positions as the next inspected signal. For that designated signal to be positive and for the final overview to remain equal to a, two things must happen:

1. the designated signal must be positive, which occurs with probability p;
2. the remaining $N-m-1$ uninspected signals must still complete the signal set in a way consistent with overview a.

If the designated next signal is positive, the consumer moves from (m,y) to $(m+1,y+1)$, so the second probability is exactly

$$H_a(m+1, y+1; p). \tag{C.48}$$

The denominator that normalizes this conditional probability is the probability that the current history itself is overview-consistent:

$$H_a(m, y; p). \tag{C.49}$$

Hence

$$\psi_a(m,y;p) = \mathbb{P}(r_{m+1} = 1 \mid a, m, y;\, p) = p\,\frac{H_a(m+1,y+1;p)}{H_a(m,y;p)}, \tag{C.50}$$

which is equation (3.8).

Finally, averaging over the two latent-quality states gives the unconditional AI transition probability:

$$\begin{aligned}
\pi^+(a,m,y) &= \mathbb{P}(r_{m+1}=1 \mid a,m,y) && \text{(C.51)}\\
&= \mathbb{P}(r_{m+1}=1 \mid a,m,y,\theta=1)\,\mathbb{P}(\theta=1 \mid a,m,y)\\
&\quad + \mathbb{P}(r_{m+1}=1 \mid a,m,y,\theta=0)\,\mathbb{P}(\theta=0 \mid a,m,y) && \text{(C.52)}\\
&= \widetilde{\mu}^{AI}(a,m,y)\psi_a(m,y;\rho) + \Big(1-\widetilde{\mu}^{AI}(a,m,y)\Big)\psi_a(m,y;1-\rho), && \text{(C.53)}
\end{aligned}$$

which is equation (3.9).

The likelihood-normalized version (3.10) follows by substituting the posterior formula into (3.9) and cancelling the common binomial coefficient. Equivalently, the numerator is the joint probability of the current history and a favorable next inspected signal, integrated over the two latent-quality states, and the denominator is the joint probability of the current history:

$$\begin{aligned}
\pi^+(a,m,y) &= \frac{A_1(a,m,y)+A_0(a,m,y)}{D_1(a,m,y)+D_0(a,m,y)},\\
A_1(a,m,y) &= \mu_0\rho^y(1-\rho)^{m-y}\rho H_a(m+1,y+1;\rho),\\
A_0(a,m,y) &= (1-\mu_0)(1-\rho)^y\rho^{m-y}(1-\rho)H_a(m+1,y+1;1-\rho),\\
D_1(a,m,y) &= \mu_0\rho^y(1-\rho)^{m-y}H_a(m,y;\rho),\\
D_0(a,m,y) &= (1-\mu_0)(1-\rho)^y\rho^{m-y}H_a(m,y;1-\rho).
\end{aligned} \tag{C.54}$$

This expression is well-defined at every reachable history because the denominator is positive by the definition of reachability. It also clarifies that no arbitrary value of $\psi_a(m,y;p)$ is needed at latent states that have zero posterior probability.

C.3.4 No-summary Bellman benchmark

The no-summary dynamic program is the benchmark used for ex ante comparisons and for the AI-induced-verification comparison in the body. Let $\widetilde{W}(m,y;\bar{M})$ be the value after inspecting m raw signals and observing y positive signals in the environment without a free summary. The terminal condition is

$$\widetilde{W}(N,y;\bar{M}) = S(\widetilde{\mu}(N,y);\bar{M}), \tag{C.55}$$

and for $m<N$,

$$\widetilde{W}(m,y;\bar{M}) = \max\Big\{S(\widetilde{\mu}(m,y);\bar{M}),\ \widetilde{\mathcal{K}}(m,y;\bar{M})\Big\}, \tag{C.56}$$

where

$$\begin{aligned}
\widetilde{\mathcal{K}}(m,y;\bar{M}) = -c^W(m+1)&\\
&+\widetilde{\pi}^+(m,y)\widetilde{W}(m+1,y+1;\bar{M})\\
&+(1-\widetilde{\pi}^+(m,y))\widetilde{W}(m+1,y;\bar{M}).
\end{aligned} \tag{C.57}$$

The ex ante no-summary value used in Proposition 6 is $V^N(\bar{M}) = \widetilde{W}(0, 0; \bar{M})$.

C.4 Aggregation-Rule Extensions and Scope

C.4.1 Three-cell summaries and aggregation-rule tests

The binary majority overview is the benchmark used for the analytical results in the main text. In empirical applications, however, the summary output may be richer. Many review interfaces report not only favorable and unfavorable attribute summaries, but also a middle category indicating that the underlying evidence is mixed. Under the maintained authentic-pool assumption, such disagreement is interpreted as truthful reporting of real split experience among genuine reviewers, not as evidence that the reviews are unreliable. This can be modeled without changing the finite-pool logic of the paper.

Let

$$Y = \sum_{k=1}^{N} r_k$$

be the number of positive raw signals in the finite pool. Fix integer thresholds

$$0 \le T^{lo} < T^{hi} \le N, \qquad T^{lo} + 1 < T^{hi},$$

and let the overview take values in $\mathcal{A}^3 = \{-, \mathsf{m}, +\}$, where m denotes a mixed summary:

$$a(Y) = \begin{cases} +, & Y \ge T^{hi}, \\ \mathsf{m}, & T^{lo} < Y < T^{hi}, \\ -, & Y \le T^{lo}. \end{cases} \tag{C.58}$$

Equivalently, each summary cell $a \in \mathcal{A}^3$ is an interval

$$C_- = \{0, \ldots, T^{lo}\}, \qquad C_{\mathsf{m}} = \{T^{lo} + 1, \ldots, T^{hi} - 1\}, \qquad C_+ = \{T^{hi}, \ldots, N\}.$$

The only object that changes is the consistency probability. For $a \in \mathcal{A}^3$, define

$$H_a^3(m, y; p) = \mathbb{P}\big(y + B \in C_a\big), \qquad B \sim \mathrm{Bin}(N - m, p). \tag{C.59}$$

Writing ℓ_a and u_a for the lower and upper endpoints of C_a, this is

$$H_a^3(m, y; p) = \sum_{b=\max\{0, \ell_a - y\}}^{\min\{N-m, u_a - y\}} \binom{N-m}{b} p^b (1-p)^{N-m-b}, \tag{C.60}$$

with the convention that an empty sum is zero.

It is useful to distinguish two interpretations of the middle cell. The three-cell rule in (C.58) allows a platform to impose a genuine middle band of sentiment shares. If instead "mixed" is interpreted narrowly as an exact majority-rule tie, then the mixed cell has positive probability only when the relevant number of summarized signals is even. In the unidimensional model this requires even N; in the multi-dimensional model it requires an even realized coverage count $C_d(M)$ for the displayed dimension. With odd signal counts and strict majority aggregation, there is no tied evidence state, so a mixed cell must come from a deliberately wider threshold band.

Proposition 1 (Three-cell coarsening). Replace the binary overview by the three-cell rule in (C.58).

(i) The exact posterior, transition probability, martingale property, and one-step verification cutoff are obtained from the binary model by replacing H_a with H_a^3 and letting $a \in \{-, \mathsf{m}, +\}$.

(ii) If a binary overview is obtained by merging adjacent cells of the three-cell overview, then the three-cell overview Blackwell-dominates that binary overview. Hence a consumer who observes the three-cell overview has weakly higher ex ante continuation value than a consumer who observes only the coarsened binary overview.

(iii) If the mixed cell is symmetric, $T^{lo} + T^{hi} = N$, then

$$H^3_{\mathsf{m}}(0, 0; \rho) = H^3_{\mathsf{m}}(0, 0; 1 - \rho)$$

and therefore

$$\widetilde{\mu}^{AI}(\mathsf{m}, 0, 0) = \mu_0.$$

As $\rho \to 1$, $\mathbb{P}(a = \mathsf{m} \mid \theta) \to 0$, while $\widetilde{\mu}^{AI}(+, 0, 0) \to 1$ and $\widetilde{\mu}^{AI}(-, 0, 0) \to 0$ for any interior prior.

Proof in Appendix C.6.

This extension gives the model a direct aggregation-rule test. If the researcher observes both the summary output and the review corpus, the thresholds in (C.58) can be checked by comparing the reported chip to the corpus-implied count of positive and negative mentions. In the multi-dimensional version, the test is about set membership. Since the overview is free to consume in the benchmark, the ordering of reported chips carries no economic content unless consumers separately respond to salience. The empirical audit has two parts. First, recover a candidate set of dimensions from the review corpus and test whether the displayed summary set equals the top K dimensions by prevalence, up to ties and the platform's display cap. Second, conditional on a displayed dimension, test whether the reported sentiment cell falls into the positive, mixed, or negative interval implied by the estimated thresholds.

The extension also sharpens the behavioral prediction, but only with an important finite-pool qualification. Decisive positive and negative summaries should substitute most strongly for review reading. Under symmetric thresholds, a mixed summary leaves the posterior at the prior; it can therefore identify a state in which verification remains decision-changing when the cell contains several feasible signal counts or when other dimensions remain unresolved. A mixed cell that is only an exact scalar tie, however, identifies the total count and has no residual scalar verification value by Lemma 10. Thus the state-contingent prediction is that clicks migrate toward nondegenerate mixed bands or mixed products with residual uncovered dimensions, not mechanically toward every chip labeled mixed. The high-precision logic is preserved in ex ante terms: decisive cells become nearly sufficient as $\rho \to 1$, and the probability of a nondegenerate mixed cell vanishes. Conditional on a rare broad mixed cell, posterior uncertainty can remain high, so the all-histories version of the high-precision no-verification result should be read as applying to decisive cells, while expected verification still vanishes because the mixed state itself becomes rare.

C.4.2 General finite-output summary kernels

The binary majority rule is the paper's benchmark because it delivers transparent closed forms and a direct information-aggregation interpretation. The finite-pool logic is broader. Let a known summary technology report a finite output $\omega \in \Omega$ according to a kernel $Q_\phi(\omega \mid R, z)$, where $R \in \{0,1\}^N$ is the realized signal pool and z collects public product or corpus covariates. For any observed output ω and inspected history $h_m = (m, y)$, define the integrated likelihood

$$\mathcal{L}_\theta^\phi(\omega, h_m) = \sum_{R \in \mathcal{R}(h_m)} \mathbb{P}(R \mid \theta)\, Q_\phi(\omega \mid R, z), \tag{C.61}$$

where $\mathcal{R}(h_m)$ is the set of full signal pools consistent with the opened signals. The posterior is then the likelihood ratio

$$\mu^\phi(\omega, h_m) = \frac{\mu_0 \mathcal{L}_1^\phi(\omega, h_m)}{\mu_0 \mathcal{L}_1^\phi(\omega, h_m) + (1-\mu_0)\mathcal{L}_0^\phi(\omega, h_m)}. \tag{C.62}$$

The next-signal distribution is obtained by the same integrated-likelihood logic with one additional restriction. For $x \in \{0,1\}$, let $\mathcal{R}(h_m, x) \subseteq \mathcal{R}(h_m)$ be the set of full signal pools that are consistent with the opened history h_m and in which the next inspected signal equals x. Define

$$\mathcal{L}_\theta^\phi(\omega, h_m, x) = \sum_{R \in \mathcal{R}(h_m, x)} \mathbb{P}(R \mid \theta)\, Q_\phi(\omega \mid R, z). \tag{C.63}$$

Then

$$\begin{aligned} \mathbb{P}(r_{m+1} = x \mid \omega, h_m) &= \sum_{\theta \in \{0,1\}} \mathbb{P}(r_{m+1} = x, \theta \mid \omega, h_m) \\ &= \frac{\mu_0 \mathcal{L}_1^\phi(\omega, h_m, x) + (1-\mu_0)\mathcal{L}_0^\phi(\omega, h_m, x)}{\mu_0 \mathcal{L}_1^\phi(\omega, h_m) + (1-\mu_0)\mathcal{L}_0^\phi(\omega, h_m)}. \end{aligned} \tag{C.64}$$

This is the general finite-kernel analogue of (3.10). When Q_ϕ is a majority tail, the numerator and denominator collapse to the corresponding H^τ terms.

Proposition 2 (Aggregation-rule scope). Fix a finite signal pool, a known finite-output summary kernel Q_ϕ, and a reachable history with positive integrated likelihood. Then the Bayesian posterior, the predictive distribution over the next inspected signal, the posterior martingale property, and the exact within-option Bellman cutoff are finite likelihood-ratio objects under Q_ϕ. If the output cells are count intervals, these likelihoods reduce to binomial interval probabilities; if the cells are binary majority tails, they reduce to the H-function in (3.4). Majority-specific claims, including Condorcet improvement and the simple high-precision tail in Lemma 8, require the maintained majority or endpoint-separating threshold assumptions. A recovered rule with sentiment-dependent selection remains an exact finite-state problem, but it need not admit the componentwise binomial closed forms.

The scope result is a guardrail for the model. The decision-margin and Bellman logic do not require literal majority rule; they require a known finite summary technology so that the consumer can form integrated likelihoods. Majority delivers closed-form binomial tails and Condorcet comparative statics, while a calibrated empirical rule remains tractable as a finite-state object even when the neat closed forms disappear.

Proof. The posterior in (C.62) is Bayes' rule after summing over all unobserved full signal pools consistent with the opened history. The predictive probability of any next signal realization is the same integrated likelihood with the additional next-signal restriction in the numerator and the current integrated likelihood in the denominator, so it is well defined whenever the current history is reachable. The posterior martingale property follows from the law of iterated expectations. The Bellman cutoff follows from comparing the stopping payoff with the expected continuation value net of the next inspection cost. Count-interval and majority-tail rules are the special cases in which the integrated likelihood collapses to a binomial interval or tail. The Condorcet and simple-tail statements use the algebra of majority aggregation and therefore do not follow for arbitrary Q_ϕ. □

C.5 Supplementary Scalar Propositions and Examples

This subsection collects scalar results that are useful for applications and calibration but are not needed for the main theoretical spine.

Proposition 3 (AI-induced one-step verification). Consider the one-step environment of Main Paper Proposition 1. Let the no-summary first-inspection posteriors be $\widetilde{\mu}(1,0)$ and $\widetilde{\mu}(1,1)$. If

$$b(\bar{M}) \notin \left(\widetilde{\mu}(1,0), \widetilde{\mu}(1,1)\right),$$

then the no-summary one-step gross value is zero. If, for some positive- probability summary realization a,

$$\mu^{a-} < b(\bar{M}) < \mu^{a+} \quad \text{and} \quad 0 < c^W(1) \leq \mathcal{I}^a(\bar{M}),$$

then one-step verification is weakly optimal after observing a, and is strictly optimal when the cost inequality is strict. Thus a free summary can create review opening by moving the consumer from a no-reading state into the decision-margin region.

Corollary 1 (Exact-count revelation exhausts scalar sign learning). In the scalar finite-pool model, suppose the free summary reveals the exact total favorable count $Y = \sum_{n=1}^{N} r_n$. Maintain exchangeable conditionally i.i.d. signs and passive inspection independent of signal realizations. At every reachable history with m opened signs, of which y are favorable,

$$\mathbb{P}(\theta = 1 \mid Y, m, y) = \mathbb{P}(\theta = 1 \mid Y), \qquad \mathbb{P}(r_{m+1} = 1 \mid Y, m, y, \theta) = \frac{Y - y}{N - m}. \tag{C.65}$$

Hence any sequence of additional passively inspected scalar signs is conditionally independent of θ given Y, its gross verification value is zero, and strictly positive inspection costs imply immediate stopping when scalar sign learning is the only possible benefit. For a folded aspect, exact pooled counts exhaust learning about the group total but not the within-fold allocation; rounded or capped reports reveal only an interval and need not exhaust scalar learning.

Theorem 2 (Scalar dynamic reachability and absorption). At a reachable scalar state (a, m, y), let

$$\mathcal{Y}(a, m, y) = \big\{ Y \in \{y, \dots, y + N - m\} : H_a^{\tau}(N, Y; \rho) > 0 \text{ or } H_a^{\tau}(N, Y; 1 - \rho) > 0 \big\},$$

and define

$$\mu^{\min}(a, m, y) = \min_{Y \in \mathcal{Y}(a,m,y)} \widetilde{\mu}^{AI}(a, N, Y), \qquad \mu^{\max}(a, m, y) = \max_{Y \in \mathcal{Y}(a,m,y)} \widetilde{\mu}^{AI}(a, N, Y).$$

Suppose every remaining inspection cost is strictly positive. If

$$b(\bar{M}) \notin \big(\mu^{\min}(a, m, y), \mu^{\max}(a, m, y)\big), \tag{A.R}$$

immediate stopping is strictly optimal. The reachable terminal interval is nested after every inspection, so once the margin leaves the interval it cannot re-enter along a continuation history. Under neutral majority aggregation and $\rho > 1/2$, the endpoints are attained at the smallest and largest feasible terminal favorable counts and have the primitive formulas derived below.

Corollary 2 (Scalar one-step/dynamic sandwich). Let $C^1(\bar{M})$ be the set of reachable scalar states at which opening one signal and then stopping covers the next inspection cost, let $C^{\infty}(\bar{M})$ be the unrestricted dynamic continuation region, and let

$$\mathcal{B}^R(\bar{M}) = \{(a, m, y) : \mu^{\min}(a, m, y) < b(\bar{M}) < \mu^{\max}(a, m, y)\}.$$

Then

$$C^1(\bar{M}) \subseteq C^{\infty}(\bar{M}) \subseteq \mathcal{B}^R(\bar{M}).$$

Corollary 3 (Tent-shaped one-step verification value). Let $x = \mu^{a-}$, $z = \mu^{a+}$, $q = \pi^a$, and suppose $0 < q < 1$, $x < z$, and $\mu^a = qz + (1-q)x$. As a function of the stopping margin $b = b(\bar{M})$, the one-step verification value is

$$\frac{\mathcal{I}^a(b)}{\gamma} = \begin{cases} 0, & b \leq x, \\ (1-q)(b-x), & x < b \leq \mu^a, \\ q(z-b), & \mu^a < b < z, \\ 0, & b \geq z. \end{cases} \tag{C.66}$$

It reaches its unique peak at $b = \mu^a$, with height

$$\mathcal{I}^a(\mu^a) = \gamma q(1-q)(z-x). \tag{C.67}$$

The tent formula gives the cleanest intuition for post-summary reading. Review inspection is valuable only when the stopping margin lies between the two beliefs that could follow the next signal. The value is highest exactly at the current posterior, where a small favorable or unfavorable movement is most decision-relevant, and it falls to zero once both possible next beliefs imply the same stopping action.

Figure 2 gives the same geometry graphically.

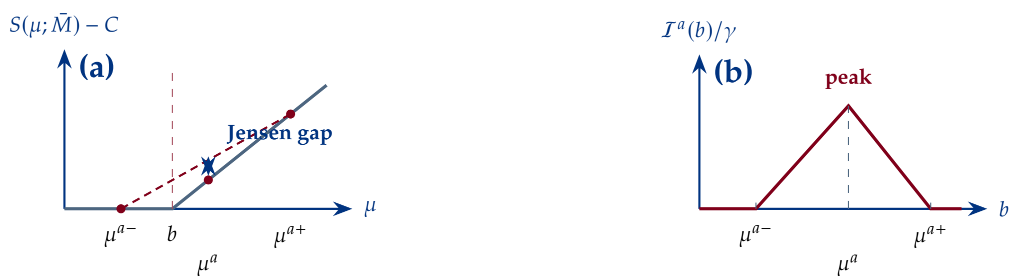

Figure 2: Decision-margin representation of one-step verification. Panel (a) shows the value of opening one more raw signal as the Jensen gap of the kinked stopping payoff. Panel (b) shows the equivalent tent-shaped value as the outside margin varies.

Example 1 (Strict one-step/dynamic sandwich)**.** The inclusions in Corollary 2 can be strict. Consider the scalar majority model with $N = 5$, $\rho = 0.53$, $\mu_0 = 0.20$, unfavorable overview $a = 0$, and state $(m, y) = (3, 0)$. Normalize $\gamma = 1$, set the stopping margin to $b(\bar{M}) = 0.13$, and use the stopping payoff $S(\mu) = \max\{\mu, 0.13\}$. Let the remaining marginal inspection costs be constant:

$$c^W(4) = c^W(5) = 0.001.$$

The exact finite-state dynamic program gives:

Object at state $(a, m, y) = (0, 3, 0)$	Value	Implication
$\widetilde{\mu}^{AI}(0, 3, 0)$	0.148460	current belief above b
$\mathcal{I}^1(0, 3, 0; \bar{M})$	0.000000	one-step opening is not optimal
$\mathbb{E}[W(0, 4, Y_4) \mid 0, 3, 0] - S(0, 3, 0)$	0.002043	dynamic gross gain
$c^W(4)$	0.001000	dynamic continuation is optimal

Thus the state lies in $C^\infty(\bar{M})$ but not in $C^1(\bar{M})$. Economically, strictness arises because each individual signal is too weak to create a one-step Jensen gain at the current margin, but a sequence of weak signals has option value.

Corollary 4 (Mixed-cell verification requires a nondegenerate margin). Consider the passive scalar one-step inspection environment with a three-cell threshold summary, and focus on the middle cell m. If the middle cell reveals a singleton total count of positive signals, including the exact majority tie $Y = N/2$ when N is even, then $\mu^{\mathsf{m}+} = \mu^{\mathsf{m}-} = \mu^{\mathsf{m}}$ and $\mathcal{I}^{\mathsf{m}}(\bar{M}) = 0$. If instead the middle cell contains multiple feasible total counts and the next posteriors are ordered $\mu^{\mathsf{m}-} < \mu^{\mathsf{m}+}$, then one-step verification after a mixed summary has positive gross value exactly when

$$\mu^{\mathsf{m}-} < b(\bar{M}) < \mu^{\mathsf{m}+}.$$

The caveat is important for interpreting "mixed" labels. A broad mixed band can leave residual scalar uncertainty because it pools several possible signal counts. An exact tie cell is different: it reveals the total count, so randomly opening more exchangeable scalar signals only reshuffles known evidence and has no decision value. This is why the model predicts extra reading after nondegenerate mixed bands or omitted dimensions, not after every label called mixed.

Proposition 4 (Post-summary action regions). In the one-step environment of Proposition 1, fix an exogenous outside value $\bar{M}$ and overview realization a. Let

$$S^a = \max\{0, \delta(\mu^a), \bar{M}\}.$$

If $\kappa^W > \mathcal{I}^a(\bar{M})$, the consumer does not initiate verification and any stopping action attaining S^a is optimal. In particular, with an arbitrary fixed tie-breaking rule, the selected post-summary action is exit if $S^a = 0$, choosing or keeping the current product if $S^a = \delta(\mu^a)$, and taking the outside opportunity if $S^a = \bar{M}$. If $\kappa^W \leq \mathcal{I}^a(\bar{M})$, verification is weakly optimal before choosing among exit, the current product, and the outside opportunity.

This proposition separates two decisions that are easy to conflate. If the verification value is below cost, the consumer stops and chooses the best available stopping action. If verification value exceeds cost, the consumer delays that terminal choice because one more review can change which stopping action is best. The summary therefore affects both what the consumer currently prefers and whether she is willing to postpone committing to it.

Proposition 5 (One-step initiation threshold). Consider a single-product within problem with constant outside option $\bar{M}$, finite $N > 1$, majority-rule aggregation, and at most one post-summary inspection at marginal cost κ^W. The scalar body uses neutral tie-breaking, $\tau = 1/2$, so the statement is parity robust. Fix overview realization $a \in \{0, 1\}$, and use the notation μ^a, μ^{a+}, μ^{a-}, and π^a from Proposition 1. Suppose the outside option can be written as

$$\bar{M} = \mathbf{x}'\beta - \alpha P + \gamma\bar{\mu}$$

and that the normalized no-purchase option is weakly dominated in the states considered below. Then:

a. If
$$\mu^{a-} < \bar{\mu} \le \mu^a < \mu^{a+},$$
so the overview leads the consumer to keep the option for now but an unfavorable follow-up signal would induce switching, then verification is optimal if and only if
$$\kappa^W \le \gamma(1-\pi^a)(\bar{\mu} - \mu^{a-}).$$

b. If
$$\mu^{a-} < \mu^a \le \bar{\mu} < \mu^{a+},$$
so the overview leads the consumer to favor the outside option for now but a favorable follow-up signal would induce retention of the option, then verification is optimal if and only if
$$\kappa^W \le \gamma\pi^a(\mu^{a+} - \bar{\mu}).$$

The threshold has a simple economic reading. In case (a), the consumer currently leans toward the product, so the value of reading is the chance that a bad follow-up review makes the outside option preferable. In case (b), the consumer currently leans toward the outside option, so the value is the chance that a good follow-up review rescues the current product. The same decision-margin logic produces substitution or AI-induced verification depending on which side of the margin the summary places the consumer.

Example 2 (One-step threshold calculation when $N = 3$)**.** To make Proposition 5(a) concrete, consider the favorable-overview $N = 3$ specialization with constant outside option $\bar{M}$ and at most one post-summary inspection at marginal cost κ^W. Let

$$\zeta(\rho) = 3\rho^2 - 2\rho^3.$$

Write

$$\mu^+ \equiv \mu^1, \qquad \mu^{++} \equiv \mu^{1+}, \qquad \mu^{+-} \equiv \mu^{1-}, \qquad \pi^+ \equiv \pi^+1.$$

Then

$$\mu^+ = \frac{\mu_0\zeta(\rho)}{\mu_0\zeta(\rho) + (1-\mu_0)(1-\zeta(\rho))}, \tag{C.68}$$

$$\mu^{++} = \frac{\mu_0\rho^2(2-\rho)}{\mu_0\rho^2(2-\rho) + (1-\mu_0)(1-\rho)^2(1+\rho)}, \tag{C.69}$$

$$\mu^{+-} = \frac{\mu_0\rho}{\mu_0\rho + (1-\mu_0)(1-\rho)}. \tag{C.70}$$

Let

$$\pi^+ = \frac{\mu_0\rho^2(2-\rho) + (1-\mu_0)(1-\rho)^2(1+\rho)}{\mu_0\zeta(\rho) + (1-\mu_0)(1-\zeta(\rho))}. \tag{C.71}$$

If

$$\mu^{+-} < \bar{\mu} \le \mu^+ < \mu^{++}, \tag{C.72}$$

then after a favorable overview the consumer continues to inspect if and only if

$$\kappa^W \le \gamma(1-\pi^+)(\bar{\mu} - \mu^{+-}). \tag{C.73}$$

Equivalently,

$$\bar{\mu}^* = \mu^{+-} + \frac{\kappa^W}{\gamma(1-\pi^+)}$$

is the outside-option belief threshold for one-step verification.

The $N = 3$ example is the smallest signal pool in which the mechanism is visible without degeneracy. It shows how the abstract posterior objects become closed-form functions of the primitive precision, prior, taste scale, and inspection cost. The threshold $\bar{\mu}^*$ is the outside opportunity at which a favorable summary is not enough by itself: the consumer keeps reading because an unfavorable raw review would change her action.

Remark 1 (Outside the threshold region). If $\bar{\mu} \le \mu^{+-}$, the outside option dominates after either inspected signal, so one-step verification has no decision-changing value in this special case. If instead $\mu^+ < \bar{\mu} < \mu^{++}$, the consumer prefers the outside option immediately after the favorable overview, but verification can still be valuable because a favorable follow-up signal would overturn that provisional ranking. If $\bar{\mu} \ge \mu^{++}$, the outside option dominates even after a favorable inspected signal, so verification again has no decision-changing value.

Lemma 7 (Posterior martingale under one-step inspection). Fix any finite $N > 1$, the neutral majority rule, and any overview realization $a \in \{0, 1\}$. Using the notation of Proposition 5, posterior beliefs satisfy

$$\mu^a = \pi^a \mu^{a+} + (1-\pi^a)\mu^{a-}. \tag{C.74}$$

The martingale lemma is the Bayesian-learning discipline behind the value of information calculations. Opening a review can move beliefs up or down, but before seeing the signal the expected posterior equals the current posterior. Positive verification value therefore comes only from the kink in the stopping payoff, not from belief drift.

Lemma 8 (Simple high-precision tail under neutral majority). Suppose $N > 1$ and the overview is generated by neutral majority rule. Let

$$g^+(\rho) = 1 - \pi^+(1, 0, 0), \qquad g^-(\rho) = \pi^+(0, 0, 0)$$

be the probabilities of a summary-disconfirming raw signal after favorable and unfavorable overviews. Then there exist finite positive constants $C_N^+(\mu_0)$ and $C_N^-(\mu_0)$ such that, as $\rho \to 1$,

$$g^+(\rho) \sim C_N^+(\mu_0)(1-\rho), \qquad g^-(\rho) \sim C_N^-(\mu_0)(1-\rho).$$

In particular,

$$\lim_{\rho\to 1} g^+(\rho) = \lim_{\rho\to 1} g^-(\rho) = 0. \tag{C.75}$$

For all $N \geq 3$, $C_N^+(\mu_0) = C_N^-(\mu_0) = 1$. For $N = 2$, the tie state contributes at first order, so the constants remain positive and finite but depend on the prior.

This tail lemma quantifies why highly precise summaries crowd out scalar review reading. After a favorable majority summary, the probability that the next raw signal contradicts the summary is first order in $1-\rho$; the analogous positive signal after an unfavorable summary is equally rare. As precision approaches one, the review that could overturn the provisional decision becomes too unlikely to justify a positive reading cost.

Corollary 5 (Sufficiently precise signals eliminate one-step verification). Fix any finite $N > 1$, $\mu_0 \in (0,1)$, $\kappa^W > 0$, and $\gamma > 0$. Consider the one-step post-summary inspection problem under neutral majority-rule aggregation with constant outside option $\bar{M}$. Define the summary-disconfirming probability

$$d^a(\rho) = \begin{cases} 1-\pi^+(1,0,0), & a = 1, \\ \pi^+(0,0,0), & a = 0. \end{cases} \tag{C.76}$$

Then $d^a(\rho) \to 0$ for each $a \in \{0,1\}$ as $\rho \to 1$, and the primitive one-step value $\mathcal{I}^a(\bar{M})$ converges to zero uniformly over feasible outside-option beliefs $\bar{\mu} \in [0,1]$. Hence there exists $\rho^* < 1$ such that for all $\rho > \rho^*$ and all feasible outside-option beliefs, continuation after the overview is not optimal in the one-step problem.

The one-step high-precision corollary is the finite-pool version of the same idea. A near-perfect summary leaves almost no chance that a single additional review will be decision-changing. Because the cost is strictly positive, the gross option value eventually falls below cost for every outside-option belief.

Corollary 6 (Asymptotic verification substitution). Fix any finite $N > 1$, $\mu_0 \in (0,1)$, $\kappa^W > 0$, and $\gamma > 0$ in the one-step post-summary inspection problem under neutral majority aggregation. Then for a favorable overview,

$$\frac{\partial \bar{\mu}^*}{\partial \rho} > 0$$

for all ρ sufficiently close to 1. Equivalently, for any fixed outside-option belief $\bar{\mu}$, the set of post-summary states in which verification is optimal shrinks as raw-signal precision rises near 1.

This comparative static is the local substitution prediction. Near perfect precision, making raw signals even cleaner primarily makes the free summary more decisive and follow-up costly inspection less attractive. The continuation region therefore contracts near the high-precision endpoint.

Corollary 7 (High precision eliminates finite verification). Fix any finite $N > 1$ and any finite inspection horizon $\bar{m} \leq N$. Suppose within-product marginal inspection costs satisfy $c^W(m) > 0$ for every $m \in \{1, \ldots, \bar{m}\}$, and use neutral majority aggregation, $\tau = 1/2$. Then there exists $\rho^* < 1$ such that, for every $\rho > \rho^*$, both overview realizations $a \in \{0, 1\}$, and every feasible outside-option belief $\bar{\mu} \in [0, 1]$, immediate stopping is optimal at the post-overview entry state. Equivalently, any outside-option cutoff for post-overview continuation lies outside the feasible unit interval; after a favorable overview the corresponding threshold satisfies $\bar{\mu}^* > 1$.

This is the unrestricted finite-horizon counterpart of the one-step endpoint comparative static. As $\rho \to 1$, the summary becomes nearly decisive and the probability that any finite continuation plan encounters a summary-disconfirming signal converges uniformly to zero. The full proof below makes the uniformity over feasible continuation plans and outside beliefs explicit.

Figure 3 illustrates the one-step geometry under $\mu_0 = 0.5$ and $\gamma = 1$. Appendix C.1 gives the arbitrary-tie-breaking comparison.

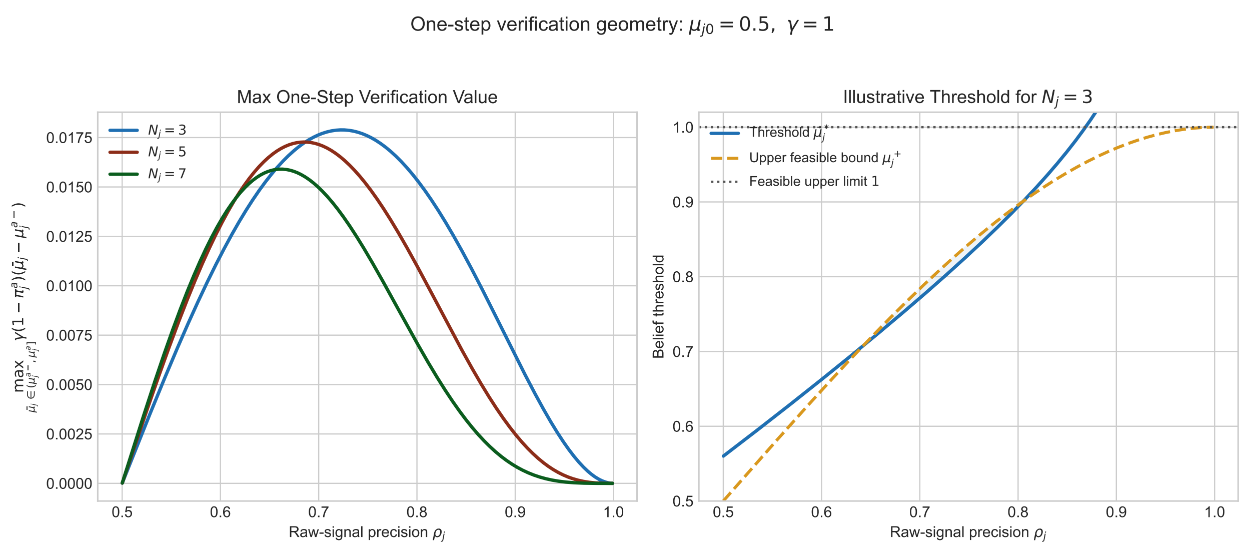

Figure 3: One-step verification geometry under $\mu_0 = 0.5$ and $\gamma = 1$. The left panel plots the maximal one-step verification value as a function of raw-signal precision for illustrative signal-set sizes $N = 3$, $N = 5$, and $N = 7$. The right panel specializes to $N = 3$ and translates that value object into a policy threshold for illustrative inspection cost $\kappa^W = 0.015$. Verification is possible only where the outside-option threshold lies below the upper feasible belief.

Proposition 6 (Free summary value). Fix a product j, a finite signal pool, an outside value $\bar{M}$, and the same within-product inspection cost schedule in the AI and no-summary environments. Let $V^A(\bar{M})$ be the ex ante value before observing the summary in the AI environment, and let $V^N(\bar{M})$ be the corresponding value without a summary. If the summary is costless and the consumer updates correctly, then

$$V^A(\bar{M}) \geq V^N(\bar{M}).$$

The free-summary value result is an ex ante information result, not a statement that consumers always read less. A costless summary weakly expands the consumer's information set, so expected value cannot fall under correct updating. The realized effect on review opening can still be positive or negative because the summary may move a consumer either away from or toward the decision margin.

Proposition 7 (State-specific within-product cutoff). Fix the product, outside value $\bar{M}$, and reachable AI state (a, m, y). Holding the continuation value function $W(\cdot, \cdot, \cdot; \bar{M})$ fixed, define

$$\bar{c}(a, m, y; \bar{M}) = \Delta^A(a, m, y; \bar{M}).$$

Then continuation at state (a, m, y) is optimal if and only if

$$c^W(m+1) \leq \bar{c}(a, m, y; \bar{M}).$$

In particular, a higher current marginal within-product cost weakly lowers the incentive to continue at that state.

