Free Information, Costly Search: AI Summaries, Consideration, and Verification

Dheer Avashia

June 18, 2026

Latest version of the paper can be found here.

Abstract

This paper develops a demand-side model of search in which entering an option (product) reveals a free AI-generated compression of binary signals (reviews) and deeper individual inspection of those signals is costly. I embed a compression-based information structure into a dynamic search environment with both across-option consideration and within-option verification. I show analytically, for finite signal-set sizes, that better underlying reviews can lead to *less* review-reading when AI summaries are present: the primitive one-step value of verification vanishes as raw-signal precision approaches complete certainty, the favorable-overview continuation region exits the feasible belief space, and under standard nondegenerate majority rules the set of states in which costly inspection is optimal shrinks near perfect precision. The mechanism is that AI summaries compress the same evidence consumers would unpack at a cost, so improving raw-signal precision simultaneously strengthens the summary and devalues the underlying signals. By numerically solving the model with and without AI, I find that the costly search freed by AI-driven verification substitution can expand consideration sets, creating a reallocation from search depth to search breadth whose net effect on choice-quality outcomes is parameter-dependent. I then extend the framework to multi-dimensional quality, in which latent quality is a vector, each review is a sparse D-dimensional object with missing entries on dimensions it does not discuss, and the AI summary reports up to the K most prevalent positive-coverage dimensions in the finite review pool. The consumer therefore updates jointly over latent quality and over what the remaining unopened reviews are likely to contain. This yields two key objects—a *coverage share* measuring how much of the consumer's preference vector the summary addresses, and a *residual learnability* term measuring the posterior probability that one more unopened review will reveal utility-relevant information outside the summary. At

one endpoint, if the summary omits every dimension the consumer values, it cannot directly update utility-relevant quality beliefs, and any remaining AI effect works only through changed beliefs about what unopened reviews still discuss. At the other, if the consumer values a single covered dimension and every review covers that dimension, the problem collapses exactly to the single-dimensional benchmark. Between those endpoints, a qualified one-step ordering shows why AI reduces verification most when coverage is high and residual learnability is low, provided omitted dimensions have comparable decision-changing payoff wedges.

Contents

1 Introduction

AI-generated summaries are changing the way consumers search. On shopping, travel, local services, and digital-content platforms, consumers increasingly encounter a short synthetic overview before they inspect reviews, ratings, product descriptions, or expert commentary. The same pattern is now visible in search engines, where an AI overview appears before the user decides whether to click into underlying links and sources. In both cases, the summary is effectively a default signal: it is negligible in cost and often consumed before deeper search. At the same time, the information summarized by the AI is not new. It is a compressed representation of raw content that consumers could inspect directly, but only by dedicating more attention and incurring time costs.

This paper takes a standard consumer-search model and changes one primitive: the information structure at product entry. In canonical sequential-search models, consumers pay to sample alternatives and decide when to stop (McCall, 1970; Stigler, 1961; Weitzman, 1979). In costly product-information search models, consumers pay to learn about product quality before making a choice (Branco, Sun, and Villas-Boas, 2012; Ke, Shen, and Villas-Boas, 2016). I build on that tradition by letting entry into an option reveal a free AI-generated summary before the consumer decides whether to acquire costly underlying information. Thus the model remains a familiar dynamic search problem: the consumer chooses whether to stop, inspect more information within the current option, or move to another option. The new element is that the first signal at entry is a compressed overview of the same finite evidence pool that later inspection can unpack.

That modification matters because an AI summary is not simply an additional independent signal. If the summary is built from the same reviews that the consumer can subsequently read, then the summary and the later inspected reviews must be statistically consistent. This is the paper's core modeling departure from standard Bayesian learning with conditionally independent signals. I introduce a compression-consistent information structure—the H-function—that tracks the dependence between the observed summary and partial inspection of the same finite review pool. This object delivers exact posteriors and a closed-form one-step threshold for verification initiation: the maximum inspection cost at which a consumer who has already seen the summary still finds it worthwhile to open one more underlying review.

The economic question is how this modified search technology changes behavior. Because the AI summary compresses evidence that consumers could otherwise unpack at a cost, it shifts search across two standard margins. Within an option, the summary can substitute for costly verification: a consumer who reads a sufficiently informative overview may rationally skip the underlying reviews. Across options, the attention saved on verification can expand consideration: the consumer may enter more options because the effective cost of screening each option falls. This links the model to consideration-set approaches in which evaluation costs restrict the set of products that receive serious attention (Hauser and Wernerfelt, 1990), and to evidence that reviews and rankings affect the consideration stage of online search (Gavilan, Avello, and Martinez-Navarro, 2018; Hu and Yang, 2020; Ursu, 2018).

The within-option margin is where the analytical contribution is sharpest. In the high-precision limit, the primitive one-step value of verification vanishes for both summary realizations and the favorable-overview continuation region leaves the feasible belief space. At intermediate precision, a stronger evidence base simultaneously improves the overview and raises the value of each individual review if inspected. Those competing forces make the verification threshold hump-shaped in signal precision. I then solve the nested finite-state model numerically to show how the within-option tension interacts with across-option entry. Three regimes emerge: verification substitution, consideration expansion, and search preservation. The numerical analysis also identifies a parameter region in which the depth-to-breadth reallocation reduces choice accuracy even as it raises ex ante welfare.

This framing also places the paper in the literature on staged information acquisition and free information provision. The timing resembles two-stage information-acquisition models, in which consumers first receive coarse information and later decide whether to acquire more precise information (Gibbard, 2022). It also connects to work on how much information firms or platforms disclose before consumers incur search costs (Anderson and Renault, 2006). The key difference is that the first-stage AI overview is a compression of the same finite evidence set that second-stage inspection reveals. The paper therefore does not replace the traditional search framework; it embeds an AI-specific information technology inside it.

The paper then develops a multi-dimensional extension that changes the underlying economics rather than only the comparative statics. Most products have multi-attribute quality: a car that is safe and reliable may nonetheless be uncomfortable for tall drivers; a restaurant that is excellent overall may be mediocre for vegan diners; a laptop that is fast and well-built may have weak battery life. In that environment, a realistic AI summary is itself lower-dimensional: it highlights up to the K most prevalent positive-coverage dimensions in the finite review pool and leaves the rest to costly inspection. Here K is a display capacity, not a claim that the product has only K economically relevant attributes. Each underlying review is modeled as a sparse D-dimensional object, with missing entries on dimensions it does not discuss. The central question is therefore no longer just whether the summary is informative, but whether it is informative about the dimensions the consumer values and whether one more unopened review is likely to speak to the uncovered dimensions she still cares about. This yields two natural objects: a coverage share, measuring how much of the consumer's preference vector the summary addresses, and a residual learnability term, measuring the posterior probability that the next unopened review will reveal preference-weighted information outside the summary. At one endpoint, if the summary omits every dimension the consumer values, it cannot directly update utility-relevant quality beliefs, though it may still change search through beliefs about what unopened reviews discuss. At the other, if the consumer values a single covered dimension and every review covers that dimension, the problem collapses back to the single-dimensional benchmark. Between those endpoints, the paper derives a qualified one-step coverage-residual ordering and treats the unrestricted dynamic version as a disciplined comparative-static prediction.

The resulting empirical implications are stated in the language of standard search data. All else equal, and in markets where users attend to the summary before deciding whether to inspect reviews, the model predicts the largest decline in within-option click-depth for products whose summaries become most decisive. Consideration sets should expand most where the summary makes initial screening cheapest. If summary outputs and the underlying corpus are both observed, the aggregation rule itself becomes testable by comparing the dimensions and sentiment cells reported by the summary to the dimensions and sentiment counts inferred from reviews. In multi-dimensional categories, consumers whose preferences align with summarized dimensions should reduce review-reading more, while reviews discussing omitted but salient attributes should retain or gain attention. When summaries include a mixed category, verification should concentrate on mixed products or attributes because those states leave posterior beliefs closest to the region where one more raw signal can change the stopping decision. These predictions connect the theory to empirical models of web browsing, rankings, and consumer search costs (De los Santos, Hortaçsu, and Wildenbeest, 2012; Honka, 2014; Ursu, Seiler, and Honka, 2025), as well as recent work on platform information provision and digital advice (Fang et al., 2024; Bundorf, Polyakova, and Tai-Seale, 2024; Dietvorst, Simmons, and Massey, 2015; Logg, Minson, and Moore, 2019).

Part I

Unidimensional Model

2 Model

2.1 Environment

Each option $j \in \{1, \ldots, J\}$ has latent binary quality

$$\theta_j \in \{0, 1\}. \tag{2.1}$$

Before any costly search, the consumer observes the vector of option characteristics $\mathbf{x}_j$ and price or effort cost p_j. Based on this freely available information, the consumer holds prior belief

$$\mu_{j0} = \mathbb{P}(\theta_j = 1 \mid \mathbf{x}_j, p_j), \tag{2.2}$$

where the prior is allowed to differ across options. If option j is consumed and the realized quality state is θ_j, ex post utility is

$$u_j(\mathbf{x}_j, p_j, \theta_j) = \mathbf{x}_j' \beta - \alpha p_j + \gamma \theta_j. \tag{2.3}$$

The preference parameters $\beta \in \mathbb{R}^K$, $\alpha > 0$, and $\gamma > 0$ are known to the consumer and held common across consumers in the baseline model.

Because $\theta_j \in \{0, 1\}$, a posterior belief μ implies

$$\mathbb{E}[\theta_j \mid \mu] = \mu. \tag{2.4}$$

It is therefore convenient to define the consumer's *expected utility at belief* μ as

$$\delta_j(\mu) \equiv \mathbb{E}[u_j(\mathbf{x}_j, p_j, \theta_j) \mid \mu] = \underbrace{\mathbf{x}_j' \beta}_{\text{observable utility}} - \underbrace{\alpha p_j}_{\text{price disutility}} + \underbrace{\gamma \mu}_{\text{expected latent-quality payoff}}. \tag{2.5}$$

In particular, before the consumer observes any AI or raw search signals, ex ante expected utility is

$$\delta_j(\mu_{j0}) = \mathbf{x}_j' \beta - \alpha p_j + \gamma \mu_{j0}. \tag{2.6}$$

Without loss of generality, $\theta_j = 1$ can be interpreted as the high-quality state and $\theta_j = 0$ as the low-quality state. The key simplifying assumption is that price and observable characteristics are freely available, while latent quality is learned through AI summaries and costly underlying search.[1]

Each option j is associated with an underlying signal set

$$\mathcal{R}_j = \{r_{j1}, \ldots, r_{jN_j}\},$$

so products are allowed to differ in the number of underlying signals they generate through the product-specific signal count N_j, and signals follow

$$r_{jk} \in \{0,1\}, \qquad \mathbb{P}(r_{jk} = 1 \mid \theta_j = 1) = \mathbb{P}(r_{jk} = 0 \mid \theta_j = 0) = \rho_j, \qquad \rho_j \in \left(\frac{1}{2}, 1\right).^{2} \tag{2.7}$$

Equivalently, the binary raw signal matches the true latent-quality state with probability $\mathbb{P}(r_{jk} = \theta_j) = \rho_j$. Conditional on θ_j, the signals are i.i.d draws from the signal set. Throughout the paper, a "positive" raw signal means a realization indicating high quality, i.e. $r_{jk} = 1$, while a "negative" raw signal means a realization indicating low quality, i.e. $r_{jk} = 0$. After seeing an AI overview, the consumer may pay to inspect these human signals one by one. Inspection order is independent of signal realizations; equivalently, each inspection draws uniformly from the unopened signals in the finite pool. The marginal cost of unpacking those signals is specified below through the within-option cost function $c_j^W(m+1)$.

2.2 Stage structure: consideration, verification, and choice

The consumer's problem involves distinct choices over across-option consideration, within-option verification, and final selection. While these stages are highly interdependent in the full recursive problem, conceptualizing them separately clarifies exactly how AI alters the search experience.

The timing for search around a single option is:

[1] One can think of $\mathbf{x}_j$ and p_j as visible on platform UI. Treating $\{\mathbf{x}_j, p_j\}_{j=1}^J$ as known at the initial node is equivalent to assuming they can be searched for free before any costly learning about latent quality begins. If $\mathbf{x}_j$ and p_j are already known, there is no need to search for them. Conversely, if they can be inspected at zero cost and enter utility directly without changing the later information structure, then inspecting them is weakly dominant, so they can be normalized into the initial information set and priors can be written as $\mu_{j0} = \mathbb{P}(\theta_j = 1 \mid \mathbf{x}_j, p_j)$ without changing the continuation problem.

[2] The restriction $\rho_j > \frac{1}{2}$ is without loss of generality once signal orientation is fixed. Because the AI overview later aggregates the binary raw signals directly, the coding must first be oriented so that $r_{jk} = 1$ is the high-quality-indicating realization and $r_{jk} = 0$ is the low-quality-indicating realization. If an original binary signal were instead negatively correlated with quality, so that $\mathbb{P}(r_{jk} = \theta_j) = \rho_j < \frac{1}{2}$, one would relabel it by defining $r'_{jk} = 1 - r_{jk}$. The relabeled signal then satisfies $\mathbb{P}(r'_{jk} = \theta_j) = 1 - \rho_j > \frac{1}{2}$, so the same informational environment can be represented using positively oriented signals only. The benchmark maintains that this orientation is known.

1. In the consideration stage, the consumer pays the across-option cost of entry to bring option j into consideration.

2. In the AI environment, entry immediately reveals the free summary a_j; in the no-summary benchmark, no such overview arrives.

3. In the verification stage, the consumer either stops using the information currently held or pays $c_j^W(1)$ to inspect the first underlying human signal.

4. After each inspected human signal, the consumer updates beliefs and either stops, inspects another signal from the same underlying set at cost $c_j^W(m+1)$, or leaves option j and pays the next across-option cost to enter another option.

Summaries are therefore free conditional on entry, but entry itself is costly. The AI environment differs from the no-summary benchmark only in what happens immediately after entry: with AI, the consumer begins the verification problem after first observing a compressed summary of the underlying evidence set.

AI environment: observe $(\mathbf{x}_j, p_j) \to$ pay $c^A(1)$ to enter option $j \to$ observe free summary $a_j \to$ choose between stop, inspect one raw signal at cost $c_j^W(1)$, or switch to another option

No-AI environment: observe $(\mathbf{x}_j, p_j) \to$ pay $c^A(1)$ to enter option $j \to$ no summary arrives $\to$ choose between stop, inspect one raw signal at cost $c_j^W(1)$, or switch to another option

Figure 1: Timing comparison between the AI and no-AI environments. The environments differ only in the arrival of the free overview at entry; all later search and stopping choices are otherwise parallel.

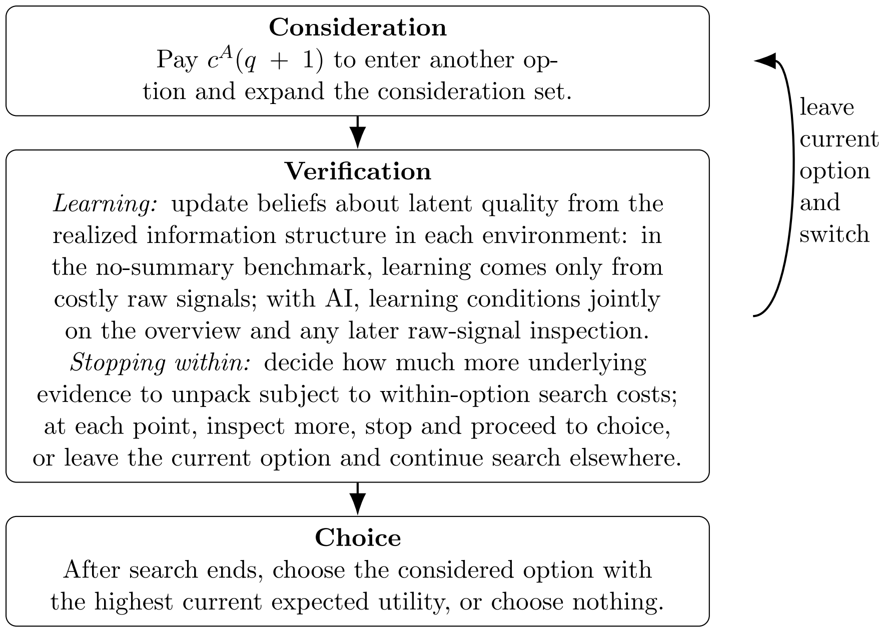

Figure 2: High-level structure of the model. Consideration determines which options enter the set, verification governs learning and within-option evidence unpacking, and choice occurs only after search stops.

2.3 Information structure

When the consumer enters option j, the platform or search interface reveals a free AI summary $a_j \in \{0,1\}$. The overview is a product-dependent aggregation rule

$$\phi_j : \{0,1\}^{N_j} \to \{0,1\},$$

so the overview is a deterministic aggregation of the full within-option signal set:

$$a_j = \phi_j(r_1, \ldots, r_{N_j}), \tag{2.8}$$

In this paper, each ϕ_j is taken as exogenous and fixed. The natural benchmark is that ϕ_j is an unbiased compression of the underlying signal environment, in the sense that it is designed to summarize evidence rather than to strategically slant it. A natural extension would let ϕ_j be chosen strategically or optimally rather than taken as fixed. That would turn the present information structure into an information-design problem in the spirit of Kamenica and Gentzkow (2011), and it is also closely related to the classic question of how much information should be disclosed for free before consumers decide whether to incur further search costs (Anderson and Renault, 2006). That problem is best treated separately from the present demand-side benchmark.

For concreteness, the benchmark overview is a majority-rule summary. In that benchmark, the aggregation family is the same across products,[3] and majority is applied to the oriented binary signals defined above. If the underlying signals were systematically misoriented and the platform aggregated them without correcting orientation, majority rule would amplify that mistake rather than improve accuracy; that non-benchmark case is excluded by the maintained assumption of known signal orientation. If N_j is odd,[4] a canonical formulation is

$$a_j = \mathbf{1}\left\{ \sum_{k=1}^{N_j} r_k \geq \frac{N_j + 1}{2} \right\}. \tag{2.9}$$

Under the assumptions used here–conditionally independent binary signals of equal precision, known orientation, and symmetric binary classification objectives–simple majority is the canonical neutral aggregation rule; richer environments with heterogeneous signal quality would instead imply weighted or threshold-shifted rules (Nitzan and Paroush, 1982). The summary should therefore be interpreted as a neutral, exogenous, costless aggregation of human signals that the consumer could otherwise process only by reading dispersed content.