This cutoff is the recursive counterpart of the one-step threshold. At each history, all future informational possibilities are summarized by the current continuation gain $\bar{c}$. The consumer opens the next review exactly when the current marginal cost is below that state-specific reservation cost.

C.6 Scalar Proofs

Lemma 9 (Weighted binomial-tail dampening). Let $X_r \sim \text{Bernoulli}(r)$ and $Z_r \sim \text{Bin}(n-1, r)$ be independent, and let

$$K_{s,\omega}(x, z) = \mathbf{1}\{x + z \geq s\} \;+\; \omega\,\mathbf{1}\{x + z = s - 1\}, \qquad 1 \leq s \leq n-1, \quad \omega \in \{0, 1/2\}.$$

Let $E_{s,\omega}$ be the event accepted by this kernel, equivalently by adding an independent randomization that accepts states with probability $K_{s,\omega}(x, z)$. The restriction $s \leq n-1$ is the strict-interior condition: after the next signal, the residual pool can still affect whether the upper-tail summary is rationalized. For $r \in (1/2, 1)$,

$$\frac{1-r}{r} < \frac{\mathbb{P}(X_r = 0 \mid E_{s,\omega})}{\mathbb{P}(X_{1-r} = 0 \mid E_{s,\omega})} < 1. \tag{C.77}$$

The corresponding likelihood-ratio bound for a one realization is

$$1 < \frac{\mathbb{P}(X_r = 1 \mid E_{s,\omega})}{\mathbb{P}(X_{1-r} = 1 \mid E_{s,\omega})} < \frac{r}{1-r}. \tag{C.78}$$

Thus conditioning on a strict upper tail or a neutral-tie adjacent upper tail makes a zero realization less likely and a one realization more likely at the higher success probability, while attenuating the likelihood-ratio content of either realization relative to an unconditional Bernoulli draw.

Proof. Write $t = r/(1-r)$, and define weights

$$w_k = \mathbf{1}\{k \geq s\} + \omega\, \mathbf{1}\{k = s-1\}, \qquad k = 0, \ldots, n,$$

where terms with indices outside their feasible ranges are omitted. After multiplying numerator and denominator by the common factor $(1+t)^n$, the conditional probability of an accepted zero is

$$\chi_{s,\omega}(t) \equiv \mathbb{P}(X_r = 0 \mid E_{s,\omega}) = \frac{D_{s,\omega}(t)}{D_{s,\omega}(t) + M_{s,\omega}(t)},$$

where

$$D_{s,\omega}(t) = \sum_{k=0}^{n-1} w_k \binom{n-1}{k} t^k, \qquad M_{s,\omega}(t) = \sum_{k=1}^{n} w_k \binom{n-1}{k-1} t^k.$$

The right inequality follows from a direct monotonicity argument. Since $\chi_{s,\omega}(t) = [1 + M_{s,\omega}(t)/D_{s,\omega}(t)]^{-1}$, it is enough to show that

$$A_{s,\omega}(t) \equiv \frac{M_{s,\omega}(t)}{D_{s,\omega}(t)}$$

is strictly increasing. Write the numerator and denominator as polynomials in the common powers t^j. Whenever both coefficients at a common power are positive, the coefficient ratio is

$$\frac{\binom{n-1}{j-1}}{\binom{n-1}{j}} = \frac{j}{n-j},$$

which is strictly increasing in j; the numerator may also have the additional highest-order term t^n, corresponding to an infinite final coefficient ratio. By Lemma 3, $A_{s,\omega}(t)$ is strictly increasing in $t > 0$. Hence $\chi_{s,\omega}(t)$ is strictly decreasing in t, which gives the right inequality in (C.77).

For the left inequality, it is enough to show that

$$\frac{\chi_{s,\omega}(t)}{1-r} > \frac{\chi_{s,\omega}(1/t)}{r}, \qquad t = \frac{r}{1-r} > 1.$$

Compare the conditional probability of a zero to its unconditional probability:

$$\begin{aligned}\frac{\chi_{s,\omega}(t)}{1-r} &= \frac{(1+t)D_{s,\omega}(t)}{D_{s,\omega}(t)+M_{s,\omega}(t)} \\ &= \left[1+\frac{M_{s,\omega}(t)-tD_{s,\omega}(t)}{(1+t)D_{s,\omega}(t)}\right]^{-1}. \end{aligned} \tag{C.79}$$

The numerator in the fraction is the two-term polynomial

$$M_{s,\omega}(t)-tD_{s,\omega}(t)=\omega\binom{n-1}{s-2}t^{s-1}+(1-\omega)\binom{n-1}{s-1}t^{s},$$

with infeasible binomial coefficients omitted. For $\omega=0$, this is the strict-tail case, and the displayed fraction is decreasing in t by Lemma 3. For $\omega=1/2$ and $s\geq 2$, the two nonzero numerator coefficients have ratios to the corresponding denominator coefficients

$$\frac{s-1}{n-s+1} \quad \text{and} \quad \frac{s}{2n-s},$$

and the second is no larger than the first because

$$\frac{s}{2n-s}\leq\frac{s-1}{n-s+1} \quad\Longleftrightarrow\quad n(s-2)\geq 0.$$

All higher-power numerator coefficients are zero. Lemma 3 again implies that the displayed fraction is weakly decreasing in t. For the remaining boundary $s=1$, the fraction equals

$$\frac{\frac{1}{2}t}{(1+t)\big((1+t)^{n-1}-\frac{1}{2}\big)}.$$

A direct comparison at reciprocal arguments gives, for $t>1$ and $n\geq 2$,

$$\frac{\frac{1}{2}t}{(1+t)\big((1+t)^{n-1}-\frac{1}{2}\big)}<\frac{\frac{1}{2}t^{n-1}}{(1+t)\big((1+t)^{n-1}-\frac{1}{2}t^{n-1}\big)}.$$

Thus the bracketed term in the preceding display is smaller at t than at $1/t$, so $\chi_{s,\omega}(t)/(1-r)>\chi_{s,\omega}(1/t)/r$. Therefore

$$\frac{\chi_{s,\omega}(t)}{\chi_{s,\omega}(1/t)}>\frac{1-r}{r},$$

which is the left inequality in (C.77).

It remains to prove the upper bound for a one realization. Define

$$\varphi_{s,\omega}(t)\equiv\mathbb{P}(X_r=1\mid E_{s,\omega})=\frac{M_{s,\omega}(t)}{D_{s,\omega}(t)+M_{s,\omega}(t)}.$$

Because $M_{s,\omega}(t)/D_{s,\omega}(t)$ is strictly increasing, $\varphi_{s,\omega}(t)$ is strictly increasing. Hence, for $t > 1$,

$$\frac{\varphi_{s,\omega}(t)}{\varphi_{s,\omega}(1/t)} > 1.$$

For the other inequality, let

$$R_{s,\omega}(t) = M_{s,\omega}(t) - tD_{s,\omega}(t) = \omega\binom{n-1}{s-2}t^{s-1} + (1-\omega)\binom{n-1}{s-1}t^{s}.$$

Then

$$\begin{aligned}\frac{\varphi_{s,\omega}(t)}{r} &= \frac{(1+t)M_{s,\omega}(t)}{t[D_{s,\omega}(t) + M_{s,\omega}(t)]} \\ &= \left[1 - \frac{R_{s,\omega}(t)}{(1+t)M_{s,\omega}(t)}\right]^{-1}.\end{aligned}$$

The coefficient ratios of $R_{s,\omega}(t)$ to $M_{s,\omega}(t)$ are weakly decreasing in the power of t: for $\omega = 0$, the ratio equals one at the lowest feasible power and zero thereafter; for $\omega = 1/2$, it equals one and then one-half at the first two feasible powers and zero thereafter. The reverse form of Lemma 3 therefore implies that $R_{s,\omega}(t)/M_{s,\omega}(t)$ is weakly decreasing. Since $1/(1+t)$ is strictly decreasing,

$$U_{s,\omega}(t) \equiv \frac{R_{s,\omega}(t)}{(1+t)M_{s,\omega}(t)}$$

is strictly decreasing. Thus, for $t > 1$,

$$\frac{\varphi_{s,\omega}(t)}{r} < \frac{\varphi_{s,\omega}(1/t)}{1-r},$$

which rearranges to

$$\frac{\varphi_{s,\omega}(t)}{\varphi_{s,\omega}(1/t)} < \frac{r}{1-r}.$$

Together with the preceding lower bound, this proves (C.78). □

Proof of Theorem 1. Part (i): matched independent-summary benchmark. In the matched independent-summary benchmark, the summary has likelihood

$$\mathbb{P}(a \mid \theta = 1) = \begin{cases} \zeta(\rho), & a = 1, \\ 1-\zeta(\rho), & a = 0, \end{cases} \qquad \mathbb{P}(a \mid \theta = 0) = \begin{cases} 1-\zeta(\rho), & a = 1, \\ \zeta(\rho), & a = 0. \end{cases}$$

Equivalently, the summary likelihood ratio is

$$\frac{\mathbb{P}(a \mid \theta = 1)}{\mathbb{P}(a \mid \theta = 0)} = \left(\frac{\zeta(\rho)}{1 - \zeta(\rho)}\right)^{2a-1}.$$

Conditional on θ, the summary is independent of the raw signals that are later inspected. If m raw signals are opened and y of them are positive, the raw-history likelihood ratio is

$$\frac{\mathbb{P}(y \mid m, \theta = 1)}{\mathbb{P}(y \mid m, \theta = 0)} = \frac{\binom{m}{y}\rho^y(1-\rho)^{m-y}}{\binom{m}{y}(1-\rho)^y\rho^{m-y}} = \left(\frac{\rho}{1-\rho}\right)^{2y-m}.$$

Multiplying the prior odds by the two independent likelihood ratios gives

$$\frac{\mu^I(a, m, y)}{1 - \mu^I(a, m, y)} = \frac{\mu_0}{1 - \mu_0}\left(\frac{\zeta(\rho)}{1 - \zeta(\rho)}\right)^{2a-1}\left(\frac{\rho}{1-\rho}\right)^{2y-m}.$$

Taking logs yields

$$\ell^I(a, m, y) = \log\frac{\mu_0}{1-\mu_0} + (2a-1)\log\frac{\zeta(\rho)}{1-\zeta(\rho)} + (2y-m)\log\frac{\rho}{1-\rho},$$

which is equation (4.1).

For a generic next realization $x \in \{0, 1\}$, the raw-history term changes from $(2y - m)\lambda^\rho$ to $[2(y + x) - (m + 1)]\lambda^\rho$. Therefore

$$\ell^I(a, m+1, y+x) - \ell^I(a, m, y) = \big[2(y+x) - (m+1) - (2y-m)\big]\lambda^\rho = (2x-1)\lambda^\rho.$$

This proves the full-weight statement for both raw realizations. The next two displays record the behaviorally salient disconfirming cases explicitly.

Now fix a favorable summary $a = 1$. A negative raw signal changes the raw count from (m, y) to $(m + 1, y)$, so

$$\ell^I(1, m+1, y) - \ell^I(1, m, y) = \big(2y - (m+1)\big)\lambda^\rho - \big(2y - m\big)\lambda^\rho = -\lambda^\rho.$$

Similarly, after an unfavorable summary $a = 0$, a positive raw signal changes (m, y) to $(m + 1, y + 1)$, so

$$\ell^I(0, m+1, y+1) - \ell^I(0, m, y) = \big(2(y+1) - (m+1)\big)\lambda^\rho - \big(2y - m\big)\lambda^\rho = \lambda^\rho.$$

Thus disconfirming raw signals receive exactly the full raw-signal log-likelihood weight in the independent-summary benchmark.

Finally, the no-summary log odds after the same raw history are

$$\ell^N(m, y) = \log\frac{\mu_0}{1-\mu_0} + (2y-m)\lambda^\rho.$$

Subtracting gives

$$\ell^I(a, m, y) - \ell^N(m, y) = (2a - 1)\log\frac{\zeta(\rho)}{1 - \zeta(\rho)}$$

for every raw history (m, y). This term does not depend on how many reviews have been opened or on whether the opened reviews themselves would imply the same majority cell. Hence the independent summary is never spent.

Part (ii): shared-pool attenuation. The favorable-summary case contains the argument; the unfavorable-summary case follows by reversing zeros and ones. Let $p = \rho$, $q = 1 - \rho$, and let

$$B_{n,s,\omega}(p) = \mathbb{P}(\mathrm{Bin}(n, p) \geq s) + \omega\,\mathbb{P}(\mathrm{Bin}(n, p) = s - 1).$$

At a favorable-summary history $(1, m, y)$, write

$$n = N - m, \qquad s = \lfloor N/2 \rfloor + 1 - y, \qquad \omega = \iota_N \tau.$$

The theorem fixes neutral tie-breaking, $\tau = 1/2$, so $\omega \in \{0, 1/2\}$. Hence $H_1^\tau(m, y; p) = B_{n,s,\omega}(p)$. The theorem's condition $1 \leq s \leq n - 1$ is precisely the strict-interior condition needed for Lemma 9: a negative next signal can be accepted, and the remaining residual pool still has room to affect whether the favorable summary is rationalized.

From (3.7), compression-consistent log odds are

$$\ell^C(1, m, y) = \log\frac{\mu_0}{1 - \mu_0} + (2y - m)\log\frac{p}{q} \; + \; \log\frac{B_{n,s,\omega}(p)}{B_{n,s,\omega}(q)}. \tag{C.80}$$

After a negative inspected signal, the log-posterior change is the log likelihood ratio of that negative signal conditional on the event that the full pool still generates the favorable majority. In the notation of Lemma 9, $r_1 = p$, $r_0 = q$, and

$$E_{s,\omega} \quad \text{is accepted by} \quad K_{s,\omega}(X, Z).$$

The conditional likelihood ratio of a negative signal is therefore

$$\frac{\mathbb{P}(X_p = 0 \mid E_{s,\omega})}{\mathbb{P}(X_q = 0 \mid E_{s,\omega})}.$$

Since $q = 1 - p$, Lemma 9 gives

$$\frac{q}{p} < \frac{\mathbb{P}(X_p = 0 \mid E_{s,\omega})}{\mathbb{P}(X_q = 0 \mid E_{s,\omega})} < 1.$$

Taking logs yields

$$-\log(p/q) < \ell^C(1, m + 1, y) - \ell^C(1, m, y) < 0,$$

which is (4.3).

In the independent-summary benchmark, the summary contributes a fixed likelihood-ratio term to posterior odds and imposes no restriction on later raw signals. Therefore a negative raw signal multiplies posterior odds by q/p, changing log odds by exactly $-\log(p/q)$. This proves the dampening comparison.

The same conditional-likelihood representation applies to a positive next signal:

$$\ell^C(1, m+1, y+1) - \ell^C(1, m, y) = \log \frac{\mathbb{P}(X_p = 1 \mid E_{s,\omega})}{\mathbb{P}(X_q = 1 \mid E_{s,\omega})}.$$

Equation (C.78) gives

$$1 < \frac{\mathbb{P}(X_p = 1 \mid E_{s,\omega})}{\mathbb{P}(X_q = 1 \mid E_{s,\omega})} < \frac{p}{q}.$$

Taking logs yields

$$0 < \ell^C(1, m+1, y+1) - \ell^C(1, m, y) < \log(p/q).$$

Combining this inequality with the negative-signal inequality gives (4.2) for both $x \in \{0, 1\}$. Exchanging zeros and ones gives the unfavorable-summary statement. The strict-interior condition is retained because at boundary histories the summary may be spent, a realization may be mechanically forced by summary consistency, or tie resolution may make the strict inequalities fail.

Part (iii): depletion. If $y > N/2$, then $H_1^\tau(m, y; p) = H_1^\tau(m, y; q) = 1$. Substituting these values into (3.7) gives

$$\ell^C(1, m, y) = \log \frac{\mu_0}{1-\mu_0} + (2y - m) \log \frac{p}{q},$$

the no-summary log odds after the same raw-signal history. The unfavorable summary case is identical after exchanging positives and negatives. If N is even and $y = N/2$, then

$$H_1^\tau(m, N/2; p) = 1 - (1-\tau)(1-p)^{N-m},$$

which is strictly below one for $m < N$; the tie-adjacent state therefore retains residual summary content, as claimed. □

Derivation of remaining summary content. This paragraph verifies Remark 1. Let $q = 1 - \rho$. From the integrated posterior formula (3.7), compression-consistent posterior odds at a nondegenerate AI history are

$$\frac{\widetilde{\mu}^{AI}(a, m, y)}{1 - \widetilde{\mu}^{AI}(a, m, y)} = \frac{\mu_0}{1-\mu_0} \frac{\rho^y q^{m-y} H_a^\tau(m, y; \rho)}{q^y \rho^{m-y} H_a^\tau(m, y; q)}.$$

Taking logs gives

$$\ell^C(a,m,y) = \log\frac{\mu_0}{1-\mu_0} + (2y-m)\lambda^\rho + \log\frac{H_a^\tau(m,y;\rho)}{H_a^\tau(m,y;q)}.$$

The last term is $\mathcal{G}_a(m,y)$. Subtracting this expression at (a,m,y) from the expression at $(a,m+1,y+x)$, for $x\in\{0,1\}$, gives

$$\ell^C(a,m+1,y+x) - \ell^C(a,m,y) = (2x-1)\lambda^\rho + \mathcal{G}_a(m+1,y+x) - \mathcal{G}_a(m,y).$$

If opened signals alone imply the summary cell, then $H_a^\tau(m,y;\rho) = H_a^\tau(m,y;q) = 1$, so $\mathcal{G}_a(m,y) = 0$. This is the spent-summary case.

Proof of Proposition 1. Let $Y = \sum_{k=1}^N r_k$ and let C_a be the count interval associated with summary cell $a \in \{-,\mathsf{m},+\}$. The event that the summary equals a is exactly the event $\{Y \in C_a\}$. After m inspected signals with y positives, the uninspected count of positives is binomial with parameters $(N-m,p)$ under latent positive-signal probability p. Hence the consistency probability is precisely

$$H_a^3(m,y;p) = \mathbb{P}(y+B\in C_a), \qquad B\sim \mathrm{Bin}(N-m,p),$$

which gives (C.59)–(C.60).

The binomial recursion follows by separating the next uninspected signal from the remaining pool:

$$H_a^3(m,y;p) = pH_a^3(m+1,y+1;p) + (1-p)H_a^3(m+1,y;p).$$

This is the only probabilistic identity used to construct ψ_a and the unconditional transition probability. Therefore all posterior and transition formulas from (3.7)–(3.10) apply after replacing H_a by H_a^3 and allowing a to range over the three cells. The martingale property follows from the law of iterated expectations exactly as in Lemma 7, and the one-step cutoff in Proposition 1 depends only on the posterior after the summary and the two possible post-inspection posteriors. It therefore also carries over cell by cell.

For the Blackwell statement, suppose a binary overview b is generated from the three-cell overview by a deterministic map $g:\{-,\mathsf{m},+\}\to\{0,1\}$ that merges adjacent cells. Then observing the three-cell summary and applying g reproduces exactly the binary summary. Any policy feasible after observing only b is therefore feasible after observing a by first garbling a into b and then following the binary-optimal policy. Since the three-cell consumer can always ignore the extra cell information, her ex ante continuation value is weakly higher.

For the symmetric mixed-cell claim, let $C_{\mathsf{m}} = \{T^{lo}+1,\dots,T^{hi}-1\}$. If $T^{lo}+T^{hi}=N$, then $y\in C_{\mathsf{m}}$ if and only if $N-y\in C_{\mathsf{m}}$. If $Y\sim\mathrm{Bin}(N,\rho)$, then $N-Y\sim\mathrm{Bin}(N,1-\rho)$. Hence

$$H_{\mathsf{m}}^3(0,0;\rho) = \mathbb{P}(Y\in C_{\mathsf{m}}) = \mathbb{P}(N-Y\in C_{\mathsf{m}}) = H_{\mathsf{m}}^3(0,0;1-\rho).$$

Substituting this equality into Bayes' rule gives $\widetilde{\mu}^{AI}(\mathsf{m},0,0) = \mu_0$ for any interior prior. Finally, as $\rho \longrightarrow 1$, the high-quality distribution of Y converges to a point mass at N and the low-quality distribution converges to a point mass at zero. Since the mixed interval excludes both endpoints, $\mathbb{P}(a = \mathsf{m} \mid \theta) \longrightarrow 0$. The positive cell contains N and the negative cell contains zero, so Bayes' rule gives $\widetilde{\mu}^{AI}(+,0,0) \longrightarrow 1$ and $\widetilde{\mu}^{AI}(-,0,0) \longrightarrow 0$ for any $\mu_0 \in (0,1)$. □

Proof of Proposition 1. Step 1: primitive cutoff. Fix overview realization a. With at most one post-summary inspection, the consumer either stops immediately and receives

$$S(\mu^a; \bar{M}) = \max\{0, \delta(\mu^a), \bar{M}\},$$

or pays κ^W to inspect one raw signal. Conditional on the overview, the next inspected signal is favorable with probability π^a and unfavorable with probability $1 - \pi^a$. After these two events, the posteriors are μ^{a+} and μ^{a-}, respectively. Hence the expected payoff from inspection is

$$-\kappa^W + \pi^a S(\mu^{a+}; \bar{M}) + (1 - \pi^a) S(\mu^{a-}; \bar{M}).$$

Inspection is optimal if and only if this expression weakly exceeds $S(\mu^a; \bar{M})$. Rearranging gives exactly

$$\kappa^W \leq \pi^a S(\mu^{a+}; \bar{M}) + (1 - \pi^a) S(\mu^{a-}; \bar{M}) - S(\mu^a; \bar{M}) = \mathcal{I}^a(\bar{M}),$$

which proves the primitive threshold characterization.

Step 2: decision-margin representation. Write

$$\delta(\mu) = \mathbf{x}'\beta - \alpha P + \gamma\mu.$$

Let $C(\bar{M}) = \max\{0, \bar{M}\}$. Since the two non-product stopping payoffs are constants, the stopping payoff can be written as

$$S(\mu; \bar{M}) = \max\{C(\bar{M}), \mathbf{x}'\beta - \alpha P + \gamma\mu\}.$$

Using the definition

$$b(\bar{M}) = \frac{C(\bar{M}) - (\mathbf{x}'\beta - \alpha P)}{\gamma},$$

this becomes

$$S(\mu; \bar{M}) = C(\bar{M}) + \gamma(\mu - b(\bar{M}))_+,$$

which proves (4.10). The value of one more signal depends on the random posterior after that signal. Bayes' rule and the law of iterated expectations imply

$$\mu^a = \mathbb{E}[\theta \mid a] = \mathbb{E}[\mathbb{E}[\theta \mid a, X_1] \mid a] = \pi^a \mu^{a+} + (1 - \pi^a)\mu^{a-},$$

where X_1 is the next inspected signal. Substituting the kinked stopping payoff into the primitive value in (4.7) gives

$$\mathcal{I}^a(\bar{M}) = \gamma\left[\pi^a(\mu^{a+} - b)_+ + (1-\pi^a)(\mu^{a-} - b)_+ - (\mu^a - b)_+\right].$$

The expression is nonnegative because $x \mapsto (x-b)_+$ is convex and μ^a is the conditional expectation of the next posterior.

If the next-posterior distribution has two ordered support points $\mu^{a-} < \mu^{a+}$ with both probabilities strictly positive, the convex function $x \mapsto (x-b)_+$ is affine on the entire support unless the kink b lies strictly between the two support points. Hence the Jensen gap is strictly positive exactly when $\mu^{a-} < b(\bar{M}) < \mu^{a+}$, and it is zero when the stopping margin is weakly outside that interval.

Step 3: finite-N precision representation. Fix the favorable overview $a = 1$, neutral majority aggregation, and write

$$q = 1 - \rho, \qquad H_1(m, y; p) \equiv H_1^{1/2}(m, y; p).$$

This final step is a direct likelihood normalization of the one-step value just characterized.

The probability of the favorable overview. At the post-summary entry state, the probability of a favorable overview under the consumer's prior is

$$D_N(\rho) = \mathbb{P}(a = 1) = \mu_0 H_1(0, 0; \rho) + (1-\mu_0) H_1(0, 0; q).$$

The first term is the probability of high quality times the probability that a high-quality product's finite signal pool generates a favorable overview. The second term is the analogous low-quality probability. Since $\mu_0 \in (0, 1)$ and the favorable overview is reachable, $D_N(\rho) > 0$.

Joint likelihoods for the next raw signal. Designate one unopened raw signal as the next inspected signal. Conditional on $\theta = 1$, this signal is positive with probability ρ and negative with probability q. Conditional on $\theta = 0$, it is positive with probability q and negative with probability ρ. If the next signal is positive, the remaining $N-1$ signals must be consistent with state $(a, m, y) = (1, 1, 1)$; if it is negative, they must be consistent with $(1, 1, 0)$. Thus the joint prior probabilities of the four events are

	$X = 1$	$X = 0$
$\theta = 1$	$L_+^H = \mu_0 \rho H_1(1, 1; \rho)$	$L_-^H = \mu_0 q H_1(1, 0; \rho)$
$\theta = 0$	$L_+^L = (1-\mu_0) q H_1(1, 1; q)$	$L_-^L = (1-\mu_0)\rho H_1(1, 0; q)$.

Let

$$L_+ = L_+^H + L_+^L, \qquad L_- = L_-^H + L_-^L.$$

The binomial recursion for the consistency function gives

$$H_1(0, 0; p) = p H_1(1, 1; p) + (1-p) H_1(1, 0; p).$$

Applying this identity at $p = \rho$ and $p = q$ yields

$$D_N(\rho) = L_+^H + L_-^H + L_+^L + L_-^L = L_+ + L_-.$$

Current and successor posteriors. Bayes' rule gives the current posterior after the favorable overview:

$$\mu^1 = \mathbb{P}(\theta = 1 \mid a = 1) = \frac{L_+^H + L_-^H}{D_N(\rho)}.$$

The predictive probability that the next inspected signal is positive is

$$\pi^1 = \mathbb{P}(X = 1 \mid a = 1) = \frac{L_+}{D_N(\rho)}.$$

Conditional on observing a positive next signal,

$$\mu^{1+} = \mathbb{P}(\theta = 1 \mid a = 1, X = 1) = \frac{L_+^H}{L_+},$$

with the convention that zero-probability branches are omitted. Conditional on observing a negative next signal,

$$\mu^{1-} = \mathbb{P}(\theta = 1 \mid a = 1, X = 0) = \frac{L_-^H}{L_-}.$$

These formulas are exactly the entry-state specializations of the integrated posterior and transition formulas in the main text.

Substitution into the kinked stopping payoff. By Step 2,

$$\frac{\mathcal{I}^1(\bar{M})}{\gamma} = \pi^1(\mu^{1+} - b)_+ + (1 - \pi^1)(\mu^{1-} - b)_+ - (\mu^1 - b)_+,$$

where $b = b(\bar{M})$. Substitute the likelihood-normalized posteriors from Step 3. For the positive branch,

$$\pi^1(\mu^{1+} - b)_+ = \frac{L_+}{D_N(\rho)}\left(\frac{L_+^H}{L_+} - b\right)_+ = \frac{\left[L_+^H - bL_+\right]_+}{D_N(\rho)} = \frac{\left[L_+^H - b(L_+^H + L_+^L)\right]_+}{D_N(\rho)}.$$

For the negative branch,

$$(1 - \pi^1)(\mu^{1-} - b)_+ = \frac{L_-}{D_N(\rho)}\left(\frac{L_-^H}{L_-} - b\right)_+ = \frac{\left[L_-^H - bL_-\right]_+}{D_N(\rho)} = \frac{\left[L_-^H - b(L_-^H + L_-^L)\right]_+}{D_N(\rho)}.$$

For the current posterior term,

$$(\mu^1 - b)_+ = \left(\frac{L_+^H + L_-^H}{D_N(\rho)} - b\right)_+ = \frac{\left[L_+^H + L_-^H - bD_N(\rho)\right]_+}{D_N(\rho)}.$$

Combining the three terms gives

$$\frac{\mathcal{I}^1(\bar{M})}{\gamma} = \frac{\left[L_+^H - b(L_+^H + L_+^L)\right]_+ + \left[L_-^H - b(L_-^H + L_-^L)\right]_+}{D_N(\rho)} \\ - \frac{\left[L_+^H + L_-^H - bD_N(\rho)\right]_+}{D_N(\rho)} = \Phi_N(\rho; b, \mu_0),$$

which proves equation (4.14).

The precision-domain cutoff and roots. Step 1 says that one-step verification is optimal after $a = 1$ if and only if

$$\kappa^W \leq \mathcal{I}^1(\bar{M}).$$

Dividing by $\gamma > 0$ and using the identity just proved gives

$$\frac{\kappa^W}{\gamma} \leq \Phi_N(\rho; b(\bar{M}), \mu_0).$$

For fixed N, the consistency terms $H_1(0, 0; p)$, $H_1(1, 1; p)$, and $H_1(1, 0; p)$ are finite binomial tails, including the neutral tie point mass when N is even. Hence they are finite polynomials in p, and therefore finite polynomials in ρ after substituting $q = 1 - \rho$. On each region where the three positive-part terms in Φ_N have fixed signs, and holding the prior μ_0 fixed, the equality

$$\Phi_N(\rho; b, \mu_0) = \kappa^W/\gamma$$

is equivalent to a finite polynomial equation after multiplying by the positive denominator $D_N(\rho)$. Thus the decision-changing precision values are contained in the roots of these finite equations, with the realized weak reading set given exactly by the inequality above. For generic costs, where the cutoff is crossed, the boundary points of the reading set are exactly those roots. The unfavorable-overview case is identical after replacing the favorable-cell consistency function H_1 by H_0. □

Corollary 8 (One-step partial identification of inspection costs). Fix N, ρ, μ_0, and $b = b(\bar{M})$, and consider the favorable-overview one-step decision in Proposition 1. Opening implies

$$\frac{\kappa^W}{\gamma} \leq \Phi_N(\rho; b, \mu_0),$$

while stopping implies the reverse weak inequality, up to indifference. If exogenous variation in b holds the other primitives fixed and reveals a switching boundary b^*, then

$$\frac{\kappa^W}{\gamma} = \Phi_N(\rho; b^*, \mu_0).$$

The identified object is normalized by the taste scale. A laboratory design can elicit or normalize γ. With observed prices, the price coefficient supplies the money metric: κ^W/α

is the inspection cost in currency units and γ/α is willingness to pay for high quality.

Proof. Proposition 1 shows that, for the favorable overview, opening one raw signal is weakly optimal if and only if

$$\kappa^W/\gamma \leq \Phi_N(\rho; b, \mu_0).$$

Therefore observed opening reveals the upper-bound inequality, while observed non-opening reveals the complementary lower-bound inequality, up to indifference. If variation in b holds (N, ρ, μ_0) fixed and reveals a switching boundary b^*, then the consumer must be indifferent at the boundary. Substituting b^* into the cutoff equality gives

$$\kappa^W/\gamma = \Phi_N(\rho; b^*, \mu_0).$$

□

Proof of Proposition 3. In the no-summary one-step problem, the current posterior is μ_0 and the two possible next posteriors are $\widetilde{\mu}(1,0)$ and $\widetilde{\mu}(1,1)$. The posterior martingale identity holds exactly as in Lemma 7. If

$$b(\bar{M}) \notin \big(\widetilde{\mu}(1,0), \widetilde{\mu}(1,1)\big),$$

then the kink in the stopping payoff lies weakly outside the support of the next posterior. The same Jensen-gap argument used in Proposition 1 therefore gives a no-summary one-step gross verification value of zero. Any strictly positive first-inspection cost then makes no-summary one-step verification strictly suboptimal.

Now condition on the AI summary realization a. If $\mu^{a-} < b(\bar{M}) < \mu^{a+}$, Proposition 1 implies $\mathcal{I}^a(\bar{M}) > 0$. If $0 < c^W(1) \leq \mathcal{I}^a(\bar{M})$, Proposition 1 implies that opening one raw signal is weakly optimal after realization a; if the inequality is strict, it is strictly optimal. Since a occurs with positive probability, the AI environment induces verification on a positive probability event whenever continuation is selected at weak indifference. □

Proof of Corollary 3. Let $b = b(\bar{M})$. Proposition 1 gives

$$\frac{\mathcal{I}^a(b)}{\gamma} = q(z-b)_+ + (1-q)(x-b)_+ - (qz + (1-q)x - b)_+.$$

If $b \leq x$, all three positive-part terms are active, and the expression equals

$$q(z-b) + (1-q)(x-b) - [qz + (1-q)x - b] = 0.$$

If $x < b \leq \mu^a = qz + (1-q)x$, the middle term is zero while the final term is active, so

$$q(z-b) - [qz + (1-q)x - b] = (1-q)(b-x).$$

If $\mu^a < b < z$, only the first term is active, so the expression is

$$q(z-b).$$

Finally, if $b \geq z$, all positive-part terms are zero. This proves the piecewise formula in (C.66). The left branch is strictly increasing on $(x, \mu^a]$ and the right branch is strictly decreasing on $[\mu^a, z)$, so the unique peak is at $b = \mu^a$. Substituting $\mu^a = qz + (1-q)x$ into either branch gives

$$\mathcal{I}^a(\mu^a) = \gamma q(z - \mu^a) = \gamma q(1-q)(z-x).$$

□

Proof of Theorem 2 and Corollary 2. Closed-form endpoint derivation. At a terminal history (a, N, Y), the full raw count Y is known. If $Y \in \mathcal{Y}(a, m, y)$, the summary cell is feasible at the terminal count, so $H_a^\tau(N, Y; \rho)$ and $H_a^\tau(N, Y; q)$ are the same positive tie-breaking weight: 1 away from a tie, τ or $1 - \tau$ at a feasible tie. The summary-consistency weight therefore cancels from posterior odds. Using (3.7),

$$\frac{\widetilde{\mu}^{AI}(a, N, Y)}{1 - \widetilde{\mu}^{AI}(a, N, Y)} = \frac{\mu_0}{1-\mu_0} \frac{\rho^Y q^{N-Y}}{q^Y \rho^{N-Y}} = \frac{\mu_0}{1-\mu_0} \left(\frac{\rho}{q}\right)^{2Y-N}.$$

Hence

$$\widetilde{\mu}^{AI}(a, N, Y) = \left[1 + \frac{1-\mu_0}{\mu_0} \left(\frac{q}{\rho}\right)^{2Y-N}\right]^{-1}.$$

Because $\rho > q$, this posterior is strictly increasing in Y. Therefore the minimum and maximum over the reachable terminal count set $\mathcal{Y}(a, m, y)$ occur at

$$Y^-(a, m, y) = \min \mathcal{Y}(a, m, y), \qquad Y^+(a, m, y) = \max \mathcal{Y}(a, m, y),$$

so substitution gives the primitive endpoints

$$\mu^{\min}(a, m, y) = \left[1 + \frac{1-\mu_0}{\mu_0} \left(\frac{q}{\rho}\right)^{2Y^-(a,m,y)-N}\right]^{-1}, \tag{C.81}$$

and

$$\mu^{\max}(a, m, y) = \left[1 + \frac{1-\mu_0}{\mu_0} \left(\frac{q}{\rho}\right)^{2Y^+(a,m,y)-N}\right]^{-1}. \tag{C.82}$$

If the stopping margin $b(\bar{M})$ lies outside this interval, all reachable terminal posteriors lie on one affine piece of the kinked stopping payoff $C(\bar{M}) + \gamma(\mu - b(\bar{M}))_+$. The posterior is a martingale, so any feasible dynamic verification policy has zero gross Jensen value. If the margin lies strictly inside the interval, there are feasible terminal completions on both sides of the stopping margin, so the remaining pool is decision-relevant before costs.

Dynamic reachability and absorption. Fix a reachable state (a, m, y). The set $\mathcal{Y}(a, m, y)$ is nonempty by reachability. Any future history that can be generated after (a, m, y) has a set of terminal completions contained in $\mathcal{Y}(a, m, y)$. In particular, if the consumer reaches some later history h, the posterior at h can be written as a conditional expectation of terminal posteriors whose terminal counts lie in $\mathcal{Y}(a, m, y)$. Hence every posterior that can be reached before the pool is exhausted lies in the closed interval

$$\left[\mu^{\min}(a, m, y), \mu^{\max}(a, m, y)\right].$$

Let $b = b(\bar{M})$ and $C = C(\bar{M}) = \max\{0, \bar{M}\}$. By Proposition 1,

$$S(\mu; \bar{M}) = C + \gamma(\mu - b)_+.$$

If $b \leq \mu^{\min}(a, m, y)$, then

$$S(\mu; \bar{M}) = C + \gamma(\mu - b)$$

on every belief that can still be reached. If $b \geq \mu^{\max}(a, m, y)$, then

$$S(\mu; \bar{M}) = C$$

on every belief that can still be reached. Thus, whenever the reachability condition in Theorem 2 holds, the stopping payoff is affine over the entire set of beliefs reachable from the current state.

Consider any feasible continuation plan and let ν be the finite stopping time at which the plan eventually stops and chooses among exit, the current product, and the outside opportunity. The posterior process is a bounded martingale under Bayesian updating, so

$$\mathbb{E}\left[\widetilde{\mu}^{AI}(a, m_\nu, y_\nu) \mid a, m, y\right] = \widetilde{\mu}^{AI}(a, m, y).$$

Since $S(\cdot; \bar{M})$ is affine on the support of the stopped posterior,

$$\mathbb{E}\left[S\left(\widetilde{\mu}^{AI}(a, m_\nu, y_\nu); \bar{M}\right) \mid a, m, y\right] = S\left(\widetilde{\mu}^{AI}(a, m, y); \bar{M}\right).$$

The continuation action requires paying the next strictly positive inspection cost, and all later costs are weakly nonnegative. Therefore every continuation plan delivers strictly less than immediate stopping when that reachability condition holds. Immediate stopping is strictly optimal, proving the necessary reachability condition.

For nestedness, if successor $(a, m + 1, y + x)$ is reachable, every terminal count that can be reached from that successor is also a terminal count that could have been reached from (a, m, y). Therefore

$$\mathcal{Y}(a, m + 1, y + x) \subseteq \mathcal{Y}(a, m, y),$$

including when N is even: under neutral tie-breaking the exact tie count is assigned positive probability in both summary cells, but no successor history creates a terminal count outside the parent state's feasible terminal set. The tie point changes weights, not the set-inclusion argument. Taking minima and maxima over the two sets gives nested reachable intervals. The absorption statement follows immediately from this inclusion.

One-step/dynamic sandwich. Fix a reachable state (a, m, y) with $m < N$. If $(a, m, y) \in C^1(\bar{M})$, then the feasible policy that opens one additional review and stops yields payoff

$$-c^W(m+1) + \sum_{x \in \mathcal{X}(a,m,y)} \mathbb{P}(x \mid a, m, y) S\Big(\widetilde{\mu}^{AI}(a, m+1, y+x); \bar{M}\Big) \geq S\Big(\widetilde{\mu}^{AI}(a, m, y); \bar{M}\Big).$$

The unrestricted dynamic problem can mimic this policy, and can also choose optimally after the next signal. Hence continuation is optimal in the unrestricted Bellman problem, so $\mathcal{C}^1(\bar{M}) \subseteq \mathcal{C}^\infty(\bar{M})$.

The second inclusion is exactly Theorem 2: if a reachable state is outside $\mathcal{B}^R(\bar{M})$, immediate stopping is strictly optimal, so the state cannot belong to $\mathcal{C}^\infty(\bar{M})$. □

Proof of Corollary 1. Write

$$p_\theta = \mathbb{P}(r_n = 1 \mid \theta) = \begin{cases} \rho, & \theta = 1, \\ 1 - \rho, & \theta = 0. \end{cases}$$

Fix an exact total count $Y = j$. For any realized signal vector $r = (r_1, \ldots, r_N)$ satisfying $\sum_{n=1}^N r_n = j$, conditional independence gives

$$\mathbb{P}(R = r \mid \theta) = p_\theta^j (1 - p_\theta)^{N-j}. \tag{C.83}$$

Summing the same likelihood over the $\binom{N}{j}$ vectors with total j gives

$$\mathbb{P}(Y = j \mid \theta) = \binom{N}{j} p_\theta^j (1 - p_\theta)^{N-j}. \tag{C.84}$$

Dividing (C.83) by (C.84) yields

$$\mathbb{P}(R = r \mid Y = j, \theta) = \frac{1}{\binom{N}{j}}, \tag{C.85}$$

which is independent of θ. Thus, once the exact total is known, the latent state affects the likelihood of the total but not the allocation of the known signs across positions in the pool.

Now suppose passive inspection has opened m positions independently of their realizations and observed y favorable signs. Conditional on $Y = j$, the count y is hypergeometric:

$$\mathbb{P}(y \mid Y = j, m, \theta) = \frac{\binom{j}{y}\binom{N-j}{m-y}}{\binom{N}{m}}, \tag{C.86}$$

on its feasible support. Again, the expression is independent of θ. Bayes' rule therefore gives

$$\mathbb{P}(\theta = 1 \mid Y = j, m, y) = \frac{\mathbb{P}(y \mid Y = j, m, \theta = 1)\, \mathbb{P}(\theta = 1 \mid Y = j)}{\sum_{\vartheta \in \{0,1\}} \mathbb{P}(y \mid Y = j, m, \theta = \vartheta)\, \mathbb{P}(\theta = \vartheta \mid Y = j)} \tag{C.87}$$

$$= \mathbb{P}(\theta = 1 \mid Y = j). \tag{C.88}$$

Exactly $j - y$ favorable signs remain among the $N - m$ unopened positions. Uniform passive sampling consequently gives

$$\mathbb{P}(r_{m+1} = 1 \mid Y = j, m, y, \theta) = \frac{j - y}{N - m}, \tag{C.89}$$

which is also independent of θ. Applying the same argument after each successor history shows by induction that every future sequence of passively opened constituent signs is conditionally independent of θ given $Y = j$.

The posterior mean of the payoff-relevant scalar state therefore remains constant along every such inspection path. The terminal stopping payoff cannot improve through scalar sign inspection, so the gross value of every policy that only opens those signs is zero. If each inspection has strictly positive cost, immediate stopping strictly dominates every nonempty scalar-sign inspection plan. This proves all claims in the corollary. □

Scope of exact-count sufficiency. The calculation uses one latent scalar state and an exact count of the binary signs generated by that state. It therefore applies directly to a singleton displayed aspect. If a folded aspect d pools fine dimensions $q \in \varpi^{-1}(d)$, observing the pooled total

$$Y_d = \sum_{q \in \varpi^{-1}(d)} Y_q$$

does not generally reveal the allocation $(Y_q)_{q \in \varpi^{-1}(d)}$. Conditional on the known margins, passive opening induces a without-replacement allocation problem; when utility loads differently on the fine latent states, learning which fine dimension generated a known sign can retain value. Thus the corollary exhausts pooled scalar-sign learning, not within-fold allocation learning.

Exactness is also essential. If a rounded or capped report reveals only $Y \in \mathcal{J}$ for a set of at least two feasible counts, then

$$\mathbb{P}(r_1 = 1 \mid Y \in \mathcal{J}, \theta) = \frac{1}{N} \mathbb{E}[Y \mid Y \in \mathcal{J}, \theta], \tag{C.90}$$

which generally depends on θ, because the conditional distribution of Y inside $\mathcal{J}$ depends on p_θ. The interval report is then a non-singleton summary cell, and constituent signs may again move beliefs.

Proof of Corollary 4. If the middle cell reveals a singleton count $Y = y$, then conditional on that cell every binary signal vector with y positives is equally likely under either latent state after conditioning on $Y = y$. For a passively sampled next review,

$$\mathbb{P}(r_1 = 1 \mid Y = y, \theta) = \frac{y}{N},$$

which is independent of θ. More generally, the likelihood ratio after observing the next scalar sign is unchanged, so $\mu^{\mathsf{m}+} = \mu^{\mathsf{m}-} = \mu^{\mathsf{m}}$. Proposition 1 then gives $\mathcal{I}^{\mathsf{m}}(\bar{M}) = 0$.

If the middle cell contains multiple feasible counts and the two posteriors are ordered with both next-signal probabilities positive, Proposition 1 applies directly with $a = \mathsf{m}$. The one-step gross verification value is positive exactly when the stopping margin lies strictly between the favorable and unfavorable next posteriors. □

Proof of Proposition 4. By Proposition 1, verification is optimal if and only if

$$\kappa^W \le \mathcal{I}^a(\bar{M}).$$

If this inequality fails, the consumer chooses the best stopping action immediately. The stopping payoff is

$$S^a = \max\{0, \delta(\mu^a), \bar{M}\}.$$

Any action attaining this maximum is optimal. Thus exit is optimal whenever the maximum is attained by 0, keeping or choosing the current product is optimal whenever it is attained by $\delta(\mu^a)$, and taking the outside opportunity is optimal whenever it is attained by $\bar{M}$. A fixed tie-breaking rule selects among multiple maximizers. If $\kappa^W \le \mathcal{I}^a(\bar{M})$, the value of one more raw signal weakly exceeds its cost, so verification is weakly optimal before making the stopping comparison at the next posterior. □

Proof of Proposition 5. Fix overview realization a. With at most one post-summary inspection, the terminal actions are to keep the current option, take the outside option, or choose nothing. Since the normalized no-purchase option is weakly dominated, the relevant terminal comparison is between $\delta(\mu)$ and $\bar{M}$.

Closed-form posterior and transition objects. Equation (3.7) implies that the posterior after the overview alone is

$$\mu^a = \frac{\mu_0 H_a(0,0;\rho)}{\mu_0 H_a(0,0;\rho) + (1-\mu_0) H_a(0,0;1-\rho)}.$$

After one additional favorable inspected signal,

$$\mu^{a+} = \frac{\mu_0 \rho H_a(1,1;\rho)}{\mu_0 \rho H_a(1,1;\rho) + (1-\mu_0)(1-\rho) H_a(1,1;1-\rho)}.$$

After one additional unfavorable inspected signal,

$$\mu^{a-} = \frac{\mu_0 (1-\rho) H_a(1,0;\rho)}{\mu_0 (1-\rho) H_a(1,0;\rho) + (1-\mu_0)\rho H_a(1,0;1-\rho)}.$$

The predictive probability that the next inspected signal is favorable after overview realization a is

$$\pi^a = \frac{\mu_0 \rho H_a(1,1;\rho) + (1-\mu_0)(1-\rho) H_a(1,1;1-\rho)}{\mu_0 H_a(0,0;\rho) + (1-\mu_0) H_a(0,0;1-\rho)}.$$

Thus the posteriors and transition probabilities are available in closed form for every finite $N > 1$ through the parity-robust H^τ object; what remains non-closed-form in the unrestricted problem is the recursive continuation policy, not the posterior objects themselves.

Expected payoff from starting verification. If the consumer pays κ^W to inspect one signal, the next signal is favorable with probability π^a and unfavorable with probability $1 - \pi^a$. After a favorable signal the posterior rises to μ^{a+}; after an unfavorable signal it falls to μ^{a-}. Therefore the expected payoff from starting verification is

$$\mathcal{V} = -\kappa^W + \pi^a \max\{\delta(\mu^{a+}), \bar{M}\} + (1 - \pi^a) \max\{\delta(\mu^{a-}), \bar{M}\}. \tag{V}$$

Part (a). Suppose the ordering

$$\mu^{a-} < \bar{\mu} \leq \mu^a < \mu^{a+}.$$

Because $\delta(\mu) = \mathbf{x}'\beta - \alpha P + \gamma\mu$ is strictly increasing in μ and the outside option payoff satisfies $\bar{M} = \mathbf{x}'\beta - \alpha P + \gamma\bar{\mu}$, this ordering implies

$$\delta(\mu^{a-}) < \bar{M} \leq \delta(\mu^a) < \delta(\mu^{a+}).$$

Therefore: (i) stopping immediately at μ^a yields $\delta(\mu^a) \geq \bar{M}$, so the consumer keeps the option; (ii) after a favorable signal, $\delta(\mu^{a+}) > \bar{M}$, so the consumer keeps the option; (iii) after an unfavorable signal, $\delta(\mu^{a-}) < \bar{M}$, so the consumer switches to the outside option.

Resolving the max operators in display (V) under these orderings:

$$\mathcal{V} = -\kappa^W + \pi^a \delta(\mu^{a+}) + (1 - \pi^a)\bar{M}.$$

Verification is optimal if and only if $\mathcal{V} \geq \delta(\mu^a)$, i.e.,

$$-\kappa^W + \pi^a \delta(\mu^{a+}) + (1 - \pi^a)\bar{M} \geq \delta(\mu^a).$$

Now substitute the affine forms $\delta(\mu) = \mathbf{x}'\beta - \alpha P + \gamma\mu$ and $\bar{M} = \mathbf{x}'\beta - \alpha P + \gamma\bar{\mu}$. The deterministic component $\mathbf{x}'\beta - \alpha P$ appears on both sides and cancels, leaving

$$-\kappa^W + \pi^a \gamma\mu^{a+} + (1 - \pi^a)\gamma\bar{\mu} \geq \gamma\mu^a,$$

which rearranges to

$$\kappa^W \leq \gamma\left[\pi^a \mu^{a+} + (1 - \pi^a)\bar{\mu} - \mu^a\right].$$

By Lemma 7, Bayesian posteriors satisfy the martingale property

$$\mu^a = \pi^a \mu^{a+} + (1 - \pi^a)\mu^{a-}.$$

Substituting $\pi^a \mu^{a+} = \mu^a - (1 - \pi^a)\mu^{a-}$ into the right-hand side:

$$\gamma\left[\pi^a \mu^{a+} + (1 - \pi^a)\bar{\mu} - \mu^a\right] = \gamma\left[\mu^a - (1 - \pi^a)\mu^{a-} + (1 - \pi^a)\bar{\mu} - \mu^a\right]$$

$$= \gamma(1 - \pi^a)(\bar{\mu} - \mu^{a-}). \tag{C.91}$$

Therefore verification after a favorable overview is optimal if and only if

$$\kappa^W \leq \gamma(1 - \pi^a)(\bar{\mu} - \mu^{a-}),$$

which proves part (a).

Part (b). Suppose instead the ordering

$$\mu^{a-} < \mu^a \leq \bar{\mu} < \mu^{a+}.$$

Now $\delta(\mu^a) \leq \bar{M}$, so stopping immediately yields $\bar{M}$ (the consumer prefers the outside option at the current posterior). After a favorable signal, $\delta(\mu^{a+}) > \bar{M}$, so the consumer keeps the option. After an unfavorable signal, $\delta(\mu^{a-}) < \bar{M}$, so the consumer still takes the outside option.

Resolving the max operators in display (V):

$$\mathcal{V} = -\kappa^W + \pi^a \delta(\mu^{a+}) + (1 - \pi^a)\bar{M}.$$

Verification is optimal if and only if $\mathcal{V} \geq \bar{M}$, i.e.,

$$-\kappa^W + \pi^a \delta(\mu^{a+}) + (1 - \pi^a)\bar{M} \geq \bar{M}.$$

Subtracting $\bar{M}$ from both sides:

$$-\kappa^W + \pi^a \left[\delta(\mu^{a+}) - \bar{M}\right] \geq 0.$$

Substituting $\delta(\mu^{a+}) - \bar{M} = \gamma(\mu^{a+} - \bar{\mu})$:

$$\kappa^W \leq \gamma \pi^a (\mu^{a+} - \bar{\mu}),$$

which proves part (b). □

Remark 2 (The one-step problem as a lower bound)*.* Fix any state (a, m, y) and reduced-form outside value M. Let $W^1(a, m, y; M)$ denote the value of the problem in which at most one additional inspection is allowed, and let $W^\infty(a, m, y; M)$ denote the value of the unrestricted finite-horizon problem with all remaining inspections available. Because the unrestricted consumer can always mimic the one-step policy and then stop, one has

$$W^\infty(a, m, y; M) \geq W^1(a, m, y; M)$$

at every reachable history. In particular, at the post-overview entry state, the one-step verification payoff characterized in Proposition 5 is a lower bound on the unrestricted verification value. The cutoff inequalities in Proposition 5 therefore give sufficient conditions for verification to start in the unrestricted problem as well. Corollary 7 closes the other side at high precision by showing that, for sufficiently large ρ, even the unrestricted continuation value is too small to justify any further inspection.

Derivation for Example 2. In the notation of Proposition 5, the worked example sets $a = 1$ and uses the shorthand

$$\mu^{+} \equiv \mu^{1}, \qquad \mu^{++} \equiv \mu^{1+}, \qquad \mu^{+-} \equiv \mu^{1-}, \qquad \pi^{+} \equiv \pi^{+}1.$$

Write $\zeta(\rho) = \mathbb{P}(a = 1 \mid \theta = 1)$. Under majority aggregation with $N = 3$,

$$\zeta(\rho) = 3\rho^2(1-\rho) + \rho^3 = 3\rho^2 - 2\rho^3.$$

Bayes' rule immediately yields

$$\mu^{+} = \frac{\mu_0\zeta(\rho)}{\mu_0\zeta(\rho) + (1-\mu_0)(1-\zeta(\rho))}.$$

After a favorable overview, an additional favorable inspected signal is compatible with the overview if at least one of the two remaining signals is also favorable. The corresponding likelihoods are

$$L_1(1,1,1) = \rho\left[1-(1-\rho)^2\right] = \rho^2(2-\rho),$$

$$L_0(1,1,1) = (1-\rho)\left[1-\rho^2\right] = (1-\rho)^2(1+\rho),$$

which gives

$$\mu^{++} = \frac{\mu_0\rho^2(2-\rho)}{\mu_0\rho^2(2-\rho) + (1-\mu_0)(1-\rho)^2(1+\rho)}.$$

Similarly, after an unfavorable inspected signal, the two remaining signals must both be favorable. Hence

$$L_1(1,1,0) = \rho^2(1-\rho), \qquad L_0(1,1,0) = \rho(1-\rho)^2,$$

so

$$\mu^{+-} = \frac{\mu_0\rho}{\mu_0\rho + (1-\mu_0)(1-\rho)}.$$

Remark 3 (A notable cancellation in μ^{+-}). After a favorable overview ($a = 1$) and an unfavorable inspected signal, the posterior μ^{+-} equals the posterior that would result from observing a single positive raw signal from the prior alone, without any summary.

This cancellation arises because the likelihood ratio of the joint history

$$(a = 1,\ \text{one inspected negative})$$

relative to the prior equals $\rho/(1-\rho)$, exactly the likelihood ratio of one positive raw signal.

Formally, the posterior odds after the favorable overview and unfavorable inspection are

$$\frac{\mu_0 L_1(1,1,0)}{(1-\mu_0)L_0(1,1,0)} = \frac{\mu_0\rho^2(1-\rho)}{(1-\mu_0)\rho(1-\rho)^2} = \frac{\mu_0\rho}{(1-\mu_0)(1-\rho)},$$

which is the posterior odds from exactly one positive signal. The cancellation is specific to the $N = 3$ majority-rule case and does not extend to larger signal sets.

The probability of a favorable inspected signal conditional on the favorable overview is therefore

$$\pi^+ = \frac{\mu_0\rho^2(2-\rho) + (1-\mu_0)(1-\rho)^2(1+\rho)}{\mu_0\zeta(\rho) + (1-\mu_0)(1-\zeta(\rho))}.$$

Under (C.72), a favorable follow-up signal leads the consumer to keep the option, while an unfavorable follow-up signal leads the consumer to switch to the outside option. Continuing for one more inspection therefore yields

$$-\kappa^W + \pi^+\delta(\mu^{++}) + (1-\pi^+)\bar{M},$$

while stopping immediately yields $\delta(\mu^+)$. Thus continuation is optimal if and only if

$$-\kappa^W + \pi^+\delta(\mu^{++}) + (1-\pi^+)\bar{M} \geq \delta(\mu^+).$$

Substituting $\bar{M} = \mathbf{x}'\beta - \alpha P + \gamma\bar{\mu}$ and cancelling common deterministic utility terms gives

$$\kappa^W \leq \gamma\left[\pi^+\mu^{++} + (1-\pi^+)\bar{\mu} - \mu^+\right].$$

By Lemma 7,

$$\mu^+ = \pi^+\mu^{++} + (1-\pi^+)\mu^{+-},$$

so the right-hand side simplifies to

$$\gamma(1-\pi^+)(\bar{\mu} - \mu^{+-}).$$

This proves (C.73). Rearranging yields

$$\bar{\mu}^* = \mu^{+-} + \frac{\kappa^W}{\gamma(1-\pi^+)},$$

and continuation is optimal if and only if $\bar{\mu} \geq \bar{\mu}^*$.

Proof of Lemma 7. Fix any finite $N > 1$, the neutral majority rule, and any overview realization $a \in \{0,1\}$. Let $\mu^a = \widetilde{\mu}^{AI}(a,0,0)$ denote the posterior after the overview alone, and let $\mu^{a+} = \widetilde{\mu}^{AI}(a,1,1)$ and $\mu^{a-} = \widetilde{\mu}^{AI}(a,1,0)$ denote the posteriors after one additional favorable or unfavorable inspected signal. The posterior after the overview is the conditional expectation of the post-inspection posterior, taken over the two possible signal realizations.

By the law of iterated expectations,

$$\mu^a = \mathbb{E}\big[\widetilde{\mu}^{AI}(a,1,\cdot) \mid a = a\big] = \pi^a \mu^{a+} + (1-\pi^a)\mu^{a-},$$

which is exactly (C.74). This is the standard martingale property of Bayesian posteriors: it holds for any finite N and any aggregation rule, because the law of iterated expectations does not depend on the signal structure. □

Proof of Lemma 8. Let $\varepsilon = 1-\rho$ and let A denote the event that the neutral majority overview is favorable. Conditional on the high-quality state and on the event that the next inspected signal is negative, the remaining $N-1$ signals must be sufficient to preserve a favorable overview.

If $N \geq 3$, one negative signal still leaves a strict positive majority possible whenever all remaining signals are correct. Therefore

$$\mathbb{P}(r_{next} = 0, A \mid \theta = 1) = \varepsilon\{1 + o(1)\},$$

while

$$\mathbb{P}(A \mid \theta = 1) \longrightarrow 1.$$

Thus

$$\mathbb{P}(r_{next} = 0 \mid A, \theta = 1) \sim \varepsilon.$$

Conditional on the low-quality state, a favorable neutral-majority overview requires at least two erroneous positive signals when $N \geq 3$: for odd N this is a strict majority of errors, and for even $N \geq 4$ the lowest accepted state is an exact tie with $N/2 \geq 2$ erroneous positives. Hence $\mathbb{P}(A \mid \theta = 0) = O(\varepsilon^2)$, so Bayes' rule gives $\mathbb{P}(\theta = 0 \mid A) = O(\varepsilon^2) = o(\varepsilon)$. Combining the two latent states yields

$$1 - \pi^+(1,0,0) = \mathbb{P}(r_{next} = 0 \mid A) \sim \varepsilon,$$

so $C_N^+(\mu_0) = 1$ for $N \geq 3$.

For $N = 2$, a high-quality favorable overview has probability $\rho^2 + \rho(1-\rho) = \rho$, and the event that the next inspected signal is negative and the overview is favorable has probability $\frac{1}{2}\rho(1-\rho)$, because the one-positive-one-negative tie is accepted with probability $1/2$. The low-quality state generates a favorable overview with probability $1-\rho = \varepsilon$, and its joint probability of a negative designated signal and a favorable tie is also $\frac{1}{2}\rho(1-\rho)$. Integrating over the two latent states therefore gives

$$\mathbb{P}(r_{\text{next}} = 0, A) = \frac{1}{2}\rho(1-\rho), \tag{C.92}$$

$$\mathbb{P}(A) = \mu_0\rho + (1-\mu_0)(1-\rho), \tag{C.93}$$

$$1 - \pi^+(1,0,0) = \frac{\frac{1}{2}\rho(1-\rho)}{\mu_0\rho + (1-\mu_0)(1-\rho)} \sim \frac{1}{2\mu_0}(1-\rho). \tag{C.94}$$

Hence $C_2^+(\mu_0) = 1/(2\mu_0)$. Reversing states and signs gives

$$C_2^-(\mu_0) = \frac{1}{2(1-\mu_0)}.$$

Both constants are finite and strictly positive for every $\mu_0 \in (0, 1)$.

The unfavorable-overview statement follows by replacing positives with negatives and $\theta = 1$ with $\theta = 0$. This proves (C.75). □

Proof of Corollary 5. Fix any finite $N > 1$ under neutral majority. Consider first $a = 1$. Lemma 8 gives

$$1 - \pi^+(1, 0, 0) \to 0.$$

Moreover, the same lemma gives the first-order expansion

$$1 - \pi^+(1, 0, 0) \sim C^\star(1-\rho) \qquad \text{as } \rho \to 1. \tag{C.95}$$

The argument for $a = 0$ is the sign-reversed analogue: the posterior after an unfavorable overview converges to zero and the probability of a positive next inspected signal satisfies

$$\pi^+(0, 0, 0) \to 0.$$

Thus $d^a(\rho) \to 0$ for both summary realizations.

It remains to show that the one-step value converges to zero uniformly over feasible outside beliefs. The stopping payoff

$$S(\mu; \bar{M}) = \max\{0, \delta(\mu), \bar{M}\}$$

is γ-Lipschitz in μ, uniformly in $\bar{M}$, because it is the maximum of constants and an affine function with slope γ. It is also uniformly bounded on $\mu, \bar{\mu} \in [0, 1]$; let $\bar{B}$ be such a bound. For $a = 1$, $\mu^a \to 1$ and $\mu^{a+} \to 1$, while the possibly nonconvergent payoff after the disconfirming signal is multiplied by $1 - \pi^a \to 0$. Hence

$$\mathcal{I}^{a=1}(\bar{M}) \le \gamma|\mu^{a+} - \mu^a| + 2\bar{B}(1 - \pi^a) \to 0$$

uniformly over feasible outside beliefs. For $a = 0$, the symmetric bound is

$$\mathcal{I}^{a=0}(\bar{M}) \le \gamma|\mu^{a-} - \mu^a| + 2\bar{B}\pi^a \to 0.$$

Thus the primitive one-step inspection value converges to zero uniformly over feasible outside-option beliefs. Since $\kappa^W > 0$, there exists $\rho^* < 1$ such that $\mathcal{I}^a(\bar{M}) < \kappa^W$ for every $\rho > \rho^*$, every feasible $\bar{\mu}$, and both $a \in \{0, 1\}$. Proposition 1 then implies that one-step verification is not optimal. Finally, in the favorable-overview case and in the case-(a) region of Proposition 5, the affine cutoff representation is

$$\bar{\mu}^* = \mu^{a-} + \frac{\kappa^W}{\gamma(1-\pi^a)}.$$

Because $1-\pi^a \to 0$, this affine cutoff eventually exceeds one. The feasible outside-belief region for continuation is therefore empty for all sufficiently high ρ. □

Proof of Corollary 6. Work under neutral majority. From Proposition 5(a), the threshold after a favorable overview is

$$\bar{\mu}^* = \mu^{a-} + \frac{\kappa^W}{\gamma(1-\pi^a)}.$$

Differentiating gives

$$\frac{\partial \bar{\mu}^*}{\partial \rho} = \frac{\partial \mu^{a-}}{\partial \rho} - \frac{\kappa^W}{\gamma}\frac{\partial(1-\pi^a)/\partial\rho}{(1-\pi^a)^2}.$$

The first term is $\partial\mu^{a-}/\partial\rho$, which remains bounded as $\rho \to 1$ because μ^{a-} is a finite polynomial likelihood ratio with a finite limit along the maintained reachable histories.

Define $g(\rho) = 1-\pi^a$. By Lemma 8,

$$g(\rho) \sim C^\star(1-\rho) \qquad \text{as } \rho \to 1.$$

Since g is a rational function of ρ near one and has a simple zero at $\rho = 1$, write

$$g(\rho) = (1-\rho)h(\rho), \qquad h(1) = C^\star > 0,$$

with h differentiable in a neighborhood of one. Therefore $1/g(\rho) = 1/[(1-\rho)h(\rho)]$, which gives

$$\frac{g(\rho)}{1-\rho} \to C^\star \qquad \text{as } \rho \to 1,$$

and consequently

$$(1-\rho)^2\frac{\partial}{\partial\rho}\left(\frac{1}{g(\rho)}\right) \to 1/C^\star.$$

Therefore

$$(1-\rho)^2\frac{\partial}{\partial\rho}\left(\frac{\kappa^W}{\gamma(1-\pi^a)}\right) \to \frac{\kappa^W}{\gamma C^\star} > 0.$$

Multiplying the bounded term $\partial\mu^{a-}/\partial\rho$ by $(1-\rho)^2$ sends it to zero, so

$$(1-\rho)^2\frac{\partial\bar{\mu}^*}{\partial\rho} \to \frac{\kappa^W}{\gamma C^\star} > 0,$$

which implies $\partial\bar{\mu}^*/\partial\rho > 0$ for all ρ sufficiently close to 1. Since continuation after a favorable overview requires $\bar{\mu} \geq \bar{\mu}^*$, an increase in ρ near 1 shrinks the set of outside-option beliefs for which verification is optimal. □

Proof of Corollary 7. Fix an overview realization $a \in \{0,1\}$ under neutral majority. The parity-robust overview precision formula in Appendix C.2.2 implies $\zeta(\rho) \to 1$ as $\rho \to 1$. Hence the post-overview posterior $\mu^a \equiv \widetilde{\mu}^{AI}(a,0,0)$ converges to certainty:

$$\mu^{a=1} \to 1, \qquad \mu^{a=0} \to 0.$$

Because the inspection horizon is finite, the set of continuation histories following the overview is finite and the associated realized payoff differences are uniformly bounded. Let B be a uniform bound on the stopping payoff $S(\mu;\bar{M}) = \max\{0,\delta(\mu),\bar{M}\}$ over $\mu,\bar{\mu} \in [0,1]$. The function S is also γ-Lipschitz in μ uniformly in $\bar{M}$.

For any feasible continuation plan σ after overview realization a, let $E_{a,\sigma,\rho}$ be the event that at least one inspected raw signal contradicts the overview: after $a = 1$, at least one inspected signal is zero; after $a = 0$, at least one inspected signal is one. On the complement of this event, every inspected signal is summary-congruent. Since the horizon is finite and posteriors after summary-congruent histories converge to the same certainty limit as the overview itself, the change in terminal stopping payoff on $E^c_{a,\sigma,\rho}$ is uniformly $o(1)$ over all feasible outside beliefs. On $E_{a,\sigma,\rho}$, the payoff difference is bounded by $2B$. Hence the gross gain from σ satisfies

$$\text{gross gain from } \sigma \le o(1) + 2B\,\mathbb{P}(E_{a,\sigma,\rho}),$$

uniformly over feasible outside beliefs. It is therefore enough to show that $\mathbb{P}(E_{a,\sigma,\rho}) \to 0$ uniformly over feasible continuation plans.

That uniformity follows from finiteness. With finite inspection horizon $\bar{m}$ and a finite action set at each reached history, the number of pure feasible continuation plans after the overview is itself finite, and randomization only convexifies their payoff consequences. Pointwise convergence of $\mathbb{P}(E_{a,\sigma,\rho})$ to zero for each pure feasible plan therefore implies

$$\sup_{\sigma} \mathbb{P}(E_{a,\sigma,\rho}) \to 0$$

simply by taking the maximum over a finite collection of terms tending to 0.

Conditional on a favorable overview, the posterior starts arbitrarily close to one when ρ is sufficiently near one, and the probability of any inspected zero signal converges to zero. Since only finitely many signals can be inspected, the probability that any feasible continuation plan encounters such a disconfirming signal converges to zero. Conditional on an unfavorable overview, the symmetric argument implies that the probability of any inspected positive signal converges to zero. Hence

$$\sup_{\sigma} \mathbb{P}(E_{a,\sigma,\rho}) \to 0 \qquad \text{for each } a \in \{0,1\}.$$

Therefore the gross option value of any finite continuation plan converges uniformly to zero for both $a = 1$ and $a = 0$.

The first additional inspection nevertheless costs at least $c^W(1) > 0$. Hence there exists $\rho^* < 1$ such that for every $\rho > \rho^*$, every feasible outside-option belief $\bar{\mu} \in [0,1]$, and both overview realizations, the net value of continuing is negative at the post-overview entry state. Immediate stopping is therefore optimal, so no additional within-option inspection occurs. If one parameterizes this stop/continue comparison by an outside-option threshold $\bar{\mu}^*$ after a favorable overview, the same conclusion is equivalently stated as $\bar{\mu}^* > 1$ for all $\rho > \rho^*$. By symmetry, after an unfavorable overview the analogous lower continuation threshold lies below zero. □

Proof of Corollary 1. It is enough to construct one set of primitives. Fix any finite $N > 1$ and an overview realization a for which, at some precision $\rho_0 \in (1/2, 1)$, the two one-signal successor posteriors are distinct:

$$\mu^{a-}(\rho_0) < \mu^{a+}(\rho_0).$$

Such a precision exists whenever the next raw signal is informative and both successor histories are feasible. Choose a stopping margin

$$b \in \big(\mu^{a-}(\rho_0), \mu^{a+}(\rho_0)\big).$$

At ρ_0, Proposition 1 implies that the gross one-step Jensen value after realization a is strictly positive. Pick a strictly positive first-inspection cost below that value:

$$0 < c^W(1) < \mathcal{I}^a(\bar{M}; \rho_0),$$

where $\bar{M}$ is chosen to induce the stopping margin b.

All finite-pool posterior and one-step payoff objects are continuous in ρ on any neighborhood of ρ_0 on which the relevant histories have positive likelihood. Hence there is an open interval around ρ_0 on which

$$\mathcal{I}^a(\bar{M}; \rho) > c^W(1),$$

so one-step verification is strictly optimal for intermediate precisions in that interval.

At $\rho = 1/2$, both the overview and the next raw signal are uninformative: the two successor posteriors coincide with the prior, so $\mathcal{I}^a(\bar{M}; 1/2) = 0$. By continuity, for ρ sufficiently close to $1/2$, $\mathcal{I}^a(\bar{M}; \rho) < c^W(1)$, and one-step verification is not optimal. Finally, Corollary 7 implies that for all ρ sufficiently close to one, the gross value of any finite continuation plan, and therefore the one-step value, is below the same positive cost. Verification is again not optimal. This gives the claimed nonmonotonicity. □

Proof of Corollary 2. Let A_N^+ denote the event that the binary majority overview is favorable when the pool size is N, and let A_N^- denote the unfavorable event. For any fixed tie-breaking rule $\tau \in [0, 1]$, the parity-robust consistency probability at entry is

$$H_1^{\tau,N}(0, 0; p) = \mathbb{P}(\mathrm{Bin}(N, p) > N/2) + \mathbf{1}\{N \text{ is even}\}\tau\, \mathbb{P}(\mathrm{Bin}(N, p) = N/2).$$

By the law of large numbers, for fixed $\rho > 1/2$,

$$H_1^{\tau,N}(0,0;\rho) \to 1, \qquad H_1^{\tau,N}(0,0;1-\rho) \to 0,$$

because the tie point mass is $O(N^{-1/2})$ and cannot affect the limit. Symmetrically,

$$H_0^{\tau,N}(0,0;\rho) \to 0, \qquad H_0^{\tau,N}(0,0;1-\rho) \to 1.$$

Therefore the posterior after a favorable overview converges to one and the posterior after an unfavorable overview converges to zero.

Fix any $t \le \bar{m}$ and any raw-signal history h_t of length t. Because $\bar{m}$ is fixed independently of N, the likelihood ratio contribution of h_t is bounded above and below by constants depending only on ρ and $\bar{m}$. To see that the conditioning on h_t cannot undo the overview, write $H_1^{\tau,N}(t,y;p)$ for the favorable-overview consistency probability when the full pool size is N, the bounded history contains y positive signals, and the tie rule is τ. For each fixed $t \le \bar{m}$ and $y \le t$,

$$H_1^{\tau,N}(t,y;\rho) \to 1, \qquad H_1^{\tau,N}(t,y;1-\rho) \to 0$$

as $N \to \infty$, uniformly over the finitely many bounded histories. The strict-tail terms converge exponentially by Chernoff bounds, and the only parity-specific tie term is a single binomial point mass, hence $O(N^{-1/2})$. The raw-history likelihood ratio is bounded because t is bounded, so the posterior odds after A_N^+ and h_t still diverge to infinity. The unfavorable-overview case is symmetric, using $H_0^{\tau,N}$, with posterior odds converging to zero. Hence, uniformly over all histories with $t \le \bar{m}$,

$$\mathbb{P}(\theta = 1 \mid A_N^+, h_t) \to 1, \qquad \mathbb{P}(\theta = 1 \mid A_N^-, h_t) \to 0.$$

Thus every posterior reachable within $\bar{m}$ post-summary inspections converges to the same certainty limit as the overview itself.