However, what is not exogenous is the overview's informativeness, which is implied by the precision of the underlying human signals and the aggregation rule. Under majority aggregation, the induced overview precision is

$$\zeta_j = \mathbb{P}(a_j = \theta_j \mid \theta_j) = \underbrace{\sum_{m=(N_j+1)/2}^{N_j} \binom{N_j}{m} \rho_j^m (1-\rho_j)^{N_j - m}}_{\text{probability that a majority of underlying human signals are correct}}, \tag{2.10}$$

A step-by-step derivation appears in Appendix A.4. The majority rule is symmetric, so $\mathbb{P}(a_j = 1 \mid \theta_j = 1) = \mathbb{P}(a_j = 0 \mid \theta_j = 0) = \zeta_j$. The induced overview precision exceeds ρ_j whenever $\rho_j > \frac{1}{2}$ and $N_j > 1$; this is the standard Condorcet jury theorem result (Austen-Smith and Banks, 1996; Boland, 1989). The consumer is assumed to know the aggregation rule and therefore knows the true overview precision. The model also assumes that the consumer knows the size of the summarized signal set, N_j, and that later inspection unpacks

[3]Product-specific notation ϕ_j is retained only to allow for the possibility of heterogeneous rules outside the benchmark in future extensions.

[4]The odd-N_j restriction avoids tie-breaking. If N_j is even, the same derivation goes through after specifying a rule for ties. Under random tie-breaking, one adds one half of the probability mass at $m = N_j/2$. If the tie-breaking rule were instead chosen strategically by a seller or platform, tied evidence states would create a small but nontrivial information-design margin: for example, breaking ties toward the product in ambiguous states would tilt the overview toward acceptance and could lower buyer welfare by reducing verification too aggressively. That strategic tie-breaking problem is conceptually interesting, but it is best treated as part of the broader endogenous-ϕ_j agenda rather than folded into the benchmark model here.

that set.[5,6] Because ζ_j is increasing in both ρ_j and N_j (with the even-N_j case qualitatively similar once tie-breaking is fixed; see Appendix A.1), changes in the precision or abundance of underlying signals change the informativeness of the free overview and therefore the incentive for costly verification. The resulting comparative statics are central to the paper's analysis of when AI substitutes for within-option search.

2.4 Learning

This subsection isolates how consumers learn about latent quality in the no-AI and AI environments. The environments differ in what information arrives first and therefore in how the posterior is updated, but both use the same underlying human signal set and the same Bayesian updating logic. In both environments, the updating formulas also use the assumption that, conditional on latent quality, the underlying human signals are independent of observed characteristics and price. Thus observables enter learning through the prior $\mu_{j0} = \mathbb{P}(\theta_j = 1 \mid \mathbf{x}_j, p_j)$, while the signal likelihood depends only on θ_j. The next two subsections define the corresponding posteriors and favorable signal probabilities.

2.4.1 Learning without AI

The natural counterfactual removes the free AI summary altogether. In that environment, the consumer enters option j with prior μ_{j0} and learns only through costly inspection of the underlying human signals one by one.[7] Let $\widetilde{\mu}_j(m, y)$ denote the posterior after inspecting m signals and observing y positive (high-quality-indicating) realizations:

$$\widetilde{\mu}_j(m,y) = \frac{\overbrace{\mu_{j0}\rho_j^y(1-\rho_j)^{m-y}}^{\text{prior weight times likelihood under high quality}}}{\underbrace{\mu_{j0}\rho_j^y(1-\rho_j)^{m-y} + (1-\mu_{j0})(1-\rho_j)^y\rho_j^{m-y}}_{\text{total likelihood of the observed raw-signal history}}}. \tag{2.11}$$

$$\widetilde{\pi}_j^+(m,y) = \widetilde{\mu}_j(m,y)\rho_j + (1-\widetilde{\mu}_j(m,y))\,(1-\rho_j). \tag{2.12}$$

[5]In e-commerce settings, this corresponds to the review count displayed on product pages, which is visible on most major platforms. The consumer can therefore observe the size of the evidence pool being summarized before deciding whether to unpack it.

[6]Relaxing known precision by allowing consumers to hold priors over ζ_j would introduce an additional learning problem, because the overview would then update beliefs not only about quality but also about the information environment itself.

[7]This abstracts from baseline first free review in the no-AI environment, such as a visible review snippet or the first review shown without an extra click, and instead treats all raw-signal inspection as costly. Allowing a limited free preview would likely attenuate, but not eliminate, the advantage of an AI overview that compresses many underlying signals.

Step-by-step derivations of (2.11) *and* (2.12) *appear in Appendix A.5 and Appendix A.7.* This is the no-AI predictive learning rule. The probability of seeing a positive signal is the probability the signal is correct when in the high quality environment or probability of seeing a wrong positive signal in the low quality environment. The no-AI benchmark therefore isolates learning through direct costly inspection of the raw human signals.

2.4.2 Learning with AI

Learning in the AI environment tracks both the overview outcome and what fraction of the underlying evidence has already been unpacked. Let

$$T_j \equiv \frac{N_j + 1}{2}$$

for odd N_j, and let m denote the number of inspected signals so far, with y of them positive (high-quality-indicating). For $a \in \{0, 1\}$ and $p \in (0, 1)$,[8] define

$$H_{ja}(m, y; p) = \begin{cases} \mathbb{P}(\mathrm{Bin}(N_j - m, p) \geq T_j - y)\,, & a = 1, \\ \mathbb{P}(\mathrm{Bin}(N_j - m, p) < T_j - y)\,, & a = 0, \end{cases} \tag{2.13}$$

where, by convention, $H_{j1}(m, y; p) = 1$ and $H_{j0}(m, y; p) = 0$ whenever $y \geq T_j$ (the majority threshold is already met by the inspected signals alone), and symmetrically $H_{j1}(m, y; p) = 0$ and $H_{j0}(m, y; p) = 1$ whenever $m - y \geq N_j - T_j + 1$ (the inspected negatives already make a positive majority impossible). Histories with $H_{j1} = H_{j0} = 0$ are impossible and excluded from the reachable state space.[9] Intuitively, $H_{ja}(m, y; p)$ is a consistency probability: after observing overview a and partial inspection history (m, y), it measures how likely the remaining uninspected signals are to complete the finite signal pool in a way that still agrees with the overview.

Definition 1 (Reachable AI histories)**.** *For a fixed option j, a within-option AI history (a, m, y) is reachable if $a \in \{0, 1\}$, $0 \leq y \leq m \leq N_j$, and*

$$H_{ja}(m, y; \rho_j) > 0 \qquad or \qquad H_{ja}(m, y; 1 - \rho_j) > 0.$$

Equivalently, reachable histories are exactly those for which the overview realization and the

[8] Here p is only a generic Bernoulli success-probability argument used to write the binomial-consistency term compactly. In the actual model, the primitive raw-signal precision is ρ_j, so the relevant evaluations are $p = \rho_j$ under $\theta_j = 1$ and $p = 1 - \rho_j$ under $\theta_j = 0$.

[9] The majority threshold $T_j \equiv (N_j + 1)/2$ used here is distinct from the tie-breaking probability $\tau_j \in [0, 1]$ introduced in the appendix on even-N_j aggregation.

inspected signal counts can arise with positive probability under at least one latent-quality state.

The likelihood of overview outcome $a_j = a$ and inspection state (m, y) under $\theta_j = 1$ is

$$L_{j1}(a, m, y) = \binom{m}{y} \rho_j^y (1 - \rho_j)^{m-y} H_{ja}(m, y; \rho_j), \tag{2.14}$$

while under $\theta_j = 0$ it is

$$L_{j0}(a, m, y) = \binom{m}{y} (1 - \rho_j)^y \rho_j^{m-y} H_{ja}(m, y; 1 - \rho_j). \tag{2.15}$$

Hence the exact posterior is, for every history (a, m, y) with positive probability,

$$\widetilde{\mu}_j^{AI}(a, m, y) = \frac{\overbrace{\mu_{j0} L_{j1}(a, m, y)}^{\text{prior weight times likelihood under high quality}}}{\underbrace{\mu_{j0} L_{j1}(a, m, y) + (1 - \mu_{j0}) L_{j0}(a, m, y)}_{\text{total likelihood of the overview and inspection history}}}. \tag{2.16}$$

The posterior after seeing only the overview is the special case $\widetilde{\mu}_j^{AI}(a, 0, 0)$. In computation, the Bellman problem need only be evaluated on reachable histories. *A step-by-step derivation appears in Appendix A.6.*

To characterize continuation, define the probability that the next inspected signal is favorable conditional on state (a, m, y) and latent precision p:[10]

$$\psi_{ja}(m, y; p) = p \, \frac{H_{ja}(m + 1, y + 1; p)}{H_{ja}(m, y; p)}. \tag{2.17}$$

This ratio is defined whenever $H_{ja}(m, y; p) > 0$; the derivation follows from the conditional exchangeability argument in the preceding footnote. Intuitively, $\psi_{ja}(m, y; p)$ is the next-signal analogue of H: it is the probability that the next inspected signal is positive after conditioning on the fact that the remaining signal pool must still be compatible with the already observed overview. The unconditional probability that the next inspected signal is

[10]This ratio is well-defined whenever $H_{ja}(m, y; p) > 0$ and follows from conditional exchangeability of the underlying signals given θ_j: conditioning on the overview outcome and on y positives among m inspected signals, every remaining signal position is symmetric, so the probability that the $(m+1)$-th inspected signal is favorable equals the probability that one designated remaining signal is positive times the probability that the residual uninspected set of $N_j - m - 1$ signals remains consistent with the overview realizing at a. If $H_{ja}(m, y; p) = 0$, the latent state associated with that value of p has zero posterior probability at the history, so the corresponding conditional ratio is payoff irrelevant. Equation (2.19) below gives the unconditional transition probability in a form that is defined on every reachable history.

favorable is therefore

$$\pi_j^+(a,m,y) = \underbrace{\widetilde{\mu}_j^{AI}(a,m,y)\psi_{ja}(m,y;\rho_j)}_{\text{next signal favorable if quality is high}} + \underbrace{\left(1-\widetilde{\mu}_j^{AI}(a,m,y)\right)\psi_{ja}(m,y;1-\rho_j)}_{\text{next signal favorable if quality is low}}. \tag{2.18}$$

Equivalently, and without any division by a zero consistency probability, the same transition probability can be written directly as

$$\pi_j^+(a,m,y) = \frac{\begin{aligned}&\mu_{j0}\rho_j^y(1-\rho_j)^{m-y}\rho_j H_{ja}(m+1,y+1;\rho_j)\\ &\quad + (1-\mu_{j0})(1-\rho_j)^y\rho_j^{m-y}(1-\rho_j)H_{ja}(m+1,y+1;1-\rho_j)\end{aligned}}{\mu_{j0}\rho_j^y(1-\rho_j)^{m-y}H_{ja}(m,y;\rho_j) + (1-\mu_{j0})(1-\rho_j)^y\rho_j^{m-y}H_{ja}(m,y;1-\rho_j)}. \tag{2.19}$$

A step-by-step derivation of (2.17) *and* (2.18) *appears in Appendix A.7.*

2.4.3 Three-cell summaries and aggregation-rule tests

The binary majority overview is the benchmark used for the analytical and numerical results below. 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. This can be modeled without changing the finite-pool logic of the paper.

Let

$$Y_j = \sum_{k=1}^{N_j} r_{jk}$$

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

$$0 \le T_j^{lo} < T_j^{hi} \le N_j, \qquad T_j^{lo} + 1 < T_j^{hi},$$

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

$$a_j(Y_j) = \begin{cases} +, & Y_j \ge T_j^{hi}, \\ \mathsf{m}, & T_j^{lo} < Y_j < T_j^{hi}, \\ -, & Y_j \le T_j^{lo}. \end{cases} \tag{2.20}$$

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

$$\mathcal{C}_{j-} = \{0, \dots, T_j^{lo}\}, \qquad \mathcal{C}_{j\mathsf{m}} = \{T_j^{lo}+1, \dots, T_j^{hi}-1\}, \qquad \mathcal{C}_{j+} = \{T_j^{hi}, \dots, N_j\}.$$

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

$$H^3_{ja}(m, y; p) = \mathbb{P}(y + B \in \mathcal{C}_{ja}), \qquad B \sim \mathrm{Bin}(N_j - m, p). \tag{2.21}$$

Writing ℓ_{ja} and u_{ja} for the lower and upper endpoints of $\mathcal{C}_{ja}$, this is

$$H^3_{ja}(m, y; p) = \sum_{b=\max\{0, \ell_{ja}-y\}}^{\min\{N_j - m, u_{ja} - y\}} \binom{N_j - m}{b} p^b (1-p)^{N_j - m - b}, \tag{2.22}$$

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 (2.20) 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_j; in the multi-dimensional model it requires an even realized coverage count $C_{jd}(M_j)$ 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 rather than from majority rule itself.

Proposition 1 (Three-cell coarsening)**.** *Replace the binary overview by the three-cell rule in* (2.20)*.*

(i) The exact posterior, transition probability, martingale property, and one-step verification cutoff are obtained from the binary model by replacing H_{ja} with H^3_{ja} 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, $T_j^{lo} + T_j^{hi} = N_j$, then

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

and therefore

$$\widetilde{\mu}_j^{AI}(\mathsf{m}, 0, 0) = \mu_{j0}.$$

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

Proof in Appendix A.8.

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 (2.20) 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 rather than display order: 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. Decisive positive and negative summaries should substitute most strongly for review reading. Mixed summaries, by contrast, leave beliefs closest to the prior under symmetric thresholds and therefore identify states in which verification is most likely to remain decision-changing. Thus the strongest state-contingent prediction is not simply that AI lowers click depth, but that clicks should migrate toward products or attributes flagged as mixed. The high-precision logic is preserved in ex ante terms: decisive cells become nearly sufficient as $\rho_j \to 1$, and the probability of the mixed cell vanishes. Conditional on the rare mixed cell, however, 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.

2.5 Search costs

The model distinguishes between within-option search costs and across-option search costs. Let m denote the number of human signals already unpacked within option j, and let q denote the number of options already entered elsewhere in the choice set before the consumer considers leaving option j for another alternative. The marginal cost of opening the $(m+1)$-th human signal within option j is

$$c_j^W(m+1), \tag{2.23}$$

where $c_j^W(\cdot)$ is weakly increasing in signals. For convenience the affine schedule

$$c_j^W(m+1) = \kappa_j^W + \lambda_j^W m, \qquad \kappa_j^W > 0,\ \lambda_j^W \geq 0. \tag{2.24}$$

is used. This captures the idea that deeper inspection becomes progressively more burdensome as attention is depleted. The weakly increasing shape is a benchmark rather than a knife-edge restriction: one could allow flatter or even non-monotone schedules in richer applications, but the affine increasing case keeps the main consideration-versus-verification tradeoff transparent.

Similarly, the marginal cost of opening one more option after the consumer has already entered q other options is

$$c^A(q+1), \tag{2.25}$$

where $c^A(\cdot)$ is weakly increasing. The convenient affine specification is

$$c^A(q+1) = \kappa^A + \lambda^A q, \qquad \kappa^A > 0,\ \lambda^A \geq 0. \tag{2.26}$$

This captures the idea that browsing additional products, links, or sources becomes more costly as the consumer opens more alternatives.

2.6 Dynamic problem

The learning subsection defined how beliefs evolve in the no-AI and AI environments. The dynamic problem turns those beliefs into behavior. Conditional on the current posterior, the consumer chooses how much more to learn within an option, when to stop, and when to leave the current option and continue search elsewhere.

2.6.1 Within-option stopping in the no-summary benchmark

Let $M_j^N(q)$ denote the reduced-form continuation value from leaving option j in the no-summary environment:

$$M_j^N(q) = \underbrace{\max_{\ell \in \mathcal{U}_j(q)} \left\{ V_\ell^N(q+1) - c^A(q+1) \right\}}_{\text{best remaining-product value net of across-option search cost}}, \tag{2.27}$$

where $\mathcal{U}_j(q)$ is the set of unvisited products available if the consumer leaves j after q previous entries. Thus $M_j^N(q)$ is the value of moving to the best remaining product before learning about it, net of the across-option entry cost. This object is reduced-form in the narrow sense that it compresses the rest of the across-option problem into a single scalar continuation value. One interpretation is as a coarse outside-option heuristic for behavioral agents who do not fully re-solve the future search tree at every step, but the same object also serves as an analytical conditioning device: fixing the continuation value at any level isolates the within-option margin while preserving the interpretation that the outside action is moving to the best remaining product.