The stopping payoff $S(\mu;\bar{M})$ is γ-Lipschitz in μ uniformly over feasible outside-option beliefs. It is also uniformly bounded on $\mu, \bar{\mu} \in [0,1]$. Therefore the gross improvement in terminal stopping payoff from any continuation plan that uses at most $\bar{m}$ post-summary inspections is $o(1)$, uniformly over feasible outside-option beliefs and over all such plans. Since the first inspection costs at least $\underline{c} > 0$, there exists N^* such that for every $N \ge N^*$ this gross improvement is strictly smaller than the first inspection cost after either summary realization. Immediate stopping is then optimal at the post-summary entry state. □

Proof of Benchmark Observation. Fix the product, the finite signal pool, the cost schedule, and the outside value $\bar{M}$. Any no-summary strategy is a feasible AI strategy: after observing the free summary, the consumer can ignore it and choose exactly the same stopping and raw-review inspection actions as in the no-summary environment. The distribution of subsequently inspected raw signals, conditional on the underlying finite pool, is unchanged by ignoring the summary, and the summary carries no cost. Therefore the AI value is at

least the no-summary value:
$$V^A(\bar{M}) \geq V^N(\bar{M}).$$
The inequality can be strict whenever the summary changes the optimal within-product action on a set of summary realizations with positive probability. □

Proof of Proposition 7. At AI state (a, m, y), the best stopping payoff is

$$S(\widetilde{\mu}^{AI}(a, m, y); \bar{M}) = \max\{0, \delta(\widetilde{\mu}^{AI}(a, m, y)), \bar{M}\}.$$

By the within-product Bellman equation, continuation is optimal if and only if

$$\mathcal{K}(a, m, y; \bar{M}) \geq S(\widetilde{\mu}^{AI}(a, m, y); \bar{M}).$$

Adding $c^W(m+1)$ to both sides yields

$$c^W(m+1) \leq \mathcal{K}(a, m, y; \bar{M}) + c^W(m+1) - S(\widetilde{\mu}^{AI}(a, m, y); \bar{M}).$$

The right-hand side is exactly $\Delta^A(a, m, y; \bar{M}) = \bar{c}(a, m, y; \bar{M})$. Hence continuation is optimal if and only if $c^W(m+1) \leq \bar{c}(a, m, y; \bar{M})$. Holding the continuation value fixed, a higher current marginal cost raises only the left-hand side, so it weakly lowers the incentive to continue. □

D Finite-Slot Multidimensional Model

D.1 Dynamic Cutoff and Posterior Integration

This subsection expands the updating, continuation, and taste-region objects from Section 5 step by step. The goal is to show exactly where the new finite-pool topic-arrival friction enters relative to the scalar benchmark and to record the closed-form objects available at a fixed history.

D.1.1 Recovered summary rules and finite-state tractability

The majority/top-K rule in the body is the analytical benchmark, not a maintained claim about a deployed interface. Let z collect public category, product, and corpus metadata available to both the researcher and the consumer. A general summary technology is a known finite-output rule

$$g_\phi : (R, M, z) \longmapsto \omega^\phi = \left(\mathcal{S}^\phi, (a_d^\phi)_{d \in \mathcal{S}^\phi}\right), \tag{D.1}$$

where ϕ is a finite-dimensional parameter vector. The realized output ω^ϕ records both the displayed set and its sentiment cells. The benchmark above is the special case $g_{\phi^{\text{maj}}}$ that selects the top K positive coverage counts and uses the componentwise majority map in (5.4).

One transparent empirical parameterization separates dimension selection from sentiment aggregation. Let

$$Y_d(R) = \sum_{n:m_{nd}=1} r_{nd}, \qquad v_d(R, M) = \frac{2Y_d(R) - C_d(M)}{C_d(M)} \tag{D.2}$$

for every positive-prevalence dimension. A candidate selection score is

$$B_d^{\phi} = \beta_M \log(1 + C_d(M)) + \beta_V |v_d| + \beta_H \operatorname{Help}_d + \beta_R \operatorname{Rec}_d + \beta_X' X_d, \tag{D.3}$$

where Help_d, Rec_d, and X_d are, respectively, pre-specified helpfulness, recency, and other observable features. The rule can display the $K_\phi(z)$ highest-scoring dimensions, with a fixed tie rule. Conditional on selection, a three-cell aggregation rule is

$$a_d^{\phi} = \begin{cases} +, & v_d \geq \vartheta_{+,c}, \\ \mathsf{m}, & \vartheta_{-,c} < v_d < \vartheta_{+,c}, \\ -, & v_d \leq \vartheta_{-,c}, \end{cases} \tag{D.4}$$

where the thresholds may vary by pre-specified category c and, if supported by the data, review-pool size. This example is deliberately interpretable rather than exhaustive: g_ϕ can instead be any recovered finite rule whose inputs are observed in the corpus audit.

For posterior calculations, allow $Q_\phi(\omega \mid R, M, z)$ to denote the probability that the technology reports output ω from a realized corpus. For a deterministic estimated rule, Q_ϕ is simply the indicator of (D.1); the notation also accommodates randomized tie-breaking. The rule is empirically estimable from observed pairs of review corpora and displayed summaries, but it is treated as known by consumers once it has been estimated or calibrated.

There is a useful tractable subclass. Suppose selection is *sign-invariant*. Thus the selected set is $\mathcal{S}^{\phi}(M, z)$. Let selected dimensions use a componentwise count kernel $q_d^{\phi}(a \mid Y, C, z)$. After m opened signals on dimension d, y of which are positive, define

$$H_{da}^{\phi}(m, y; p, C, z) = \sum_{b=0}^{C-m} \binom{C-m}{b} p^b (1-p)^{C-m-b} q_d^{\phi}(a \mid y + b, C, z). \tag{D.5}$$

This is the probability that the remaining signals map into cell a. Majority aggregation makes q_d^{ϕ} an upper or lower tail, and the three-cell threshold rule makes it an interval. If selection itself depends on realized sentiment, as in the β_V term of (D.3), this componentwise factorization is no longer valid. The general integrated likelihood below remains exact, but the model should then be solved by finite-state enumeration rather than advertised as closed form.

Proposition 8 (Nesting and tractability under a recovered summary rule). Fix finite N

and D, a known coverage prior G, and a known output kernel Q_ϕ. Then the integrated likelihood in (5.7), the posterior, and the within-option Bellman problem are all finite-state objects. In particular, the exact dynamic cutoff in Proposition 9 applies under every such rule after conditioning on ω.

If, in addition, selection is sign-invariant and aggregation is componentwise as in (D.5), let $c_{d,t}$ and $y_{d,t}$ be the numbers of opened signals on dimension d and favorable opened signals among them, and set

$$p_d(\theta_d) = \begin{cases} \rho_d, & \theta_d = 1, \\ 1-\rho_d, & \theta_d = 0. \end{cases}$$

For compactness, define the observed-sign and summary-cell factors

$$\ell_{d,t}(\theta_d) = p_d(\theta_d)^{y_{d,t}}\big(1-p_d(\theta_d)\big)^{c_{d,t}-y_{d,t}},$$
$$\mathsf{h}^{\phi}_{d,t}(\theta_d, M) = H^{\phi}_{da_d}\big(c_{d,t}, y_{d,t}; p_d(\theta_d), C_d(M), z\big).$$

The integrated likelihood then factorizes as

$$\mathcal{L}^{\phi}(\omega, h_t \mid \theta, M, z) = \mathbb{1}\{M \text{ is consistent with } h_t\}\mathbb{1}\{\mathcal{S}^{\phi}(M,z) = \mathcal{S}^{\phi}\} \times \prod_{d=1}^{D} \ell_{d,t}(\theta_d) \prod_{d\in\mathcal{S}^{\phi}} \mathsf{h}^{\phi}_{d,t}(\theta_d, M). \tag{D.6}$$

Thus strict-majority and threshold-band rules retain the same binomial-tail or binomial-interval calculations as benchmark special cases. A rule with sentiment-dependent selection remains exactly defined by (5.7), but generally requires finite-state enumeration rather than a componentwise closed form.

The interpretation is practical. Estimating a richer empirical summary rule does not invalidate the theory; it changes how the likelihood is computed. When the recovered rule keeps sign-invariant selection and componentwise sentiment aggregation, the paper's binomial-tail machinery remains available. When the rule lets selection depend on sentiment or other corpus features, the same Bayesian object is still exact, but it must be computed by enumeration rather than written as a compact closed form. The finite-state and factorization claims are proved in the derivations below. These general-kernel claims should be distinguished from the mechanism-specific selection results: self-concealing omission and the observable count bound use top-K prevalence selection, while the current prevalence-design and slot-scarcity statements use their additional stated restrictions. They are not implications of an arbitrary finite Q_ϕ.

Operational multidimensional statements moved from the main paper

The main paper uses the following results as operational support rather than as part of the main theorem spine. They are stated here using the notation of the main text and proved in the derivations that follow.

Corollary 9 (Observable top-K omission bound). Maintain top-K prevalence selection with fixed, known tie-breaking, and suppose the summary reveals the displayed coverage counts. If the summary fills all K slots, let

$$c_*(\omega^C) = \min_{s\in\mathcal{S}^K} C_s(M).$$

Every omitted dimension $d \notin \mathcal{S}^K$ then satisfies

$$C_d(M) \le c_*(\omega^C), \qquad \eta_{d,t}(\omega^C, h_t) \le \frac{c_*(\omega^C) - c_{d,t}}{N-t}. \tag{D.7}$$

At entry this becomes $\eta_{d,0} \le c_*(\omega^C)/N$. If fewer than K dimensions have positive coverage, every omitted dimension has zero coverage and hence $\eta_{d,t} = 0$ at every reachable history. The count inequality is weak because a dimension tied at the marginal displayed count can lose the fixed tie-break.

For either model, let $\mathcal{X}(s_t)$ be the positive-probability next-realization set, let $s_t(x)$ be its successor, and write $[L(s), U(s)]$ for the terminal support hull. Define the successor-absorption layer

$$\mathcal{E}_1(s_t) = \big(L(s_t), U(s_t)\big) \setminus \bigcup_{x\in\mathcal{X}(s_t)} \big(L(s_t(x)), U(s_t(x))\big). \tag{D.8}$$

Corollary 10 (Successor absorption and exact dynamic verification). Under the affine stopping-payoff assumptions of Theorem 2 in the scalar model or Theorem 4 in the multidimensional model, suppose every remaining inspection cost is strictly positive. If $b \in \mathcal{E}_1(s_t)$, then the current state is decision-insufficient but every reachable successor $s_t(x)$ is decision-sufficient. Consequently,

$$V(s_t(x)) = S(s_t(x)) \qquad \text{for every } x \in \mathcal{X}(s_t), \tag{D.9}$$

and the Bellman equation reduces to

$$V(s_t) = \max\big\{S(s_t), -c^W(t+1) + \mathbb{E}[S(s_t(X_{t+1})) \mid s_t]\big\}. \tag{D.10}$$

Hence the exact dynamic gross cutoff equals the one-step Jensen value,

$$\bar{c}(s_t) = \mathbb{E}[V(s_t(X_{t+1})) \mid s_t] - S(s_t) = \mathcal{I}(s_t), \tag{D.11}$$

and continuation is optimal if and only if $c^W(t+1) \le \mathcal{I}(s_t)$.

Corollary 11 (Geometry of successor absorption). Use the exactness layer $\mathcal{E}_1(s_t)$ in (D.8). Under neutral scalar majority aggregation, if $m \le N-2$, then $\mathcal{E}_1(a, m, y)$ has Lebesgue measure zero and contains at most one terminal posterior value. At $m = N-1$, whenever both terminal successors are reachable,

$$\mathcal{E}_1(a, N-1, y) = \big(\mu^{\min}(a, N-1, y), \mu^{\max}(a, N-1, y)\big).$$

In the multidimensional model, the exactness layer contains a positive-length interval whenever successor terminal-index hulls are branch-separated. In particular, if there are $u < v$ such that every successor hull lies weakly below u or weakly above v, with positive-probability successors on both sides, then $(u, v) \subseteq \mathcal{E}_1(s_t)$.

Corollary 12 (Cost inequalities on and off the exactness layer). At a state with $b \in \mathcal{E}_1(s_t)$, unrestricted opening reveals $c^W(t+1) \leq \mathcal{I}(s_t)$, while stopping reveals the reverse weak inequality, up to indifference. If information primitives and cost are fixed, utility scale is known or normalized, and exogenous variation moves b continuously through an interior switching boundary b^* while remaining inside the layer, then

$$c^W(t+1) = \mathcal{I}(s_t; b^*).$$

Outside the layer, stopping still implies $c^W(t+1) \geq \bar{c}(s_t) \geq \mathcal{I}(s_t)$, whereas opening implies only $c^W(t+1) \leq \bar{c}(s_t)$. Thus an opener does not supply a one-step cost upper bound without an exactness condition.

Corollary 13 (Folding–nondisplay equivalence). In the local smoothed approximation, a folded displayed aspect and nondisplay are verification-equivalent for taste vector γ whenever they leave the same posterior-predictive covariance in the consumer's within-fold taste direction. In particular, if a broad chip leaves $\gamma_G' \Sigma^{\phi}_{G,t} \gamma_G$ unchanged relative to nondisplay, then displaying that chip does not reduce the local within-fold verification component for that consumer.

Corollary 14 (Capacity threshold for uniform large-pool verification elimination). Maintain componentwise binary majority, sign-invariant selection, independent priors across dimensions, coverage-quality independence, a correctly specified coverage prior, and the finite-D coverage-share convergence used in Main Paper Theorem 3, but allow display capacity to vary with pool size, writing K_N. Let

$$\mathcal{D}^+ = \{d : \lambda_d > 0\}, \qquad D^+ = |\mathcal{D}^+|,$$

and fix a finite inspection horizon $\bar{m}$, a positive lower bound on its inspection costs, and a bounded taste class $\Gamma \subset \mathbb{R}^D_+$. There is no folding.

(i) If $K_N \geq D^+$ eventually, the gross value of every continuation plan using at most $\bar{m}$ inspections converges to zero uniformly over $\gamma \in \Gamma$ and stopping margins. All such consumers therefore stop immediately for sufficiently large N.

(ii) If $K_N = K < D^+$ eventually and the limiting top-K cutoff is strict, there is an omitted dimension and a nonempty open set of type-margin pairs for which one-step verification has a strictly positive large-pool limit. Sufficiently low positive costs then induce verification for all sufficiently large N on that open set.

Thus review volume eliminates post-summary verification uniformly only when display capacity covers every positive-share experience dimension.

Theorem 3 (Three-cell band resolution threshold). Fix a displayed dimension d with $C_d(M) \to \infty$ in probability and prior $\mu_{d,0} \in (0,1)$. Assume that selection and displayed coverage are uninformative about θ_d. Let the displayed sentiment be generated by a symmetric share rule with

$$0 < q_\ell < \frac{1}{2} < q_h < 1, \qquad q_\ell = 1 - q_h,$$

reporting $+$ when $Z_d = Y_d/C_d(M) \geq q_h$, $-$ when $Z_d \leq q_\ell$, and m otherwise. Assume $\rho_d \neq q_h$.

(i) If $\rho_d > q_h$, the endpoint cell identifies the true state with probability approaching one, and any fixed number of additional signals about d has vanishing gross verification value.

(ii) If $\rho_d \in (1/2, q_h)$, the mixed cell occurs with probability approaching one under either state and its posterior converges to the prior. If the next review covers d with positive limiting probability, the two one-signal posteriors straddle the dimension-specific margin, and the limiting review experiment is focal-separable for d, the full limiting one-step verification value after m is strictly positive. Sufficiently low positive first-inspection costs therefore induce verification even though d is displayed.

Under binary majority $q_h = 1/2$, so the mixed-band case is impossible for $\rho_d > 1/2$.

For a post-summary state s_t, let $\mathfrak{S}(s_t)$ be the finite set of reachable future posterior states and $\mathfrak{T}(s_t) \subseteq \mathfrak{S}(s_t)$ the terminal descendants reached after revealing the remaining pool. Define

$$\Lambda_t^{\min} = \min_{s' \in \mathfrak{S}(s_t)} \sum_d \gamma_d \mu_d(s'), \qquad \Lambda_t^{\max} = \max_{s' \in \mathfrak{S}(s_t)} \sum_d \gamma_d \mu_d(s'),$$

and define $\Lambda_t^{T,\min}$ and $\Lambda_t^{T,\max}$ by replacing $\mathfrak{S}(s_t)$ with $\mathfrak{T}(s_t)$. These are the operational objects used in the next theorem.

Theorem 4 (Decision sufficiency and multidimensional support geometry). Fix any known finite truthful summary rule Q_ϕ and post-summary state s_t. Under the affine index structure of the main paper:

(i) the extrema over all reachable future decision indices equal the extrema over terminal descendants,

$$[\Lambda_t^{\min}, \Lambda_t^{\max}] = [\Lambda_t^{T,\min}, \Lambda_t^{T,\max}]; \tag{D.12}$$

(ii) define the gross value of revealing the entire remaining pool by

$$\mathcal{J}^T(s_t;\gamma) = \mathbb{E}[S(\bar{s};\gamma) \mid s_t] - S(s_t;\gamma), \qquad \bar{s} \in \mathfrak{T}(s_t). \tag{D.13}$$

If $b \notin (\Lambda_t^{T,\min}, \Lambda_t^{T,\max})$, every reachable index lies on one affine branch of the stopping payoff, $\mathcal{J}^T(s_t;\gamma) = 0$, and strictly positive inspection costs imply immediate stopping. If $\Lambda_t^{T,\min} < b < \Lambda_t^{T,\max}$, then $\mathcal{J}^T(s_t;\gamma) > 0$, and verification starts whenever

$$\sum_{\ell=t+1}^{N} c^W(\ell) < \mathcal{J}^T(s_t;\gamma). \tag{D.14}$$

When the terminal interval is nondegenerate, the strict straddle and positive gross value hold on a nonempty open set of type-margin pairs.

(iii) every successor has a weakly smaller reachable interval,

$$[\Lambda_{t+1}^{\min}, \Lambda_{t+1}^{\max}] \subseteq [\Lambda_t^{\min}, \Lambda_t^{\max}]; \tag{D.15}$$

(iv) if $C^{1,MD}$ is the one-step-and-stop continuation region, $C^{\infty,MD}$ the unrestricted region, and $\mathcal{B}^{R,MD} = \{s_t : \Lambda_t^{\min} < b < \Lambda_t^{\max}\}$, then

$$C^{1,MD} \subseteq C^{\infty,MD} \subseteq \mathcal{B}^{R,MD}. \tag{D.16}$$

Outside $\mathcal{B}^{R,MD}$, the dynamic gross cutoff is zero; inside it, $\bar{c}^{MD}(s_t;\gamma) \geq \mathcal{I}^{MD}(s_t;\gamma)$.

In the componentwise no-spillover benchmark, once no payoff-relevant dimension can arrive in a future review, the reachable interval is degenerate and this zero-value region is absorbing.

Proposition 9 (Exact multidimensional initiation cutoff). At any post-summary state $s_t = (a, h_t)$ with at least one unopened review, within-option search starts in the full dynamic problem if and only if

$$c^W(t+1) \leq \bar{c}^{MD}(s_t;\gamma). \tag{D.17}$$

If the inequality fails, immediate stopping—exit, choosing the current option, or moving to the outside option—is optimal. The cutoff is exact but generally recursive because $V(s_{t+1})$ embeds all future inspection decisions.

Proposition 10 (Primitive multidimensional verification). At any post-summary state $s_t = (a, h_t)$, if the consumer is restricted to at most one additional review before stopping, opening one more review is optimal if and only if

$$c^W(t+1) \leq \mathcal{I}^{MD}(s_t;\gamma). \tag{D.18}$$

Fix a focal uncovered dimension $d \notin \mathcal{S}^K$ with $\gamma_d > 0$. If the next review does not cover d, let $S_d^0(s_t)$ denote the expected stopping payoff after the full sparse review is observed, conditional on noncoverage of d. If the next review covers d, let $S_d^+(s_t)$ and $S_d^-(s_t)$ denote the expected stopping payoffs conditional, respectively, on the dimension-d realization being favorable or unfavorable, integrating over any other dimensions covered by the same review. The one-step payoff from continuing for the chance to learn about dimension d is

$$\underline{\mathcal{K}}_d(s_t) = -c^W(t+1)+\eta_{d,t}(s_t)\Big[\pi^+_{d,t}(s_t)S^+_d(s_t)+\big(1-\pi^+_{d,t}(s_t)\big)S^-_d(s_t)\Big]+\big(1-\eta_{d,t}(s_t)\big)S^0_d(s_t). \quad \text{(D.19)}$$

Proposition 11 (One-step verification condition). At any post-summary state $s_t = (a, h_t)$ and for any uncovered dimension d with $\gamma_d > 0$, continuing for one more review is optimal whenever

$$\begin{aligned} c^W(t+1) \le\ & \eta_{d,t}(s_t)\Big(\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big] + \big(1-\pi^+_{d,t}(s_t)\big)\big[S^-_d(s_t) - S(s_t;\gamma)\big]\Big) \\ & + \big(1-\eta_{d,t}(s_t)\big)\big[S^0_d(s_t) - S(s_t;\gamma)\big]. \end{aligned} \quad \text{(D.20)}$$

If the no-coverage branch is payoff-neutral or favorable, $S^0_d(s_t) \ge S(s_t;\gamma)$, the simpler sufficient condition that drops the last term is also valid.

Proposition 12 (Coverage-residual ordering under homogeneous omitted dimensions). Hold fixed $(\mu_{d,0}, \rho_d)_{d=1}^D$, the prior over coverage matrices G, the within-option cost schedule, and K. Consider two consumers with preference vectors γ and γ' at the same post-summary history (a, h_t). Suppose uncovered dimensions have a common one-step payoff wedge per unit of taste weight: for every $d \notin \mathcal{S}^K$ there is a common scalar $B_t(s_t) \ge 0$ such that the one-step benefit of dimension d is $\gamma_d\eta_{d,t}(s_t)B_t(s_t)$. If

$$\mathrm{RL}_t(\gamma; a, h_t) \le \mathrm{RL}_t(\gamma'; a, h_t),$$

then the residual-learnability inequality is sufficient for the omitted-dimension component of the one-step verification cutoff for γ to be weakly lower than the corresponding component for γ'. If covered-dimension one-step gains are zero or identical across the two taste vectors, the total one-step verification cutoff for γ is weakly lower than the total one-step verification cutoff for γ'.

Taste regions. When the no-coverage branch is payoff-neutral or favorable, the focal one-step gain from uncovered dimension d can be written as

$$\mathcal{B}_d(s_t;\gamma) = \eta_{d,t}(s_t)\Big(\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big] + (1-\pi^+_{d,t}(s_t))\big[S^-_d(s_t) - S(s_t;\gamma)\big]\Big). \quad \text{(D.21)}$$

The associated sufficient focal inspection region is

$$\mathcal{T}^I_d(s_t) = \big\{\gamma \in \mathbb{R}^D_+ : c^W(t+1) \le \mathcal{B}_d(s_t;\gamma)\big\}. \quad \text{(D.22)}$$

Under local separability,

$$S^+_d(s_t) - S(s_t;\gamma) = \gamma_d\Delta^+_{d,t}(s_t), \qquad S^-_d(s_t) - S(s_t;\gamma) = \gamma_d\Delta^-_{d,t}(s_t), \quad \text{(D.23)}$$

so the focal inspection region becomes

$$\mathcal{T}_d^I(s_t) = \Big\{\gamma \in \mathbb{R}_+^D : \gamma_d \eta_{d,t}(s_t)\Big[\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t)\Big] \geq c^W(t+1)\Big\}. \quad \text{(D.24)}$$

Whenever the bracketed term is strictly positive, this is equivalent to the local threshold

$$\gamma_d \geq \gamma_d^*(s_t) = \frac{c^W(t+1)}{\eta_{d,t}(s_t)\Big[\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t)\Big]}. \quad \text{(D.25)}$$

Proposition 13 (One-topic taste hyperplane). Suppose each opened review covers at most one payoff-relevant dimension and that the local separability condition (D.23) holds for every dimension. Define the per-unit taste gain

$$\Omega_{d,t}(s_t) = \eta_{d,t}(s_t)\Big[\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t)\Big]. \quad \text{(D.26)}$$

Then the exact one-step taste region is the half-space

$$\mathcal{T}^1(s_t) = \left\{\gamma \in \mathbb{R}_+^D : \sum_{d=1}^{D} \gamma_d \Omega_{d,t}(s_t) \geq c^W(t+1)\right\}. \quad \text{(D.27)}$$

On the normalized simplex $\Gamma = \{\gamma \geq 0 : \sum_d \gamma_d = 1\}$, if $c^W(t+1) > \max_d \Omega_{d,t}(s_t)$, no taste type starts one-step verification; if $c^W(t+1) \leq \min_d \Omega_{d,t}(s_t)$, every taste type starts; otherwise the boundary is the hyperplane $\sum_d \gamma_d \Omega_{d,t}(s_t) = c^W(t+1)$.

General-rule local approximation. For a known finite rule Q_ϕ, let

$$\Lambda_t = \gamma' \mu_t, \qquad O = \max\{0, \bar{M}\}, \qquad b = O - (\mathbf{x}'\beta - \alpha P).$$

For fixed smoothing scale $\varepsilon > 0$, define the smooth stopping payoff

$$S_\varepsilon(\Lambda) = O + \varepsilon \log\left(1 + \exp\left(\frac{\Lambda - b}{\varepsilon}\right)\right). \quad \text{(D.28)}$$

Let $\Sigma_t^\phi = \mathrm{Var}_\phi(\mu_{t+1} \mid s_t)$, and define the smoothed one-step value

$$\mathcal{I}_\varepsilon^\phi(s_t;\gamma) = \mathbb{E}_\phi[S_\varepsilon(\Lambda_{t+1}) \mid s_t] - S_\varepsilon(\Lambda_t).$$

Proposition 14 (General-rule local approximation). For any known finite rule Q_ϕ and fixed $\varepsilon > 0$,

$$\mathcal{I}_\varepsilon^\phi(s_t;\gamma) = \frac{1}{2\varepsilon}\sigma\left(\frac{\Lambda_t - b}{\varepsilon}\right)\left[1 - \sigma\left(\frac{\Lambda_t - b}{\varepsilon}\right)\right]\gamma'\Sigma_t^\phi\gamma + R_\varepsilon^\phi(s_t), \quad \text{(D.29)}$$

where $\sigma(z) = 1/(1+e^{-z})$ and the remainder satisfies

$$\left|R_\varepsilon^\phi(s_t)\right| \le \frac{\overline{S}_\varepsilon^{(3)}}{6}\,\mathbb{E}_\phi\big[|\Lambda_{t+1}-\Lambda_t|^3 \mid s_t\big], \qquad \overline{S}_\varepsilon^{(3)} = \sup_\lambda\left|S_\varepsilon^{(3)}(\lambda)\right|. \tag{D.30}$$

Thus, for fixed ε, the leading local verification value is the curvature of the smoothed stopping payoff at the decision margin times the taste-weighted posterior variance left by the summary rule. The approximation is local in belief movements relative to the smoothing scale; it is not uniform as $\varepsilon \to 0$.

Corollary 15 (Local residual-variance design). Fix a state s_t, a taste vector γ, and two known finite summary rules ϕ and ϕ' evaluated at the same current decision index Λ_t. Up to the third-order remainders in (D.30), rule ϕ generates weakly lower local smoothed verification value than rule ϕ' if and only if

$$\gamma'\Sigma_t^\phi\gamma \le \gamma'\Sigma_t^{\phi'}\gamma.$$

For a population of taste vectors, the corresponding local design objective is to reduce the curvature-weighted average of this taste-weighted posterior variance.

Proposition 15 (Folding decomposition). Let $G \subset \{1,\dots,D\}$ be a group of fine dimensions folded into one displayed aspect. Holding outside-group beliefs fixed, let γ_G be the vector of taste weights on the folded group and $\Sigma_{G,t}^\phi$ the posterior-predictive covariance matrix of next beliefs over the fine dimensions in G. Then the folded group's local smoothed one-step verification component is

$$\mathcal{I}_\varepsilon^{\phi,G}(s_t;\gamma) = \mathcal{W}_\varepsilon(\Lambda_t)\,\gamma_G'\Sigma_{G,t}^\phi\gamma_G + R_\varepsilon^{\phi,G}(s_t), \qquad \mathcal{W}_\varepsilon(\Lambda_t) = \frac{1}{2\varepsilon}\sigma\!\left(\frac{\Lambda_t-b}{\varepsilon}\right)\left[1-\sigma\!\left(\frac{\Lambda_t-b}{\varepsilon}\right)\right]. \tag{D.31}$$

Thus a broad displayed cell leaves no local second-order verification motive inside the fold exactly when it leaves zero posterior variance in the consumer's taste direction, $\gamma_G'\Sigma_{G,t}^\phi\gamma_G = 0$.

Consumer-relevance design. Suppose displaying dimension d changes its per-unit one-step gain from the uncovered value $\Omega_{d,t}^U(s_t)$ to the covered value $\Omega_{d,t}^C(s_t)$, with

$$0 \le \Omega_{d,t}^C(s_t) \le \Omega_{d,t}^U(s_t) \qquad \text{for every } d. \tag{D.32}$$

For a display set S and taste distribution F, define expected one-step inspection pressure by

$$\mathcal{P}(S) = \mathbb{E}_F\left[\sum_{d\in S}\gamma_d\Omega_{d,t}^C(s_t) + \sum_{d\notin S}\gamma_d\Omega_{d,t}^U(s_t)\right]. \tag{D.33}$$

Proposition 16 (Consumer-relevance slot ranking). In the one-topic local-separability environment satisfying (D.32), the display set of size K that minimizes $\mathcal{P}(S)$ consists of the K dimensions with the largest consumer-relevance indices

$$\mathcal{R}_{d,t} = \mathbb{E}_F[\gamma_d]\left(\Omega^U_{d,t}(s_t) - \Omega^C_{d,t}(s_t)\right),$$

which are the same indices defined in the main text in (5.29). A top-K prevalence rule is inspection-pressure optimal if and only if its selected dimensions coincide with the top-K dimensions under $\mathcal{R}_{d,t}$, up to ties.

Proof of the exact dynamic initiation cutoff

At a post-summary state $s_t = (a, h_t)$, the consumer can either stop immediately and receive

$$S(s_t;\gamma) = \max\{0, \delta(a, h_t), \bar{M}\},$$

or pay the current marginal inspection cost $c^W(t+1)$ and move to the posterior continuation state s_{t+1}. The Bellman equation is therefore

$$V(s_t) = \max\left\{S(s_t;\gamma),\ -c^W(t+1) + \mathbb{E}[V(s_{t+1}) \mid s_t]\right\}.$$

Continuation is optimal if and only if the second term in the maximum weakly exceeds the first:

$$-c^W(t+1) + \mathbb{E}[V(s_{t+1}) \mid s_t] \geq S(s_t;\gamma).$$

Rearranging gives

$$c^W(t+1) \leq \mathbb{E}[V(s_{t+1}) \mid s_t] - S(s_t;\gamma) \equiv \bar{c}^{MD}(s_t;\gamma),$$

which proves Proposition 9. The cutoff is recursive because $V(s_{t+1})$ is itself the value of a future stopping problem.

Joint posterior over latent quality and coverage

Fix the product. Let the post-summary history after t opened reviews be (ω, h_t), where:

- $\omega = (\mathcal{S}, (a_d)_{d\in\mathcal{S}})$ is the realized displayed set and sentiment vector,
- h_t records which t reviews have been opened and the sparse vectors observed on those reviews,
- $\theta = (\theta_1, \ldots, \theta_D) \in \{0,1\}^D$ is the latent quality vector,
- $M = \{m_{nd}\}_{n\leq N, d\leq D}$ is the unknown coverage matrix.

  Bayes' rule gives

$$\mathbb{P}(\theta, M \mid \omega, h_t, z;\phi) = \frac{\mathbb{P}(\omega, h_t \mid \theta, M, z;\phi)\,\mathbb{P}(\theta, M)}{\sum_{\theta'}\sum_{M'}\mathbb{P}(\omega, h_t \mid \theta', M', z;\phi)\,\mathbb{P}(\theta', M')}. \tag{D.34}$$

By prior independence between latent quality and coverage,

$$\mathbb{P}(\theta, M) = \mathbb{P}(\theta)\, G(M). \tag{D.35}$$

Conditional on (θ, M), the opened reviews and the summary are still random because the unopened signal signs have not been realized from the consumer's perspective. Let $\mathcal{R}(M; h_t)$ be the finite set of full sparse signal matrices consistent with coverage matrix M and the opened-review history h_t. The integrated likelihood of the observed summary and opened history is

$$\mathcal{L}^{\phi}(\omega, h_t \mid \theta, M, z) \equiv \mathbb{P}(\omega, h_t \mid \theta, M, z; \phi) = \sum_{R \in \mathcal{R}(M; h_t)} \mathbb{P}(R \mid \theta, M) Q_{\phi}(\omega \mid R, M, z). \tag{D.36}$$

The likelihood inside the sum factors across covered signs. Let

$$p_d(\theta_d) = \begin{cases} \rho_d, & \theta_d = 1, \\ 1 - \rho_d, & \theta_d = 0, \end{cases}$$

and recall that missing entries $m_{nd} = 0$ carry no sign likelihood. For any full sparse realization R consistent with M,

$$\mathbb{P}(R \mid \theta, M) = \prod_{n=1}^{N} \prod_{d=1}^{D} \left[p_d(\theta_d)^{r_{nd}} \left(1 - p_d(\theta_d)\right)^{1 - r_{nd}} \right]^{m_{nd}}. \tag{D.37}$$

Opened reviews restrict the feasible set $\mathcal{R}(M; h_t)$. Unopened signs are summed out through (D.36), which is why the posterior conditions on (ω, h_t) without conditioning on an unobserved full signal matrix. For a deterministic rule, Q_{ϕ} is one or zero; randomized tie-breaking simply makes it a nondegenerate kernel. This likelihood is the exact multi-dimensional analogue of the H-function: it integrates over the unobserved signs in the finite review pool that remain consistent with the observed summary output and the opened reviews.

Substituting (D.35) and (D.36) into (D.34) yields

$$\mathbb{P}(\theta, M \mid \omega, h_t, z; \phi) \propto \mathbb{P}(\theta)\, G(M)\, \mathcal{L}^{\phi}(\omega, h_t \mid \theta, M, z), \tag{D.38}$$

which is equation (5.6) in the main text.

The marginal posterior on quality dimension d is obtained by summing over all latent states and coverage matrices consistent with $\theta_d = 1$:

$$\mu_{d,t}(\omega, h_t; \phi) = \mathbb{P}(\theta_d = 1 \mid \omega, h_t, z; \phi) = \sum_{M} \sum_{\theta_{-d}} \mathbb{P}(\theta_d = 1, \theta_{-d}, M \mid \omega, h_t, z; \phi). \tag{D.39}$$

This is the dimension-level posterior appearing in (5.5). At a fixed history, it is a finite weighted sum over feasible (θ, M) pairs and is therefore available in closed form once the finite support of $G(M)$ is specified.

Likelihood-normalized next-review transition. The same integrated likelihood produces the exact transition kernel without conditioning on an unobserved signal matrix. Let x denote a feasible sparse realization of the designated next unopened review, including its coverage indicators and all signs on covered dimensions. Define $\mathcal{R}(M;h_t,x)$ as the full sparse matrices consistent with both the current history and that next realization, and set

$$\mathcal{L}^{\phi}(\omega,h_t,x \mid \theta,M,z) = \sum_{R\in\mathcal{R}(M;h_t,x)} \mathbb{P}(R \mid \theta,M)Q_{\phi}(\omega \mid R,M,z). \tag{D.40}$$

If the next review identity is sampled uniformly rather than fixed in advance, the right-hand side is first averaged over the $N-t$ unopened identities. In either convention, the feasible realizations partition the current continuation event:

$$\sum_{x}\mathcal{L}^{\phi}(\omega,h_t,x \mid \theta,M,z) = \mathcal{L}^{\phi}(\omega,h_t \mid \theta,M,z). \tag{D.41}$$

Applying the law of total probability over (θ,M) gives

$$\mathbb{P}_{\phi}(x \mid \omega,h_t,z) = \frac{\sum_{\theta}\sum_{M}\mathbb{P}(\theta)G(M)\mathcal{L}^{\phi}(\omega,h_t,x \mid \theta,M,z)}{\sum_{\theta}\sum_{M}\mathbb{P}(\theta)G(M)\mathcal{L}^{\phi}(\omega,h_t \mid \theta,M,z)}. \tag{D.42}$$

The denominator is positive at every reachable history, and (D.41) verifies that the probabilities in (D.42) sum to one. The successor posterior is obtained by replacing h_t with $h_{t+1}=(h_t,x)$ in (D.34). Thus the posterior, transition kernel, and Bellman recursion all use the same likelihood that integrates over the unobserved sparse review realizations.