Let $\widetilde{W}_j(m, y; M_j^N(q))$ denote the value of being at option j in the no-summary environment after inspecting m raw signals and observing y positives. The terminal condition is

$$\widetilde{W}_j(N_j, y; M_j^N(q)) = \max\{0, \delta_j(\widetilde{\mu}_j(N_j, y)), M_j^N(q)\}, \tag{2.28}$$

and for $m < N_j$,

$$\widetilde{W}_j(m, y; M_j^N(q)) = \max\Big\{0,\ \delta_j(\widetilde{\mu}_j(m, y)),\ M_j^N(q),\ \widetilde{\mathcal{K}}_j(m, y; M_j^N(q))\Big\}, \tag{2.29}$$

where

$$\begin{aligned}\widetilde{\mathcal{K}}_j(m, y; M_j^N(q)) = \quad & \underbrace{-c_j^W(m+1)}_{\text{pay current within-option cost}} \\ & + \underbrace{\widetilde{\pi}_j^+(m, y)\widetilde{W}_j(m+1, y+1; M_j^N(q))}_{\text{continuation value after a positive next signal}} \\ & + \underbrace{(1 - \widetilde{\pi}_j^+(m, y))\widetilde{W}_j(m+1, y; M_j^N(q))}_{\text{continuation value after a negative next signal}}.\end{aligned} \tag{2.30}$$

In words, $\widetilde{\mathcal{K}}_j(m, y; M_j^N(q))$ is the value of paying for one more raw signal in the no-summary environment. The consumer compares the best current stopping payoff to the value of learning more from direct inspection. The ex ante value of entering option j without AI is therefore

$$V_j^N(q) = \widetilde{W}_j(0, 0; M_j^N(q)). \tag{2.31}$$

2.6.2 Within-option stopping with AI

Fix option j and let $M_j^A(q)$ denote the reduced-form continuation value from leaving option j after q other options have already been entered in the AI environment. In the reduced-form version of the across-option problem,

$$M_j^A(q) = \underbrace{\max_{\ell \in \mathcal{U}_j(q)} \left\{V_\ell^A(q+1) - c^A(q+1)\right\}}_{\text{best remaining-product value net of across-option search cost}}, \tag{2.32}$$

where $V_\ell^A(q+1)$ is the ex ante value of entering remaining product ℓ in the AI environment when the next option visit would be the $(q+1)$-th across-option search. Thus the reduced-form continuation value is explicitly the value of switching to the best remaining product, and it is decreasing in the number of alternatives already opened whenever $c^A(\cdot)$ is increasing.

Let $W_j(a, m, y; M_j^A(q))$ be the value of being at option j after observing overview outcome a, inspecting m human signals, and observing y positive signals among those m. The terminal condition is

$$W_j(a, N_j, y; M_j^A(q)) = \max\{0, \delta_j(\widetilde{\mu}_j^{AI}(a, N_j, y)), M_j^A(q)\}. \tag{2.33}$$

For $m < N_j$,

$$W_j(a, m, y; M_j^A(q)) = \max\Big\{0,\ \delta_j(\widetilde{\mu}_j^{AI}(a, m, y)),\ M_j^A(q),\ \mathcal{K}_j(a, m, y; M_j^A(q))\Big\}, \tag{2.34}$$

where

$$\begin{aligned} \mathcal{K}_j(a, m, y; M_j^A(q)) = \quad & \underbrace{-c_j^W(m+1)}_{\text{pay current within-option cost}} \\ & + \underbrace{\pi_j^+(a, m, y) W_j(a, m+1, y+1; M_j^A(q))}_{\text{continuation value after a favorable next signal}} \\ & + \underbrace{(1 - \pi_j^+(a, m, y)) W_j(a, m+1, y; M_j^A(q))}_{\text{continuation value after an unfavorable next signal}}. \end{aligned} \tag{2.35}$$

Define the value of one additional human signal as

$$\begin{aligned} \Delta_j^A(a, m, y; q) = \quad & \underbrace{\mathcal{K}_j(a, m, y; M_j^A(q)) + c_j^W(m+1)}_{\text{expected continuation value before subtracting the current signal cost}} \\ & - \underbrace{\max\{0, \delta_j(\widetilde{\mu}_j^{AI}(a, m, y)), M_j^A(q)\}}_{\text{best stopping payoff at the current state}}. \end{aligned} \tag{2.36}$$

The consumer continues verifying within option j if and only if

$$\Delta_j^A(a, m, y; q) \geq c_j^W(m+1). \tag{2.37}$$

2.6.3 Across-option continuation and the full dynamic program

Across-option continuation determines which option to enter next and therefore which options enter the consideration set, while learning about latent quality occurs only after entry through the within-option process described above. In that sense, entering an option does not reveal utility. It reveals either the opportunity to begin learning from raw signals in the no-summary benchmark or, in the AI environment, a free overview followed by the opportunity to pay for deeper verification of the underlying human evidence. Final choice is a separate terminal stage: once the consumer stops searching, she chooses the considered option with the highest current expected utility.

For some comparative statics it is useful to summarize the outside option by the reduced-form objects $M_j^A(q)$ and $M_j^N(q)$. In the AI environment, the ex ante value of entering option j after q previous entries is

$$V_j^A(q) = \mathbb{E}_{a_j}\left[W_j(a_j, 0, 0; M_j^A(q))\right], \tag{2.38}$$

with the analogous no-AI object

$$V_j^N(q) = \widetilde{W}_j(0, 0; M_j^N(q)). \tag{2.39}$$

These objects play the role of outer-layer continuation values in the spirit of Weitzman-style search, but the nested problem might not collapse to a simple scalar rule once revisits and within-option learning are allowed. They should therefore be read primarily as analytical summaries of the richer future search problem. A myopic interpretation is possible only as a secondary behavioral approximation in which consumers compare the current option to a coarse perceived outside opportunity rather than re-solving the entire future tree at every history. In particular, because consumers observe $\mathbf{x}_j$ and p_j before any costly search begins, the first option entered in the AI environment is the option with the highest ex ante entry value $V_j^A(0)$, while the first option entered in the no-summary environment is the option with the highest ex ante entry value $V_j^N(0)$. Only in the special symmetric case in which options differ solely in prior beliefs does this reduce to inspecting first the option with the highest prior.

The full nested problem makes the consideration stage explicit. Let the state at time t be

$$s_t = \left(q_t, \mathcal{V}_t, \{a_{jt}, m_{jt}, y_{jt}\}_{j \in \mathcal{V}_t}, \mathcal{U}_t\right), \tag{2.40}$$

where t indexes decision steps in the search process rather than calendar time, q_t is the number of options already entered, $\mathcal{V}_t$ is the set of visited options, (a_{jt}, m_{jt}, y_{jt}) records the overview outcome and inspection history for visited option j, and $\mathcal{U}_t$ is the set of unvisited options. At each decision step the consumer can

1. stop and choose one visited option,
2. inspect another underlying human signal for a visited option and pay $c_j^W(m_{jt} + 1)$,
3. enter an unvisited option, pay $c^A(q_t + 1)$, and observe its AI summary.

This state space is manageable in simulation for small J and provides the structural interpretation of the reduced-form outside-option representation used in the propositions below.[11]

[11]The problem in this paper is dynamic in information and search state rather than in calendar time:

The key point is that across-option consideration, within-option learning, and stopping are distinct objects. Across-option search determines where the consumer looks next. Within-option verification determines how much the consumer learns about latent quality once an option has been entered. Stopping converts those learning histories into final choice.

For any visited option $j \in \mathcal{V}_t$, define the current posterior

$$\mu_{jt} = \widetilde{\mu}_j^{AI}(a_{jt}, m_{jt}, y_{jt}), \tag{2.41}$$

and let the stopping payoff from selecting the best currently visited option, or choosing nothing, be

$$S(s_t) = \max\left\{0, \max_{j \in \mathcal{V}_t} \delta_j(\mu_{jt})\right\}. \tag{2.42}$$

If the consumer instead inspects one more underlying signal for visited option j, the continuation value is

$$\begin{aligned} K_j^W(s_t) = \quad & \underbrace{-c_j^W(m_{jt}+1)}_{\text{pay within-option inspection cost}} \\ + \quad & \underbrace{\pi_j^+(a_{jt}, m_{jt}, y_{jt}) V(s_{t+1}^{W,j,+})}_{\text{continuation value after a favorable signal}} \\ + \quad & \underbrace{\left(1 - \pi_j^+(a_{jt}, m_{jt}, y_{jt})\right) V(s_{t+1}^{W,j,-})}_{\text{continuation value after an unfavorable signal}} \end{aligned} \tag{2.43}$$

where $s_{t+1}^{W,j,+}$ and $s_{t+1}^{W,j,-}$ denote the next states after a favorable or unfavorable additional human signal for option j.

If the consumer instead enters a new option $k \in \mathcal{U}_t$, the continuation value is

$$K_k^A(s_t) = \underbrace{-c^A(q_t+1)}_{\text{pay across-option entry cost}} + \underbrace{\mathbb{E}_{a_k}\left[V(s_{t+1}^{A,k}(a_k))\right]}_{\text{expected continuation value after entering a new option}}, \tag{2.44}$$

where $s_{t+1}^{A,k}(a_k)$ is the next state after option k is entered and its AI overview is observed. Entry therefore creates a new within-option problem; it does not reveal the option's latent quality directly.

---

what evolves is the consumer's belief and search position, not the timing of consumption itself. A standard extension would introduce a discount factor multiplying continuation values without changing the basic distinction between consideration and verification, I abstract away from that though.

The full Bellman equation is

$$V(s_t) = \max\Big\{ \underbrace{S(s_t)}_{\text{stop and choose from the current consideration set}}, \underbrace{\max_{j \in \mathcal{V}_t} K_j^W(s_t)}_{\text{inspect further within a visited option}}, \underbrace{\max_{k \in \mathcal{U}_t} K_k^A(s_t)}_{\text{enter a new option and expand consideration}} \Big\}. \tag{2.45}$$

Equation (2.45) makes clear how the model nests the consideration and choice stages. Across-option entry expands the consideration set at cost $c^A(\cdot)$, within-option inspection deepens learning inside the current consideration set at cost $c_j^W(\cdot)$, and stopping converts the resulting consideration set into a final choice by selecting the visited option with the highest current expected utility.

The no-AI version of the full dynamic program is obtained by replacing each new-entry transition $s_{t+1}^{A,k}(a_k)$ with a state in which option k is entered without an overview and begins from $(m_{kt}, y_{kt}) = (0, 0)$. In that environment, the consumer still chooses across options and within options, but the first-stage costless aggregation margin is absent. The reduced-form objects $M_j^A(q)$ and $M_j^N(q)$ used in the propositions below are environment-specific summaries of the best across-option action in this fuller problem.

To obtain tractable analytical results, the next subsection collapses across-option search into these reduced-form outside values. The full nested problem remains finite and numerically solvable, but it does not generally admit a closed-form policy rule once revisits and consideration-set expansion are allowed, since objectives evolve recursively.

3 Results

3.1 Analytical Results

3.1.1 Analytically tractable special cases

The full nested model is designed for numerical solution rather than closed-form analysis. To obtain sharper analytical results, I study simpler special cases that preserve the core tradeoff between compression and verification. The natural starting point fixes a continuation value $\bar{M}_j$ from leaving the current product, interpreted as the value of entering the best remaining product net of across-option search cost:

$$\bar{M}_j(q) = \max_{\ell \in \mathcal{U}_j(q)} \left\{ V_\ell^A(q+1) - c^A(q+1) \right\}. \tag{3.1}$$

This collapses the across-option choice problem into a single outside value, so that the problem becomes a within-product stopping problem after the overview arrives. In the full heterogeneous search problem the relevant object is the realized unvisited set in the state; the notation $\mathcal{U}_j(q)$ should therefore be read as a reduced-form or order-fixed representation of that set. The three stopping actions are transparent: exit with payoff 0, choose the current product with payoff $\delta_j(\mu)$, or move to the best remaining product with payoff $\bar{M}_j(q)$. Because every proposition below conditions on the fixed $\bar{M}_j$, the results isolate the within-option verification margin while retaining a direct across-option interpretation.

This analytical special case has a reservation-value flavor similar to Weitzman-style sequential search (Weitzman, 1979), but it is not the canonical Weitzman problem. In Weitzman, opening an option reveals the payoff-relevant object for stopping and comparison. Here, entry reveals only a compressed overview, so a separate within-option learning problem remains after entry.

The simplest nondegenerate version of that problem sets $N_j = 3$ and restricts attention to the post-summary state $(a, 0, 0)$. Then the consumer observes a majority-rule overview based on three underlying signals and decides whether to pay to inspect one of those same signals. With constant within-option marginal cost, the continuation condition can be written explicitly in terms of the current posterior, the probability of a favorable next signal, and the value of the two possible next states. That formulation is especially useful for studying how the value of additional verification varies with the primitive signal precision ρ_j.

More generally, belief updating itself remains available in closed form at every reachable history: for any (a, m, y), the posterior $\widetilde{\mu}_j^{AI}(a, m, y)$ is given by (2.16), where the H-function is a finite binomial-tail probability. What does not remain in closed form once multiple inspections are allowed is the optimal continuation value, because after one more signal the consumer must again choose between stopping, continuing within the option, and leaving for the outside opportunity. The one-step restriction is analytically useful precisely because it removes that recursive continuation margin while preserving the post-summary comparison between compression and verification.

This simplified problem highlights the two channels through which ρ_j matters. A higher ρ_j makes the overview more informative, pushing the post-summary posterior toward an extreme and reducing the option value of further search. At the same time, a higher ρ_j also makes each additional underlying signal more informative once inspected. Those two forces generally work in opposite directions, so the analytically interesting object is the post-summary continuation cutoff rather than a global closed-form policy rule for the full model. In that reduced problem, one can often characterize threshold values of ρ_j above which AI fully crowds out within-option verification and below which consumers still find it optimal to unpack the summary.[12]

Proposition 2 (Primitive one-step verification threshold)**.** *Consider a single-option version of the model with constant outside option $\bar{M}$, finite $N_j > 1$, majority-rule aggregation, and at most one post-summary inspection at marginal cost κ_j^W. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$ as in Appendix A.1. Fix overview realization $a \in \{0,1\}$ and define*

$$\mu_j^a = \widetilde{\mu}_j^{AI}(a,0,0), \qquad \mu_j^{a+} = \widetilde{\mu}_j^{AI}(a,1,1), \qquad \mu_j^{a-} = \widetilde{\mu}_j^{AI}(a,1,0),$$

and

$$\pi_j^a = \pi_j^+(a,0,0).$$

Let the current stopping payoff at posterior μ be

$$S_j(\mu;\bar{M}) = \max\{0, \delta_j(\mu), \bar{M}\}. \tag{3.2}$$

The primitive gross value of one post-summary inspection is

$$\mathcal{I}_j^a(\bar{M}) = \pi_j^a S_j(\mu_j^{a+};\bar{M}) + (1-\pi_j^a) S_j(\mu_j^{a-};\bar{M}) - S_j(\mu_j^a;\bar{M}). \tag{3.3}$$

Verification is optimal after overview realization a if and only if

$$\kappa_j^W \le \mathcal{I}_j^a(\bar{M}). \tag{3.4}$$

Proof in Appendix A.8.

Proposition 2 is the most compact analytical object in the paper. It says that post-summary verification is governed by a primitive value-of-information wedge: the expected improvement in the terminal stopping payoff generated by allowing the posterior to move one more time. The case-specific thresholds below are useful because they show how this wedge collapses to simple closed forms in economically interpretable regions.

[12] The next analytical step is to reintroduce a small across-option margin, for example with $J = 2$ symmetric options. The corresponding threshold question is then whether the increase in ex ante entry value generated by the summary is large enough to overcome the across-option entry cost. The extension can help derive sufficient conditions under which AI expands consideration for low-prior consumers even though it substitutes for within-option verification after entry.

Proposition 3 (Post-summary action regions)**.** *Fix $\bar{M}_j$ as in (3.1) and fix overview realization a. Let*

$$S_j^a = \max\{0, \delta_j(\mu_j^a), \bar{M}_j\}.$$

If $\kappa_j^W > \mathcal{I}_j^a(\bar{M}_j)$, the consumer does not initiate verification and any stopping action attaining S_j^a is optimal. In particular, with an arbitrary fixed tie-breaking rule, the selected post-summary action is:

a. exit if $S_j^a = 0$;

b. choose or keep the current product if $S_j^a = \delta_j(\mu_j^a)$;

c. move to the best remaining product if $S_j^a = \bar{M}_j$.

If $\kappa_j^W \leq \mathcal{I}_j^a(\bar{M}_j)$, verification is weakly optimal before choosing among exit, the current product, and the best remaining product.

Proof in Appendix A.8.

Proposition 3 is useful for connecting the analytical cutoff to consideration. A summary does not only determine whether the consumer reads another review; it can move the consumer directly into exit, acceptance of the current product, or switching to the next product. Verification occurs only when one more raw signal has enough chance of changing which of these stopping actions is best.

Proposition 4 (Verification initiation under one-step inspection)**.** *Consider a single-option version of the model with constant outside option $\bar{M}$, finite $N_j > 1$, majority-rule aggregation, and at most one post-summary inspection at marginal cost κ_j^W. If N_j is even, fix a tie-breaking rule $\tau_j \in [0, 1]$ as in Appendix A.1. Fix overview realization $a \in \{0, 1\}$ and define*

$$\mu_j^a = \widetilde{\mu}_j^{AI}(a, 0, 0), \qquad \mu_j^{a+} = \widetilde{\mu}_j^{AI}(a, 1, 1), \qquad \mu_j^{a-} = \widetilde{\mu}_j^{AI}(a, 1, 0),$$

and

$$\pi_j^a = \pi_j^+(a, 0, 0).$$

Suppose the outside option can be written as

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

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

a. If

$$\mu_j^{a-} < \bar{\mu}_j \le \mu_j^a < \mu_j^{a+},$$

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

$$\kappa_j^W \le \gamma(1-\pi_j^a)(\bar{\mu}_j - \mu_j^{a-}).$$

b. If

$$\mu_j^{a-} < \mu_j^a \le \bar{\mu}_j < \mu_j^{a+},$$

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

$$\kappa_j^W \le \gamma\pi_j^a(\mu_j^{a+} - \bar{\mu}_j).$$

Proof in Appendix A.8.

Proposition 4 characterizes the extensive margin of within-option verification: whether the consumer starts raw-signal inspection at all. In case (a), the value of beginning verification comes from the chance that one more signal overturns provisional acceptance of the option; in case (b), it comes from the chance that one more signal overturns provisional rejection.

Example 1 (One-step threshold calculation when $N_j = 3$)**.** *To make Proposition 4(a) concrete, consider the favorable-overview $N_j = 3$ specialization of the single-option problem with constant outside option $\bar{M}$ and at most one post-summary inspection at marginal cost κ_j^W. Let*

$$\zeta(\rho_j) = 3\rho_j^2 - 2\rho_j^3 \tag{3.5}$$

and specialize the notation from Proposition 4 to the favorable-overview state $a = 1$:

$$\mu_j^1 \equiv \mu_j^a\big|_{a=1}, \qquad \mu_j^{1+} \equiv \mu_j^{a+}\big|_{a=1}, \qquad \mu_j^{1-} \equiv \mu_j^{a-}\big|_{a=1}, \qquad \pi_j^1 \equiv \pi_j^a\big|_{a=1}.$$

For readability, write

$$\mu_j^+ \equiv \mu_j^1, \qquad \mu_j^{++} \equiv \mu_j^{1+}, \qquad \mu_j^{+-} \equiv \mu_j^{1-}, \qquad \pi_j^+ \equiv \pi_j^1.$$

Then

$$\mu_j^+ = \frac{\mu_{j0}\zeta(\rho_j)}{\mu_{j0}\zeta(\rho_j) + (1-\mu_{j0})(1-\zeta(\rho_j))}, \tag{3.6}$$

$$\mu_j^{++} = \frac{\mu_{j0}\rho_j^2(2-\rho_j)}{\mu_{j0}\rho_j^2(2-\rho_j) + (1-\mu_{j0})(1-\rho_j)^2(1+\rho_j)}, \tag{3.7}$$

$$\mu_j^{+-} = \frac{\mu_{j0}\rho_j}{\mu_{j0}\rho_j + (1-\mu_{j0})(1-\rho_j)}. \tag{3.8}$$

Let

$$\pi_j^{+} = \frac{\mu_{j0}\rho_j^2(2-\rho_j) + (1-\mu_{j0})(1-\rho_j)^2(1+\rho_j)}{\mu_{j0}\zeta(\rho_j) + (1-\mu_{j0})(1-\zeta(\rho_j))} \tag{3.9}$$

denote the probability that the next inspected signal is favorable after a favorable overview. Suppose the outside option can be written as

$$\bar{M} = \mathbf{x}_j'\beta - \alpha p_j + \gamma\bar{\mu}_j \tag{3.10}$$

for some $\bar{\mu}_j$ satisfying

$$\mu_j^{+-} < \bar{\mu}_j \le \mu_j^{+} < \mu_j^{++}. \tag{3.11}$$

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

$$\kappa_j^W \le \gamma(1-\pi_j^{+})(\bar{\mu}_j - \mu_j^{+-}). \tag{3.12}$$

Equivalently, there is a threshold outside-option belief

$$\bar{\mu}_j^{*} = \mu_j^{+-} + \frac{\kappa_j^W}{\gamma(1-\pi_j^{+})} \tag{3.13}$$

such that continuation after a favorable overview is optimal if and only if $\bar{\mu}_j \ge \bar{\mu}_j^{*}$.

Derivation in Appendix A.8.

Condition (3.11) isolates the economically interesting case in which the consumer would keep the option after a favorable follow-up signal but would switch to the outside option after an unfavorable one. In that region, the option value of further verification comes entirely from the ability to overturn the favorable first impression when the inspected signal goes against it. A useful anchor in the $N_j = 3$ majority-rule case is that

$$\mu_j^{+-} = \frac{\mu_{j0}\rho_j}{\mu_{j0}\rho_j + (1-\mu_{j0})(1-\rho_j)},$$

which is exactly the posterior that would result from observing one favorable raw signal from the prior alone. In other words, a favorable overview followed by one unfavorable inspected signal collapses to a net single positive piece of evidence in this special case. Appendix A.8 records the algebra. By symmetry, an analogous expression applies after an unfavorable overview.

Remark 1 (Outside the threshold region). If $\bar{\mu}_j \le \mu_j^{+-}$, then the outside option dominates after either inspected signal, so one-step verification has no decision-changing value in this special case. If instead $\mu_j^{+} < \bar{\mu}_j < \mu_j^{++}$, then 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; this is exactly the case-(b) region in Proposition 4. Finally, if $\bar{\mu}_j \ge \mu_j^{++}$, then the outside option dominates even after a favorable inspected signal, so verification again has no decision-changing value. The nontrivial threshold worked out in Example 1 therefore arises in the intermediate region (3.11), where the inspected signal can reverse the post-summary decision after a favorable overview.

Lemma 1 (Posterior martingale under one-step inspection)**.** *Fix any finite $N_j > 1$ and any overview realization $a \in \{0,1\}$. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$ as in Appendix A.1. Using the notation of Proposition 4, posterior beliefs satisfy*

$$\mu_j^a = \pi_j^a \mu_j^{a+} + (1 - \pi_j^a)\mu_j^{a-}. \tag{3.14}$$

Proof in Appendix A.8.

Lemma 1 records the standard martingale property of Bayesian posteriors. For $N_j = 3$ with a favorable overview, equation (3.14) reduces to $\mu_j^{+} = \pi_j^{+}\mu_j^{++} + (1-\pi_j^{+})\mu_j^{+-}$ using the closed-form posteriors in Example 1.

Lemma 2 (Simple high-precision tail under odd majority)**.** *Suppose $N_j > 1$ is odd and the overview is generated by majority rule. Let*

$$g_j(\rho_j) = 1 - \pi_j^{+}(1,0,0)$$

be the probability of a disconfirming raw signal after a favorable overview. Then

$$\lim_{\rho_j \to 1} \frac{g_j(\rho_j)}{1-\rho_j} = 1. \tag{3.15}$$

Equivalently, the endpoint extension satisfies $g_j(1) = 0$ and $g_j'(1) = -1$. By symmetry,

$$\pi_j^{+}(0,0,0) \sim 1 - \rho_j$$

after an unfavorable overview.

Proof in Appendix A.8.

Corollary 1 (Sufficiently precise signals eliminate one-step verification)**.** *Fix any finite $N_j > 1$, $\mu_{j0} \in (0,1)$, $\kappa_j^W > 0$, and $\gamma > 0$. Consider the one-step post-summary inspection problem under majority-rule aggregation with constant outside option $\bar{M}$. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$ as in Appendix A.1. Define the summary-disconfirming probability*

$$d_j^a(\rho_j) = \begin{cases} 1 - \pi_j^+(1,0,0), & a = 1, \\ \pi_j^+(0,0,0), & a = 0. \end{cases} \tag{3.16}$$

Then $d_j^a(\rho_j) \to 0$ for each $a \in \{0,1\}$ as $\rho_j \to 1$, and the primitive one-step value $\mathcal{I}_j^a(\bar{M})$ in (3.3) converges to zero uniformly over feasible outside-option beliefs $\bar{\mu}_j \in [0,1]$. Hence there exists $\rho_j^* < 1$ such that for all $\rho_j > \rho_j^*$ and all feasible outside-option beliefs $\bar{\mu}_j \in [0,1]$, continuation after the overview is not optimal. In the favorable-overview case and in the case-(a) region of Proposition 4, the feasible outside-option-belief region for continuation is empty for all ρ_j sufficiently close to one; in the affine cutoff representation, this is equivalently $\bar{\mu}_j^* > 1$ eventually.

Proof in Appendix A.8.

Corollary 2 (Asymptotic verification substitution)**.** *Fix any finite $N_j > 1$, $\mu_{j0} \in (0,1)$, $\kappa_j^W > 0$, and $\gamma > 0$ in the one-step post-summary inspection problem. Suppose either N_j is odd, so Lemma 2 applies, or N_j is even with a fixed tie-breaking rule $\tau_j \in [0,1]$ for which the favorable-overview disconfirming probability $1 - \pi_j^+(1,0,0)$ has a simple zero at $\rho_j = 1$. Then for a favorable overview ($a = 1$),*

$$\frac{\partial \bar{\mu}_j^*}{\partial \rho_j} > 0 \tag{3.17}$$

for all ρ_j sufficiently close to 1. Equivalently, for any fixed outside-option belief $\bar{\mu}_j$, the set of post-summary states in which verification is optimal shrinks as raw-signal precision rises near 1. By symmetry, an analogous monotonicity holds for the lower threshold after an unfavorable overview when the corresponding favorable-signal probability $\pi_j^+(0,0,0)$ has a simple zero at $\rho_j = 1$.

Proof in Appendix A.8.

Corollary 3 (High precision eliminates finite verification)**.** *Fix any finite $N_j > 1$ and any finite inspection horizon $\bar{m}_j \leq N_j$. Suppose within-option marginal inspection costs satisfy $c_j^W(m) > 0$ for every $m \in \{1, \ldots, \bar{m}_j\}$. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$ as in Appendix A.1. Then there exists $\rho_j^* < 1$ such that for all $\rho_j > \rho_j^*$, for both overview* *realizations $a_j \in \{0,1\}$, and for all feasible outside-option beliefs $\bar{\mu}_j \in [0,1]$, the optimal action at the post-overview entry state is immediate stopping rather than further inspection. Equivalently, any outside-option cutoff for post-overview continuation lies outside the feasible unit interval; after a favorable overview, the corresponding threshold satisfies $\bar{\mu}_j^* > 1$.*

Proof in Appendix A.8.

For even N_j, the proofs differ by the presence of a τ_j-weighted exact-tie term in the overview-consistency probabilities. That term does not change the high-precision limiting behavior in Corollaries 1 and 3. The derivative statement in Corollary 2 additionally requires the stated simple-zero condition, which holds for the usual odd-N_j benchmark and for neutral even-N_j tie-breaking. Appendix A.1 gives the corresponding tie-breaking formula and Figure 8 verifies numerically that the one-step threshold geometry is qualitatively the same for $N_j = 4$ under neutral tie-breaking.

Figure 3 makes the one-step geometry concrete under the illustrative parameters $\mu_{j0} = 0.5$ and $\gamma = 1$. The left panel plots the maximal continuation value from one additional inspection, evaluated at the most verification-friendly outside-option belief, for three illustrative signal-set sizes: $N_j = 3$, $N_j = 5$, and $N_j = 7$. For each displayed N_j, this one-step value object is available in closed form through the finite-history posterior and predictive-probability formulas. All three curves are hump-shaped. As N_j rises, the peak shifts leftward and the value of one more inspected signal falls sooner at high precision because the overview itself becomes more informative. Economically, each peak is the highest inspection cost for which one-step verification can be optimal somewhere in the state space. The right panel then specializes back to $N_j = 3$ and translates that value object into a policy threshold for an illustrative inspection cost, $\kappa_j^W = 0.015$. Verification after a favorable overview is feasible only when the threshold cutoff $\bar{\mu}_j^*$ lies below the upper bound μ_j^+, so the shaded band marks the set of outside-option beliefs for which one more inspection is optimal. As ρ_j rises, that band narrows and then disappears. Appendix Figure 8 shows that the same geometry is qualitatively similar for even N_j under neutral tie-breaking.

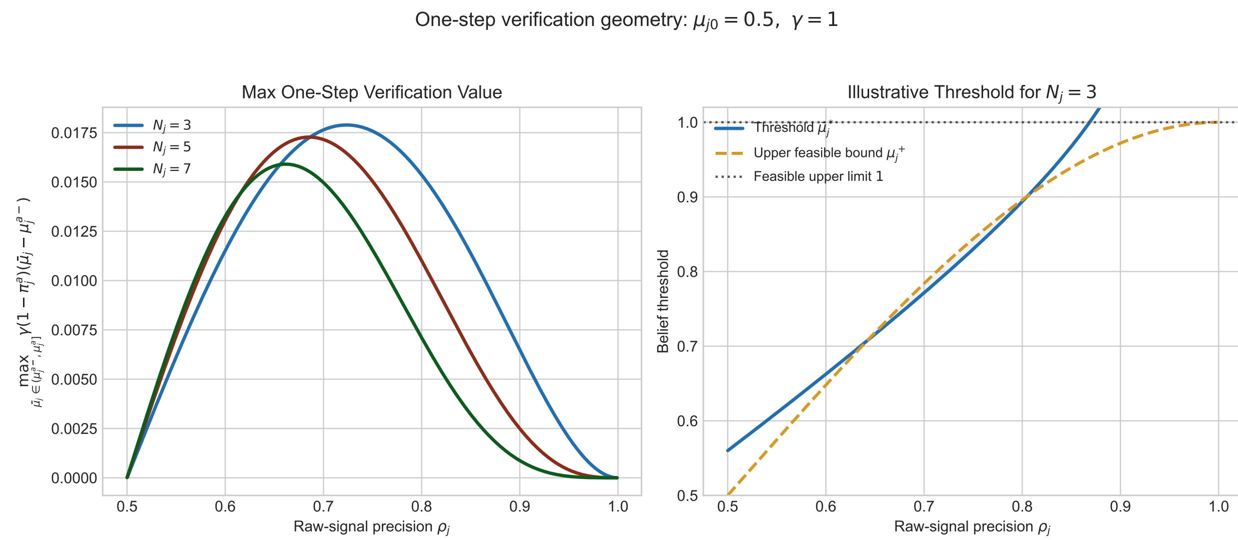

Figure 3: One-step verification geometry under $\mu_{j0} = 0.5$ and $\gamma = 1$. The left panel plots the maximal one-step verification value as a function of raw-signal precision for illustrative signal-set sizes $N_j = 3$, $N_j = 5$, and $N_j = 7$. The right panel specializes to $N_j = 3$ and translates that value object into a policy threshold for illustrative inspection cost $\kappa_j^W = 0.015$ by plotting the outside-option cutoff $\bar{\mu}_j^*$ together with the feasible upper bound μ_j^+ and the unit bound. Verification is possible only where $\bar{\mu}_j^* \leq \mu_j^+$, so the shaded band is the set of outside-option beliefs for which one more inspection is optimal. Appendix Figure 8 shows the corresponding odd-versus-even comparison under neutral tie-breaking.

3.1.2 General benchmark results

The previous subsection established closed-form results in a stripped-down special case. This section records two benchmark results that hold more generally in the rational model before turning to numerical counterfactuals. The first is a dominance result: if the AI summary is free and correctly interpreted, it cannot lower expected value relative to the no-summary environment because the consumer can always ignore it. The second is a state-specific cutoff characterization for within-option continuation. Richer comparative statics over primitives are then characterized numerically in the nested model.

Benchmark observation. Consider a common full search state, with the same visited set, unvisited set, costs, and feasible actions in the AI and no-summary environments. Let $V^A(s)$ denote the value in that state when the consumer observes a costless AI summary upon entry, and let $V^N(s)$ denote the value when no AI summary is shown. If the summary is costless and the consumer updates correctly, then

$$V^A(s) \geq V^N(s). \tag{3.18}$$

The same comparison applies to the option-entry values $V_j^A(q)$ and $V_j^N(q)$ whenever the reduced-form outside continuation value is held fixed across environments or is generated by the same full AI-enhanced search problem. The logic is simple: the AI consumer can always replicate the no-AI policy by ignoring the summary, so a free correctly interpreted signal cannot make the consumer worse off in the rational benchmark.

Proposition 5 (State-specific cutoff characterization)**.** *Fix option j, across-option state q, and state (a, m, y) in the AI environment. Holding the continuation value function $W_j(\cdot, \cdot, \cdot; M_j^A(q))$ fixed, define*

$$\bar{c}_j(a, m, y; q) = \Delta_j^A(a, m, y; q). \tag{3.19}$$

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

$$c_j^W(m+1) \leq \bar{c}_j(a, m, y; q). \tag{3.20}$$

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

Proof in Appendix A.8. This proposition isolates the local economics of verification. At any reached state, the continuation decision is governed by a one-step comparison between the cost of unpacking one more signal and the option value of allowing beliefs to move before stopping.

The benchmark observation and the cutoff proposition establish the minimal rational benchmark. A free summary weakly raises ex ante value, and continuation at any reached state is governed by a transparent local cost-benefit comparison. What those analytical results do not pin down is how summary-driven changes in within-option continuation feed back into endogenous across-option entry. In particular, the propositions above apply to a stripped-down single-option problem with a reduced-form outside value, whereas the full characterization of when AI expands consideration versus substitutes for verification comes from solving the nested model numerically.

3.2 Numerical Solutions

The full nested model is solved numerically by exact dynamic programming on a finite state space rather than by simulation or estimation. A global state records, for each option, whether it has been entered and, if so, the current within-option inspection history: (m, y) in the no-AI environment and (a, m, y) in the AI environment. At each reached state, the Bellman equation compares three action classes: stop and choose from the currently considered set, continue by opening one more underlying signal in any entered option, or enter an unvisited option by paying the next across-option cost. This makes the outside option fully endogenous in the numerical solution. Leaving the current option does not invoke an exogenous scalar continuation value; it means paying the next across-option cost and returning to the full Bellman problem on the enlarged state space. Because the state space is finite for fixed J and finite N_j, the policy and value functions can be computed exactly by backward recursion with memoization, and the reported outcome statistics are then obtained by integrating the resulting optimal policy over the latent-quality states and signal realizations implied by the parameterization. Thus "expected" outcomes in this section are exact ex ante expectations under the solved policy, not averages across Monte Carlo simulation draws.

The model is built to compare two environments:

1. **AI environment:** the consumer receives a free summary before costly search.
2. **No-AI environment:** the consumer must rely on priors and costly raw search alone.

For any parameterization, the quantitative objects of interest are

$$\begin{aligned} \mathbb{E}[\text{options entered}], &\qquad \mathbb{E}[\text{raw signals opened}], \\ \mathbb{P}(\text{stop immediately}), &\qquad \mathbb{P}(\text{choose high-quality option}), \end{aligned} \tag{3.21}$$

as well as actual expected payoff. Define the consumer's ex ante value by

$$V = \mathbb{E}[\text{actual payoff}] - \mathbb{E}[\text{within-option search cost}] - \mathbb{E}[\text{across-option search cost}]. \tag{3.22}$$

For the AI versus no-AI comparison, the main reported counterfactual objects are

$$\begin{aligned} \Delta E[\text{options}], &\qquad \Delta E[\text{signals}], \\ \Delta \mathbb{P}(\text{high quality}), &\qquad \Delta\text{payoff}, \qquad \Delta V. \end{aligned} \tag{3.23}$$

The companion tables below therefore report both levels and AI-minus-no-AI differences. They also decompose the ex ante value effect into three pieces:

$$\Delta V = \Delta\text{payoff} + \text{within-cost savings} + \text{across-cost savings}.$$

This welfare decomposition is useful because choice accuracy alone is not a welfare object: AI can improve consumer value either by improving realized choices or by economizing on costly search, and those two channels need not move together. By quantifying these distinct channels, the full Bellman solution explicitly connects the local within-option verification choice to the broader problem of consideration set formation.