Factorization under sign-invariant selection

The general likelihood need not factor because the selected set can depend on the realized signs in several dimensions. For the tractable subclass in Section D.1.1, selection is instead $\mathcal{S}^{\phi}(M,z)$ and the reported cell for a selected dimension is drawn from $q_d^{\phi}(a \mid Y_d,C_d,z)$. Let $c_{d,t}$ and $y_{d,t}$ be the number of opened non-missing signals and positive opened signals on dimension d. Conditional on θ_d, their likelihood is

$$\ell_{d,t}(\theta_d) = p_d(\theta_d)^{y_{d,t}}\left(1-p_d(\theta_d)\right)^{c_{d,t}-y_{d,t}},$$

where $p_d(1)=\rho_d$ and $p_d(0)=1-\rho_d$. Among the $C_d(M)-c_{d,t}$ unopened signals on d, the number of additional positives is binomial. Hence the probability that their total count, combined with the opened count, yields reported cell a_d is exactly

$$H^{\phi}_{da_d}\big(c_{d,t}, y_{d,t}; p_d(\theta_d), C_d(M), z\big).$$

Writing q^{ϕ}_d for the cell kernel makes this integration explicit:

$$\begin{aligned} &H^{\phi}_{da_d}\big(c_{d,t}, y_{d,t}; p, C_d(M), z\big) \\ &= \sum_{b=0}^{C_d(M)-c_{d,t}} \binom{C_d(M)-c_{d,t}}{b} p^b (1-p)^{C_d(M)-c_{d,t}-b} q^{\phi}_d\big(a_d \mid y_{d,t}+b, C_d(M), z\big). \end{aligned} \tag{D.43}$$

Conditional independence of signs across dimensions gives

$$\begin{aligned} \mathcal{L}^{\phi}(\omega, h_t \mid \theta, M, z) = \mathbb{1}\{M \text{ is consistent with } h_t\}\mathbb{1}\{\mathcal{S}^{\phi}(M,z) = \mathcal{S}\} \prod_{d=1}^{D} \ell_{d,t}(\theta_d) \\ \times \prod_{d\in\mathcal{S}} H^{\phi}_{da_d}\big(c_{d,t}, y_{d,t}; p_d(\theta_d), C_d(M), z\big), \end{aligned}$$

which is (D.6). When q^{ϕ}_d is strict majority, the last factor is the usual binomial tail; when it is a three-cell threshold rule, it is the corresponding binomial interval. This proves the factorization claim in Proposition 8. For an arbitrary known Q_{ϕ}, the finite sets of M, R, and future histories make the posterior and Bellman recursion finite-state, proving the remaining claim of that proposition.

Derivation of the topic-arrival probability $\eta_{d,t}$

Let n_{t+1} denote a uniformly chosen unopened review. Conditional on a particular coverage matrix M, exactly $C_d(M) - c_{d,t}$ unopened reviews cover dimension d, where

$$C_d(M) = \sum_{n=1}^{N} m_{nd} \qquad \text{and} \qquad c_{d,t} = \sum_{n\in\text{opened at } t} m_{nd}. \tag{D.44}$$

Since there are $N - t$ unopened reviews,

$$\mathbb{P}(m_{n_{t+1},d} = 1 \mid a, h_t, M) = \frac{C_d(M) - c_{d,t}}{N - t}. \tag{D.45}$$

Applying the law of iterated expectations over the posterior distribution of M gives

$$\begin{aligned} \eta_{d,t}(a, h_t) &= \mathbb{P}(m_{n_{t+1},d} = 1 \mid a, h_t) && \text{(D.46)} \\ &= \sum_{M} \mathbb{P}(m_{n_{t+1},d} = 1 \mid a, h_t, M)\,\mathbb{P}(M \mid a, h_t) && \text{(D.47)} \end{aligned}$$

$$= \sum_{M} \frac{C_d(M) - c_{d,t}}{N - t} \, \mathbb{P}(M \mid a, h_t) \tag{D.48}$$

$$= \mathbb{E}\left[\frac{C_d(M) - c_{d,t}}{N - t} \,\middle|\, a, h_t \right], \tag{D.49}$$

which is equation (5.9).

Proof of Theorem 2. At the entry history $t = 0$, equation (5.9) gives

$$\eta_{d,0} = \mathbb{E}\left[\frac{C_d(M)}{N} \,\middle|\, \omega \right].$$

For the first claim, condition only on the event $O_d = \{d \notin \mathcal{S}^K(M)\}$. Since no reviews have been opened,

$$\mathbb{E}[\eta_{d,0} \mid O_d] = \frac{1}{N} \, \mathbb{E}[C_d(M) \mid O_d].$$

We next establish the distributional ordering. Write $C = C_d(M)$ and $U = C_{-d}(M)$ for the vector of competing coverage counts. For any fixed realization u, define

$$I_d(c, u) = \mathbf{1}\{d \notin \mathcal{S}^K(c, u)\}.$$

Under top-K selection with fixed tie-breaking, increasing c can only move dimension d weakly upward in the ranking. Therefore

$$c' < c'' \quad \Longrightarrow \quad I_d(c'', u) \le I_d(c', u) \qquad \text{for every } u.$$

Independence of C and U makes the conditional law of U invariant to c. Hence

$$\begin{aligned} g(c) &= \mathbb{P}(O_d \mid C = c) \\ &= \sum_{u} I_d(c, u) \, \mathbb{P}(U = u \mid C = c) \\ &= \sum_{u} I_d(c, u) \, \mathbb{P}(U = u), \end{aligned}$$

which is weakly decreasing in c. It is nonconstant by assumption.

Let $f(c) = \mathbb{P}(C = c)$ and $f_O(c) = \mathbb{P}(C = c \mid O_d)$. At every point with $f(c) > 0$, Bayes' rule gives

$$f_O(c) = \frac{\mathbb{P}(O_d \mid C = c) \, \mathbb{P}(C = c)}{\mathbb{P}(O_d)} = \frac{g(c) f(c)}{\mathbb{P}(O_d)}, \tag{D.50}$$

$$\frac{f_O(c)}{f(c)} = \frac{g(c)}{\mathbb{P}(O_d)}. \tag{D.51}$$

The last ratio is weakly decreasing and nonconstant. This is precisely the stated monotone-likelihood-ratio dominance of $C \mid O_d$ by C.

For completeness, the increasing-function and first-order stochastic-dominance implications can be derived directly. Let h be weakly increasing. Using $\mathbb{P}(O_d) = \mathbb{E}[g(C)]$,

$$\mathbb{E}[h(C) \mid O_d] - \mathbb{E}[h(C)] = \frac{\mathbb{E}[h(C)g(C)]}{\mathbb{E}[g(C)]} - \mathbb{E}[h(C)] \tag{D.52}$$

$$= \frac{\mathrm{Cov}(h(C), g(C))}{\mathbb{P}(O_d)}. \tag{D.53}$$

If C' is an independent copy of C, then

$$2\,\mathrm{Cov}(h(C), g(C)) = \mathbb{E}\Big[\big(h(C) - h(C')\big)\big(g(C) - g(C')\big)\Big] \leq 0, \tag{D.54}$$

because h is increasing and g is decreasing. This proves (5.18). If h is strictly increasing on the support, nonconstancy of g makes the inequality strict.

Taking $h_k(c) = \mathbf{1}\{c \geq k\}$ gives

$$\mathbb{P}(C \geq k \mid O_d) \leq \mathbb{P}(C \geq k)$$

for every cutoff k, which is (5.17). Since the two distributions differ, the tail inequality is strict for at least one cutoff. Finally, taking $h(c) = c/N$ gives

$$\mathbb{E}[C_d \mid O_d] < \mathbb{E}[C_d]$$

and hence (5.19).

For the matched exogenous-display benchmark, let E_d denote a nondisplay event that is independent of C_d. Bayes' rule gives

$$\mathbb{P}(C_d = c \mid E_d) = \mathbb{P}(C_d = c) \qquad \text{for every } c.$$

Hence the entry topic-arrival probabilities in the endogenous and exogenous nondisplay states are, respectively,

$$\eta^O_{d,0} \equiv \mathbb{P}(m_{n_1 d} = 1 \mid O_d) = \frac{1}{N}\,\mathbb{E}[C_d \mid O_d], \tag{D.55}$$

$$\eta^E_{d,0} \equiv \mathbb{P}(m_{n_1 d} = 1 \mid E_d) = \frac{1}{N}\,\mathbb{E}[C_d]. \tag{D.56}$$

The strict mean ordering already proved yields $\eta^O_{d,0} < \eta^E_{d,0}$. Thus every increasing expected gain is weakly lower under top-K omission than under exogenous nondisplay by (5.18). In particular, when the focal one-step omitted-dimension gain is $\gamma_d \eta_{d,0} B_d$ with $B_d \geq 0$, the conditional expected gross one-step value is weakly lower, and it is strictly lower when $\gamma_d B_d > 0$. Corollary 3 states the resulting behavioral interval explicitly. □

Proof of Corollary 3. The proof has four steps.

Step 1: topic arrival under the two nondisplay rules. At entry, no reviews have been opened, so conditional on a realized coverage count $C_d = c$, uniform passive sampling gives

$$\mathbb{P}(m_{n_1 d} = 1 \mid C_d = c) = \frac{c}{N}. \tag{D.57}$$

Integrating (D.57) against the posterior coverage distribution in the top-K omission state gives

$$\eta^O_{d,0} = \sum_c \frac{c}{N}\,\mathbb{P}(C_d = c \mid O_d) \tag{D.58}$$
$$= \frac{1}{N}\,\mathbb{E}[C_d \mid O_d]. \tag{D.59}$$

In the exogenous state, nondisplay is independent of C_d, so

$$\eta^E_{d,0} = \sum_c \frac{c}{N}\,\mathbb{P}(C_d = c \mid E_d) \tag{D.60}$$
$$= \sum_c \frac{c}{N}\,\mathbb{P}(C_d = c) \tag{D.61}$$
$$= \frac{1}{N}\,\mathbb{E}[C_d]. \tag{D.62}$$

Theorem 2 establishes

$$\mathbb{E}[C_d \mid O_d] < \mathbb{E}[C_d], \tag{D.63}$$

and therefore $\eta^O_{d,0} < \eta^E_{d,0}$.

Step 2: exact focal one-step values. Under the one-topic and local-separability assumptions, the next review either covers dimension d or does not. Conditional on coverage, the fixed gross gain is $\gamma_d B_d$. Conditional on no coverage, the focal gain is zero, and all non-d branches are held fixed across the matched states. Thus, for $q \in \{O, E\}$, the law of total expectation gives

$$\mathcal{I}^q_d = \mathbb{P}(m_{n_1 d} = 1 \mid q)\,\gamma_d B_d + \mathbb{P}(m_{n_1 d} = 0 \mid q)\,0 \tag{D.64}$$
$$= \gamma_d B_d \eta^q_{d,0}. \tag{D.65}$$

Substituting (D.59) and (D.62) yields

$$\mathcal{I}^O_d = \frac{\gamma_d B_d}{N}\,\mathbb{E}[C_d \mid O_d], \tag{D.66}$$
$$\mathcal{I}^E_d = \frac{\gamma_d B_d}{N}\,\mathbb{E}[C_d]. \tag{D.67}$$

Because $\gamma_d > 0$, $B_d > 0$, and (D.63) holds, subtraction gives

$$\mathcal{I}_d^E - \mathcal{I}_d^O = \frac{\gamma_d B_d}{N}(\mathbb{E}[C_d] - \mathbb{E}[C_d \mid O_d]) \tag{D.68}$$
$$> 0. \tag{D.69}$$

This proves Main Paper equations (5.20) and (5.22) and shows that the interval has strictly positive length.

Step 3: behavior on the cost interval. In the one-step problem, inspection is optimal exactly when its gross value is weakly at least the first-inspection cost. Therefore, if

$$\mathcal{I}_d^O < c^W(1) \le \mathcal{I}_d^E, \tag{D.70}$$

inspection is strictly suboptimal after top-K omission but is optimal under exogenous nondisplay. If costs are heterogeneous with cumulative distribution F_c, and either the distribution is atomless or indifferent consumers inspect, the corresponding difference in one-step opening probabilities is

$$\mathbb{P}(\text{open} \mid E_d) - \mathbb{P}(\text{open} \mid O_d) = F_c(\mathcal{I}_d^E) - F_c(\mathcal{I}_d^O) \ge 0, \tag{D.71}$$

with strict inequality whenever the cost distribution places positive mass on $(\mathcal{I}_d^O, \mathcal{I}_d^E]$.

Step 4: scope of the dynamic and no-summary interpretations. The comparison above is exact for opening once and then stopping. If both matched states satisfy the successor-absorption conditions in Section D.1.2, the one-step cutoff equals the unrestricted dynamic cutoff, so the same interval orders dynamic initiation. Finally, exogenous nondisplay is a local no-summary benchmark for dimension d only when all other displayed information leaves the current stopping margin, B_d, and every non-d branch unchanged. Without those conditions, the corollary does not sign total AI-versus-no-AI search. □

Proof of Corollary 9. Fix a realized count-revealing output

$$\omega^C = \left(\mathcal{S}^K, \left(a_s, C_s(M)\right)_{s \in \mathcal{S}^K}\right)$$

with $|\mathcal{S}^K| = K$, and write

$$c_* = \min_{s \in \mathcal{S}^K} C_s(M).$$

Choose a displayed dimension $s_* \in \mathcal{S}^K$ attaining the minimum, so $C_{s_*}(M) = c_*$. Consider any omitted dimension $d \notin \mathcal{S}^K$ and any coverage matrix M consistent with the observed output. If $C_d(M) > c_*$, then dimension d has strictly higher coverage than the displayed dimension s_*. A rule that selects the top K coverage counts could not then display s_* while omitting d, a contradiction. Therefore, pointwise over the posterior support,

$$C_d(M) \le c_*. \tag{D.72}$$

Equality is possible when several dimensions tie at the marginal displayed count and the fixed tie-breaking rule favors s_*. This is why the result is a weak rather than strict inequality.

Because (D.72) holds for every coverage matrix assigned positive posterior probability after ω^C,

$$\mathbb{P}(C_d(M) \le c_* \mid \omega^C) = 1. \tag{D.73}$$

At entry, equation (5.9) and iterated expectations give

$$\eta_{d,0}(\omega^C) = \mathbb{E}\left[\frac{C_d(M)}{N} \,\middle|\, \omega^C\right] \tag{D.74}$$

$$\le \mathbb{E}\left[\frac{c_*}{N} \,\middle|\, \omega^C\right] \tag{D.75}$$

$$= \frac{c_*}{N}, \tag{D.76}$$

which proves the entry bound.

Next fix any reachable history h_t with $t < N$. The history reveals that $c_{d,t}$ of the opened reviews covered dimension d. For every posterior-feasible M, reachability and (D.72) imply

$$0 \le C_d(M) - c_{d,t} \le c_* - c_{d,t}.$$

Dividing by the positive number $N - t$ and taking the posterior expectation conditional on (ω^C, h_t) gives

$$\eta_{d,t}(\omega^C, h_t) = \mathbb{E}\left[\frac{C_d(M) - c_{d,t}}{N - t} \,\middle|\, \omega^C, h_t\right] \tag{D.77}$$

$$\le \frac{c_* - c_{d,t}}{N - t}, \tag{D.78}$$

which is the dynamic bound.

If fewer than K dimensions have positive coverage, top-K selection displays every positive-coverage dimension. Every omitted dimension therefore has $C_d(M) = 0$, so its topic-arrival probability is zero at entry and after every reachable history.

No independence restriction was used: the argument first bounds every posterior-feasible coverage matrix and only then integrates. To clarify the posterior interpretation, Bayes' rule can be written as

$$\mathbb{P}(C_d = c \mid \omega^C) \propto \mathbb{P}(\omega^C \mid C_d = c)\,\mathbb{P}(C_d = c). \tag{D.79}$$

The first factor is zero for $c > c_*$, establishing the support truncation, but it need not be constant for $c \le c_*$. Correlation with displayed counts can reweight all interior masses, and fixed tie-breaking can reweight the boundary mass at $c = c_*$. Hence the corollary asserts the support and topic-arrival bounds, not that the posterior must equal an unweighted truncation of the marginal prior. $\square$

Predictive sign probability and unconditional relevance

Now condition on the event that the next unopened review covers dimension d. The exact object is the posterior predictive probability

$$\pi^{+}_{d,t}(\omega, h_t) = \mathbb{P}(r_{n_{t+1},d} = 1 \mid m_{n_{t+1},d} = 1, \omega, h_t).$$

It must average over both latent quality and feasible coverage matrices. Given (θ, M), define the conditional probability that a covered next signal on dimension d is favorable by

$$\chi^{+}_{d,t}(\theta, M; \omega, h_t) = \mathbb{P}(r_{n_{t+1},d} = 1 \mid m_{n_{t+1},d} = 1, \theta, M, \omega, h_t).$$

In the componentwise sign-invariant subclass, if d is displayed and its cell is a_d, then

$$\chi^{+}_{d,t}(\theta, M; \omega, h_t) = p_d(\theta_d) \frac{H^{\phi}_{da_d}(c_{d,t} + 1, y_{d,t} + 1; p_d(\theta_d), C_d(M), z)}{H^{\phi}_{da_d}(c_{d,t}, y_{d,t}; p_d(\theta_d), C_d(M), z)},$$

whenever the denominator is positive. If d is not subject to any summary cell that restricts its unobserved signs, this reduces to $\chi^{+}_{d,t}(\theta, M; \omega, h_t) = p_d(\theta_d)$.

Using the posterior over (θ, M), the exact conditional sign probability is

$$\pi^{+}_{d,t}(\omega, h_t) = \frac{\sum_{\theta,M} \mathbb{P}(\theta, M \mid \omega, h_t, z; \phi) \frac{C_d(M) - c_{d,t}}{N - t} \chi^{+}_{d,t}(\theta, M; \omega, h_t)}{\sum_{\theta,M} \mathbb{P}(\theta, M \mid \omega, h_t, z; \phi) \frac{C_d(M) - c_{d,t}}{N - t}}, \tag{D.80}$$

whenever the denominator is positive. If the denominator is zero, the next review cannot cover d, so the value assigned to $\pi^{+}_{d,t}$ is irrelevant and the unconditional probability below is zero.

The unconditional predictive probability that the next unopened review both covers dimension d and delivers a favorable sign is

$$\psi^{+}_{d,t}(\omega, h_t) = \mathbb{P}(m_{n_{t+1},d} = 1, r_{n_{t+1},d} = 1 \mid \omega, h_t) \tag{D.81}$$

$$= \eta_{d,t}(\omega, h_t)\, \pi^{+}_{d,t}(\omega, h_t), \tag{D.82}$$

which is equation (5.11). This identity is just the product rule. The simple formula $\rho_d \mu_{d,t} + (1 - \rho_d)(1 - \mu_{d,t})$ is recovered only when coverage of the next review is not additionally informative about θ_d and the summary imposes no count restriction on dimension d's unobserved signs.

Proof of the primitive one-step threshold

Fix state $s_t = (a, h_t)$. If the consumer stops immediately, she receives $S(s_t; \gamma)$. If she opens one more review and then stops, she pays $c^W(t+1)$ and receives expected stopping payoff

$$\mathbb{E}[S(s_t(X_{t+1}); \gamma) \mid s_t].$$

The one-step review is optimal if and only if

$$-c^W(t+1) + \mathbb{E}[S(s_t(X_{t+1}); \gamma) \mid s_t] \geq S(s_t; \gamma),$$

which rearranges to

$$c^W(t+1) \leq \mathbb{E}[S(s_t(X_{t+1}); \gamma) \mid s_t] - S(s_t; \gamma) = \mathcal{I}^{MD}(s_t; \gamma).$$

This proves Proposition 10. Since the full dynamic problem can always mimic the one-step policy and then stop, the primitive one-step value is a lower bound on the exact dynamic cutoff:

$$\mathcal{I}^{MD}(s_t; \gamma) \leq \bar{c}^{MD}(s_t; \gamma).$$

Derivation of the one-step continuation bound

Fix a post-summary state $s_t = (a, h_t)$ and an uncovered focal dimension d. Let

$$S(s_t; \gamma) = \max\{0, \delta(a, h_t), \bar{M}\} \tag{D.83}$$

denote the current stopping payoff. For the focal lower-bound policy, there are three relevant cases under the one-step deviation:

1. with probability $1 - \eta_{d,t}(s_t)$, the next review does not cover dimension d, and the expected stopping payoff after observing the full sparse review is denoted by $S_d^0(s_t)$;

2. with probability $\eta_{d,t}(s_t)\pi_{d,t}^+(s_t)$, the next review covers d and gives a favorable signal, so the consumer stops at $S_d^+(s_t)$;

3. with probability $\eta_{d,t}(s_t)(1 - \pi_{d,t}^+(s_t))$, the next review covers d and gives an unfavorable signal, so the consumer stops at $S_d^-(s_t)$.

Subtracting the inspection cost $c^W(t+1)$, the expected payoff from this one-step deviation is

$$\begin{aligned}\underline{\mathcal{K}}_d(s_t) = {} & -c^W(t+1) + (1 - \eta_{d,t}(s_t))S_d^0(s_t) \\ & + \eta_{d,t}(s_t)\pi_{d,t}^+(s_t)S_d^+(s_t) + \eta_{d,t}(s_t)(1 - \pi_{d,t}^+(s_t))S_d^-(s_t),\end{aligned} \tag{D.84}$$

which is equation (D.19). Subtracting $S(s_t;\gamma)$ from both sides gives

$$\begin{aligned}\underline{\mathcal{K}}_d(s_t) - S(s_t;\gamma) = &-c^W(t+1) + \eta_{d,t}(s_t)\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big]\\ &+ \eta_{d,t}(s_t)(1-\pi^+_{d,t}(s_t))\big[S^-_d(s_t) - S(s_t;\gamma)\big]\\ &+ (1-\eta_{d,t}(s_t))\big[S^0_d(s_t) - S(s_t;\gamma)\big],\end{aligned} \tag{D.85}$$

so a sufficient condition for continuation is

$$\begin{aligned}c^W(t+1) \le\ &\eta_{d,t}(s_t)\Big(\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big] + (1-\pi^+_{d,t}(s_t))\big[S^-_d(s_t) - S(s_t;\gamma)\big]\Big)\\ &+ (1-\eta_{d,t}(s_t))\big[S^0_d(s_t) - S(s_t;\gamma)\big],\end{aligned} \tag{D.86}$$

which is exactly (D.20). If this inequality holds, the focal one-step deviation weakly dominates immediate stopping. Because the full dynamic problem can implement this deviation, continuation is optimal in the unrestricted Bellman problem as well. If $S^0_d(s_t) \ge S(s_t;\gamma)$, the last term is nonnegative and can be dropped to obtain the simpler sufficient condition stated after the proposition. One sufficient primitive condition for this inequality is that the no-coverage event does not change payoff-relevant beliefs in expectation; then posterior beliefs are a martingale over the remaining realized signs, and convexity of the stopping payoff gives the Jensen inequality. This proves Proposition 11.

Analytical taste-region characterization

When the no-coverage branch is payoff-neutral or favorable, the focal one-step sufficient taste region in (D.22) is the simplified sufficient condition obtained from (D.86) by dropping the nonnegative no-coverage term. Define

$$\mathcal{B}_d(s_t;\gamma) = \eta_{d,t}(s_t)\Big(\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big] + (1-\pi^+_{d,t}(s_t))\big[S^-_d(s_t) - S(s_t;\gamma)\big]\Big). \tag{D.87}$$

Then the inspection region is

$$\mathcal{T}^I_d(s_t) = \big\{\gamma \in \mathbb{R}^D_+ : c^W(t+1) \le \mathcal{B}_d(s_t;\gamma)\big\}, \tag{D.88}$$

which is a sufficient focal-dimension region for the unrestricted Bellman problem. It is generally implicit because the stopping-payoff wedges can depend on the full taste vector through the max operator.

Under the local separability condition

$$S^+_d(s_t) - S(s_t;\gamma) = \gamma_d\Delta^+_{d,t}(s_t), \qquad S^-_d(s_t) - S(s_t;\gamma) = \gamma_d\Delta^-_{d,t}(s_t), \tag{D.89}$$

substituting (D.89) into (D.87) yields

$$\mathcal{B}_d(s_t;\gamma) = \gamma_d\eta_{d,t}(s_t)\Big(\pi^+_{d,t}(s_t)\Delta^+_{d,t}(s_t) + (1-\pi^+_{d,t}(s_t))\Delta^-_{d,t}(s_t)\Big). \tag{D.90}$$

Hence, whenever the bracketed term is positive,

$$\gamma_d \geq \gamma_d^*(s_t) = \frac{c^W(t+1)}{\eta_{d,t}(s_t)\Big(\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t)\Big)}. \tag{D.91}$$

This is the closed-form local taste threshold reported in (D.25).

Exact one-topic taste regions

The focal threshold above gives a sufficient condition dimension by dimension. It becomes an exact one-step taste-region characterization under the one-topic-per-review specialization. Suppose that each opened review covers at most one payoff-relevant dimension and that the local separability condition holds for every dimension $d = 1, \ldots, D$. Conditional on dimension d being the topic of the next review, the expected stopping-payoff gain per unit of taste is

$$\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t).$$

Multiplying by the probability $\eta_{d,t}(s_t)$ that the next review covers dimension d gives the per-unit taste gain

$$\Omega_{d,t}(s_t) = \eta_{d,t}(s_t)\Big[\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t)\Big].$$

Since at most one payoff-relevant dimension arrives in the next review, the exact one-step gross value of inspection is the sum of these dimension-specific gains:

$$\mathcal{I}^{MD}(s_t;\gamma) = \sum_{d=1}^{D} \gamma_d \Omega_{d,t}(s_t).$$

Therefore one-step inspection starts if and only if

$$\sum_{d=1}^{D} \gamma_d \Omega_{d,t}(s_t) \geq c^W(t+1),$$

which is the half-space in (D.27).

On the normalized simplex $\Gamma = \{\gamma \geq 0 : \sum_d \gamma_d = 1\}$, the left-hand side is a convex combination of the values $\{\Omega_{d,t}(s_t)\}_{d=1}^D$. Hence it lies between $\Omega^{\min}(s_t)$ and $\Omega^{\max}(s_t)$. If $c^W(t+1) > \Omega^{\max}(s_t)$, even the highest-gain pure taste type will not start one-step search. If $c^W(t+1) \leq \Omega^{\min}(s_t)$, even the lowest-gain pure taste type starts. The intermediate case is the selective region cut out by the boundary hyperplane

$$\sum_{d=1}^{D} \gamma_d \Omega_{d,t}(s_t) = c^W(t+1).$$

For a two-dimensional taste rotation between dimensions r and d,

$$\gamma(\lambda) = (1-\lambda)e_r + \lambda e_d, \qquad \lambda \in [0,1],$$

one-step search starts iff

$$(1-\lambda)\Omega_{r,t}(s_t) + \lambda\Omega_{d,t}(s_t) \geq c^W(t+1).$$

If $\Omega_{d,t}(s_t) > \Omega_{r,t}(s_t)$, this is equivalent to

$$\lambda \geq \lambda^*_{rd}(s_t) = \frac{c^W(t+1) - \Omega_{r,t}(s_t)}{\Omega_{d,t}(s_t) - \Omega_{r,t}(s_t)},$$

with the threshold truncated to the unit interval for interpretation. If $\Omega_{d,t}(s_t) < \Omega_{r,t}(s_t)$ the inequality reverses. If the two per-unit gains are equal, either all values of λ start or none do, depending on whether the common gain exceeds the inspection cost.

Proof of decision sufficiency, reachability, and the sandwich

Proof of Theorem 4. Let $\bar{s}$ denote the terminal state reached by opening every remaining review. The terminal decision index is

$$\Lambda(\bar{s}) = \sum_{d=1}^{D} \gamma_d \mu_d(\bar{s}).$$

Because posterior beliefs are martingales, the current decision index satisfies

$$\Lambda_t = \mathbb{E}[\Lambda(\bar{s}) \mid s_t].$$

Support geometry, part (i): equality of terminal and all-reachable extrema. Let $s_u \in \mathfrak{S}(s_t)$ be any reachable intermediate state and let $\mathfrak{T}(s_u) \subseteq \mathfrak{T}(s_t)$ be its positive-probability descendant terminal states. Write Λ_N for the random terminal index. The tower property gives

$$\Lambda(s_u) = \mathbb{E}[\Lambda_N \mid s_u].$$

Because the terminal support is finite, a conditional expectation lies between the smallest and largest values in its conditional support:

$$\min_{\bar{s}\in\mathfrak{T}(s_u)} \Lambda(\bar{s}) \leq \Lambda(s_u) \leq \max_{\bar{s}\in\mathfrak{T}(s_u)} \Lambda(\bar{s}).$$

Since $\mathfrak{T}(s_u) \subseteq \mathfrak{T}(s_t)$, every reachable intermediate index lies in $[\Lambda_t^{T,\min}, \Lambda_t^{T,\max}]$. Hence

$$\Lambda_t^{\min} \geq \Lambda_t^{T,\min}, \qquad \Lambda_t^{\max} \leq \Lambda_t^{T,\max}.$$

Conversely, terminal states belong to $\mathfrak{S}(s_t)$, so minimizing and maximizing over the larger set gives the reverse inequalities. Therefore

$$[\Lambda_t^{\min}, \Lambda_t^{\max}] = [\Lambda_t^{T,\min}, \Lambda_t^{T,\max}],$$

which proves (D.12).

Part (ii): decision sufficiency. Suppose first that the state is decision-sufficient. If $b \le \Lambda_t^{T,\min}$, then every terminal index lies weakly above the stopping margin. Any posterior index reached by any feasible future history is a conditional expectation of terminal indices and therefore also lies weakly above b. On this support,

$$S(s;\gamma) = O + \Lambda(s) - b$$

is affine. For any finite stopping time ν,

$$\mathbb{E}[S(s_\nu;\gamma) \mid s_t] = O + \mathbb{E}[\Lambda_\nu \mid s_t] - b = O + \Lambda_t - b = S(s_t;\gamma).$$

If instead $b \ge \Lambda_t^{T,\max}$, every terminal and intermediate posterior index lies weakly below b, so

$$S(s;\gamma) = O$$

throughout the reachable support and again $\mathbb{E}[S(s_\nu;\gamma) \mid s_t] = S(s_t;\gamma)$. Hence no feasible policy can raise gross expected stopping payoff above immediate stopping. In particular, full revelation has zero gross value: $\mathcal{J}^T(s_t;\gamma) = 0$.

Conversely, if $\Lambda_t^{T,\min} < b < \Lambda_t^{T,\max}$, the terminal index distribution places positive probability on both sides of b. The function $x \mapsto O + (x-b)_+$ is convex and is affine only on each side of the kink. Since the terminal support crosses the kink, Jensen's inequality is strict:

$$\mathbb{E}[O + (\Lambda(\bar{s}) - b)_+ \mid s_t] > O + (\mathbb{E}[\Lambda(\bar{s}) \mid s_t] - b)_+ = S(s_t;\gamma).$$

Thus $\mathcal{J}^T(s_t;\gamma) > 0$. This proves the equivalence between decision insufficiency and strictly positive gross full-revelation value.

The preceding argument also proves the zero-gross-value statement for every finite stopping policy, because every intermediate index lies on the same affine branch of the stopping payoff.

When the state is decision-sufficient, no feasible policy can raise gross expected stopping payoff, and strictly positive inspection costs make every nontrivial policy strictly worse than stopping. If the state is decision-insufficient and the total cost of opening the remaining pool is below $\mathcal{J}^T(s_t;\gamma)$, the feasible policy that opens all remaining reviews and then stops gives strictly higher net value than immediate stopping. The optimal dynamic policy must therefore start verification.

If $\Lambda_t^{T,\min} < \Lambda_t^{T,\max}$, then every $b \in (\Lambda_t^{T,\min}, \Lambda_t^{T,\max})$ makes the state decision-insufficient. Part (ii) then implies $\mathcal{J}^T(s_t;\gamma) > 0$. Choosing positive inspection costs with total remaining cost below this value yields verification. Therefore any finite summary rule that leaves a nondegenerate terminal decision index fails to eliminate gross residual value uniformly over stopping margins for that type.

To verify the openness claim, choose positive-probability terminal states $\bar{s}^-$ and $\bar{s}^+$ on opposite sides of the maintained margin. For each terminal state $\bar{s}$, the terminal index $\Lambda(\bar{s};\gamma) = \sum_d \gamma_d \mu_d(\bar{s})$ is continuous in γ, and the inequalities

$$\Lambda(\bar{s}^-;\gamma) < b < \Lambda(\bar{s}^+;\gamma)$$

are strict at the maintained type-margin pair. By continuity, there is an open neighborhood of (γ, b) such that the same two strict inequalities hold for every nearby (γ', b'):

$$\Lambda(\bar{s}^-;\gamma') < b' < \Lambda(\bar{s}^+;\gamma').$$

Because the two terminal states retain positive posterior-predictive probability at the fixed history s_t, the terminal decision-index support continues to straddle b' throughout this neighborhood. Hence s_t is decision-insufficient for every nearby type-margin pair. The omission and folding statement is an application: omitted dimensions and non-aligned folded cells generate distinct terminal indices in the relevant taste direction precisely when they leave two positive-probability terminal states satisfying the strict straddle inequalities.

Support geometry, part (ii): absorption and topic arrival. Reachability is transitive. Conditional on any realized successor s_{t+1}, every future state reachable from s_{t+1} is also reachable from s_t through that successor. Hence

$$\mathfrak{S}(s_{t+1}) \subseteq \mathfrak{S}(s_t)$$

on every path. Taking the minimum and maximum of the linear index over nested sets gives

$$[\Lambda_{t+1}^{\min}, \Lambda_{t+1}^{\max}] \subseteq [\Lambda_t^{\min}, \Lambda_t^{\max}].$$

Thus if b is outside the open interval at t, it remains outside the open interval at every later state.

The spread $\Lambda_t^{\max} - \Lambda_t^{\min}$ is non-increasing by the preceding nestedness result. In the componentwise topic-arrival benchmark, future movement in dimension d's posterior can occur only through future signals that cover dimension d. Dimensions with $\gamma_d = 0$ do not enter the index, and dimensions with no future coverage cannot move the index through direct topic arrival. If $\mathcal{R}_t \cap \{d : \gamma_d > 0\} = \emptyset$ and there are no cross-dimensional informational spillovers, then every payoff-relevant belief is fixed for all future histories, so $\Lambda_t^{\min} = \Lambda_t^{\max} = \Lambda_t$. The open interval is empty, and part (ii) applies. Absorption follows from nestedness.

Support geometry, part (iii): continuation-value bracket. Let

$$C^{1,MD} = \left\{s_t : c^W(t+1) \leq \mathcal{I}^{MD}(s_t;\gamma)\right\}$$

and

$$C^{\infty,MD} = \{s_t : c^W(t+1) \le \bar{c}^{MD}(s_t;\gamma)\}.$$

The first inclusion in (D.16) holds because the unrestricted dynamic problem can always mimic the feasible policy that opens one review and then stops. The second inclusion is the contrapositive of part (ii) of Theorem 4. If $s_t \notin \mathcal{B}^{R,MD}$, then each successor is also outside $\mathcal{B}^{R,MD}$ by absorption, so $V(s_{t+1}) = S(s_{t+1};\gamma)$. Part (ii) then gives

$$\bar{c}^{MD}(s_t;\gamma) = \mathbb{E}[V(s_{t+1}) \mid s_t] - S(s_t;\gamma) = \mathbb{E}[S(s_{t+1};\gamma) \mid s_t] - S(s_t;\gamma) = 0.$$

On $\mathcal{B}^{R,MD}$, optimal continuation weakly dominates the feasible one-step-and-stop policy, so

$$\bar{c}^{MD}(s_t;\gamma) = \mathbb{E}[V(s_{t+1}) \mid s_t] - S(s_t;\gamma) \ge \mathbb{E}[S(s_t(X_{t+1});\gamma) \mid s_t] - S(s_t;\gamma) = \mathcal{I}^{MD}(s_t;\gamma).$$

In the one-topic-per-review local-separability specialization, the exact one-step value is

$$\sum_d \gamma_d \Omega_{d,t}(s_t),$$

as shown in the proof of Proposition 13. □

D.1.2 Successor absorption and exact one-step cutoffs

Proof of Corollary 10. Fix a reachable state s_t. Let $\mathfrak{T}(s_t)$ denote its finite set of positive-probability terminal descendants, and let $Z_T(\bar{s})$ be the terminal decision index at $\bar{s} \in \mathfrak{T}(s_t)$. For every feasible next realization $x \in \mathcal{X}(s_t)$, write $\mathfrak{T}(s_t(x))$ for the corresponding descendant terminal set. The next realization partitions terminal histories, so

$$\mathfrak{T}(s_t) = \dot{\bigcup_{x\in\mathcal{X}(s_t)}} \mathfrak{T}(s_t(x)). \tag{D.92}$$

Consequently,

$$L(s_t) = \min_{x\in\mathcal{X}(s_t)} L(s_t(x)), \qquad U(s_t) = \max_{x\in\mathcal{X}(s_t)} U(s_t(x)). \tag{D.93}$$

Now suppose $b \in \mathcal{E}_1(s_t)$. By (D.8),

$$L(s_t) < b < U(s_t),$$

so the current terminal support has positive-probability realizations strictly below and strictly above b. At the same time, for every successor x,

$$b \notin \big(L(s_t(x)), U(s_t(x))\big).$$

Because the successor hull is an interval, this exclusion is equivalent to one of the two weak orderings

$$U(s_t(x)) \le b \qquad \text{or} \qquad b \le L(s_t(x)). \tag{D.94}$$

Thus the stopping payoff

$$S(s) = O + g(Z(s) - b)_+, \qquad g > 0,$$

is affine on the entire terminal support below each successor. Posterior martingality gives, for every finite stopping time ν following successor $s_t(x)$,

$$\mathbb{E}[Z_\nu \mid s_t(x)] = Z(s_t(x)).$$

If the left inequality in (D.94) holds, then

$$\mathbb{E}[S(s_\nu) \mid s_t(x)] = O = S(s_t(x));$$

if the right inequality holds, then

$$\mathbb{E}[S(s_\nu) \mid s_t(x)] = O + g\big(\mathbb{E}[Z_\nu \mid s_t(x)] - b\big) = S(s_t(x)).$$

Hence every future policy from $s_t(x)$ has zero gross information value. Any policy that opens another review pays a strictly positive cost, so immediate stopping is strictly optimal and

$$V(s_t(x)) = S(s_t(x)) \qquad \text{for every } x \in \mathcal{X}(s_t). \tag{D.95}$$

It remains to show that the one-step value at the parent is strictly positive. Because $L(s_t) < b < U(s_t)$, choose positive-probability terminal histories $\bar{s}^-$ and $\bar{s}^+$ satisfying

$$Z_T(\bar{s}^-) < b < Z_T(\bar{s}^+).$$

Let x^- and x^+ be their respective next-period branches. The two histories cannot descend through the same successor, because that successor's terminal hull would then straddle b. By (D.94),

$$U(s_t(x^-)) \le b, \qquad b \le L(s_t(x^+)).$$

Moreover, each relevant conditional support contains a strict realization on its indicated side. Therefore its conditional mean is strict:

$$Z(s_t(x^-)) < b < Z(s_t(x^+)).$$

Both branches have positive posterior-predictive probability. The next-period index distribution consequently crosses the kink, and strict Jensen gives

$$\begin{aligned}\mathcal{I}(s_t) &= \mathbb{E}[S(s_t(X_{t+1})) \mid s_t] - S(s_t) \\ &= g\{\mathbb{E}[(Z_{t+1} - b)_+ \mid s_t] - (Z_t - b)_+\} > 0.\end{aligned} \tag{D.96}$$

Substituting (D.95) into the Bellman equation at s_t yields

$$V(s_t) = \max\left\{S(s_t), -c^W(t+1) + \sum_{x \in \mathcal{X}(s_t)} \mathbb{P}(x \mid s_t) S(s_t(x))\right\},$$

which is (D.10). Before paying the current inspection cost, the exact gross dynamic cutoff is

$$\begin{aligned}
\bar{c}(s_t) &= \sum_{x \in \mathcal{X}(s_t)} \mathbb{P}(x \mid s_t) V(s_t(x)) - S(s_t) \\
&= \sum_{x \in \mathcal{X}(s_t)} \mathbb{P}(x \mid s_t) S(s_t(x)) - S(s_t) \\
&= \mathcal{I}(s_t).
\end{aligned}$$

Thus continuation is optimal exactly when $c^W(t+1) \le \mathcal{I}(s_t)$. In the scalar model this identity reads $\Delta^A = \mathcal{I}^1$; in the multidimensional model it reads $\bar{c}^{MD} = \mathcal{I}^{MD}$. □

Proof of the geometry claims in Corollary 10. Part (i): parity-robust scalar majority. Under neutral tie-breaking, the terminal positive counts consistent with each binary summary cell form the contiguous sets

$$\mathcal{A}_1 = \{\lceil N/2 \rceil, \dots, N\}, \qquad \mathcal{A}_0 = \{0, \dots, \lfloor N/2 \rfloor\}. \tag{D.97}$$

For odd N, the two sets partition the count support. For even N, the tie count $N/2$ belongs to both sets because neutral randomization assigns it positive probability in each summary cell. This changes likelihood weights but does not destroy contiguity.

Let $n = N - m$. At reachable state (a, m, y), the parent count support and the two successor supports are

$$\mathcal{Y}(a, m, y) = \mathcal{A}_a \cap \{y, \dots, y + n\}, \tag{D.98}$$
$$\mathcal{Y}(a, m+1, y) = \mathcal{A}_a \cap \{y, \dots, y + n - 1\}, \tag{D.99}$$
$$\mathcal{Y}(a, m+1, y+1) = \mathcal{A}_a \cap \{y + 1, \dots, y + n\}. \tag{D.100}$$

Each nonempty set is an integer interval. At any feasible terminal count Y, the summary-cell weight cancels from posterior odds, giving

$$\nu(Y) \equiv \widetilde{\mu}^{AI}(a, N, Y) = \left[1 + \frac{1-\mu_0}{\mu_0}\left(\frac{1-\rho}{\rho}\right)^{2Y-N}\right]^{-1}. \tag{D.101}$$

Since $\rho > 1/2$, $\nu(Y)$ is strictly increasing in Y.

Write the parent count interval as $\{L_Y, \dots, U_Y\}$. If the summary cell truncates the lower raw-count endpoint, so $L_Y > y$, then the positive successor support in (D.100) equals the parent support. If the cell truncates the upper endpoint, so $U_Y < y + n$, then the negative successor support in (D.99) equals the parent support. In either case, provided the parent interval is nondegenerate, one successor open hull equals the parent open hull and $\mathcal{E}_1(a, m, y) = \varnothing$. If the parent support is degenerate, its open hull is already empty.

It remains to consider the untrimmed case $\mathcal{Y}(a, m, y) = \{y, \ldots, y + n\}$. When $n \geq 3$, the two successor open posterior hulls are

$$(\nu(y), \nu(y + n - 1)) \qquad \text{and} \qquad (\nu(y + 1), \nu(y + n)).$$

They overlap and cover the parent open hull $(\nu(y), \nu(y + n))$. Hence the exactness layer is empty. When $n = 2$, the successor hulls are

$$(\nu(y), \nu(y + 1)) \qquad \text{and} \qquad (\nu(y + 1), \nu(y + 2)),$$

so the parent open hull is covered except for the single point $\nu(y + 1)$. Finally, when $n = 1$, both successors are terminal and their open hulls are empty. Whenever both are reachable,

$$\mathcal{E}_1(a, N - 1, y) = (\nu(y), \nu(y + 1)) = (\mu^{\min}(a, N - 1, y), \mu^{\max}(a, N - 1, y)).$$

Since $m \leq N - 2$ is equivalent to $n \geq 2$, the pre-penultimate layer has measure zero and contains at most one terminal posterior point for both parities of N.

Part (ii): multidimensional branch separation. Suppose $u < v$ and every successor satisfies either

$$U(s_t(x)) \leq u \qquad \text{or} \qquad v \leq L(s_t(x)),$$

with positive-probability successors of both types. The terminal partition (D.92) implies

$$L(s_t) \leq u < v \leq U(s_t).$$

For every $b \in (u, v)$, the parent hull therefore straddles b, while no successor open hull contains b. By (D.8),

$$(u, v) \subseteq \mathcal{E}_1(s_t).$$

The stated pivotal-topic configuration is an application: requiring the next review to cover the pivotal dimension removes the noncoverage branch, and the assumed separation of the complete descendant hulls verifies the two displayed inequalities. By contrast, if exactly one covering review remains among several unopened reviews, a next noncover leaves that review in the residual pool. If its position is exchangeable, its next-period arrival probability changes from

$$\frac{1}{N - t} \quad \text{to} \quad \frac{1}{N - t - 1},$$

rather than falling to zero. That weaker count restriction therefore does not establish successor absorption. □

Proof of Corollary 12. On the exactness layer, Corollary 10 reduces the unrestricted Bellman comparison to

$$\text{open at } s_t \quad \Longleftrightarrow \quad c^W(t+1) \leq \mathcal{I}(s_t; b). \tag{D.102}$$

Therefore an observed opening gives the upper bound $c^W(t+1) \leq \mathcal{I}(s_t; b)$, while an observed stop gives the reverse weak inequality, up to the consumer's behavior at equality.

Now hold the information kernel, posterior state apart from the exogenous margin movement, inspection cost, tastes, and utility scale fixed. Because the state space is finite and $S = O + g(Z - b)_+$, the function

$$\mathcal{I}(s_t; b) = \sum_x \mathbb{P}(x \mid s_t)\big[O + g(Z(s_t(x)) - b)_+\big] - \big[O + g(Z(s_t) - b)_+\big]$$

is continuous and piecewise linear in b. If exogenous variation moves b through an interior switching boundary b^* while layer membership continues to hold, observations on opposite sides of the boundary imply

$$c^W(t+1) \leq \mathcal{I}(s_t; b^*) \quad \text{and} \quad c^W(t+1) \geq \mathcal{I}(s_t; b^*).$$

Hence

$$c^W(t+1) = \mathcal{I}(s_t; b^*).$$

This is point identification in utility units under a common deterministic cost. If the utility scale is not observed, only the corresponding normalized cost is identified. If costs are heterogeneous, opening probabilities trace the cost distribution evaluated at $\mathcal{I}(s_t; b)$, subject to the usual support and independence conditions, rather than identifying one common cost.

Finally, an outer support band can certify, but need not exactly characterize, layer membership. If the current exact terminal hull straddles b and a computable outer band for every successor lies entirely weakly on one side of b, then the successor's exact hull, being contained in that outer band, also excludes b. This is sufficient for $b \in \mathcal{E}_1(s_t)$. Failure of an outer band to certify the exclusion does not imply that the exact successor hull straddles b, which is why a top-K truncation bound cannot generally recover the exact layer by itself. □

Off-layer revealed-preference inequalities. At an arbitrary reachable state, write the exact dynamic gross gain and the one-step-and-stop gain as

$$\bar{c}(s_t) = \mathbb{E}[V(s_{t+1}) \mid s_t] - S(s_t), \qquad \mathcal{I}(s_t) = \mathbb{E}[S(s_{t+1}) \mid s_t] - S(s_t).$$

Since $V(s_{t+1}) \geq S(s_{t+1})$ state by state, $\bar{c}(s_t) \geq \mathcal{I}(s_t)$. The Bellman comparison therefore implies, up to behavior at equality,

$$\text{stop} \implies c^W(t+1) \geq \bar{c}(s_t) \geq \mathcal{I}(s_t),$$

whereas

$$\text{open} \implies c^W(t+1) \le \bar{c}(s_t),$$

with no general comparison between $c^W(t+1)$ and $\mathcal{I}(s_t)$. Thus the one-step value remains a valid cost lower bound for stoppers outside the exactness layer, but it is an opener cost upper bound only when an exactness condition establishes $\bar{c}(s_t) = \mathcal{I}(s_t)$.

Proof of the general-rule local approximation

Proof of Proposition 14. For any known finite rule Q_ϕ, the integrated likelihood in (5.7) defines a posterior predictive distribution over next-review histories. Under Bayesian updating,

$$\mathbb{E}_\phi[\mu_{t+1} \mid s_t] = \mu_t.$$

Multiplying by γ gives the decision-index martingale

$$\mathbb{E}_\phi[\Lambda_{t+1} \mid s_t] = \Lambda_t, \qquad \Lambda_t = \gamma' \mu_t.$$

Let

$$\Delta_{t+1} = \Lambda_{t+1} - \Lambda_t.$$

The smoothed stopping payoff in (D.28) is three times continuously differentiable. Writing $z = (\Lambda - b)/\varepsilon$, direct differentiation gives

$$S'_\varepsilon(\Lambda) = \sigma(z), \tag{D.103}$$

$$S''_\varepsilon(\Lambda) = \frac{1}{\varepsilon}\sigma(z)[1-\sigma(z)], \tag{D.104}$$

$$S'''_\varepsilon(\Lambda) = \frac{1}{\varepsilon^2}\sigma(z)[1-\sigma(z)][1-2\sigma(z)]. \tag{D.105}$$

In particular,

$$\overline{S}^{(3)}_\varepsilon = \sup_\Lambda |S'''_\varepsilon(\Lambda)| \le \frac{1}{4\varepsilon^2} < \infty.$$

Equivalently applying Lemma 6 to the martingale movement Δ_{t+1}, Taylor's theorem around Λ_t gives, for each successor realization,

$$S_\varepsilon(\Lambda_{t+1}) = S_\varepsilon(\Lambda_t) + S'_\varepsilon(\Lambda_t)\Delta_{t+1} + \frac{1}{2}S''_\varepsilon(\Lambda_t)\Delta^2_{t+1} + r_{t+1},$$

where

$$|r_{t+1}| \le \frac{\overline{S}^{(3)}_\varepsilon}{6}|\Delta_{t+1}|^3, \qquad \overline{S}^{(3)}_\varepsilon = \sup_\lambda |S^{(3)}_\varepsilon(\lambda)| < \infty.$$

Taking conditional expectations and using $\mathbb{E}_\phi[\Delta_{t+1} \mid s_t] = 0$ yields

$$\mathcal{I}^\phi_\varepsilon(s_t;\gamma) = \frac{1}{2}S''_\varepsilon(\Lambda_t)\,\mathrm{Var}_\phi(\Lambda_{t+1} \mid s_t) + R^\phi_\varepsilon(s_t),$$

with the remainder bound in (D.30). Finally,

$$\mathrm{Var}_\phi(\Lambda_{t+1} \mid s_t) = \mathrm{Var}_\phi(\gamma' \mu_{t+1} \mid s_t) = \gamma' \Sigma_t^\phi \gamma.$$

Substituting the expression for $S''_\varepsilon(\Lambda_t)$ proves (D.29). The stated local approximation follows immediately when the third absolute moment is lower order than the taste-weighted posterior variance. □

Corollary 16 (Full local residual-variance design formula). Fix a state s_t, a taste vector γ, and two known finite summary rules ϕ and ϕ' that are evaluated at the same current decision index Λ_t. Let

$$\mathcal{W}_\varepsilon(\Lambda_t) = \frac{1}{2\varepsilon}\sigma\left(\frac{\Lambda_t - b}{\varepsilon}\right)\left[1 - \sigma\left(\frac{\Lambda_t - b}{\varepsilon}\right)\right].$$

Then their smoothed one-step verification values satisfy

$$\mathcal{I}_\varepsilon^\phi(s_t;\gamma) - \mathcal{I}_\varepsilon^{\phi'}(s_t;\gamma) = \mathcal{W}_\varepsilon(\Lambda_t)\,\gamma'\left(\Sigma_t^\phi - \Sigma_t^{\phi'}\right)\gamma + R_\varepsilon^\phi(s_t) - R_\varepsilon^{\phi'}(s_t). \tag{D.106}$$

Therefore, up to the third-order remainders in (D.30), a rule reduces local verification pressure relative to another rule if and only if it lowers the taste-weighted posterior variance $\gamma' \Sigma_t^\phi \gamma$. For a population of taste vectors F, the corresponding local design objective is to minimize

$$\mathbb{E}_F\left[\mathcal{W}_\varepsilon(\Lambda_t(\gamma))\,\gamma' \Sigma_t^\phi(\gamma)\gamma\right]. \tag{D.107}$$

Proof of Corollary 15. Apply Proposition 14 to the two rules ϕ and ϕ' at the same current decision index Λ_t. The curvature term

$$\mathcal{W}_\varepsilon(\Lambda_t) = \frac{1}{2} S''_\varepsilon(\Lambda_t)$$

is common across the two rules. Subtracting the two expansions gives (D.106). The population objective in (D.107) follows by taking expectations over $\gamma \sim F$ and dropping terms that do not depend on the rule. The qualification about third-order remainders is exactly the remainder bound in (D.30) applied to each rule. □

Proof of Proposition 15. Holding outside-group beliefs fixed, the decision-index movement generated by the next review through the folded group is

$$\Delta_{G,t+1} = \gamma_G'(\mu_{G,t+1} - \mu_{G,t}).$$

Bayesian updating gives

$$\mathbb{E}_\phi[\mu_{G,t+1} \mid s_t] = \mu_{G,t},$$

and therefore

$$\mathbb{E}_\phi[\Delta_{G,t+1} \mid s_t] = 0, \qquad \mathrm{Var}_\phi(\Delta_{G,t+1} \mid s_t) = \gamma_G' \Sigma_{G,t}^\phi \gamma_G.$$

Applying the same Taylor expansion used in Proposition 14, but with $\Delta_{G,t+1}$ replacing $\Lambda_{t+1} - \Lambda_t$, gives

$$\mathcal{I}_\varepsilon^{\phi,G}(s_t;\gamma) = \frac{1}{2} S_\varepsilon''(\Lambda_t)\gamma_G' \Sigma_{G,t}^{\phi} \gamma_G + R_\varepsilon^{\phi,G}(s_t).$$

Since $\frac{1}{2} S_\varepsilon''(\Lambda_t) = \mathcal{W}_\varepsilon(\Lambda_t)$, this is (D.31). The remainder bound is identical to (D.30) with $|\Delta_{G,t+1}|^3$ in place of $|\Lambda_{t+1} - \Lambda_t|^3$.

If $\gamma_G' \Sigma_{G,t}^{\phi} \gamma_G = 0$, the folded group produces no local second-order posterior variance in the consumer's taste direction, so it does not generate a local verification component apart from higher-order terms. If $\gamma_G' \Sigma_{G,t}^{\phi} \gamma_G > 0$, then for finite ε the curvature factor is strictly positive. Near the decision margin the second-order term is therefore strictly positive and dominates whenever the third-order remainder is small relative to the variance term. This proves the stated local folding criterion. □

Proof of Corollary 13. Let ϕ^F denote the folded displayed rule and ϕ^0 the rule that does not display the folded group. Applying Proposition 15 to each rule gives

$$\mathcal{I}_\varepsilon^{\phi^F,G}(s_t;\gamma) = \mathcal{W}_\varepsilon(\Lambda_t)\,\gamma_G' \Sigma_{G,t}^{\phi^F} \gamma_G + R_\varepsilon^{\phi^F,G}(s_t)$$

and

$$\mathcal{I}_\varepsilon^{\phi^0,G}(s_t;\gamma) = \mathcal{W}_\varepsilon(\Lambda_t)\,\gamma_G' \Sigma_{G,t}^{\phi^0} \gamma_G + R_\varepsilon^{\phi^0,G}(s_t).$$

In the local smoothed approximation, the leading difference is therefore

$$\mathcal{W}_\varepsilon(\Lambda_t)\,\gamma_G' \Big(\Sigma_{G,t}^{\phi^F} - \Sigma_{G,t}^{\phi^0} \Big) \gamma_G.$$

If the folded chip leaves the same posterior-predictive covariance in the consumer's within-fold taste direction as nondisplay, then this leading difference is zero. In particular, if

$$\gamma_G' \Sigma_{G,t}^{\phi^F} \gamma_G = \gamma_G' \Sigma_{G,t}^{\phi^0} \gamma_G,$$

then the broad chip does not reduce the local second-order verification component relative to nondisplay. If the two rules also induce the same third-order distribution of $\gamma_G'(\mu_{G,t+1} - \mu_{G,t})$, the remainder terms are identical and the smoothed local components coincide exactly. □

Proof of the consumer-relevance ranking

Proof of Proposition 16. Under the componentwise coverage effect in (D.32), the expected one-step inspection pressure for display set S is

$$\mathcal{P}(S) = \mathbb{E}_F \left[\sum_{d \in S} \gamma_d \Omega_{d,t}^{C}(s_t) + \sum_{d \notin S} \gamma_d \Omega_{d,t}^{U}(s_t) \right] \tag{D.108}$$

$$= \sum_{d=1}^{D} \mathbb{E}_F[\gamma_d]\Omega^U_{d,t}(s_t) - \sum_{d\in S} \mathbb{E}_F[\gamma_d]\Big(\Omega^U_{d,t}(s_t) - \Omega^C_{d,t}(s_t)\Big) \tag{D.109}$$

$$= \sum_{d=1}^{D} \mathbb{E}_F[\gamma_d]\Omega^U_{d,t}(s_t) - \sum_{d\in S} \mathcal{R}_{d,t}. \tag{D.110}$$

The first term is independent of S. Therefore minimizing $\mathcal{P}(S)$ over all display sets of size K is equivalent to maximizing $\sum_{d\in S} \mathcal{R}_{d,t}$. The solution is any set containing the K largest relevance indices, with arbitrary tie-breaking among equal indices. If a selected dimension r has lower relevance than an omitted dimension q, replacing r by q changes pressure by

$$\mathcal{P}((S \setminus \{r\}) \cup \{q\}) - \mathcal{P}(S) = -\mathcal{R}_{q,t} + \mathcal{R}_{r,t} = -(\mathcal{R}_{q,t} - \mathcal{R}_{r,t}) < 0.$$

Thus the replacement reduces expected one-step inspection pressure by $\mathcal{R}_{q,t} - \mathcal{R}_{r,t}$. A top-$K$ prevalence rule is optimal for this objective if and only if its chosen dimensions coincide with the top-K dimensions under $\mathcal{R}_{d,t}$, up to ties. □

Proof of Theorem 4. Write

$$B(S) = \sum_{d\in S} \mathcal{R}_{d,t}$$

for the inspection-pressure reduction delivered by display set S. By definition of w_d,

$$\mathcal{R}_{d,t} = \mathrm{Prev}_d w_d.$$

Part (i): approximation guarantee. Because S^P maximizes total prevalence among size-K sets,

$$\sum_{d\in S^P} \mathrm{Prev}_d \geq \sum_{d\in S^R} \mathrm{Prev}_d. \tag{D.111}$$

The lower alignment bound gives

$$B(S^P) = \sum_{d\in S^P} \mathrm{Prev}_d w_d \tag{D.112}$$

$$\geq \underline{w} \sum_{d\in S^P} \mathrm{Prev}_d \tag{D.113}$$

$$\geq \underline{w} \sum_{d\in S^R} \mathrm{Prev}_d, \tag{D.114}$$

where the last inequality is (D.111). The upper alignment bound implies

$$\mathcal{R}_{d,t} = \mathrm{Prev}_d w_d \leq \overline{w}\, \mathrm{Prev}_d,$$

or equivalently

$$\mathrm{Prev}_d \geq \frac{\mathcal{R}_{d,t}}{\overline{w}}.$$

Summing this inequality over S^R and substituting into (D.114) yields

$$B(S^P) \geq \frac{\underline{w}}{\overline{w}} \sum_{d \in S^R} \mathcal{R}_{d,t} = \frac{\underline{w}}{\overline{w}} B(S^R).$$

Since $B(S^R) > 0$, division gives (5.31). If $w_d = w$ for every dimension, then $B(S) = w \sum_{d \in S} \mathrm{Prev}_d$, so every prevalence-maximizing set is relevance-maximizing and the bound holds with equality at one.

Part (ii): sharpness. Fix $0 < \underline{w} < \overline{w}$, take $K = 1$ and $D = 2$, and let

$$\mathrm{Prev}_1 = q(1+\delta), \qquad \mathrm{Prev}_2 = q, \qquad w_1 = \underline{w}, \qquad w_2 = \overline{w},$$

where $q > 0$ is a common scale and

$$0 < \delta < \frac{\overline{w}}{\underline{w}} - 1.$$

Prevalence uniquely selects dimension one, while $(1+\delta)\underline{w} < \overline{w}$ makes dimension two uniquely optimal under consumer relevance. These scores can be implemented inside the normalized model. Let F be degenerate at $(\gamma_1, \gamma_2) = (1/2, 1/2)$, set $\Omega^C_{d,t} = 0$, and choose

$$\Omega^U_{1,t} = 2q(1+\delta)\underline{w}, \qquad \Omega^U_{2,t} = 2q\overline{w}.$$

Choosing

$$0 < q \leq \min\left\{\frac{1}{1+\delta}, \frac{1}{2\max\{(1+\delta)\underline{w}, \overline{w}\}}\right\}$$

ensures both $0 < \mathrm{Prev}_d \leq 1$ and $0 \leq \Omega^C_{d,t} \leq \Omega^U_{d,t} \leq 1$. Equation (5.29) delivers exactly the stated relevance scores. The achieved ratio is

$$\frac{B(S^P)}{B(S^R)} = \frac{q(1+\delta)\underline{w}}{q\overline{w}} = \frac{(1+\delta)\underline{w}}{\overline{w}}.$$

As $\delta \downarrow 0$, this converges to $\underline{w}/\overline{w}$ from above. Hence no larger constant can be a dimension-free guarantee. When $\underline{w} = \overline{w}$, part (i) already gives exact optimality.

Part (iii): no-alignment converse. Fix any $K < D$ and $\varepsilon > 0$. Choose positive prevalence scores with

$$\mathrm{Prev}_d = \frac{3}{4} \quad (d = 1, \ldots, K), \qquad \mathrm{Prev}_{K+1} = \frac{1}{2}, \qquad 0 < \mathrm{Prev}_d < \frac{1}{2} \quad (d > K+1). \tag{D.115}$$

There are exactly K highest-prevalence dimensions, so the top-K rule uniquely selects the set

$$S^P = \{1, \ldots, K\}.$$

We now construct the relevance scores from admissible model objects rather than postulating them. Let F be degenerate at the normalized uniform taste vector, so $\mathbb{E}_F[\gamma_d] = 1/D$ for every d. Fix $r_0 \in (0, 1/D]$, choose $\delta \in (0, 1)$, set $\Omega^C_{d,t} = 0$ for every d, and set

$$\Omega^U_{d,t} = \begin{cases} Dr_0\delta, & d = 1, \dots, K, \\ Dr_0, & d = K+1, \\ 0, & d > K+1. \end{cases} \tag{D.116}$$

These objects satisfy $0 \le \Omega^C_{d,t} \le \Omega^U_{d,t} \le 1$. By the definition of the consumer-relevance index, they imply

$$\mathcal{R}_{d,t} = r_0\delta \quad \text{for } d = 1, \dots, K, \qquad \mathcal{R}_{K+1,t} = r_0,$$

with all remaining relevance indices equal to zero. The prevalence-selected reduction is $Kr_0\delta$, whereas the relevance-optimal size-K reduction is $r_0 + (K-1)r_0\delta$: it selects dimension $K+1$ and any $K-1$ of the dimensions in S^P. Choose $\delta > 0$ so that

$$\frac{Kr_0\delta}{r_0 + (K-1)r_0\delta} = \frac{K\delta}{1 + (K-1)\delta} < \varepsilon.$$

Then

$$\frac{B(S^P)}{B(S^R)} < \varepsilon,$$

which proves (5.32). Replacing any displayed dimension $r \in S^P$ by dimension $K+1$ lowers expected inspection pressure by

$$\mathcal{R}_{K+1,t} - \mathcal{R}_{r,t} = r_0(1-\delta) > 0.$$

This completes the converse. □

D.1.3 Welfare representation and click nonordering

Proof of Main Paper Theorem 5. The proof has five steps. The first two derive the unrestricted dynamic welfare representation. The third connects the one-review deviation to the dynamic initiation gain. The last two construct strict open-set examples delivering the two opposite welfare orderings for the same ordering of review-opening rates.

Step 1: unrestricted fixed-cost welfare. Fix a realized entry output ω. The first-inspection cost is $c = c^W(1)$, while $V_\phi(s_1)$ contains the optimal continuation policy and all costs from the second inspection onward. The unrestricted entry Bellman equation is

$$\begin{aligned} V_\phi(\omega) &= \max\{S_\phi(\omega), -c + \mathbb{E}_\phi[V_\phi(s_1) \mid \omega]\} && \text{(D.117)} \\ &= S_\phi(\omega) + \max\{0, \mathbb{E}_\phi[V_\phi(s_1) \mid \omega] - S_\phi(\omega) - c\} && \text{(D.118)} \\ &= S_\phi(\omega) + \left(\mathcal{I}^\infty_\phi(\omega) - c\right)_+. && \text{(D.119)} \end{aligned}$$

Taking expectations over the common prior, corpus technology, and output kernel gives

$$\mathbb{E}_\phi[V_\phi(\omega)] = \mathbb{E}_\phi\Big[S_\phi(\omega) + \big(\mathcal{I}_\phi^\infty(\omega) - c\big)_+\Big], \tag{D.120}$$

which is Main Paper equation (5.34). At output ω, opening is optimal exactly when

$$\begin{aligned} -c + \mathbb{E}_\phi[V_\phi(s_1) \mid \omega] &\geq S_\phi(\omega) && \text{(D.121)} \\ \iff \qquad \mathcal{I}_\phi^\infty(\omega) &\geq c. && \text{(D.122)} \end{aligned}$$

Integrating the indicator of this event gives Main Paper equation (5.35).

Step 2: heterogeneous first-inspection costs. Let $\widetilde{c} \geq 0$ be independent of the information environment, have cdf F_c, and be observed before the opening choice. Only the first cost is heterogeneous; the future cost schedule is fixed, so $\mathcal{I}_\phi^\infty(\omega)$ does not depend on $\widetilde{c}$. For any fixed gross value $I \geq 0$, write the positive part as

$$(I - \widetilde{c})_+ = \int_0^I \mathbf{1}\{\widetilde{c} \leq u\}\, du. \tag{D.123}$$

Indeed, the integral is zero when $\widetilde{c} > I$, and otherwise has length $I - \widetilde{c}$. Since the integrand is nonnegative, Tonelli's theorem gives

$$\begin{aligned} \mathbb{E}_{\widetilde{c}}[(I - \widetilde{c})_+] &= \mathbb{E}_{\widetilde{c}}\left[\int_0^I \mathbf{1}\{\widetilde{c} \leq u\}\, du\right] && \text{(D.124)} \\ &= \int_0^I \mathbb{P}(\widetilde{c} \leq u)\, du && \text{(D.125)} \\ &= \int_0^I F_c(u)\, du. && \text{(D.126)} \end{aligned}$$

Applying (D.126) output by output to (D.119), and then integrating over ω, proves Main Paper equation (5.36). Conditional on ω, the opening probability is

$$\mathbb{P}(\widetilde{c} \leq \mathcal{I}_\phi^\infty(\omega) \mid \omega) = F_c(\mathcal{I}_\phi^\infty(\omega)), \tag{D.127}$$

when costs are atomless or indifference is resolved toward inspection. Taking expectations proves Main Paper equation (5.37). For the matched states in Main Paper Corollary 3, the stopping payoff is common and $\mathcal{I}_d^O < \mathcal{I}_d^E$. On the one-step-exact layer these values equal the corresponding dynamic initiation gains. Subtracting the two applications of (D.126) therefore gives

$$\left[S + \int_0^{\mathcal{I}_d^E} F_c(u)\, du\right] - \left[S + \int_0^{\mathcal{I}_d^O} F_c(u)\, du\right] \tag{D.128}$$

$$= \int_{\mathcal{I}_d^O}^{\mathcal{I}_d^E} F_c(u)\,du\,, \tag{D.129}$$

which proves Main Paper equation (5.41). This subtraction is conditional on the two matched nondisplay states; it is not an ex ante ranking of the two selection rules.

Step 3: one-review lower bound and exactness layer. At every first-review successor, immediate stopping remains feasible, so

$$V_\phi(s_1) \geq S(s_1;\gamma). \tag{D.130}$$

Taking conditional expectations and subtracting $S_\phi(\omega)$ gives

$$\mathcal{I}_\phi^\infty(\omega) - \mathcal{I}_\phi^1(\omega) = \mathbb{E}_\phi\big[V_\phi(s_1) - S(s_1;\gamma) \mid \omega\big] \tag{D.131}$$

$$\geq 0. \tag{D.132}$$

Moreover, posterior beliefs are a martingale and the stopping payoff is the maximum of affine actions, hence convex. Conditional Jensen therefore gives

$$\mathcal{I}_\phi^1(\omega) = \mathbb{E}_\phi[S(s_1;\gamma) \mid \omega] - S_\phi(\omega) \geq 0. \tag{D.133}$$

Equations (D.132) and (D.133) prove the hierarchy in Main Paper equation (5.38). Because the integrand in (D.132) is nonnegative, equality holds if and only if $V_\phi(s_1) = S(s_1;\gamma)$ at every positive-probability successor. This is exactly one-step exactness. Thus the one-review condition is always sufficient for starting search and is necessary precisely on that layer.

Step 4: fewer openings and higher welfare through information substitution. Consider $D = 2$, $K = 1$, and $N = 2$. Reviews are one-topic: one review covers dimension 1 and one covers dimension 2, so $C_1 = C_2 = 1$. Qualities are independent with $\mathbb{P}(\theta_1 = 1) = \mathbb{P}(\theta_2 = 1) = 1/2$, and covered signs have common precision $\rho \in (1/2, 1)$. The consumer has $\gamma_1 = 0$, $\gamma_2 = 1$, observed utility intercept $A = 0$, and outside payoff $O = 1/2$. Before receiving information about dimension 2, her stopping payoff is therefore

$$S_0 = \max\{O, \mathbb{E}[\theta_2]\} = \frac{1}{2}. \tag{D.134}$$

A single dimension-2 signal is positive with predictive probability $1/2$. The posterior is ρ after a positive signal and $1-\rho$ after a negative signal. Hence its gross decision value is

$$B(\rho) = \frac{1}{2}\max\left\{\frac{1}{2}, \rho\right\} + \frac{1}{2}\max\left\{\frac{1}{2}, 1-\rho\right\} - \frac{1}{2} \tag{D.135}$$

$$= \frac{1}{2}\rho + \frac{1}{4} - \frac{1}{2} \tag{D.136}$$

$$= \frac{2\rho - 1}{4} > 0. \tag{D.137}$$

Let rule Q^{rel} always display dimension 2 and report its singleton majority cell. The free cell reveals the only dimension-2 sign in the corpus; opening either review afterward cannot provide further payoff-relevant information. Thus

$$\text{Read}(Q^{\text{rel}}; c) = 0, \qquad \mathcal{W}(Q^{\text{rel}}; c) = \frac{1}{2} + B(\rho). \tag{D.138}$$

Let rule Q^{irr} instead always display dimension 1. Its cell is payoff irrelevant and independent of θ_2. A uniformly opened review covers dimension 2 with probability 1/2, so

$$\mathcal{I}_{\text{irr}} = \frac{1}{2} B(\rho). \tag{D.139}$$

Choose

$$0 < c < \frac{1}{2} B(\rho) \tag{D.140}$$

and choose all later inspection costs above the maximal remaining full-information gain. Then the dynamic problem is one-step exact,

$$\text{Read}(Q^{\text{irr}}; c) = 1, \qquad \mathcal{W}(Q^{\text{irr}}; c) = \frac{1}{2} + \frac{1}{2} B(\rho) - c. \tag{D.141}$$

Subtracting gives

$$\mathcal{W}(Q^{\text{rel}}; c) - \mathcal{W}(Q^{\text{irr}}; c) = \frac{1}{2} B(\rho) + c > 0. \tag{D.142}$$

Thus the rule with fewer openings has strictly higher welfare. All behavioral and welfare inequalities are strict and the finite-state values are continuous in (ρ, c, O, A, γ), so they persist on an open neighborhood of these primitives.