Table 1 collects the symmetric benchmark parameterization used in the $J = 3$ numerical exercises.

Table 1: Benchmark Numerical Parameterization

Parameter	Value	Interpretation
J	3	Number of options in the market
N_j	5	Human signals summarized within each option
μ_{j0}	0.5	Prior probability of high quality
ρ_j	0.7	Precision of each raw human signal
$\mathbf{x}_j'\beta$	0.2	Deterministic utility from observables
p_j	1.0	Posted price or deterministic cost component
α	0.5	Price sensitivity
γ	1.0	Utility weight on latent quality
No-purchase utility	0.0	Normalized terminal payoff from choosing nothing
$c^A(q+1)$	$0.03 + 0.02q$	Across-option marginal search cost
$c_j^W(m+1)$	$0.05 + 0.02m$	Within-option marginal search cost

3.2.1 Homogeneous benchmark results

The numerical exercises are used in a comparative-static spirit: the goal is to study how shifts in economically central parameters change search behavior, the probability of choosing a high-quality option, and welfare, rather than to claim a fully estimated quantitative fit. The benchmark row in Tables 2–4 uses the parameterization in Table 1. The remaining rows vary one economically central object at a time: prior beliefs, across-option search costs, and raw-signal precision. First, AI strongly substitutes for within-option verification in the baseline calibration. Table 2 shows the levels directly: expected raw-signal openings fall to 0.0 from 1.5 in the no-AI environment, while expected options entered remain at 1.75 in both worlds. At the same time, ex ante value rises from 0.2025 to 0.3973, actual expected payoff rises from 0.35 to 0.4698, and the probability of choosing a high-quality option rises from 0.65 to 0.7323. This zero-verification outcome in the AI arm is a corner solution for the baseline calibration rather than a claim of generality; with $\rho_j = 0.7$ and $N_j = 5$, the implied overview precision is already high enough that the free summary often pushes posterior beliefs directly beyond the continuation region, so no additional raw signals are opened.

Second, when priors are weak, AI expands consideration rather than merely reducing verification. Setting $\mu_{j0} = 0.2$ and leaving the remaining parameters unchanged, the no-AI consumer stops immediately and chooses nothing, whereas the AI consumer enters 2.20

options on average and chooses a market option with probability about 0.654. The levels table is useful here because the difference $\Delta E[\text{options}] = 2.20$ is not a shift from one active search policy to another; it is a shift from zero entry in the no-AI environment to substantial screening in the AI environment.

Third, high across-option costs create a region in which AI preserves broader consideration. Raising the across-option cost schedule to $c^A(q+1) = 0.10 + 0.02q$, the no-AI consumer enters only one option on average, while the AI consumer still enters 1.75 options on average and retains the same within-option no-search pattern as in the baseline. This is the clearest numerical example of the paper's main mechanism: the analytical verification-substitution force remains active within options, but in the full nested model it also preserves entry into additional options that would otherwise be screened out by high across-option costs.

Table 2: Numerical Comparative Statics: Levels by Environment

	E[options]		E[signals]		Pr(high quality)		Actual payoff	
Case	No AI	AI	No AI	AI	No AI	AI	No AI	AI
Baseline	1.75	1.75	1.50	0.00	0.650	0.732	0.350	0.470
Low prior ($\mu_0 = 0.2$)	0.00	2.20	0.00	0.00	0.000	0.367	0.000	0.171
High across cost	1.00	1.75	0.00	0.00	0.500	0.732	0.200	0.470
Low raw-signal precision ($\rho = 0.6$)	1.00	1.75	0.00	0.00	0.500	0.637	0.200	0.337
Near-uninformative signals ($\rho = 0.5001$)	1.00	1.00	0.00	0.00	0.500	0.500	0.200	0.200

Notes: Entries report the exact finite-state Bellman solution in each environment. Reporting levels alongside differences is useful because several rows involve regime shifts—for example, the low-prior row moves from no entry at all in the no-AI environment to active screening under AI.

Table 3: Numerical Comparative Statics: AI Minus No-AI Differences

Case	ΔE[options]	ΔE[signals]	Δ Pr(high quality)	Δ payoff	Main margin
Baseline	0.00	−1.50	0.082	0.120	Verification substitution
Low prior ($\mu_0 = 0.2$)	2.20	0.00	0.367	0.171	Consideration expansion
High across cost	0.75	0.00	0.232	0.270	Search-preserving consideration
Low raw-signal precision ($\rho = 0.6$)	0.75	0.00	0.137	0.137	Consideration expansion
Near-uninformative signals ($\rho = 0.5001$)	0.00	0.00	0.000	0.000	No informational role

Notes: Entries report AI minus no-AI outcomes from the nested rational model. Under the normalized no-purchase utility, changes in actual expected payoff reflect both changes in the probability of making any purchase and changes in the probability that the chosen option is truly high quality.

Table 4: Welfare Decomposition in the Homogeneous Numerical Exercises

Case	ΔV	Δ payoff	Quality margin	Purchase margin	Within-cost savings	Across-cost savings
Baseline	0.195	0.120	0.082	0.037	0.075	0.000
Low prior ($\mu_0 = 0.2$)	0.072	0.171	0.367	-0.196	0.000	-0.100
High across cost	0.175	0.270	0.232	0.037	0.000	-0.095
Low raw-signal precision ($\rho = 0.6$)	0.094	0.137	0.137	0.000	0.000	-0.042
Near-uninformative signals ($\rho = 0.5001$)	0.000	0.000	0.000	0.000	0.000	0.000

Notes: ΔV is the change in ex ante value. The quality margin is the change in the latent-quality component of actual payoff, while the purchase margin is the deterministic utility index times the change in the purchase probability. Positive cost savings mean AI reduces expected search costs; negative entries mean AI induces additional search spending on that margin. By construction, $\Delta V = \Delta\text{payoff} + \text{within-cost savings} + \text{across-cost savings}$.

Table 5: Testable Cross-Sectional Predictions

Environment	High ρ_j	Low μ_{j0}	High c^A
AI effect on within-option search	Large negative	Small or ambiguous	Small
AI effect on options entered	Near zero	Positive	Positive
AI effect on probability of choosing a high-quality option	Positive or weakly positive	Positive	Positive
Main mechanism	Verification substitution	Consideration expansion	Search-preserving consideration

Three patterns emerge from Tables 2 and 3. First, AI can reduce total search by eliminating within-option verification while leaving consideration unchanged, as in the baseline. Second, AI can expand consideration when the no-AI consumer would otherwise stay near the outside option, as in the low-prior and high-across-cost environments. Third, the decomposition between across-option and within-option search is essential. Looking only at total clicks would miss that AI sometimes lowers search by substituting for verification and sometimes raises search by making additional options worth entering.

The raw-signal precision counterfactual is especially informative about the role of compression. The overview is most valuable when the underlying human-signal environment is itself strong, because majority aggregation then produces a highly accurate compressed signal. This is consistent with the force isolated analytically in Corollaries 1–3: as raw signals become more precise, the continuation region shrinks and verification becomes less attractive. By contrast, when ρ_j is lower, the overview is a weaker object and the model predicts less scope for AI to substitute for within-option verification. In that sense, holding other primitives fixed, stronger signal environments should exhibit larger AI-driven reductions in within-option search, while weaker signal environments should exhibit weaker substitution away from verification.

At the boundary, this logic becomes exact. In a near-uninformative robustness case with $\rho_j = 0.5001$, the AI and no-AI environments are numerically indistinguishable in the nested model: expected options entered, expected raw-signal openings, the probability of choosing a high-quality option, and expected payoff all coincide. Once the underlying human-signal environment becomes essentially pure noise, the overview ceases to play any informational role.

Table 4 makes the welfare comparison more transparent. In the baseline, AI raises ex ante value by about 0.195. Roughly 0.120 of that comes from higher actual payoff and the remaining 0.075 comes from reduced within-option search costs. Within the actual-payoff term, the probability of buying falls slightly, which mechanically raises payoff by 0.0375 because the deterministic utility index net of price is negative in the benchmark; the larger contribution is the 0.082 increase in the probability of ending with a high-quality option. In the low-prior case, by contrast, AI raises actual payoff by 0.171 but only raises ex ante value by 0.072 because the consumer now incurs about 0.10 of additional across-option search cost in order to screen options that the no-AI consumer would never open. The high-across-cost and low-precision rows have the same flavor: most of the quality gain comes from broader consideration, but part of that gain is spent on the additional across-option search needed to realize it.

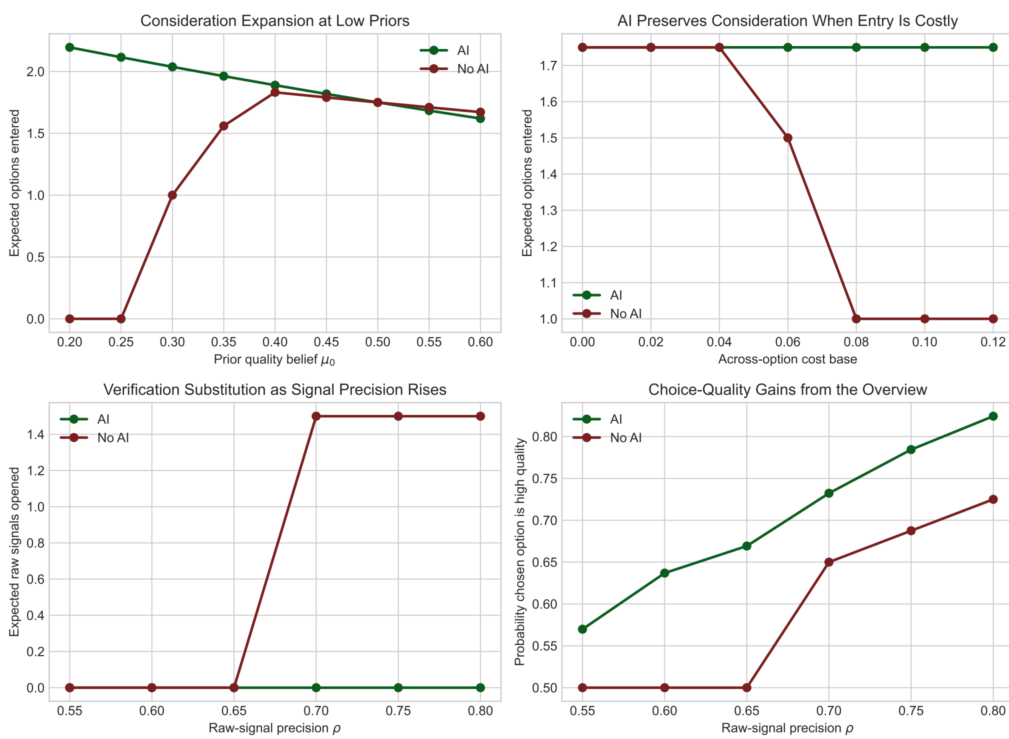

Figure 4: Preliminary numerical illustrations from the nested $J = 3$ model. The top-left panel shows that AI expands consideration when prior quality beliefs are weak. The top-right panel shows that AI preserves consideration when across-option entry becomes costly. The bottom-left panel shows that AI crowds out within-option verification as raw-signal precision rises. The bottom-right panel shows the associated gains in the probability of choosing a high-quality option.

Figure 4 reports the baseline sweeps used to motivate the main numerical regimes. Figure 5 adds two further checks that matter for interpretation. First, the top-right panel shows that the conclusion's non-monotonicity claim is not a feature of the baseline calibration. In the baseline, the AI effect on choice quality remains positive throughout the plotted ρ_j range. The non-monotonicity appears when within-option verification is made cheaper, using the alternative schedule $c_j^W(m+1) = 0.01 + 0.005m$: in that case the no-AI consumer optimally keeps reading at intermediate precisions, so the AI summary can crowd out informative depth strongly enough that the probability of choosing a high-quality option is lower under AI around $\rho_j \in [0.65, 0.70]$ before turning positive again at higher precision. Second, the bottom panels sweep the homogeneous signal count N_j. In the benchmark calibration, increasing N_j from 3 to 11 raises the AI overview precision enough that the AI-versus-no-AI choice-quality gap grows from about 0.036 to about 0.157. The search margins stay at their corner values in this calibration—AI keeps opening zero signals and no-AI keeps opening about 1.5—so the N_j sweep isolates the direct role of a larger evidence pool in strengthening the compressed overview.

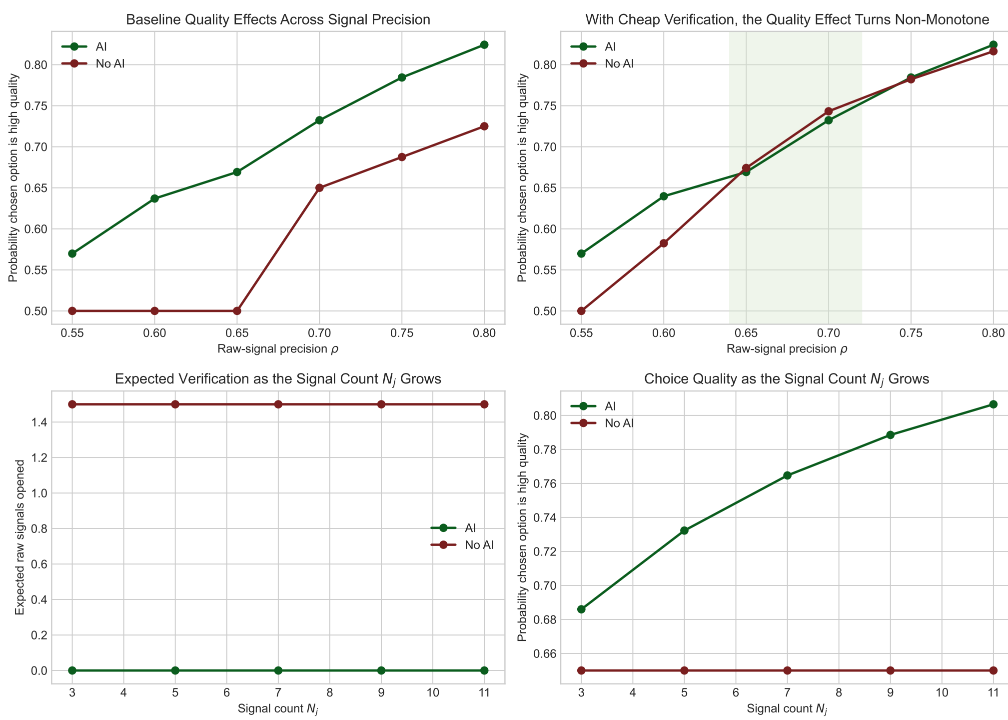

Figure 5: Additional numerical checks. The top-left panel reports the baseline quality sweep across raw-signal precision. The top-right panel repeats that sweep under a lower within-option cost schedule, $c_j^W(m+1) = 0.01 + 0.005m$, and shows that the AI effect on choice quality can become non-monotone when verification is cheap. The bottom-left panel sweeps the homogeneous signal count N_j and reports expected raw-signal openings. The bottom-right panel reports the associated probability of choosing a high-quality option.

3.2.2 Heterogeneous option results

The symmetric exercises isolate the paper's baseline mechanism. The heterogeneous results show how that same mechanism plays out when options differ in three economically distinct dimensions: ex ante attractiveness μ_{j0}, the amount of evidence available for compression N_j, and the strength of the underlying evidence ρ_j. Table 6 shows that each margin creates a selection effect in consideration. With heterogeneous priors alone, AI expands consideration onto the middle-prior option even though both environments always enter the high-prior option first. When options also differ in signal precision, AI assigns positive entry probability to a low-prior option whose underlying evidence is especially informative. When options differ in signal-set size, AI likewise screens a low-prior option backed by a large evidence pool. Economically, these cases are the across-option counterpart to the analytical cutoff logic: the free overview raises the option value of entry most for options that start with weak priors but can nonetheless generate informative compressed summaries because their underlying evidence is strong or abundant. In all three cases, the no-AI consumer enters only the highest-prior option, while AI broadens the consideration set and raises the probability of choosing a high-quality option. Ex ante value also rises in all three cases, though by less than actual payoff, because some of the quality gain is spent on the additional across-option search needed to screen the newly viable options.

The second row of Table 6 is especially informative about the selection mechanism. The low-prior option starts at $\mu_{j0} = 0.25$, so without AI it is never worth entering. But its raw signals are highly precise, with $\rho_j = 0.85$, so with $N_j = 5$ the implied overview precision is about 0.973. A favorable overview therefore moves the posterior on that option all the way to roughly 0.924, while an unfavorable overview drives it down to about 0.009. The AI consumer consequently treats the option as a useful screen: it is entered with probability 0.204 and chosen with probability 0.054, even though the no-AI consumer never enters it. The key point is that the summary does not rescue the option by changing its prior. It rescues it by revealing, at low cost, that this low-prior option sits on an unusually informative evidence base. That is the heterogeneous-option analogue of consideration expansion in the symmetric model.

Table 6: Heterogeneous-Option Numerical Extension

Case	Option profiles (μ_{j0}, ρ_j, N_j)	ΔE[options]	ΔE[signals]	Δ Pr(high quality)	Δ payoff	ΔV	AI option-entry probabilities
Heterogeneous priors	$(0.20, 0.70, 5), (0.50, 0.70, 5), (0.80, 0.70, 5)$	0.298	0.000	0.059	0.059	0.045	(0.000, 0.298, 1.000)
Heterogeneous priors and signal precision	$(0.25, 0.85, 5), (0.50, 0.70, 5), (0.75, 0.60, 5)$	0.613	0.000	0.070	0.070	0.036	(0.204, 0.409, 1.000)
Heterogeneous priors and signal counts	$(0.25, 0.70, 11), (0.50, 0.70, 5), (0.75, 0.70, 3)$	0.537	0.000	0.087	0.087	0.056	(0.179, 0.358, 1.000)

Notes: Entries report AI minus no-AI outcomes from the exact finite-state Bellman solution. In all three cases the no-AI consumer enters only the highest-prior option, while AI additionally assigns positive entry probability to lower-prior options whose summaries make them worth screening. The welfare gains are smaller than the payoff gains because consideration expansion requires additional across-option search expenditure.