Step 5: fewer openings and lower welfare through tail-dimension suppression. Now take $D = 2$, $K = 1$, and $N = 4$, retaining the same independent quality priors, tastes, and stopping margin. Reviews remain one-topic. The coverage matrix has two possible count vectors:

$$(C_1, C_2) = \begin{cases} (3, 1), & \text{state } L \text{ with probability } 1 - p, \\ (1, 3), & \text{state } H \text{ with probability } p, \end{cases} \qquad 0 < p < \frac{1}{2}. \tag{D.143}$$

Conditional on each count vector, review order is uniformly random. Compare two truthful one-slot kernels that use the same componentwise majority cell. Rule Q^P selects the dimension with the larger realized coverage count. Rule Q^R selects either dimension with probability 1/2, independently of the coverage matrix, and then reports the selected dimension's majority cell.

First calculate at the boundary $\rho = 1$; continuity will return the construction to the maintained interior precision set. Equation (D.137) gives $B(1) = 1/4$. Under prevalence selection, display of dimension 1 reveals state L, so an opened review reaches the omitted relevant dimension with probability 1/4. Its gross one-step value is therefore

$$\mathcal{I}_L^P = \frac{1}{4}B(1) = \frac{1}{16}. \tag{D.144}$$

In state H, prevalence selection displays dimension 2 and its majority cell reveals θ_2 perfectly at $\rho = 1$, so no review is opened.

Under randomized display, selection of dimension 1 carries no information about the coverage state. Its sentiment cell is also independent of the coverage state and of θ_2: with a symmetric prior on θ_1, an odd-count majority cell is favorable with unconditional probability 1/2 whether $C_1 = 1$ or $C_1 = 3$. Consequently,

$$\eta_2^R = \mathbb{E}\left[\frac{C_2}{4}\right] = \frac{(1-p)1 + p3}{4} = \frac{1+2p}{4}, \tag{D.145}$$

$$\mathcal{I}_{\text{omit}}^R = \eta_2^R B(1) = \frac{1+2p}{16}. \tag{D.146}$$

Choose the first-inspection cost in the nonempty interval

$$\frac{1}{16} < c < \frac{1+2p}{16}, \tag{D.147}$$

and again choose later costs high enough to make the first inspection the only potentially optimal one. Prevalence selection then induces no opening in either coverage state, whereas randomized display induces opening exactly when it selects dimension 1. Hence

$$\text{Read}(Q^P; c) = 0 < \frac{1}{2} = \text{Read}(Q^R; c). \tag{D.148}$$

At $\rho = 1$, a displayed dimension-2 majority cell reveals θ_2, so its expected stopping payoff is 3/4. Prevalence selection displays dimension 2 only in state H, giving

$$\mathcal{W}(Q^P; c) = (1-p)\frac{1}{2} + p\frac{3}{4} = \frac{1}{2} + \frac{p}{4}. \tag{D.149}$$

Randomized display selects dimension 2 with probability 1/2, yielding payoff 3/4. When it selects dimension 1, the consumer opens and obtains net gain $\mathcal{I}_{\text{omit}}^R - c$. Therefore

$$\mathcal{W}(Q^R; c) = \frac{1}{2}\cdot\frac{3}{4} + \frac{1}{2}\left(\frac{1}{2} + \mathcal{I}_{\text{omit}}^R - c\right) = \frac{5}{8} + \frac{1}{2}\left(\mathcal{I}_{\text{omit}}^R - c\right). \tag{D.150}$$

Subtracting (D.149) gives

$$\mathcal{W}(Q^R; c) - \mathcal{W}(Q^P; c) = \frac{1}{8} - \frac{p}{4} + \frac{1}{2}\left(\mathcal{I}_{\text{omit}}^R - c\right) \tag{D.151}$$

$$> 0, \tag{D.152}$$

because $p < 1/2$ and $c < \mathcal{I}^R_{\text{omit}}$. Thus the prevalence rule has both fewer openings and lower ex ante consumer welfare under the common data-generating distribution.

The boundary calculation delivers an admissible interior open set. All finite-state likelihoods, posteriors, stopping payoffs, and one-step values are continuous in ρ. The inequalities in (D.147), (D.148), and (D.152) are strict. Moreover, after a displayed dimension-2 majority cell, residual verification value converges to zero as $\rho \uparrow 1$. Hence there exists $\varepsilon > 0$ such that every $\rho \in (1-\varepsilon, 1)$ sufficiently close to one, together with nearby values of (p, c, O, A, γ), preserves the same opening and welfare orderings. These precisions satisfy the maintained restriction $\rho \in (1/2, 1)$. This proves both nonordering claims on open sets of admissible primitives. □

Proof of Theorem 3. Let $S^*(\lambda)$ denote the set of dimensions with the K largest limiting positive coverage shares. Since the K-th and $(K+1)$-st largest positive shares are distinct, convergence $C_d(M)/N \to \lambda_d$ in probability implies

$$\mathbb{P}(\mathcal{S}^K(M) = S^*(\lambda)) \to 1.$$

This proves the first display-set claim.

Consider any $d \in S^*(\lambda)$ with $\lambda_d > 0$. Then $C_d(M) \to \infty$ in probability. Conditional on $\theta_d = 1$, the share of favorable covered signs on d converges in probability to $\rho_d > 1/2$. Conditional on $\theta_d = 0$, it converges in probability to $1-\rho_d < 1/2$. Hence the componentwise majority cell on dimension d equals θ_d with probability approaching one. With an interior prior, Bayes' rule implies

$$\mu_{d,0}(a) \to \theta_d$$

in probability under the data-generating state and along the high-probability realization of the displayed majority cell. The stopping payoff is Lipschitz in the taste-weighted belief index. Therefore, after the displayed posterior has converged to a degenerate belief, any fixed number of additional raw signs on the same displayed dimension can change expected stopping payoff by only $o(1)$. To see this without treating the summary and later signs as independent, fix any finite opened-sign history on d. Its count changes the remaining majority threshold by only a constant. Conditional on the true state, the resulting majority-consistency probability still converges to one; conditional on the false state, it still converges to zero. The exact posterior-odds expression, including the ratio of those consistency probabilities, therefore converges to the same certainty limit after every fixed history. Finitely many reviews cannot undo the degenerate posterior; Theorem 1 reinforces this conclusion because summary-disconfirming signals are additionally dampened before the summary cell is spent. This proves part (i).

For part (ii), use the topic-arrival formula at entry:

$$\eta_{d,0}(a, \emptyset) = \mathbb{E}\left[\frac{C_d(M)}{N} \,\middle|\, a, \mathcal{S}^K(M)\right].$$

For $d \notin S^*(\lambda)$, the event $\mathcal{S}^K(M) = S^*(\lambda)$ occurs with probability approaching one and

$C_d(M)/N \to \lambda_d$ in probability. Since $C_d(M)/N \in [0,1]$, bounded convergence applied on this high-probability event gives

$$\eta_{d,0}(a,\emptyset) \to \lambda_d$$

in probability. If the realized displayed-sign cell is conditioned on as well, sign-invariance and prior independence imply that it does not change the omitted dimension's sign posterior in the limit:

$$\mathbb{P}(\theta_d = 1 \mid a, \mathcal{S}^K(M) = S^*(\lambda)) - \mathbb{P}(\theta_d = 1 \mid \mathcal{S}^K(M) = S^*(\lambda)) \to 0.$$

Because selection is sign-invariant and coverage is independent of latent quality, the event that d is omitted is a function of M, not of the realized signs on dimension d. Under prior independence across dimensions, sentiment cells displayed on other dimensions do not directly update θ_d. Thus omission affects dimension d's verification role through the limiting topic-arrival probability λ_d, not through a direct sentiment posterior.

For part (iii), fix an omitted dimension d^* satisfying (5.27)–(5.28). Let

$$\Lambda^{\infty}_{-d^*} = \sum_{r \in S^*(\lambda)} \gamma_r \theta_r + \sum_{\substack{q \notin S^*(\lambda) \\ q \neq d^*}} \gamma_q \mu_{q,0}.$$

The dimension-specific margin is

$$\widehat{b}_{d^*} = \frac{b - \Lambda^{\infty}_{-d^*}}{\gamma_{d^*}}.$$

Let

$$\pi^+_{d^*} = \mu_{d^*,0}\rho_{d^*} + (1 - \mu_{d^*,0})(1 - \rho_{d^*})$$

be the predictive probability of a favorable covered sign, and let

$$\mu^+_{d^*} = \mathbb{P}(\theta_{d^*} = 1 \mid r_{nd^*} = 1, m_{nd^*} = 1), \qquad \mu^-_{d^*} = \mathbb{P}(\theta_{d^*} = 1 \mid r_{nd^*} = 0, m_{nd^*} = 1)$$

be the one-signal posteriors from the omitted dimension's prior. Since $\rho_{d^*} > 1/2$ and $\mu_{d^*,0} \in (0,1)$, $\mu^-_{d^*} < \mu_{d^*,0} < \mu^+_{d^*}$. The one-step Jensen gap from a covered signal on d^* is

$$\begin{aligned} J^{\infty}_{d^*} = {} & \pi^+_{d^*}\Big[O + \big(\Lambda^{\infty}_{-d^*} + \gamma_{d^*}\mu^+_{d^*} - b\big)_+\Big] \\ & + (1 - \pi^+_{d^*})\Big[O + \big(\Lambda^{\infty}_{-d^*} + \gamma_{d^*}\mu^-_{d^*} - b\big)_+\Big] \\ & - \Big[O + \big(\Lambda^{\infty}_{-d^*} + \gamma_{d^*}\mu_{d^*,0} - b\big)_+\Big]. \end{aligned}$$

The one-step straddle condition (5.28) says exactly that the kink lies strictly between the two possible one-signal posterior indices. By the scalar decision-margin argument, equivalently by strict convexity of the kink over a two-point support that straddles it, $J^{\infty}_{d^*} > 0$.

Since $\eta_{d^*,0} \to \lambda_{d^*} > 0$, the limiting gross value contributed by the event that the passive review covers this omitted dimension is

$$\lambda_{d^*} J_{d^*}^{\infty} > 0.$$

The focal-separability condition is what turns this focal component into a lower bound for the full review experiment. On the event that the next review covers d^*, all other payoff-relevant posterior components are fixed, so the conditional gross gain is exactly $J_{d^*}^{\infty}$. On the complementary event, focal separability gives

$$\mathbb{E}\big[S(s_1;\gamma) \mid s_0,\, m_{n_1 d^*} = 0\big] \geq S(s_0;\gamma).$$

Consequently,

$$\begin{aligned} \mathcal{I}^{MD,\infty}(s_0;\gamma) &= \lambda_{d^*}\big(\mathbb{E}\big[S(s_1;\gamma) \mid s_0,\, m_{n_1 d^*} = 1\big] - S(s_0;\gamma)\big) \\ &\quad + (1-\lambda_{d^*})\big(\mathbb{E}\big[S(s_1;\gamma) \mid s_0,\, m_{n_1 d^*} = 0\big] - S(s_0;\gamma)\big) \\ &\geq \lambda_{d^*} J_{d^*}^{\infty} > 0. \end{aligned}$$

By continuity of the finite-state posterior and payoff objects, the finite-N one-step value is eventually above any cost strictly below this positive lower bound. Therefore opening a review and then stopping dominates immediate stopping for all sufficiently large N. Since the unrestricted Bellman problem can mimic that feasible one-step policy, verification also starts in the full dynamic problem.

Finally, consider adding d^* as one additional free summary slot. Since $\lambda_{d^*} > 0$, the coverage count for d^* diverges and the dimension-level majority summary reveals θ_{d^*} with probability approaching one. The limiting gross value of the extra slot is therefore

$$\mathcal{V}_{d^*}^{slot} = \mathbb{E}\big[O + \big(\Lambda_{-d^*}^{\infty} + \gamma_{d^*}\theta_{d^*} - b\big)_+\big] - \big[O + \big(\Lambda_{-d^*}^{\infty} + \gamma_{d^*}\mu_{d^*,0} - b\big)_+\big],$$

where the expectation is over the posterior prior on θ_{d^*}. Because the weaker full-information pivotality condition $0 < \widehat{b}_{d^*} < 1$ puts the kink strictly between the low and high quality states of dimension d^*, and because $\mu_{d^*,0} \in (0,1)$, the convexity of $x \mapsto (x-b)_+$ implies $\mathcal{V}_{d^*}^{slot} > 0$. This proves that the value of the additional slot remains bounded away from zero while scalar verification value on displayed dimensions vanishes. □

Proof of Corollary 14. Write

$$\bar{\gamma} = \sup_{\gamma \in \Gamma} \sum_{d=1}^{D} \gamma_d < \infty.$$

The finiteness of D and the componentwise majority rule are used in both directions of the argument.

Part (i). Suppose $K_N \geq D^+$ for all sufficiently large N. Every dimension in $\mathcal{D}^+$ has coverage count diverging in probability. Since dimensions outside $\mathcal{D}^+$ have coverage $o_p(N)$, every positive-share dimension belongs to the top-K_N set with probability approaching one. Conditional on that event, the majority cell of each $d \in \mathcal{D}^+$ identifies θ_d with probability approaching one by the same law-of-large-numbers argument used in Theorem 3(i). A fixed number of subsequently opened signals cannot undo this posterior concentration.

For a dimension $d \notin \mathcal{D}^+$, the entry topic-arrival probability converges to zero. Because there are only finitely many dimensions and at most $\bar{m}$ reviews can be opened, the posterior-predictive probability that a continuation plan encounters any zero-share dimension is $o(1)$, uniformly over feasible adaptive plans. Thus, uniformly over $\gamma \in \Gamma$, every posterior movement in the decision index

$$\Lambda_t = \sum_{d=1}^{D} \gamma_d \mu_{d,t}$$

after at most $\bar{m}$ inspections is $o_p(1)$, and boundedness gives

$$\sup_{\gamma \in \Gamma} \sup_{\sigma : \nu_\sigma \le \bar{m}} \mathbb{E}[|\Lambda_{\nu_\sigma} - \Lambda_0|] \longrightarrow 0.$$

For every stopping margin b, the function $x \mapsto O + (x - b)_+$ is one-Lipschitz. Hence the gross payoff improvement under any such policy σ is at most the preceding expected absolute index movement. This proves uniform convergence of gross continuation value to zero. A positive lower bound on the first inspection cost then implies immediate stopping for all sufficiently large N.

Part (ii). Let $K_N = K < D^+$ eventually. The strict limiting cutoff selects a deterministic set $S^*(\lambda)$ with probability approaching one, and it omits some $d^* \in \mathcal{D}^+$. Consider first the focal type $\gamma^0 = g e_{d^*}$ for any $g > 0$. Since d^* is omitted, sign-invariant selection and coverage-quality independence leave its limiting posterior at $\mu_{d^*,0}$. Choose a normalized margin $\widehat{b} \in (\mu^-_{d^*}, \mu^+_{d^*})$ and set $b^0 = g\widehat{b}$. The next review covers d^* with limiting probability $\lambda_{d^*} > 0$. When it does not cover d^*, no payoff-relevant posterior changes for the focal type. When it does cover d^*, the two successor decision indices straddle b^0. Let

$$\pi^+_{d^*} = \mu_{d^*,0}\rho_{d^*} + (1 - \mu_{d^*,0})(1 - \rho_{d^*})$$

be the limiting predictive probability of a favorable covered signal. Therefore

$$\begin{aligned}
\mathcal{I}^{MD,\infty}(s_0; \gamma^0) = \lambda_{d^*}\Big[&\pi^+_{d^*}\big(O + (g\mu^+_{d^*} - b^0)_+\big) \\
&+ (1 - \pi^+_{d^*})\big(O + (g\mu^-_{d^*} - b^0)_+\big) \\
&- \big(O + (g\mu_{d^*,0} - b^0)_+\big)\Big] > 0.
\end{aligned}$$

The inequality is strict. By continuity of this finite-state one-step value in (γ, b), it remains strict on a neighborhood of (γ^0, b^0). That neighborhood contains a pair with every taste weight strictly positive; a smaller Euclidean neighborhood around such an interior pair remains inside $\mathbb{R}^D_{++} \times \mathbb{R}$ and preserves the strict inequality. Convergence of the finite-N posterior objects then implies that the one-step value is bounded away from zero on a possibly smaller open neighborhood for all sufficiently large N. Any positive first inspection cost below that bound makes one-step verification, and hence dynamic verification, optimal. This proves the converse. □

Proof of Theorem 3. Fix dimension d. Conditional on $\theta_d = 1$ and on the set of covered signals, covered signs are independent Bernoulli draws with success probability ρ_d. Conditional on $\theta_d = 0$, they are independent Bernoulli draws with success probability $1 - \rho_d$. Since $C_d(M) \to \infty$ in probability, the random-sample-size law of large numbers gives

$$Z_d \equiv \frac{Y_d}{C_d(M)} \xrightarrow{p} \rho_d \quad \text{under } \theta_d = 1, \qquad Z_d \xrightarrow{p} 1 - \rho_d \quad \text{under } \theta_d = 0.$$

First suppose $\rho_d > q_h$. Then $1 - \rho_d < 1 - q_h = q_\ell$. Hence

$$\mathbb{P}(a_d = + \mid \theta_d = 1) \to 1, \qquad \mathbb{P}(a_d = + \mid \theta_d = 0) \to 0,$$

and symmetrically

$$\mathbb{P}(a_d = - \mid \theta_d = 0) \to 1, \qquad \mathbb{P}(a_d = - \mid \theta_d = 1) \to 0.$$

Bayes' rule with an interior prior implies that the posterior after a favorable displayed cell converges to one and the posterior after an unfavorable displayed cell converges to zero. A fixed number L of later covered signals changes the residual count and each threshold by at most L. More explicitly, after $\ell \le L$ covered signals containing y positives, the favorable endpoint consistency probability is

$$H_+(\ell, y; p, C_d) = \mathbb{P}\big(\mathrm{Bin}(C_d - \ell, p) \ge \lceil q_h C_d \rceil - y\big).$$

Because ℓ and y are fixed while $C_d \to \infty$, this probability converges to one at $p = \rho_d > q_h$ and to zero at $p = 1 - \rho_d < q_\ell < q_h$. The unfavorable endpoint is symmetric. Substitution into the exact integrated posterior therefore shows that the posterior remains at the same certainty limit after every fixed history. Since the stopping payoff is Lipschitz in beliefs, the gross value of any fixed finite verification plan about dimension d converges to zero.

Now suppose $\rho_d \in (1/2, q_h)$. Because $q_\ell = 1 - q_h$,

$$q_\ell < 1 - \rho_d < \frac{1}{2} < \rho_d < q_h.$$

Both state-contingent limits of Z_d therefore lie inside the mixed band. The law of large numbers implies

$$\mathbb{P}(a_d = \mathsf{m} \mid \theta_d = 1) \to 1, \qquad \mathbb{P}(a_d = \mathsf{m} \mid \theta_d = 0) \to 1.$$

Bayes' rule then gives

$$\mathbb{P}(\theta_d = 1 \mid a_d = \mathsf{m}) = \frac{\mu_{d,0}\, \mathbb{P}(a_d = \mathsf{m} \mid \theta_d = 1)}{\mu_{d,0}\, \mathbb{P}(a_d = \mathsf{m} \mid \theta_d = 1) + (1 - \mu_{d,0})\, \mathbb{P}(a_d = \mathsf{m} \mid \theta_d = 0)} \longrightarrow \mu_{d,0}.$$

Let $\Lambda^{\mathsf{m}}_{-d}$ denote the limiting contribution of all other dimensions to the buy payoff after the mixed cell, and define the dimension-specific margin

$$\widehat{b}^{\mathsf{m}}_d = \frac{b - \Lambda^{\mathsf{m}}_{-d}}{\gamma_d}.$$

From the mixed-cell posterior limit, the predictive probability of a positive covered signal on d converges to

$$\pi^+_d = \mu_{d,0}\rho_d + (1 - \mu_{d,0})(1 - \rho_d),$$

and the two one-signal posteriors converge to

$$\mu^+_d = \frac{\mu_{d,0}\rho_d}{\mu_{d,0}\rho_d + (1 - \mu_{d,0})(1 - \rho_d)}, \qquad \mu^-_d = \frac{\mu_{d,0}(1 - \rho_d)}{\mu_{d,0}(1 - \rho_d) + (1 - \mu_{d,0})\rho_d}.$$

Because $\rho_d > 1/2$, $\mu^-_d < \mu_{d,0} < \mu^+_d$. Under the stated straddle condition,

$$\mu^-_d < \widehat{b}^{\mathsf{m}}_d < \mu^+_d.$$

The limiting one-signal Jensen gap from a covered signal on d is therefore

$$\begin{aligned} J^{\mathsf{m}}_d = {} & \pi^+_d\Big[O + \big(\Lambda^{\mathsf{m}}_{-d} + \gamma_d\mu^+_d - b\big)_+\Big] \\ & + (1 - \pi^+_d)\Big[O + \big(\Lambda^{\mathsf{m}}_{-d} + \gamma_d\mu^-_d - b\big)_+\Big] \\ & - \Big[O + \big(\Lambda^{\mathsf{m}}_{-d} + \gamma_d\mu_{d,0} - b\big)_+\Big]. \end{aligned}$$

The kink lies strictly between the two possible posterior indices, so the same decision-margin/Jensen-gap argument used in the scalar benchmark gives $J^{\mathsf{m}}_d > 0$. If the next review covers d with limiting probability $\lambda_d > 0$, focal separability implies that the covered branch changes no other payoff-relevant posterior component and that the complementary branch has weakly nonnegative gross value. The limiting one-step gross verification value is therefore bounded below by

$$\lambda_d J^{\mathsf{m}}_d > 0.$$

Hence any first-inspection cost strictly below this positive lower bound is eventually dominated by the feasible policy that opens one review and then stops. The full dynamic Bellman problem can always mimic that feasible one-step policy, so verification starts for all sufficiently large pools.

Finally, binary majority is the limiting special case $q_h = 1/2$. Since the model maintains $\rho_d > 1/2$, the interval $(1/2, q_h)$ is empty. Thus the mixed-band case cannot arise under the binary majority benchmark; it is a property of a genuine three-cell band rule. □

D.2 Quantitative Illustration: Formulas and Scope

This subsection derives the calculations reported in Main Paper Section 7. They are theory diagnostics evaluated at descriptive audit inputs, not structural estimates of consumer primitives.

Scalar two-point calculation

Fix a summary cell and write its current posterior as μ, its two next-review posteriors as $\mu^- < \mu^+$, and the predictive probability of the favorable successor as π. The martingale identity gives

$$\mu = \pi\mu^+ + (1-\pi)\mu^-.$$

Normalize the payoff scale to one and write the stopping payoff as $s(\mu; b) = (\mu - b)_+$. Inside the straddle region, direct substitution gives

$$\frac{\mathcal{I}(b)}{\gamma} = \begin{cases} (1-\pi)(b-\mu^-), & \mu^- < b \le \mu, \\ \pi(\mu^+ - b), & \mu \le b < \mu^+, \\ 0, & b \notin (\mu^-, \mu^+). \end{cases} \tag{D.153}$$

For example, when $b \le \mu$,

$$\begin{aligned} \frac{\mathcal{I}(b)}{\gamma} &= \pi(\mu^+ - b) - (\mu - b) \\ &= \pi\mu^+ - \pi b - \pi\mu^+ - (1-\pi)\mu^- + b \\ &= (1-\pi)(b-\mu^-). \end{aligned}$$

The other branch follows symmetrically. The two branches meet at $b = \mu$, so

$$\max_b \frac{\mathcal{I}(b)}{\gamma} = \pi(1-\pi)(\mu^+ - \mu^-). \tag{D.154}$$

For normalized cost $k = \kappa^W/\gamma$, the one-step verification region is

$$\left[\mu^- + \frac{k}{1-\pi},\ \mu^+ - \frac{k}{\pi}\right], \tag{D.155}$$

provided the lower endpoint does not exceed the upper endpoint. The values in Main Paper Table 2 evaluate μ, μ^-, μ^+, π with the finite binomial-interval likelihood (C.60), the recovered descriptive thresholds $q_\ell = 0.2707$ and $q_h = 0.6667$, prior 1/2, and the representative cell-specific (N, ρ) inputs shown in the table.

For a realized next sign x, let Δ_x be the shared-pool log-odds change and let $\lambda =$

$\log(\rho/(1-\rho))$ be the absolute raw-signal log-likelihood ratio. The reported attenuation is

$$\text{Atten}_x = 1 - \frac{|\Delta_x|}{\lambda}. \tag{D.156}$$

It is evaluated at the signal that contradicts the displayed endpoint cell. The modal-sign share used as ρ is a transparent agreement sensitivity calibration. It is not truth-validated signal precision, and conditioning that input on the realized cell means these rows cannot be used to reproduce unconditional cell frequencies.

Multidimensional count bound

Suppose a top-K summary fills all slots and reveals displayed counts. Let c_* be the smallest displayed count and $c_{\max}$ the largest. With fixed tie-breaking, every omitted dimension obeys the weak inequality $C_d \le c_*$; equality is possible when it loses a tie at the K-th slot. Because every displayed count is drawn from the common pool, $N \ge c_{\max}$. At entry,

$$\eta_{d,0} = \mathbb{E}\left[\frac{C_d}{N}\,\middle|\,\omega\right] \le \frac{c_*}{c_{\max}}. \tag{D.157}$$

Now let N^{written} be a manually recorded written-review count. If this is the same eligible pool used by the summary, so that $N^{\text{written}} = N$, and each review contributes at most one mention to any given dimension count, then conditioning on the observed summary gives the weakly sharper observable bound

$$\eta_{d,0} = \mathbb{E}\left[\frac{C_d}{N}\,\middle|\,\omega\right] \tag{D.158}$$

$$= \frac{1}{N^{\text{written}}}\,\mathbb{E}[C_d \mid \omega] \tag{D.159}$$

$$\le \frac{c_*}{N^{\text{written}}} \le \frac{c_*}{c_{\max}}. \tag{D.160}$$

The first equality uses the definition of the next-review topic-arrival probability under uniform sampling without replacement; the second uses the same-pool restriction, the first inequality uses top-K omission, and the last uses $c_{\max} \le N^{\text{written}}$. If Amazon computes chips from a smaller eligibility-filtered pool, if the displayed count is rounded or capped, or if one review can contribute multiple units to the same displayed dimension count, equation (D.160) is only a sensitivity calculation rather than a valid bound. The final inequality is strict whenever $N^{\text{written}} > c_{\max}$. If the display is under-filled because fewer than K dimensions have positive eligible coverage, exact top-K selection instead implies $C_d = \eta_{d,0} = 0$ for every omitted dimension.

For the transparent one-topic sensitivity calculation, let $\Gamma = \sum_r \gamma_r$, $w = \gamma_d/\Gamma$, and suppose the omitted dimension has prior $1/2$, precision ρ_d, and topic-arrival probability η_d. Conditional on arrival, the favorable and unfavorable posteriors are ρ_d and $1-\rho_d$, each with probability 1/2. Conditional on no arrival, the focal index does not move. Applying (D.154) and multiplying by the arrival and taste weights gives

$$\max_b \frac{\mathcal{I}_d^{MD}(b)}{\Gamma} = \frac{w\eta_d(2\rho_d - 1)}{4}. \tag{D.161}$$

For normalized cost $k = \kappa^W/\Gamma$, the focal verification region in the dimension-specific margin is

$$\left[1 - \rho_d + \frac{2k}{w\eta_d},\ \rho_d - \frac{2k}{w\eta_d}\right], \tag{D.162}$$

again provided it is nonempty. If η_d is replaced by the upper bound in (D.157), equations (D.161)–(D.162) are upper-bound diagnostics: a peak below cost rules verification out, while a peak above cost only leaves it possible.

The contemporaneous depth archive contains 1,296 products with summaries: 842 fill all eight slots and 454 are under-filled. All 842 full-slot products have count-bearing chips. Their median denominator-free bound in (D.157) is 0.164. Setting $\rho_d = 0.70$ and $w = 0.50$ produces a median peak bound of 0.0082; at normalized costs $k = 0.002$ and $k = 0.005$, 96.1% and 75.3%, respectively, are not ruled out.

The manual written-count tracker supplies a positive denominator for 279 of the 842 full-slot products. Every one satisfies $c_{\max} \le N^{\text{written}}$. For this nonrandom denominator subsample, the quartiles of c_*/N^{written} are

$$Q_{.25} = 0.0160, \qquad Q_{.50} = 0.0420, \qquad Q_{.75} = 0.0718.$$

The baseline median peak is therefore

$$\text{Med}\left[\max_b \frac{\mathcal{I}_d^{MD}(b)}{\Gamma}\right] = \frac{(0.50)(0.0420)\{2(0.70) - 1\}}{4} \tag{D.163}$$

$$= 0.00210. \tag{D.164}$$

At $k = 0.002$, 52.0% of the product-level upper bounds weakly exceed cost; at $k = 0.005$, 4.7% do; and at $k = 0.010$, none do. To check repeated ASINs that share a review medley, I collapse identical tuples of the written count, ratings count, and complete eight-chip count/sentiment vector. This leaves 212 distinct review-pool signatures, with median arrival bound 0.0417, median peak 0.00209, and corresponding shares 52.4% and 6.1%.

For the product-level median $\eta = 0.0419847$, equation (D.162) gives, at $\rho_d = 0.70$ and $w = 0.50$,

$$\text{at } k = 0.001: \quad \left[0.30 + \frac{0.002}{(0.50)(0.0419847)}, 0.70 - \frac{0.002}{(0.50)(0.0419847)}\right] = [0.3953, 0.6047],$$

$$\text{at } k = 0.002: \quad \left[0.30 + \frac{0.004}{(0.50)(0.0419847)}, 0.70 - \frac{0.004}{(0.50)(0.0419847)}\right] = [0.4905, 0.5095].$$

The interval collapses at the median peak cost $k = 0.00210$. Raising precision to $\rho_d = 0.80$, holding $w = 0.50$, raises the median peak to 0.00315, and raising the taste share to $w = 1$, holding $\rho_d = 0.70$, raises it to 0.00420. These are transparent sensitivity movements, not estimated heterogeneity.

The written-denominator subsample is concentrated in categories for which the manual tracker was filled and is not population-representative. Moreover, the manual count has not yet been shown to equal Amazon's eligibility-filtered summary pool. I therefore retain both rows in Main Paper Table 3: the displayed-count bound is denominator-free but loose, while the written-count row is sharper but depends on an explicit measurement assumption. The public ratings count is never used as the denominator. Finally, the assistant was actively queried for content beyond the displayed themes; its retrieved evidence is consequently query-conditioned rather than a random draw from the unopened review pool. Semantic omission and folding classifications are also awaiting blinded human adjudication and a discovery-recall audit. Neither object is used to estimate η_d here.

D.3 Experimental Predictions and Mixed-Cell Caveats

Supplementary statements

Proposition 17 (Coverage-alignment effect). Consider the one-topic, local-separability, one-step environment. Let two assigned rules ϕ and ϕ' generate the same state, except for the per-unit verification gain on dimension r. Then

$$\mathcal{I}^{MD,\phi'}(s_t;\gamma) - \mathcal{I}^{MD,\phi}(s_t;\gamma) = \gamma_r\Big[\Omega^{\phi'}_{r,t}(s_t) - \Omega^{\phi}_{r,t}(s_t)\Big]. \tag{D.165}$$

If the rule that covers dimension r lowers its residual gain, the reduction in the gross value of review reading is exactly proportional to the pre-treatment need weight γ_r.

This is the cleanest lab interaction implied by the model. A summary treatment should have the largest effect on review opening for participants who assign high need weight to the dimension that the treatment newly covers or resolves. The prediction is therefore not just a treatment main effect; it is a treatment-by-need-weight gradient.

Proposition 18 (Free information refinement). Fix the review corpus, inspection technology, and consumer primitives. If the output generated by rule ϕ is a garbling of the output generated by rule ϕ', then ex ante consumer value under ϕ' is weakly higher than under ϕ:

$$\mathbb{E}\big[V^{\phi'}\big] \geq \mathbb{E}\big[V^{\phi}\big]. \tag{D.166}$$

This comparison does not order review opening: a refinement can reduce verification by resolving a decision, or increase it by identifying a residual uncertainty worth resolving.

The refinement result separates welfare from click depth. A more informative summary cannot hurt a Bayesian consumer ex ante, but it need not reduce review opening in every state. Refinement can save time by resolving the decision, or it can reveal that one more raw review is worth reading because the consumer is now close to the margin.

Lemma 10 (Singleton-cell caveat for mixed sentiment). In the scalar finite-pool model with exchangeable signals and passive inspection, suppose a summary cell reveals one exact total count $Y = y$. Conditional on that cell, the signs of any sequence of inspected reviews are independent of θ. Hence positive-cost inspection of those scalar signs has zero value. In particular, a "mixed" cell defined solely as the exact majority tie $Y = N/2$ has no residual within-dimension verification value in this scalar benchmark.

The singleton caveat disciplines the mixed-chip prediction. Extra clicks after "mixed" summaries require either a nondegenerate mixed band, omitted dimensions, or multidimensional folding. A scalar cell that reveals exactly one total count leaves no residual scalar information for passive review opening to recover.

Laboratory treatment predictions and singleton mixed cells

Under the one-topic and local-separability assumptions, the exact one-step value is

$$\mathcal{I}^{MD,\phi}(s_t;\gamma) = \sum_{d=1}^{D} \gamma_d \Omega_{d,t}^{\phi}(s_t).$$

If two experimental rules differ only in their per-unit gain for dimension r, subtracting the two expressions gives

$$\mathcal{I}^{MD,\phi'}(s_t;\gamma) - \mathcal{I}^{MD,\phi}(s_t;\gamma) = \gamma_r \left[\Omega_{r,t}^{\phi'}(s_t) - \Omega_{r,t}^{\phi}(s_t)\right],$$

which proves Proposition 17. Combining the one-step cutoff with an additive cost shock gives

$$\mathbb{P}(\text{open} \mid \gamma, s_t, \phi) = F_{\xi}\left(\sum_{d} \gamma_d \Omega_{d,t}^{\phi}(s_t) - \kappa\right),$$

and the deterministic version follows when the shock is degenerate at zero.

For Proposition 18, let g be the garbling that maps the more informative output $\omega^{\phi'}$ into ω^{ϕ}. A consumer who observes $\omega^{\phi'}$ can apply g herself and then follow any policy feasible under ϕ, including all later inspection and stopping choices. That reproduces the distribution of payoffs available under ϕ. Optimizing over the larger policy set available after $\omega^{\phi'}$ proves

$$\mathbb{E}[V^{\phi'}] \geq \mathbb{E}[V^{\phi}].$$

The argument compares feasible information sets only and says nothing about the monotonicity of later review opening.

Finally, consider the scalar finite-pool environment in Lemma 10. Conditional on $Y = y$, every binary vector with exactly y positive entries is equally likely under either latent state. For any k passively sampled, distinct reviews that produce a particular ordered sign pattern $x = (x_1, \ldots, x_k)$ with $s = \sum_{\ell=1}^{k} x_\ell$ positives,

$$\mathbb{P}(x \mid Y = y, \theta) = \frac{\binom{N-k}{y-s}}{\binom{N}{y}},$$

with the probability understood to be zero when $y - s \notin \{0, \ldots, N-k\}$. This expression is independent of θ. Bayes' rule therefore leaves the posterior over θ unchanged after every such inspection history. Since the scalar signs have no direct payoff apart from information about θ, they cannot change the stopping payoff; a strictly positive inspection cost then makes inspection strictly suboptimal. Taking $y = N/2$ yields the exact-tie claim in the lemma.