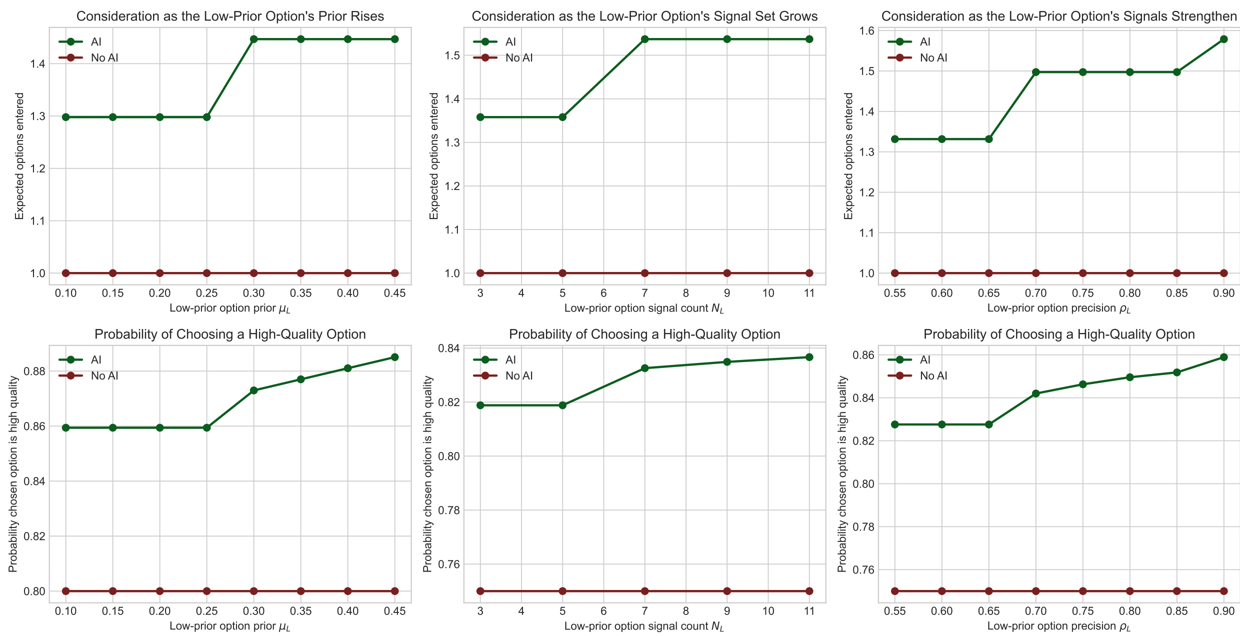

Figure 6: Heterogeneous comparative statics. The left column varies the low-prior option's prior while holding the other options fixed. The middle column varies the low-prior option's signal-set size. The right column varies the low-prior option's raw-signal precision. The top row reports expected options entered and the bottom row reports the probability of choosing a high-quality option. Across all three margins, AI expands consideration most when a low-prior option can generate a more informative overview.

Part II

Multi-Dimensional Extension

4 Model

Part I treats latent quality as a scalar. This section instead studies the case in which product quality is multi-attribute, consumers care unevenly across attributes, and the AI summary is itself lower-dimensional. The key change is simple: the summary no longer compresses *all* of the evidence. It compresses only the dimensions that are most prevalent in the finite review pool, leaving the rest to costly inspection. The parameter K should be read as a display capacity: a platform may show at most six chips, for example, even when review text contains many more economically meaningful dimensions. The model therefore concerns which dimensions enter the displayed set, not the order in which those dimensions are arranged on the page.

That formulation is closer to how many summaries are actually consumed. A restaurant summary highlights food, service, and ambiance but may omit vegan options; a laptop summary highlights speed, battery, and build quality but may omit left-handed ergonomics; a car summary highlights safety and reliability but may omit legroom for unusually tall drivers. The economically relevant question therefore becomes: after the consumer has seen the summary, is one more unopened review likely to teach her something about a dimension she still cares about? To answer that question precisely, the consumer must update not only beliefs about latent quality, but also beliefs about what the remaining unopened reviews still contain.

4.1 Environment

Fix an option j and a consumer with preference vector $\gamma = (\gamma_1, \dots, \gamma_D)$, where $\gamma_d \geq 0$ and $\sum_{d=1}^{D} \gamma_d > 0$. Latent quality is

$$\theta_j = (\theta_{j1}, \dots, \theta_{jD}) \in \{0,1\}^D, \tag{4.1}$$

and flow utility is additively separable across dimensions:

$$u_j(\theta_j) = \mathbf{x}_j'\beta - \alpha p_j + \sum_{d=1}^{D} \gamma_d \theta_{jd}. \tag{4.2}$$

The prior on each dimension is

$$\mu_{jd,0} = \mathbb{P}(\theta_{jd} = 1 \mid \mathbf{x}_j, p_j), \qquad d = 1, \ldots, D, \tag{4.3}$$

and the dimensions are independent under the prior.

Option j has a finite pool of underlying reviews indexed by $n = 1, \ldots, N_j$, where N_j is known to the consumer. Review n is a sparse D-dimensional object

$$r_{jn} = (r_{jn1}, \ldots, r_{jnD}), \qquad r_{jnd} \in \{0, 1, \varnothing\}, \tag{4.4}$$

where $r_{jnd} = \varnothing$ means that review n contains no information about dimension d. Define the coverage indicator

$$m_{jnd} = \mathbf{1}\{r_{jnd} \neq \varnothing\}. \tag{4.5}$$

The finite review pool can therefore be summarized by a sparse signal matrix $R_j = \{r_{jnd}\}_{n,d}$ together with its coverage matrix $M_j = \{m_{jnd}\}_{n,d}$.

Conditional on latent quality and coverage, signs are dimension-specific:

$$\mathbb{P}(r_{jnd} = \theta_{jd} \mid \theta_{jd}, m_{jnd} = 1) = \rho_{jd}, \qquad \mathbb{P}(r_{jnd} = 1 - \theta_{jd} \mid \theta_{jd}, m_{jnd} = 1) = 1 - \rho_{jd}, \tag{4.6}$$

with $\rho_{jd} \in (1/2, 1)$. If $m_{jnd} = 0$, then $r_{jnd} = \varnothing$ deterministically. Across reviews and dimensions, signs are conditionally independent given (θ_j, M_j).

The consumer need not know the exact sparse matrix M_j ex ante, but she knows the finite review count N_j and a prior distribution $G_j(M)$ over coverage matrices. This is what makes the updating problem nontrivial: after observing the AI summary and some opened reviews, the consumer updates jointly over latent quality and over what the remaining unopened reviews are likely to discuss.

4.2 A K-Dimensional Summary over the Most Prevalent Dimensions

Let $K < D$. The AI summary does not attempt to summarize the full D-dimensional review environment. Instead, it reports only the most prevalent dimensions in the finite review pool, up to a maximum of K. For any coverage matrix M_j, define the dimension-specific prevalence counts

$$C_{jd}(M_j) = \sum_{n=1}^{N_j} m_{jnd}, \qquad d = 1, \ldots, D. \tag{4.7}$$

Let $\mathcal{D}_j^+(M_j) = \{d : C_{jd}(M_j) > 0\}$ be the set of dimensions that appear at least once in the finite pool, and let $L_j(M_j) = \min\{K, |\mathcal{D}_j^+(M_j)|\}$. Define

$$\mathcal{S}_j^K(M_j) \subset \mathcal{D}_j^+(M_j), \qquad |\mathcal{S}_j^K(M_j)| = L_j(M_j), \tag{4.8}$$

to be the set of indices corresponding to the $L_j(M_j)$ largest positive values of $C_{jd}(M_j)$, with ties broken by a fixed rule known to the consumer. When no confusion arises, I suppress the dependence on M_j and write $\mathcal{S}_j^K$. Upon entry, the consumer observes the effective summary vector

$$a_j = (a_{jd})_{d \in \mathcal{S}_j^K}, \qquad a_{jd} \in \{0, 1\}, \tag{4.9}$$

where a_{jd} is a compressed signal about dimension d. A natural benchmark is that each a_{jd} is the majority-rule overview of the non-missing dimension-d signals:

$$\mathbb{P}(a_{jd} = 1 \mid R_j, M_j) = \begin{cases} 1, & \sum_{n:\, m_{jnd}=1} r_{jnd} > C_{jd}(M_j)/2, \\ \tau_{jd}, & \sum_{n:\, m_{jnd}=1} r_{jnd} = C_{jd}(M_j)/2, \\ 0, & \sum_{n:\, m_{jnd}=1} r_{jnd} < C_{jd}(M_j)/2, \end{cases} \qquad d \in \mathcal{S}_j^K, \tag{4.10}$$

where $\tau_{jd} \in [0, 1]$ is a fixed dimension-level tie-breaking probability known to the consumer. Thus every reported summary component has positive realized prevalence. If the finite pool contains at least K positive-prevalence dimensions, then $|\mathcal{S}_j^K(M_j)| = K$; if it contains fewer, the unreported zero-prevalence dimensions are uninformative null components and are dropped from the effective summary. The consumer knows the aggregation rule, knows that at most the K most prevalent dimensions are summarized, and knows that any posterior about the remaining unopened reviews must be consistent with both of those facts.

Under this strict majority benchmark, an exact mixed sentiment state can arise only through a tie in the non-missing signals for a displayed dimension. Hence it has positive probability only when $C_{jd}(M_j)$ is even. If $C_{jd}(M_j)$ is odd, majority rule maps every realized dimension-level signal pool into either a favorable or unfavorable summary. A platform can still report a mixed category for odd coverage counts, but doing so corresponds to the wider three-cell threshold rule in Section 2.4.3, not to a pure majority tie.

This specification is intentionally asymmetric. Covered dimensions are summarized for free; uncovered dimensions are not. The summary is therefore informative about common dimensions of discussion, not necessarily about the dimensions a particular consumer values.

4.3 Beliefs after the summary and the value of another review

Let h_t denote the history after t opened reviews: it records which reviews have been opened and the full sparse vectors observed on those reviews. For any covered dimension $d \in \mathcal{S}_j^K$, let

$$\mu_{jd,t}(a_j, h_t) = \mathbb{P}(\theta_{jd} = 1 \mid a_j, h_t) \tag{4.11}$$

denote the posterior on latent quality dimension d after observing the summary and the first t opened reviews. At entry, before any sparse reviews are opened, write $\mu_{jd,0}(a_j) \equiv \mu_{jd,0}(a_j, \emptyset)$. For uncovered dimensions, the summary does not directly report a compressed sign, but it still carries indirect information through the fact that those dimensions failed to make the top-K set. Formally, the consumer updates jointly over (θ_j, M_j) using the integrated likelihood of the observed summary and opened reviews:

$$\mathbb{P}(\theta_j, M_j \mid a_j, h_t) \propto \mathbb{P}(\theta_j)\, G_j(M_j)\, \mathbf{1}\big\{\mathcal{S}_j^K(M_j) = \mathcal{S}_j^K\big\}\, \mathcal{L}_j(a_j, h_t \mid \theta_j, M_j), \tag{4.12}$$

where

$$\mathcal{L}_j(a_j, h_t \mid \theta_j, M_j) = \sum_{R_j \in \mathcal{R}_j(M_j; h_t)} \mathbb{P}(R_j \mid \theta_j, M_j)\, \mathbb{P}(a_j \mid R_j, M_j) \tag{4.13}$$

and $\mathcal{R}_j(M_j; h_t)$ is the finite set of full sparse signal matrices that are consistent with coverage matrix M_j and the opened-review history h_t. Equation (4.13) is the multi-dimensional analogue of the H-function in Part I: it integrates over the unobserved signs in the finite review pool rather than conditioning on an unobserved full signal matrix. Equation (4.12) is therefore the precise updating object behind the extension: the summary restricts posterior mass through the integrated probability that the unseen review pool could have generated both the observed top-K summary and the opened sparse reviews. Appendix B.1 derives (4.12) and the related updating objects step by step.

Marginalizing (4.12) over M_j gives the quality posterior in (4.11). The consumer's expected utility at history (a_j, h_t) is

$$\delta_j(a_j, h_t) = \mathbf{x}_j'\beta - \alpha p_j + \sum_{d=1}^{D} \gamma_d \mu_{jd,t}(a_j, h_t). \tag{4.14}$$

Inspection technology is passive: opening one more review means drawing one review uniformly from the unopened finite pool rather than targeting a review by topic. Let

$$\eta_{jd,t}(a_j, h_t) = \mathbb{P}\big(m_{j,n_{t+1},d} = 1 \mid a_j, h_t\big)\,, \tag{4.15}$$

where n_{t+1} denotes a uniformly chosen unopened review. Because the consumer knows the summary rule and the finite review count, she can compute this probability from the joint posterior:

$$\eta_{jd,t}(a_j, h_t) = \mathbb{E}\left[\frac{C_{jd}(M_j) - c_{jd,t}}{N_j - t} \;\middle|\; a_j, h_t\right], \tag{4.16}$$

where $c_{jd,t}$ is the number of opened reviews up to time t that covered dimension d. This is the exact formal version of the intuition that the consumer keeps reading when there is still a meaningful chance that the next unopened review will speak to an uncovered dimension she cares about.

For the sign calculation, maintain the tractable separability assumption

$$m_{j,n_{t+1},d} \perp \theta_{jd} \mid a_j, h_t. \tag{4.17}$$

This says that once the consumer conditions on the realized summary and the opened-review history, the event that the next unopened review covers dimension d carries no additional information about the latent quality on that same dimension. Under (4.17), if the next unopened review covers dimension d, the predictive probability of a favorable signal on that dimension is

$$\begin{aligned} \pi^{+}_{jd,t}(a_j, h_t) &= \mathbb{P}(r_{j,n_{t+1},d} = 1 \mid m_{j,n_{t+1},d} = 1, a_j, h_t) \\ &= \mu_{jd,t}(a_j, h_t)\rho_{jd} + \big(1 - \mu_{jd,t}(a_j, h_t)\big)(1 - \rho_{jd}). \end{aligned} \tag{4.18}$$

The unconditional predictive probability that the next unopened review delivers a favorable signal on dimension d is therefore

$$\psi^{+}_{jd,t}(a_j, h_t) = \eta_{jd,t}(a_j, h_t)\, \pi^{+}_{jd,t}(a_j, h_t). \tag{4.19}$$

The role of (4.17) should be interpreted narrowly. It is a tractability restriction that separates uncertainty about whether the next unopened review discusses dimension d from uncertainty about the sign of that discussion once it appears. Economically, it rules out environments in which topic incidence is itself informative about latent quality beyond the realized summary and the opened-review history—for example, settings in which very poor products generate more discussion of defects or very strong products generate more discussion of standout attributes. Without (4.17), the fact that the next review covers dimension d would itself shift beliefs about θ_{jd}, so the clean decomposition in (4.19) would be replaced by a more general predictive term in which topic arrival is directly informative about quality.

Equation (4.19) is the heart of the extension. In the single-dimensional model every next review is automatically payoff-relevant. Here, continuation depends not only on how informative a review would be if opened, but also on the posterior probability that the next unopened review will address a utility-relevant dimension at all.

4.4 The within-option dynamic program

To connect the updating objects above to search behavior, let the within-option state after the summary be $s_t = (a_j, h_t)$, where h_t records the first t opened sparse reviews. Let $\bar{M}_j$ denote the continuation value of leaving the current option for the best outside opportunity. The current stopping payoff is

$$S_j(s_t; \gamma) = \max\left[0, \delta_j(a_j, h_t), \bar{M}_j\right]. \tag{4.20}$$

The within-option Bellman equation is

$$V_j(s_t) = \max\left\{S_j(s_t; \gamma),\ -c_j^W(t+1) + \mathbb{E}[V_j(s_{t+1}) \mid s_t]\right\}. \tag{4.21}$$

The stopping value is the best of exit, keeping the current option, and taking the outside opportunity. Continuation is costly because another review must be opened before its sparse content is known.

The expectation in (4.21) is over all sparse next-review realizations consistent with the posterior in (4.12). That is precisely where the finite-pool logic matters: after seeing the summary and some opened reviews, the consumer updates not only latent-quality beliefs but also the probability distribution over what the unopened reviews still discuss.

Define the exact dynamic gross value of starting within-option search at state s_t as

$$\bar{c}_j^{MD}(s_t; \gamma) = \mathbb{E}[V_j(s_{t+1}) \mid s_t] - S_j(s_t; \gamma). \tag{4.22}$$

Proposition 6 (Exact multidimensional initiation cutoff)**.** *At any post-summary state $s_t = (a_j, h_t)$ with at least one unopened review, within-option search starts in the full dynamic problem if and only if*

$$c_j^W(t+1) \leq \bar{c}_j^{MD}(s_t; \gamma). \tag{4.23}$$

If the inequality fails, immediate stopping—exit, choosing the current option, or moving to the outside option—is optimal. The cutoff in (4.23) is exact but generally recursive because $V_j(s_{t+1})$ embeds all future inspection decisions.

Proof in Appendix B.1.

4.5 A one-step continuation bound

The full Bellman problem in (4.21) is high-dimensional. To obtain a sharp analytical object analogous to Part I, fix a post-summary state $s_t = (a_j, h_t)$ and consider the one-step deviation in which the consumer opens exactly one more review and then stops.

Let $\mathcal{X}_j(s_t)$ denote the finite set of sparse next-review realizations that can arrive at state s_t, and let $s_t(x)$ be the post-review state after realization $x \in \mathcal{X}_j(s_t)$. The exact primitive one-step value of opening one more review and then stopping is

$$\mathcal{I}_j^{MD}(s_t;\gamma) = \mathbb{E}[S_j(s_t(X_{t+1});\gamma) \mid s_t] - S_j(s_t;\gamma), \tag{4.24}$$

where the expectation is taken under the posterior predictive distribution implied by (4.12).

Proposition 7 (Multidimensional primitive verification threshold)**.** *At any post-summary state $s_t = (a_j, h_t)$, if the consumer is restricted to at most one additional review before stopping, opening one more review is optimal if and only if*

$$c_j^W(t+1) \le \mathcal{I}_j^{MD}(s_t;\gamma). \tag{4.25}$$

Proof in Appendix B.1.

Fix a focal uncovered dimension $d \notin \mathcal{S}_j^K$ with $\gamma_d > 0$. Consider the feasible one-step policy that opens one more review and then stops. If the next review does not cover d, the policy ignores whatever other information arrives and stops at $S_j(s_t;\gamma)$, so the construction is a lower bound on the full one-step value. If the next review covers d, let $S_{jd}^+(s_t)$ and $S_{jd}^-(s_t)$ denote the expected stopping payoffs after the full sparse review is observed, conditional respectively on the dimension-d realization being favorable or unfavorable; these expectations integrate over any other dimensions covered by the same review. The one-step value of continuing for the chance to learn about dimension d is then bounded below by

$$\underline{\mathcal{K}}_{jd}(s_t) = -c_j^W(t+1) + \eta_{jd,t}(s_t)\Big[\pi_{jd,t}^+(s_t)S_{jd}^+(s_t) + \big(1-\pi_{jd,t}^+(s_t)\big)S_{jd}^-(s_t)\Big] + \big(1-\eta_{jd,t}(s_t)\big)S_j(s_t;\gamma). \tag{4.26}$$

Proposition 8 (One-step verification condition)**.** *At any post-summary state $s_t = (a_j, h_t)$ and for any uncovered dimension d with $\gamma_d > 0$, define $S_{jd}^+(s_t)$ and $S_{jd}^-(s_t)$ as the conditional expected stopping payoffs described above. If*

$$c_j^W(t+1) \le \eta_{jd,t}(s_t)\left(\pi_{jd,t}^+(s_t)\left[S_{jd}^+(s_t) - S_j(s_t;\gamma)\right] + \left(1-\pi_{jd,t}^+(s_t)\right)\left[S_{jd}^-(s_t) - S_j(s_t;\gamma)\right]\right), \tag{4.27}$$

then continuing for one more review is optimal at state s_t.

Proof in Appendix B.1.

Proposition 8 isolates a sufficient one-step residual-learnability channel. Search survives in this focal deviation not merely because the summary omitted something the consumer values, but because the omitted dimension is still likely to appear in the next unopened review and is valuable enough to overturn the current stopping payoff. The object $\eta_{jd,t}(s_t)$ is what scales this component of further search relative to the single-dimensional benchmark: if $\eta_{jd,t}(s_t) = 0$, the focal omitted dimension cannot be recovered from the remaining finite review pool and this omitted-information motive for continuation disappears. Appendix B.1 also expands the algebra behind the one-step bound in (4.26)–(4.27).

4.6 Coverage and residual learnability

Two scalars summarize the consumer-facing content of the summary.

Definition 2 (Coverage share)**.** *For consumer preference vector γ and summary coverage set $\mathcal{S}_j^K$, define*

$$\mathrm{Cove}_j(\gamma; K) = \frac{\sum_{d \in \mathcal{S}_j^K} \gamma_d}{\sum_{d=1}^{D} \gamma_d} \in [0, 1]. \tag{4.28}$$

The coverage share measures how much of the consumer's preference mass lies on dimensions explicitly summarized by AI.

Definition 3 (Residual learnability)**.** *For consumer preference vector γ and history (a_j, h_t), define*

$$\mathrm{RL}_{j,t}(\gamma; a_j, h_t) = \sum_{d \notin \mathcal{S}_j^K} \gamma_d \eta_{jd,t}(a_j, h_t). \tag{4.29}$$

Residual learnability is the preference-weighted posterior probability that one additional unopened review will speak to an uncovered dimension. It is the simplest scalar object capturing the new friction in (4.19). A consumer may care about an uncovered dimension, but if the remaining unopened reviews are unlikely to discuss it then the incentive to keep reading is small; conversely, a dimension can be omitted by the summary yet still sustain verification if it is both salient to the consumer and sufficiently likely to appear in the remaining finite review pool.

5 Results

The multi-dimensional extension is designed to formalize a new mechanism rather than to duplicate the theorem-by-theorem structure of Part I. The strongest global statements are endpoint results, but the one-step problem also yields an exact multidimensional verification threshold and a qualified coverage-residual ordering.

5.1 Endpoint propositions

Proposition 9 (Zero-direct-coverage benchmark). *Suppose* $\{d : \gamma_d > 0\} \cap \mathcal{S}_j^K = \emptyset$*. Under prior independence across latent-quality dimensions, the observed summary vector* a_j *does not directly update any utility-relevant quality posterior: for every* d *with* $\gamma_d > 0$*,*

$$\mu_{jd,0}(a_j) = \mu_{jd,0}.$$

Hence any AI effect on within-option verification must operate through the coverage probabilities $\eta_{jd,t}(a_j, h_t)$ *rather than through direct updating of utility-relevant quality beliefs. If, in addition, the coverage counts* $\{C_{jd}(M_j) : \gamma_d > 0\}$ *are known ex ante, then the within-option problem coincides with the no-AI benchmark.*

Proof in Appendix B.1.

Proposition 9 is the right zero-coverage benchmark for the top-K model. Omission no longer implies an automatic full null, because failing to make the top-K set can itself be informative about whether unopened reviews are likely to discuss the omitted dimension. What omission does imply is that AI cannot directly move utility-relevant quality beliefs. Exact coincidence with no-AI requires the stronger condition that coverage on those utility-relevant dimensions is already known.

Proposition 10 (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* $d^* \in \mathcal{S}_j^K$*, every review covers that dimension, i.e.* $m_{jnd^*} = 1$ *for all* $n = 1, \ldots, N_j$*, and the summary component* a_{jd^*} *is the same compression of the dimension-*d^* *signal pool studied in Part I, then the consumer's within-option problem on dimension* d^* *is exactly the single-dimensional problem with primitives* $(\mu_{jd^*,0}, \rho_{jd^*}, N_j, \gamma_{d^*})$*.*

Proof in Appendix B.1.

Proposition 10 shows exactly where the new model nests the original one. If a single covered dimension drives utility and every review speaks to that dimension, then Part II reduces to Part I. If coverage on that dimension is incomplete, the multidimensional model differs from Part I through the additional topic-arrival risk captured by $\eta_{jd^*,t}(a_j, h_t)$.

5.2 Interior coverage-residual ordering

The remaining cases are neither exact nulls nor exact reductions. When some preference mass lies on covered dimensions and some lies on uncovered dimensions, the summary partially resolves relevant uncertainty but leaves a residual incentive to keep reading. Coverage and residual learnability alone do not generically order verification across all possible taste vectors, because omitted dimensions can differ in posterior uncertainty, signal precision, and stopping-payoff wedges. The sharp statement is therefore a one-step ordering under a homogeneous omitted-dimension condition; outside that condition the same logic should be read as an empirical and numerical prediction.

Proposition 11 (Coverage-residual ordering under homogeneous omitted dimensions)**.** *Hold fixed $(\mu_{jd,0}, \rho_{jd})_{d=1}^{D}$, the prior over coverage matrices G_j, the within-option cost schedule, and K. Consider two consumers with preference vectors γ and γ' at the same post-summary history (a_j, h_t). Suppose that uncovered dimensions have a common one-step payoff wedge per unit of taste weight: for every $d \notin \mathcal{S}_j^K$ there is a common scalar $B_t(s_t) \geq 0$ such that the one-step benefit of dimension d is $\gamma_d \eta_{jd,t}(s_t) B_t(s_t)$. If*

$$\mathrm{Cove}_j(\gamma; K) \geq \mathrm{Cove}_j(\gamma'; K) \qquad \text{and} \qquad \mathrm{RL}_{j,t}(\gamma; a_j, h_t) \leq \mathrm{RL}_{j,t}(\gamma'; a_j, h_t),$$

then the one-step verification cutoff for γ is weakly lower than the one-step verification cutoff for γ'. In the unrestricted dynamic problem, the same monotonicity is the maintained comparative-static prediction when continuation values preserve this one-step ordering.

Derivation in Appendix B.1.

Proposition 11 says that AI crowds out verification most when two conditions are jointly satisfied: the summary covers the dimensions the consumer values, and the remaining unopened review pool is unlikely to deliver utility-relevant information outside that coverage set. The homogeneity qualification matters. Without it, two consumers with the same coverage share can differ because their omitted taste weight may sit on dimensions with different precision, posterior variance, or decision-changing potential. With it, residual learnability is the sufficient statistic for the omitted-dimension search motive.

The proposition yields a simple three-regime reading of the model:

- *Zero-direct-coverage regime:* $\mathrm{Cove}_j(\gamma; K) = 0$. By Proposition 9, AI cannot directly move utility-relevant quality beliefs; any remaining within-option effect must come from coverage inference.

- *Selective-persistence regime:* $0 < \mathrm{Cove}_j(\gamma; K) < 1$. The summary resolves some but not all relevant uncertainty, so verification persists selectively.

- *High-coverage regime:* $\mathrm{Cove}_j(\gamma; K) = 1$ and $\mathrm{RL}_{j,t}(\gamma; a_j, h_t)$ is close to zero. The consumer is pushed back toward the single-dimensional benchmark of Part I.

5.3 Analytical taste regions

The ordering can be sharpened analytically at the one-step level. Fix a post-summary state $s_t = (a_j, h_t)$ and an uncovered dimension $d \notin \mathcal{S}_j^K$. Define the focal-dimension lower-bound gain from opening one more review for the chance to learn about dimension d:

$$\mathcal{B}_{jd}(s_t; \gamma) = \eta_{jd,t}(s_t) \left(\pi^+_{jd,t}(s_t) \left[S^+_{jd}(s_t) - S_j(s_t; \gamma) \right] + \left(1 - \pi^+_{jd,t}(s_t)\right) \left[S^-_{jd}(s_t) - S_j(s_t; \gamma) \right] \right). \tag{5.1}$$

Proposition 8 then implies the sufficient one-step inspection region for dimension d:

$$\mathcal{T}^I_{jd}(s_t) = \left\{ \gamma \in \mathbb{R}^D_+ : c^W_j(t+1) \le \mathcal{B}_{jd}(s_t; \gamma) \right\}. \tag{5.2}$$

Equation (5.2) is the analytical counterpart to the numerical taste-type exercise: taste enters continuation through the stopping-payoff wedges on the uncovered dimension, scaled by the probability that the next unopened review actually discusses that dimension.

The region becomes especially transparent under a local separability condition. Suppose that, for the focal one-step deviation, the next review changes stopping payoffs only through dimension d, so there exist state objects $\Delta^+_{jd,t}(s_t)$ and $\Delta^-_{jd,t}(s_t)$ such that

$$S^+_{jd}(s_t) - S_j(s_t; \gamma) = \gamma_d \Delta^+_{jd,t}(s_t), \qquad S^-_{jd}(s_t) - S_j(s_t; \gamma) = \gamma_d \Delta^-_{jd,t}(s_t). \tag{5.3}$$

Then the one-step inspection region collapses to a half-space in taste space:

$$\mathcal{T}^I_{jd}(s_t) = \left\{ \gamma \in \mathbb{R}^D_+ : \gamma_d \ge \gamma^*_{jd}(s_t) \right\}, \tag{5.4}$$

where

$$\gamma^*_{jd}(s_t) = \frac{c^W_j(t+1)}{\eta_{jd,t}(s_t) \left(\pi^+_{jd,t}(s_t) \Delta^+_{jd,t}(s_t) + \left(1 - \pi^+_{jd,t}(s_t)\right) \Delta^-_{jd,t}(s_t) \right)} \tag{5.5}$$

whenever the denominator is positive. Thus uncovered-dimension salience matters only insofar as the omitted dimension is still recoverable from the remaining finite review pool.

This characterization clarifies the distinction between taste regions and simple heterogeneous matching. Matching depends only on whether taste weight lies on covered or uncovered dimensions. By contrast, (5.5) shows that two consumers with the same coverage share can face different continuation incentives if their omitted taste weight falls on dimensions with different residual learnability or different payoff consequences after one additional review.

The same logic gives an exact one-step taste-region classification in the common one-topic-per-review specialization used in the numerical section. Suppose each opened review covers at most one payoff-relevant dimension and the local separability condition in (5.3) holds for every dimension d. Define the per-unit taste gain

$$\Omega_{jd,t}(s_t) = \eta_{jd,t}(s_t)\left(\pi^+_{jd,t}(s_t)\Delta^+_{jd,t}(s_t) + \left(1-\pi^+_{jd,t}(s_t)\right)\Delta^-_{jd,t}(s_t)\right). \tag{5.6}$$

Then the one-step initiation region is the half-space

$$\mathcal{T}^1_j(s_t) = \left\{\gamma \in \mathbb{R}^D_+ : \sum_{d=1}^{D} \gamma_d \Omega_{jd,t}(s_t) \geq c^W_j(t+1)\right\}. \tag{5.7}$$

On the normalized taste simplex $\Gamma = \{\gamma \geq 0 : \sum_d \gamma_d = 1\}$, let

$$\Omega^{\min}_j(s_t) = \min_d \Omega_{jd,t}(s_t), \qquad \Omega^{\max}_j(s_t) = \max_d \Omega_{jd,t}(s_t).$$

There are three one-step regimes:

(i) if $c^W_j(t+1) > \Omega^{\max}_j(s_t)$, no normalized taste type starts one-step within-option search;

(ii) if $c^W_j(t+1) \leq \Omega^{\min}_j(s_t)$, every normalized taste type starts one-step within-option search;

(iii) if $\Omega^{\min}_j(s_t) < c^W_j(t+1) \leq \Omega^{\max}_j(s_t)$, one-step search starts only for tastes on the high-gain side of the hyperplane in (5.7).

Appendix B.1 records the proof, the two-dimensional rotation thresholds, and the relationship between this exact one-step region and the sufficient focal-dimension region in (5.2).

5.4 Numerical illustration

To give Part II the same quantitative discipline as Part I, I solve the within-option Bellman problem in (4.21) exactly under a tractable numerical specialization of the sparse-review environment. Each review covers exactly one dimension, topic arrival is i.i.d. across the finite pool with shares $(q_1, q_2, q_3, q_4) = (0.40, 0.30, 0.18, 0.12)$, and the AI interface reports up to the top $K = 2$ positive-prevalence dimensions in a pool of $N_j = 3$ reviews. Realized coverage-count ties are broken by the fixed dimension order, and dimension-level sign ties use $\tau_{jd} = 0$, so tied signs break negative. The latent-quality and sign-precision primitives are $(\mu_{jd,0})^4_{d=1} = (0.55, 0.50, 0.48, 0.47)$ and $(\rho_{jd})^4_{d=1} = (0.82, 0.76, 0.72, 0.68)$, the outside opportunity is $\bar{M}_j =$

0.02, and the per-review verification cost is constant at $c_j^W = 0.015$. This one-topic-per-review specialization is narrower than the full model in Section 4.1, but it preserves the key finite-pool topic-arrival friction through the posterior objects $\eta_{jd,t}$ and $\psi_{jd,t}^{+}$.

Table 7 reports five benchmark consumer types. The aligned type loads all taste weight on a mainstream covered dimension and exhibits the clean substitution result from Part I: AI cuts expected review reading from 2.395 to 0.667 while raising ex ante value by 0.026. The mixed and uncovered types instead illustrate selective persistence. In both cases the no-AI consumer stops immediately, while AI induces positive review reading because the summary cheaply resolves covered dimensions and leaves utility-relevant uncertainty on an omitted one. The covered-split and tail types are low-search corners, but even there AI weakly raises value by improving the stopping decision.

Table 7: Multi-Dimensional Benchmark Consumer Types

Type	γ	$\text{Cove}_j(\gamma; K)$	No AI E[reviews]	AI E[reviews]	ΔV
Aligned	$(1, 0, 0, 0)$	1.00	2.395	0.667	0.026
Covered split	$(0.5, 0.5, 0, 0)$	1.00	0.000	0.000	0.010
Mixed	$(0.6, 0.2, 0.2, 0)$	0.80	0.000	0.377	0.008
Uncovered	$(0, 0, 1, 0)$	0.00	0.000	0.130	0.010
Tail	$(0, 0, 0.6, 0.4)$	0.00	0.000	0.010	0.000

Notes: Entries report exact within-option Bellman outcomes under the benchmark specialization described in Section 5.4. ΔV denotes AI minus no-AI ex ante value. The AI summary reports the realized effective summary set, whose size can be less than K when fewer than K dimensions appear in the finite pool. The aligned type illustrates verification substitution; the mixed and uncovered types illustrate selective persistence because omitted but recoverable dimensions keep search alive under AI.

The more structural object is the taste-region map in Figure 7. For that figure I switch to a symmetric calibration in which all four dimensions share a common prior and signal precision, so the only primitive asymmetry comes from topic prevalence and the preference vector itself. The left panel rotates taste mass from a covered dimension toward an uncovered one; the right panel rotates taste mass within the covered set. For each taste vector, the plotted object is the largest grid value, up to 0.03, at which initial inspection is still optimal. Two patterns stand out. First, AI sustains inspection for a weakly larger set of grid costs than no-AI in both families. Second, the gap is generally largest near concentrated taste vectors and smallest at interior mixtures. Economically, that is exactly the coverage-residual logic of Proposition 11: concentration on one covered dimension makes the summary highly substitutable for raw reviews, while concentration on one omitted dimension makes residual learnability especially valuable. Interior mixtures dilute the marginal value of any one additional dimension-specific review and therefore shrink the inspection region in both environments.