D.4 Coverage-Residual Ordering and Endpoint Results

Coverage-residual ordering under homogeneous omitted dimensions

The qualified ordering in Proposition 12 follows directly from the one-step taste gain when omitted dimensions are homogeneous. Suppose that for every uncovered dimension $d \notin \mathcal{S}^K$,

$$\pi^+_{d,t}(s_t)\Delta^+_{d,t}(s_t) + (1 - \pi^+_{d,t}(s_t))\Delta^-_{d,t}(s_t) = B_t(s_t) \geq 0.$$

Then the total omitted-dimension one-step benefit for taste vector γ is

$$\mathcal{B}^U(s_t; \gamma) = \sum_{d \notin \mathcal{S}^K} \gamma_d \eta_{d,t}(s_t) B_t(s_t) \tag{D.167}$$

$$= B_t(s_t)\mathrm{RL}_t(\gamma; a, h_t). \tag{D.168}$$

Thus, under homogeneous omitted dimensions, residual learnability is a sufficient statistic for the omitted-dimension search motive. If $\mathrm{RL}_t(\gamma; a, h_t) \leq \mathrm{RL}_t(\gamma'; a, h_t)$, then $\mathcal{B}^U(s_t; \gamma) \leq \mathcal{B}^U(s_t; \gamma')$, so the omitted-dimension component of the largest cost at which one-step verification is optimal is weakly lower for γ. If covered-dimension one-step gains are zero or identical across the two taste vectors, this same inequality orders the total one-step cutoff. The coverage-share condition gives the complementary interpretation: higher coverage means less preference mass is left to be recovered through costly review reading. Without the homogeneous-wedge and covered-gain restrictions, the same scalar RL need not be sufficient because dimensions may differ in precision, posterior uncertainty, and decision-changing payoff wedges.

D.4.1 Endpoint statements and proofs

Proposition 19 (Zero-direct-coverage benchmark). Under the top-K prevalence-selection benchmark, suppose $\{d : \gamma_d > 0\} \cap \mathcal{S}^K = \emptyset$. Under prior independence across latent-quality dimensions, the observed summary vector a does not directly update any utility-relevant quality posterior: for every d with $\gamma_d > 0$,

$$\mu_{d,0}(a) = \mu_{d,0}.$$

Hence any AI effect on within-option verification must operate through the coverage probabilities $\eta_{d,t}(a, h_t)$ not through direct updating of utility-relevant quality beliefs. If, in addition, the coverage counts $\{C_d(M) : \gamma_d > 0\}$ are known ex ante, then the within-option problem coincides with the no-AI benchmark.

Proposition 20 (Reduction to the single-dimensional benchmark). Suppose the consumer cares about one dimension only, so $\gamma_{d^*} > 0$ and $\gamma_d = 0$ for $d \neq d^*$. If $d^* \in \mathcal{S}^K$, every review covers that dimension, i.e. $m_{nd^*} = 1$ for all $n = 1, \ldots, N$, and the summary component a_{d^*} is the same compression of the dimension-d^* signal pool studied in the scalar benchmark, then the consumer's within-option problem on dimension d^* is exactly the single-dimensional problem with primitives $(\mu_{d^*,0}, \rho_{d^*}, N, \gamma_{d^*})$.

Proof of Proposition 19. Let $\mathcal{D}^\gamma = \{d : \gamma_d > 0\}$ denote the utility-relevant dimensions. If $\mathcal{D}^\gamma \cap \mathcal{S}^K = \emptyset$, the realized summary signs a are functions only of sparse-review realizations on dimensions outside $\mathcal{D}^\gamma$. The event that dimensions in $\mathcal{D}^\gamma$ failed to enter the top-K summary set is a function of the coverage matrix M alone. Because latent quality is independent across dimensions and independent of M under the prior, the pair consisting of summary signs on non-utility dimensions and the top-K selection event is independent of $(\theta_d)_{d\in\mathcal{D}^\gamma}$. Therefore, for every utility-relevant dimension d,

$$\mu_{d,0}(a) = \mathbb{P}(\theta_d = 1 \mid a) = \mathbb{P}(\theta_d = 1) = \mu_{d,0}.$$

Thus the summary cannot directly update utility-relevant quality beliefs.

The omission event can still affect beliefs about coverage. In particular, the posterior distribution of M conditional on the event $d \notin \mathcal{S}^K(M)$ may differ from the prior distribution of M, so the topic-arrival probability $\eta_{d,t}(a, h_t)$ can differ from its no-summary counterpart. If the relevant coverage counts $\{C_d(M) : d \in \mathcal{D}^\gamma\}$ are known ex ante, then this residual coverage-inference channel is absent. The quality beliefs and topic-arrival probabilities relevant for utility are then the same as in the no-AI benchmark, so the within-option problem coincides with the no-AI benchmark. □

Proof of Proposition 20. Suppose $\gamma_{d^*} > 0$ and $\gamma_d = 0$ for $d \neq d^*$. All dimensions other than d^* are payoff irrelevant. If $d^* \in \mathcal{S}^K$ and every review covers d^*, then $\eta_{d^*,t}(a, h_t) = 1$ at every history with unopened reviews. There is therefore no topic-arrival risk on the only payoff-relevant dimension. Conditional on the summary component a_{d^*} and any subsequent opened signals on d^*, the posterior over θ_{d^*} is exactly the finite-pool posterior from the scalar benchmark with primitives $(\mu_{d^*,0}, \rho_{d^*}, N, \gamma_{d^*})$. Since all other dimensions have zero taste weight, they do not enter either the stopping payoff or the value of additional inspection. The Bellman problem therefore reduces exactly to the single-dimensional benchmark. □

E Notation Compendium

This table collects the main notation used in the scalar benchmark, the aggregation-rule extensions, and the finite-slot multidimensional model. A few letters are intentionally context-specific: for example, $C(\bar{M})$ is a scalar stopping-payoff constant, whereas $C_d(M)$ is a multidimensional coverage count.

Table 4: Notation Compendium

Symbol	Name or object	Meaning
General primitives and timing		
j	Product or option index	The product whose within-option information problem is being solved.
$d = 1, \dots, D$	Attribute or dimension index	A payoff-relevant product dimension in the multidimensional model.
k	Scalar signal index	Indexes the raw binary signals in the scalar benchmark.
n	Review index	Indexes sparse reviews in the multidimensional model.
N	Finite review-pool size	Number of scalar signals in the benchmark and sparse reviews in the fixed-product multidimensional model.
m	Scalar inspection count	Number of scalar raw signals already inspected.
y	Scalar positive count	Number of positive scalar signals among the m inspected signals.
t	Multidimensional time/history index	Number of sparse reviews opened after the summary.
$\mathbf{x}, P$	Observed listing attributes and price	Search-good information observed before costly review inspection in the scalar benchmark.
β, α	Utility coefficients	β loads observed characteristics; $\alpha > 0$ is the price coefficient.
γ	Scalar taste scale	Payoff gain from high latent quality in the scalar benchmark.
$\gamma = (\gamma_1, \dots, \gamma_D)$	Preference or need vector	Dimension-specific taste weights in the multidimensional model, with $\gamma_d \geq 0$.
$\bar{M}$	Outside opportunity value	Exogenous value of exit, the next product, or the best outside opportunity in the scalar within-product theory.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$c^W(\ell)$	Marginal within-product inspection cost	Cost of opening the ℓ-th raw signal or review in the scalar benchmark.
κ^W, λ^W	Cost parameters	Intercept and slope in $c^W(m+1) = \kappa^W + \lambda^W m$.
$\mathbb{P}(\cdot), \mathbb{E}[\cdot]$	Probability and expectation	Probabilities and expectations under the consumer's Bayesian model.
$(x)_+$	Positive part	$\max\{x, 0\}$, used to write kinked stopping payoffs.
Scalar benchmark primitives		
θ	Scalar latent quality	Binary product quality, $\theta \in \{0, 1\}$.
μ_0	Prior quality belief	$\mathbb{P}(\theta = 1 \mid \mathbf{x}, P)$.
$\mathcal{R}$	Scalar signal pool	Finite raw-signal pool $\{r_1, \ldots, r_N\}$.
r_k	Scalar raw signal	Binary review signal; 1 is high-quality-indicating and 0 is low-quality-indicating.
ρ	Oriented scalar signal precision	Probability that an authentic scalar signal agrees with θ, with $\rho \in (1/2, 1)$.
a	Scalar AI overview	Free binary summary of the finite scalar signal pool.
ϕ	Scalar aggregation rule	Map from the scalar signal pool into the overview a.
τ, ι_N	Tie-breaking objects	τ is the probability that an exact tie is favorable; $\iota_N = \mathbf{1}\{N \text{ is even}\}$.
$\zeta(\rho)$	Overview precision	Probability that the majority-rule overview is correct.
$\delta(\mu)$	Expected current-product utility	$\mathbf{x}'\beta - \alpha P + \gamma\mu$.
$S(\mu; \bar{M})$	Scalar stopping payoff	Best of exit, current product, and outside opportunity: $\max\{0, \delta(\mu), \bar{M}\}$.
$C(\bar{M})$	Scalar outside constant	$\max\{0, \bar{M}\}$; not a coverage count.
$b(\bar{M})$	Scalar stopping margin	Belief at which the current product ties the best non-product stopping payoff.
$\bar{\mu}$	Outside-option belief equivalent	Value satisfying $\bar{M} = \mathbf{x}'\beta - \alpha P + \gamma\bar{\mu}$ in threshold examples.
$\bar{\mu}^*$	Outside-option threshold	Cutoff outside-option belief at which one-step verification begins in the $N = 3$ example.
Scalar learning and dynamic objects		

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$\widetilde{\mu}(m, y)$	No-summary posterior	Posterior after inspecting m scalar signals and observing y positives.
$\widetilde{\pi}^+(m, y)$	No-summary predictive probability	Probability that the next scalar signal is positive after history (m, y).
$H_a^\tau(m, y; p)$	Scalar consistency probability	Probability that the uninspected scalar signals can complete the pool in a way consistent with overview cell a.
$L_\theta(a, m, y)$	Scalar likelihoods	Likelihood of overview and inspected history under high and low quality, with $\theta \in \{0, 1\}$.
$\widetilde{\mu}^{AI}(a, m, y)$	AI posterior	Posterior after overview a, m inspected signals, and y positives.
$\psi_a(m, y; p)$	Latent-state next-signal probability	Probability of a positive next scalar signal conditional on overview consistency and Bernoulli probability p.
$\pi^+(a, m, y)$	AI predictive probability	Unconditional probability that the next scalar signal is positive at AI state (a, m, y).
λ^ρ	Raw-signal log-likelihood ratio	$\log(\rho/(1-\rho))$, used to measure raw-review belief movements.
$\ell^C(a, m, y)$	Compression-consistent log odds	Log posterior odds after summary a and opened scalar history (m, y).
$\mathcal{G}_a(m, y)$	Remaining summary content	Log likelihood ratio in the residual consistency term, $\log[H_a^\tau(m, y; \rho)/H_a^\tau(m, y; 1-\rho)]$.
$\widetilde{W}(m, y; \bar{M})$	No-summary value function	Scalar Bellman value without a free summary.
$W(a, m, y; \bar{M})$	AI value function	Scalar Bellman value after summary a and inspected history (m, y).
$\widetilde{\mathcal{K}}(m, y; \bar{M})$	No-summary continuation payoff	Expected payoff from opening one more scalar signal before the max operator.
$\mathcal{K}(a, m, y; \bar{M})$	AI continuation payoff	Expected payoff from opening one more scalar signal in the AI environment before the max operator.
$\Delta^A(a, m, y; \bar{M})$	Dynamic gross value of another signal	Gross continuation gain at scalar AI state (a, m, y).
μ^a	Post-summary posterior	Shorthand for $\widetilde{\mu}^{AI}(a, 0, 0)$.
μ^{a+}, μ^{a-}	One-step successor posteriors	Posteriors after summary a and one favorable or unfavorable inspected signal.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
π^a	One-step favorable probability	Shorthand for $\pi^+(a,0,0)$.
$\mathcal{I}^a(\bar{M})$	Primitive one-step value	Gross value of opening one scalar signal after summary a and then stopping.
$\Phi_N(\rho;b,\mu_0)$	Finite-N one-step trigger	Belief-unit value of one post-summary inspection after a favorable overview, holding the prior fixed.
$\mathcal{Y}(a,m,y)$	Reachable terminal counts	Terminal positive counts still possible after scalar AI state (a,m,y).
Y^-,Y^+	Extreme reachable counts	Smallest and largest terminal positive counts in $\mathcal{Y}(a,m,y)$.
$\mu^{\min},\mu^{\max}$	Reachable terminal belief interval	Minimum and maximum terminal posteriors attainable from a scalar state.
$\mathcal{X}(a,m,y)$	Reachable next scalar signals	Feasible next-signal realizations $x\in\{0,1\}$ from state (a,m,y).
$\mathcal{I}^1(a,m,y;\bar{M})$	Local one-step value	Gross one-step value from a generic reachable scalar AI state.
$\mathcal{C}^1,\mathcal{C}^\infty$	One-step and dynamic continuation regions	Scalar states where one-step or unrestricted dynamic continuation is optimal.
$\mathcal{B}^R$	Scalar reachable-margin region	States where the reachable belief interval straddles the scalar stopping margin.
$d^a(\rho)$	Summary-disconfirming probability	Probability of a raw signal contradicting summary realization a.
$\bar{m},\underline{c}$	Large-pool proof constants	Fixed post-summary inspection horizon and positive lower bound on marginal cost.
Aggregation-rule extensions		
τ,τ_d	Tie-breaking probabilities	Scalar and dimension-level probabilities of resolving an exact tie in favor of a positive summary.
Y	Scalar positive total	Total number of positive scalar signals in the finite pool.
$\mathcal{A}^3$	Three-cell output space	Set $\{-,\mathtt{m},+\}$, where $\mathtt{m}$ denotes a mixed summary cell.
T^{lo},T^{hi}	Three-cell thresholds	Lower and upper thresholds defining negative, mixed, and positive summary cells.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
C_a	Scalar count cell	Set of total positive counts mapped into three-cell summary output a.
$H_a^3(m, y; p)$	Three-cell consistency probability	Binomial interval probability replacing H_a under a three-cell rule.
$\omega \in \Omega$	General scalar summary output	Finite output of a known scalar summary technology.
$Q_\phi(\omega \mid R, z)$	Scalar summary kernel	Probability that signal pool R and covariates z generate output ω.
$\mathcal{L}_\theta^\phi$	Scalar integrated likelihood	Likelihood of summary output and opened scalar history after integrating over unobserved signal pools.
$\mu^\phi(\omega, h_m)$	General-rule scalar posterior	Posterior induced by a known finite-output scalar summary kernel.
Multidimensional primitives and summary technology		
$\theta = (\theta_1, \ldots, \theta_D)$	Latent quality vector	Binary product quality by dimension.
θ_d	Dimension-level quality	Latent binary quality of the fixed product on dimension d.
$u(\theta)$	Realized utility index	$\mathbf{x}'\beta - \alpha P + \sum_d \gamma_d \theta_d$.
$\mu_{d,0}$	Dimension-level prior	$\mathbb{P}(\theta_d = 1 \mid \mathbf{x}, P)$.
$\mu_{d,t}(a, h_t)$	Dimension-level posterior	Posterior probability that $\theta_d = 1$ after summary and opened sparse reviews.
r_{nd}	Sparse review signal	Review n's signal on dimension d; equals 0, 1, or $\varnothing$.
m_{nd}	Coverage indicator	Equals one when review n covers dimension d.
R	Sparse signal matrix	Full matrix of sparse review signs.
M	Coverage matrix	Matrix of coverage indicators m_{nd}.
ρ_d	Dimension-specific signal precision	Probability that a covered sign on dimension d agrees with θ_d.
$G(M)$	Coverage prior	Consumer prior over feasible coverage matrices.
K	Summary display capacity	Maximum number of dimensions displayed in the finite-slot summary.
$C_d(M)$	Prevalence count	Number of reviews covering dimension d under coverage matrix M.
λ_d	Limiting coverage share	Large-pool limit of $C_d(M)/N$.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$\mathcal{D}^+(M)$	Positive-prevalence dimensions	Dimensions appearing at least once in the finite review pool.
$L(M)$	Effective number displayed	$\min\{K, \lvert\mathcal{D}^+(M)\rvert\}$.
$\mathcal{S}^K(M)$	Top-K summary set	Displayed dimensions with the largest positive prevalence counts.
$S^*(\lambda)$	Limiting top-K set	Dimensions with the K largest limiting positive coverage shares.
O_d	Omission event	Event that dimension d is not selected into $\mathcal{S}^K(M)$.
a_d	Displayed dimension summary	Summary sign or sentiment cell for dimension d.
a	Multidimensional summary vector	Collection of displayed cells $(a_d)_{d \in \mathcal{S}^K}$.
$\omega = (\mathcal{S}^\phi, a)$	Full summary output	Displayed set together with the displayed sentiment cells.
z	Public product/corpus covariates	Covariates available to the summary rule and, in the recovered-rule interpretation, to the researcher.
g_ϕ	Deterministic recovered rule	Map from realized sparse corpus and covariates into summary output.
$Q_\phi(\omega \mid R, M, z)$	Multidimensional output kernel	Known finite-output summary technology; includes deterministic and randomized tie-breaking rules.
Multidimensional beliefs, continuation, and coverage		
h_t	Opened-review history	Records which t sparse reviews were opened and their observed vectors.
$s_t = (a, h_t)$	Within-option state	Post-summary multidimensional state used in the Bellman problem.
$\mathcal{R}(M; h_t)$	Feasible sparse matrices	Full sparse signal matrices consistent with M and opened history h_t.
$\mathcal{L}^\phi$	Multidimensional integrated likelihood	Likelihood of summary output and opened history after summing over unobserved sparse signal realizations.
$\delta(a, h_t)$	Expected current-product utility	Utility from choosing the fixed product at the current multidimensional posterior.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$S(s_t;\gamma)$	Multidimensional stopping payoff	Best of exit, choosing the current product, and taking the outside opportunity.
$V(s_t)$	Multidimensional value function	Exact within-option Bellman value at state s_t.
$\bar{c}^{MD}(s_t;\gamma)$	Dynamic initiation cutoff	Gross dynamic value of starting one more review at state s_t.
$\eta_{d,t}(a,h_t)$	Topic-arrival probability	Posterior probability that the next unopened review covers dimension d.
$c_{d,t}$	Opened coverage count	Number of opened reviews up to t that covered dimension d.
$y_{d,t}$	Opened favorable count	Number of favorable opened signs on dimension d.
$\pi^+_{d,t}(a,h_t)$	Conditional favorable-sign probability	Probability a covered next signal on d is favorable.
$\psi^+_{d,t}(a,h_t)$	Unconditional favorable topic probability	Probability the next review both covers d and gives a favorable sign.
$\mathcal{X}(s_t)$	Sparse next-review support	Finite set of sparse next-review realizations possible from state s_t.
X_{t+1}	Next sparse review realization	Random next review drawn from $\mathcal{X}(s_t)$.
$\mathcal{I}^{MD}(s_t;\gamma)$	Multidimensional one-step value	Gross value of opening one more sparse review and then stopping.
S^0_d, S^+_d, S^-_d	Focal-dimension conditional payoffs	Expected stopping payoffs when the next review does not cover d, covers d favorably, or covers d unfavorably.
$\underline{\mathcal{K}}_d(s_t)$	Focal one-step continuation payoff	Payoff from the feasible policy that opens one review for possible information about dimension d.
$\mathrm{Cove}(\gamma;K)$	Coverage share	Share of the consumer's taste weight on displayed dimensions; named Cove to avoid covariance notation.
$\mathrm{RL}_t(\gamma;a,h_t)$	Residual learnability	Preference-weighted probability that the next review covers an omitted dimension.
Multidimensional taste regions, reachability, and design		
$\mathcal{B}_d(s_t;\gamma)$	Focal-dimension taste gain	One-step gain from the chance to learn about uncovered dimension d.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$\mathcal{T}_d^I(s_t)$	Focal inspection region	Taste vectors for which the focal one-step sufficient condition holds.
$\Delta_{d,t}^+, \Delta_{d,t}^-$	Local payoff wedges	Per-unit taste gains after favorable and unfavorable signals under local separability.
$\gamma_d^*(s_t)$	Local taste threshold	Minimum weight on dimension d needed for focal one-step inspection under local separability.
$\Omega_{d,t}(s_t)$	Per-unit taste gain	One-topic, local-separability gain from dimension d, inclusive of topic-arrival probability.
$\mathcal{T}^1(s_t)$	One-step taste region	Set of taste vectors for which one-step multidimensional inspection starts.
Γ	Normalized taste simplex	$\{\gamma \geq 0 : \sum_d \gamma_d = 1\}$.
$\Omega^{\min}, \Omega^{\max}$	Extreme per-unit gains	Minimum and maximum of $\Omega_{d,t}(s_t)$ across dimensions.
Λ_t	Decision index	Taste-weighted posterior index $\sum_d \gamma_d \mu_{d,t}$.
$\sigma(\cdot)$	Logistic link	Smooth CDF used in the soft-max approximation to the stopping payoff.
ε	Smoothing parameter	Positive parameter controlling how closely S_ε approximates the kinked stopping payoff.
$S_\varepsilon(\Lambda)$	Smooth stopping payoff	Soft-max approximation $O + \varepsilon \log(1 + \exp((\Lambda - b)/\varepsilon))$.
$\mathcal{W}_\varepsilon(\Lambda_t)$	Curvature weight	One-half of the local second derivative of the smoothed stopping payoff at Λ_t.
Σ_t^ϕ	Posterior-predictive belief covariance	Covariance matrix of μ_{t+1} under known finite rule Q_ϕ.
$G, \gamma_G, \mu_{G,t}$	Folded-group objects	Set of fine dimensions inside a folded aspect, with restricted taste and belief vectors.
$\Sigma_{G,t}^\phi$	Within-fold covariance	Posterior-predictive covariance of next beliefs on the fine dimensions inside folded group G.
$\mathcal{I}_\varepsilon^{\phi,G}$	Within-fold local value	Smoothed local verification component generated by belief movements inside folded group G.
$R_\varepsilon^\phi(s_t)$	Local-approximation remainder	Third-order Taylor remainder in the general-rule verification approximation.
O, A, b	Index-collapse constants	$O = \max\{0, \bar{M}\}$, $A = \mathbf{x}'\beta - \alpha P$, and $b = O - A$.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$\mathfrak{S}(s_t)$	Reachable future states	Finite set of future posterior states reachable from s_t.
$\Lambda_t^{\min}, \Lambda_t^{\max}$	Reachable index interval	Minimum and maximum decision-index values over $\mathfrak{S}(s_t)$.
$\mathcal{R}_t$	Coverage-reachable topic set	Dimensions that can still appear in at least one unopened review with positive posterior probability.
$C^{1,MD}, C^{\infty,MD}$	Multidimensional continuation regions	One-step and unrestricted dynamic continuation regions.
$\mathcal{B}^{R,MD}$	Multidimensional reachable-margin region	States where the reachable decision-index interval straddles b.
$\Omega^C_{d,t}, \Omega^U_{d,t}$	Covered and uncovered per-unit gains	Per-unit gains when dimension d is covered or uncovered by a candidate display set.
$\mathcal{P}(S)$	Inspection pressure	Population expected one-step inspection pressure under display set S.
$\mathcal{R}_{d,t}$	Consumer-relevance index	Taste-weighted reduction in residual verification pressure from covering dimension d.
$S_\phi(\omega), \mathcal{I}^1_\phi(\omega), \mathcal{I}^\infty_\phi(\omega)$	Output-level welfare objects	Immediate stopping payoff, gross one-review value, and unrestricted dynamic initiation gain under kernel Q_ϕ.
$\mathcal{W}(\phi; c)$	Ex ante consumer welfare	Expected unrestricted optimal within-product value under kernel Q_ϕ.
$\mathrm{Read}(\phi; c)$	Review-opening rate	Ex ante probability that the unrestricted dynamic initiation gain weakly exceeds the first-inspection cost.
F_c	Inspection-cost distribution	Cdf of independently heterogeneous first-inspection cost, with the future cost schedule held fixed.
Prev_d	Prevalence score	Positive count or share used by the prevalence ranking for dimension d.
w_d	Relevance per unit of prevalence	Ratio $\mathcal{R}_{d,t}/\mathrm{Prev}_d$ governing the sharp prevalence-ranking guarantee.
$\widehat{b}_d$	Dimension-specific omitted margin	Large-pool stopping margin measured in units of omitted dimension d's taste weight.
J_d^∞	Limiting omitted one-step gap	One-step Jensen gap from a covered signal on omitted dimension d in the slot-scarcity limit.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$\mathcal{V}_d^{slot}$	Marginal slot value	Gross value of adding omitted dimension d as an additional free summary slot in the large-pool limit.
F	Taste distribution	Population distribution of normalized taste vectors in the slot-ranking result.
$\mathbb{E}_F[\gamma_d]$	Average taste weight	Mean population weight on dimension d.
κ, ξ, F_ξ	Lab cost objects	Experimental inspection cost, additive decision shock, and its distribution.
$\mathcal{D}^\gamma$	Utility-relevant dimensions	Set $\{d : \gamma_d > 0\}$.
Recovered-rule and computational notation		
ϕ	Rule parameter vector	Parameters governing a recovered finite-output summary rule.
$Y_d(R)$	Dimension-level favorable count	Number of positive signs among reviews that cover dimension d.
$v_d(R, M)$	Net sentiment share	Scaled sentiment balance $(2Y_d - C_d)/C_d$ for positive-prevalence dimensions.
B_d^ϕ	Selection score	Score used by an empirical rule to rank dimensions for display.
$\text{Help}_d, \text{Rec}_d, X_d$	Selection covariates	Helpfulness, recency, and other dimension-level observable features.
$K_\phi(z)$	Rule-specific display cap	Number of displayed dimensions allowed by the recovered rule.
$\vartheta_{+,c}, \vartheta_{-,c}$	Category sentiment thresholds	Positive and negative threshold-band cutoffs for product category c.
q_d^ϕ	Componentwise count kernel	Probability of reporting a sentiment cell given total favorable count and coverage count.
H_{da}^ϕ	Dimension-level consistency probability	Binomial-tail or interval object for componentwise recovered rules.
$p_d(\theta_d)$	State-contingent positive-sign probability	Equals ρ_d if $\theta_d = 1$ and $1 - \rho_d$ if $\theta_d = 0$.
$\ell_{d,t}(\theta_d)$	Opened-sign likelihood factor	Likelihood contribution of opened signs on dimension d.
$\mathsf{h}_{d,t}^\phi$	Summary-cell likelihood factor	Probability that unopened signs on a displayed dimension remain consistent with its reported cell.

Continued on next page

Table 4: Notation Compendium (continued)

Symbol	Name or object	Meaning
$\chi^{+}_{d,t}$	Conditional favorable-sign probability	Probability that a covered next signal on dimension d is favorable conditional on (θ, M) and the current history.
$\varpi^{-1}(d)$	Folded subdimension set	Fine subdimensions pooled into displayed aspect d under aspect folding.

References

Agarwal, Saharsh, and Ananya Sen. 2026. "The Impact of Google AI Overviews on Publisher Traffic and User Experience: Evidence from a Field Experiment." Working paper, SSRN No. 6513059. https://ssrn.com/abstract=6513059.

Bundorf, M. Kate, Maria Polyakova, and Ming Tai-Seale. 2024. "How Do Consumers Interact with Digital Expert Advice? Experimental Evidence from Health Insurance." *Management Science* 70(11): 7617–7643.

Austen-Smith, David, and Jeffrey S. Banks. 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review* 90(1): 34–45.

Boland, Philip J. 1989. "Majority Systems and the Condorcet Jury Theorem." *Journal of the Royal Statistical Society: Series D (The Statistician)* 38(3): 181–189.

Caplin, Andrew, and Mark Dean. 2015. "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review* 105(7): 2183–2203.

Anderson, Simon P., and Régis Renault. 2006. "Advertising Content." *American Economic Review* 96(1): 93–113.

Armstrong, Mark. 2017. "Ordered Consumer Search." *Journal of the European Economic Association* 15(5): 989–1024.

Aybas, Yunus C., and Eray Turkel. 2026. "Persuasion with Coarse Communication." *Economic Theory*, advance online publication.

Benjamin, Daniel J. 2019. "Errors in Probabilistic Reasoning and Judgment Biases." In *Handbook of Behavioral Economics: Applications and Foundations 1*, Vol. 2, edited by B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, 69–186. Elsevier.

Branco, Fernando, Monic Sun, and J. Miguel Villas-Boas. 2012. "Optimal Search for Product Information." *Management Science* 58(11): 2037–2056.

Bizzotto, Jacopo, Jesper Rüdiger, and Adrien Vigier. 2020. "Testing, Disclosure and Approval." *Journal of Economic Theory* 187: 105002.

Choi, Michael, Anovia Yifan Dai, and Kyungmin Kim. 2018. "Consumer Search and Price Competition." *Econometrica* 86(4): 1257–1281.

Cheng, Lu, Ruocheng Guo, and Huan Liu. 2021. "Estimating Causal Effects of Multi-Aspect Online Reviews with Multi-Modal Proxies." *arXiv:2112.10274*.

Chahrour, Ryan. 2014. "Public Communication and Information Acquisition." *American Economic Journal: Macroeconomics* 6(3): 73–101.

Cheng, Xienan. 2023. "Deterrence, Diversion, and Encouragement: How Freely Provided Information May Distort Learning." Working paper, SSRN No. 4382692.

Dai, Weijia, Ginger Z. Jin, Jungmin Lee, and Michael Luca. 2018. "Aggregation of Consumer Ratings: An Application to Yelp.com." *Quantitative Marketing and Economics* 16(3): 289–339.

Chade, Hector, and Edward Schlee. 2002. "Another Look at the Radner-Stiglitz Nonconcavity in the Value of Information." *Journal of Economic Theory* 107(2): 421–452.

Chevalier, Judith A., and Dina Mayzlin. 2006. "The Effect of Word of Mouth on Sales: Online Book Reviews." *Journal of Marketing Research* 43(3): 345–354.

De los Santos, Babur, Ali Hortaçsu, and Matthijs Wildenbeest. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review* 102(6): 2955–2980.

Dietvorst, Berkeley J., Joseph P. Simmons, and Cade Massey. 2015. "Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err." *Journal of Experimental Psychology: General* 144(1): 114–126.

Doval, Laura. 2018. "Whether or Not to Open Pandora's Box." *Journal of Economic Theory* 175: 127–158.

Dziuda, Wioletta. 2011. "Strategic Argumentation." *Journal of Economic Theory* 146(4): 1362–1397.

Dye, Ronald A. 1985. "Disclosure of Nonproprietary Information." *Journal of Accounting Research* 23(1): 123–145.

Fang, Lu, Yanyou Chen, Chiara Farronato, Zhe Yuan, and Yitong Wang. 2024. "Platform Information Provision and Consumer Search: A Field Experiment." *NBER Working Paper* 32099.

Gavilan, Diana, Maria Avello, and Gema Martinez-Navarro. 2018. "The Influence of Online Ratings and Reviews on Hotel Booking Consideration." *Tourism Management* 66: 53–61.

Gibbard, Peter. 2022. "A Model of Search with Two Stages of Information Acquisition and Additive Learning." *Management Science* 68(2): 1212–1217.

Giovannoni, Francesco, and Daniel J. Seidmann. 2007. "Secrecy, Two-Sided Bias and the Value of Evidence." *Games and Economic Behavior* 59(2): 296–315.

Greminger, Rafael P. 2022. "Optimal Search and Discovery." *Management Science* 68(5): 3904–3924.

Gu, Chris, and Yike Wang. 2022. "Consumer Online Search with Partially Revealed Information." *Management Science* 68(6): 4215–4235.

Grossman, Sanford J. 1981. "The Informational Role of Warranties and Private Disclosure about Product Quality." *Journal of Law and Economics* 24(3): 461–483.

Hauser, John R., and Birger Wernerfelt. 1990. "An Evaluation Cost Model of Consideration Sets." *Journal of Consumer Research* 16(4): 393–408.

Honka, Elisabeth. 2014. "Quantifying Search and Switching Costs in the U.S. Auto Insurance Industry." *RAND Journal of Economics* 45(4): 847–884.

Hu, Xingbao (Simon), and Yang Yang. 2020. "Determinants of Consumers' Choices in Hotel Online Searches: A Comparison of Consideration and Booking Stages." *International Journal of Hospitality Management* 86: 102370.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101(6): 2590–2615.

Ke, T. Tony, Zuo-Jun Max Shen, and J. Miguel Villas-Boas. 2016. "Search for Information on Multiple Products." *Management Science* 62(12): 3576–3603.

Ke, T. Tony, and J. Miguel Villas-Boas. 2019. "Optimal Learning Before Choice." *Journal of Economic Theory* 180: 383–437.

Keane, Michael P., and Kenneth I. Wolpin. 1994. "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence." *Review of Economics and Statistics* 76(4): 648–672.

Kristensen, Dennis, Patrick K. Mogensen, John M. Moon, and Bertel Schjerning. 2021. "Solving Dynamic Discrete Choice Models Using Smoothing and Sieve Methods." *Journal of Econometrics* 223(2): 328–360.

Le Treust, Maël, and Tristan Tomala. 2019. "Persuasion with Limited Communication Capacity." *Journal of Economic Theory* 184: 104940.

Liao, Xiaoye, and Michal Szkup. 2026. "Coordination with Sequential Information Acquisition." *Theoretical Economics* 21(1): 132–166.

Logg, Jennifer M., Julia A. Minson, and Don A. Moore. 2019. "Algorithm Appreciation: People Prefer Algorithmic to Human Judgment." *Organizational Behavior and Human Decision Processes* 151: 90–103.

Matějka, Filip, and Alisdair McKay. 2015. "Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model." *American Economic Review* 105(1): 272–298.

Matysková, Ludmila, and Alfonso Montes. 2023. "Bayesian Persuasion with Costly Information Acquisition." *Journal of Economic Theory* 211: 105678.

Alavi, Soogand, and Salar Nozari. 2025. "AI Product Review Summaries and Content Homogenization: The Changing Landscape of UGC." Working paper, University of Iowa.

McCall, J. J. 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics* 84(1): 113–126.

Milgrom, Paul R. 1981. "Good News and Bad News: Representation Theorems and Applications." *Bell Journal of Economics* 12(2): 380–391.

Nitzan, Shmuel, and Jacob Paroush. 1982. "Optimal Decision Rules in Uncertain Dichotomous Choice Situations." *International Economic Review* 23(2): 289–297.

Olszewski, Wojciech, and Richard Weber. 2015. "A More General Pandora Rule?" *Journal of Economic Theory* 160: 429–437.

Saeedi, Maryam, and Ali Shourideh. 2026. "Optimal Rating Design under Moral Hazard." *arXiv:2008.09529*.

Seidmann, Daniel J., and Eyal Winter. 1997. "Strategic Information Transmission with Verifiable Messages." *Econometrica* 65(1): 163–169.

Shin, Hyun Song. 2003. "Disclosures and Asset Returns." *Econometrica* 71(1): 105–133.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy* 69(3): 213–225.

Titova, Maria, and Kun Zhang. 2025. "Persuasion with Verifiable Information." *Journal of Economic Theory* 230: 106102.

Ursu, Raluca M. 2018. "The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions." *Marketing Science* 37(4): 530–552.

Ursu, Raluca M., Stephan Seiler, and Elisabeth Honka. 2025. "The Sequential Search Model: A Framework for Empirical Research." *Quantitative Marketing and Economics* 23(1): 165–213.

Vana, Prasad, and Anja Lambrecht. 2021. "The Effect of Individual Online Reviews on Purchase Likelihood." *Marketing Science* 40(4): 708–730.

Wei, Dong. 2021. "Persuasion under Costly Learning." *Journal of Mathematical Economics* 94: 102451.

Yoon, Young-Ro. 2015. "Strategic Behavior in Acquiring and Revealing Costly Private Information." *International Review of Economics & Finance* 39: 133–148.

Weitzman, Martin L. 1979. "Optimal Search for the Best Alternative." *Econometrica* 47(3): 641–654.

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics* 24(2): 265–272.

Hopenhayn, Hugo, and Maryam Saeedi. 2026. "Optimal Simple Ratings." *Journal of Industrial Economics*, advance online publication.

Diamond, Peter A. 1971. "A Model of Price Adjustment." *Journal of Economic Theory* 3(2): 156–168.

Nelson, Phillip. 1970. "Information and Consumer Behavior." *Journal of Political Economy* 78(2): 311–329.

Nelson, Phillip. 1974. "Advertising as Information." *Journal of Political Economy* 82(4): 729–754.

Roberts, John H., and James M. Lattin. 1991. "Development and Testing of a Model of Consideration Set Composition." *Journal of Marketing Research* 28(4): 429–440.

Stahl, Dale O. 1989. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review* 79(4): 700–712.

Wolinsky, Asher. 1986. "True Monopolistic Competition as a Result of Imperfect Information." *Quarterly Journal of Economics* 101(3): 493–511.