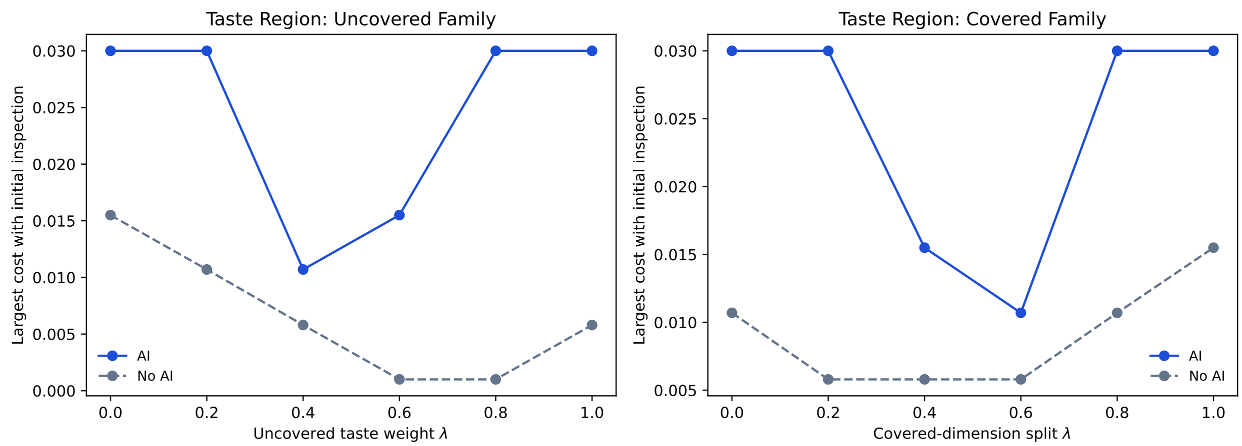

Figure 7: Taste-region analysis in the multi-dimensional extension under a symmetric calibration with common priors and signal precision across dimensions. The left panel rotates taste mass from a covered mainstream dimension toward an uncovered niche dimension. The right panel rotates taste mass across two covered dimensions. Each point reports the largest plotted per-review verification cost, up to 0.03, for which initial inspection is optimal. In both families, AI supports inspection over a weakly larger cost region than no-AI, but the size of that region depends sharply on where taste mass is concentrated.

Taken together, Table 7 and Figure 7 give numerical backing to the three-regime interpretation above. High coverage can crowd out verification sharply, omitted but still recoverable dimensions can preserve review reading under AI, and heterogeneous taste vectors sort consumers across those regimes even when they face the same finite review pool and the same summary rule.

5.5 Testable predictions

The new framework yields four sharp empirical predictions.

1. *Within-category consumer heterogeneity.* Consumers whose preferences load more heavily on the summary's covered dimensions should reduce review-reading more after AI rollout.

2. *Cross-category heterogeneity.* Categories whose review discourse is concentrated on a few widely shared dimensions should exhibit larger AI-induced click-depth reductions than categories with diffuse or niche-specific review topics.

3. *Attribute-specific click reallocation.* After AI rollout, the relative share of clicks should shift toward reviews discussing uncovered but salient attributes, since those are precisely the dimensions the summary leaves unresolved.

4. *Across-option implications.* Options whose covered dimensions overlap more with the consumer's preference vector become cheaper to screen after AI rollout and should therefore be more likely to enter the consideration set.

The last prediction links Part II back to Part I. In the single-dimensional model, AI expands consideration by freeing attention from within-option verification. In the multi-dimensional model, the same force should be strongest for options whose summaries cover the dimensions the consumer values most.

6 Discussion and extensions

Three extensions are especially natural.

A full nested multi-dimensional search model. Part II is intentionally focused on the within-option margin. The same logic should also reshape across-option search because an option's ex ante entry value depends on whether its covered dimensions line up with the consumer's preference vector. A full nested model would therefore endogenize which options are even worth opening under AI.

Estimation. The objects in (4.28) and (4.29) suggest a direct empirical path. Review text can be topic-labeled into dimensions, the platform summary can be coded for which dimensions it covers, and consumer heterogeneity can be summarized by a preference vector inferred from browsing or purchase behavior. The central empirical question is then whether AI-induced click-depth reductions are strongest when estimated coverage is high and residual learnability is low. Section 2.4.3 adds a complementary first-stage test: when the researcher observes both the summary output and the review corpus, one can estimate or audit whether the platform is using a prevalence-and-threshold rule. The key evidence is not the order of displayed chips, but whether the review corpus contains more candidate dimensions than the displayed summary and whether the displayed set is the high-prevalence subset. If a platform always displays six dimensions whenever at least six corpus dimensions are present, the natural interpretation is $K = 6$; the top-K restriction is then tested by comparing the displayed set with the six highest-prevalence corpus dimensions, and the sentiment thresholds are tested by comparing each displayed cell with the dimension-specific positive and negative mention counts.

Uncertain summary coverage. The current benchmark assumes consumers know which dimensions the summary covers. A natural next step is to let consumers be uncertain about $\mathcal{S}_j^K$ itself. That would soften Proposition 9: omission would no longer imply an exact null, but rather an approximate one whose strength depends on how likely the consumer believes it is that the summary covers her salient dimensions.

7 Conclusion

AI summaries do not simply add information to consumer search—they reallocate it. This paper develops a demand-side model of sequential search in which entering an option reveals a free compressed overview, and deeper inspection of the underlying evidence is costly. By comparing an AI environment to a no-summary benchmark, the model isolates three regimes—verification substitution, consideration expansion, and search preservation—and characterizes the conditions under which each arises.

The paper's main single-dimensional analytical results concern the within-option verification margin. For any finite signal-set size $N_j > 1$ and one-step post-summary inspection, the primitive value of verification vanishes as raw-signal precision rises and the favorable-overview continuation region exits the feasible outside-belief space in the relevant case-(a) region (Corollary 1). Under the nondegenerate majority rules in Corollary 2, the set of states in which verification is optimal also shrinks monotonically near perfect precision. When the underlying human signals are sufficiently informative, the AI summary compresses them so well that a rational consumer skips the reviews entirely. Corollary 3 extends this logic to any finite inspection horizon: no multi-step verification plan survives at sufficiently high signal precision. The implication is that, in the single-dimensional high-precision region, better reviews lead to *less* review-reading when AI summaries are present—the opposite of what one would expect in a standard search model where free information and further search are complements.

The multi-dimensional extension in Section 4 reframes the mechanism around partial coverage rather than around a scalar summary of all product attributes. Quality is D-dimensional, each review is a sparse D-dimensional object with missing entries on the dimensions it does not discuss, and the AI interface reports up to the $K < D$ most prevalent positive-coverage dimensions in the finite review pool. The consumer therefore updates jointly over latent quality and over what the remaining unopened reviews are likely to contain. That motivates two new objects: a coverage share, measuring how much of the consumer's preference vector the summary addresses, and a residual-learnability term, measuring the posterior probability that one more unopened review will reveal preference-weighted information outside the summary. Proposition 9 establishes the first endpoint: if the summary omits every dimension the consumer values, it cannot directly move utility-relevant quality beliefs, though it may still change search by altering beliefs about what unopened reviews discuss. Proposition 10 gives the opposite nesting result: if the consumer effectively cares about one covered dimension and every review covers that dimension, the problem collapses back to the single-dimensional benchmark. Proposition 11 states the qualified interior claim:

under homogeneous omitted-dimension payoff wedges, one-step verification falls as coverage rises and residual learnability falls; beyond that case, the same object is the paper's empirical comparative-static prediction. Exact within-option Bellman numerics for a $D = 4$, $K = 2$ benchmark support that logic: aligned consumers display sharp verification substitution, while consumers whose tastes load on uncovered but recoverable dimensions continue to inspect under AI.

Those within-option predictions interact naturally with the across-option margin. In Part I, AI expands consideration by freeing attention from verification. In the multi-dimensional extension, the same force is predicted to be strongest, holding other primitives fixed, for options whose summaries cover the dimensions the consumer values most. The extension therefore points toward a richer across-option model in which summary coverage affects not only whether the consumer keeps reading within an option, but whether the option is worth entering at all.

Several testable predictions follow directly from this framing. First, holding other primitives fixed, platforms that introduce AI summaries should observe the largest decline in within-option click-depth where summaries become most decisive and where consumer preferences are concentrated on covered dimensions. Second, consideration sets should expand most where AI summaries make initial screening cheapest—low-prior or unfamiliar product categories with enough evidence to generate informative summaries. Third, within a category, consumers with atypical tastes should show smaller AI-induced click-depth reductions because the omitted dimensions they care about still require raw review reading. Fourth, reviews discussing uncovered attributes should gain click share after AI rollout, since those are exactly the dimensions the summary leaves unresolved. Fifth, when the summary technology reports positive, mixed, and negative cells, raw-review clicks should concentrate on mixed products or attributes: those are the states in which the free summary is least decisive and one more raw signal is most likely to change the stopping decision. These predictions differ from those of standard sequential-search models and suggest empirical tests using platform experiments that randomize the availability of AI overviews, potentially paired with cross-sectional heterogeneity in product categories, consumer types, and summary-cell realizations.

The model is deliberately silent on two fronts. First, it takes the summary design as exogenous. Making the coverage set endogenous would open a persuasion problem (Kamenica and Gentzkow, 2011) in which the platform chooses which dimensions to summarize, not just how to aggregate them. Second, the paper abstracts from supply-side responses. If sellers can observe whether AI summaries are shown, they may adjust pricing, quality provision, or review solicitation in response, potentially undoing some of the demand-side gains. A full welfare analysis would require embedding the demand model into an equilibrium framework with strategic sellers.

Two extensions are especially natural. First, one can replace the binary latent-quality state with a continuous latent-quality model, provided the overview remains a lossy compression rather than a sufficient statistic for the underlying signal set. Second, one can allow the consumer to be uncertain about the summary's coverage set itself, which would turn the exact zero-coverage null into an approximate prediction. Both extensions preserve the paper's central reallocation mechanism while making the model more empirically flexible.

A Appendix to Part I: Unidimensional Model

A.1 Appendix Note on Strategic Tie-Breaking Under Even N_j

The benchmark analysis uses odd N_j so that majority aggregation is uniquely defined. If instead N_j 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

$$\mathbb{P}(a_j = 1 \mid \text{tie}) = \tau_j, \qquad \tau_j \in [0,1], \tag{A.1}$$

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

Under this parameterization, the high-state overview accuracy becomes

$$\begin{aligned}\mathbb{P}(a_j = 1 \mid \theta_j = 1) = &\sum_{m=N_j/2+1}^{N_j} \binom{N_j}{m} \rho_j^m (1-\rho_j)^{N_j-m} \\ &+ \tau_j \binom{N_j}{N_j/2} \rho_j^{N_j/2} (1-\rho_j)^{N_j/2}.\end{aligned} \tag{A.2}$$

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,

$$\mathbb{P}(a_j = 0 \mid \theta_j = 0) = \sum_{m=N_j/2+1}^{N_j} \binom{N_j}{m} \rho_j^m (1-\rho_j)^{N_j-m} + (1-\tau_j) \binom{N_j}{N_j/2} \rho_j^{N_j/2} (1-\rho_j)^{N_j/2}, \tag{A.3}$$

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

This tie-breaking margin is small but economically meaningful. Increasing τ_j shifts probability mass toward positive overviews precisely in the most ambiguous states, which tends to expand consideration and reduce verification by making the summary appear more decisive than the underlying evidence warrants. If consumers trust the summary, seller- or platform-chosen tie-breaking can therefore lower buyer welfare by increasing false positives and inducing premature stopping. For that reason, the main text fixes odd N_j and treats strategic tie-breaking as a simple illustration of the broader endogenous-ϕ_j agenda rather than as part of the benchmark demand-side model.

For the high-precision elimination results in Corollaries 1 and 3, however, even-N_j tie-breaking does not introduce a qualitatively new force: favorable overviews still become nearly sufficient statistics as $\rho_j \to 1$. The monotone derivative result in Corollary 2 additionally uses the simple-zero condition stated there. Figure 8 illustrates the neutral even-N_j case numerically by comparing the one-step verification geometry for $N_j = 3$, $N_j = 4$ with neutral tie-breaking $\tau_j = 1/2$, and $N_j = 5$. 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.

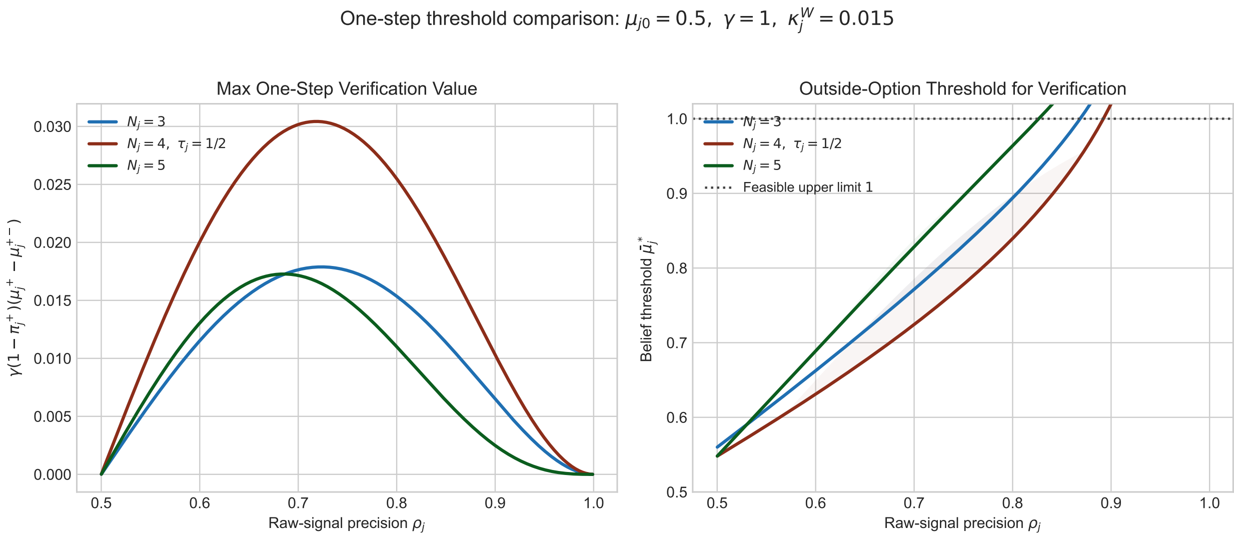

Figure 8: One-step verification thresholds with odd and even signal counts. The left panel plots the maximal one-step verification value after a favorable overview for $N_j = 3$, $N_j = 4$ with neutral tie-breaking, and $N_j = 5$ under $\mu_{j0} = 0.5$ and $\gamma = 1$. The right panel translates those value objects into outside-option thresholds $\bar{\mu}_j^*$ for illustrative inspection cost $\kappa_j^W = 0.015$. Verification is possible only where $\bar{\mu}_j^* \le \mu_j^+$, so the shaded bands mark the outside-option beliefs for which one more inspection is optimal. The even-N_j case is qualitatively similar: the verification region shrinks and eventually exits the feasible belief space as ρ_j approaches one.

A.2 Low-Cost Search Robustness

The benchmark numerical calibration is chosen so that the no-AI consumer performs meaningful but not exhaustive verification. A natural robustness check is to lower within-option search costs enough that raw-signal search clearly begins in the no-AI environment. To do so, keep the benchmark calibration from Table 1 but replace the within-option schedule by

$$c_j^W(m+1) = 0.01 + 0.005m.$$

Under this lower-cost schedule, the exact finite-state Bellman solution implies that the no-AI consumer enters an option and then initiates verification immediately; the post-entry action at state $(m, y) = (0, 0)$ is continuation rather than stopping, and the probability of stopping without search remains zero. The AI consumer, by contrast, still finds the free overview sufficiently informative that expected within-option verification remains zero in this calibration. Table 8 summarizes the implied behavior.

Notably, this is also a calibration in which AI slightly reduces choice accuracy: the no-AI consumer reads enough raw signals to identify the high-quality option more often, while the AI consumer saves on search costs and therefore attains higher ex ante value despite choosing the high-quality option slightly less often.

Table 8: Robustness to Lower Within-Option Search Costs

Outcome	No AI	AI
Ex ante value	0.304	0.397
Actual expected payoff	0.456	0.470
Expected options entered	1.905	1.750
Expected raw signals opened	4.734	0.000
Probability chosen option is high quality	0.743	0.732
Probability of stopping without search	0.000	0.000

Notes: The only change relative to the benchmark numerical calibration is the within-option cost schedule $c_j^W(m+1) = 0.01 + 0.005m$. In this calibration, the no-AI consumer starts within-option verification immediately after entry, while the AI consumer still stops after the overview. Choice quality is slightly higher in the no-AI environment because deeper verification is informative, but ex ante value remains higher with AI because the summary saves substantial search cost.

A.3 Derivations

A.4 Majority-rule overview precision

This section derives equation (2.10). Conditional on the latent quality state θ_j, define the indicator that signal k is correct:

$$X_{jk} = 1\{r_{jk} = \theta_j\}. \tag{A.4}$$

By assumption,

$$\mathbb{P}(X_{jk} = 1 \mid \theta_j) = \rho_j, \tag{A.5}$$

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

$$C_j = \sum_{k=1}^{N_j} X_{jk}, \tag{A.6}$$

is binomial:

$$C_j \mid \theta_j \sim \text{Binomial}(N_j, \rho_j). \tag{A.7}$$

Under majority aggregation with odd N_j, the AI overview is correct if and only if a strict majority of the underlying signals are correct, namely if

$$C_j \geq \frac{N_j + 1}{2}. \tag{A.8}$$

Therefore

$$\zeta_j = \mathbb{P}(a_j = \theta_j \mid \theta_j) = \mathbb{P}\left(C_j \geq \frac{N_j + 1}{2} \mid \theta_j\right). \tag{A.9}$$

Using the binomial probability mass function,

$$\mathbb{P}(C_j = m \mid \theta_j) = \binom{N_j}{m} \rho_j^m (1 - \rho_j)^{N_j - m}, \tag{A.10}$$

and summing over all values of m for which the majority is correct yields

$$\zeta_j = \sum_{m=(N_j+1)/2}^{N_j} \binom{N_j}{m} \rho_j^m (1 - \rho_j)^{N_j - m}, \tag{A.11}$$

which is equation (2.10) in the main text.

Theorem 1 (Condorcet improvement under majority aggregation)**.** *Let $N_j > 1$ be odd and suppose the underlying signals are conditionally independent, with each signal correct with probability $\rho_j \in (1/2, 1)$. Let ζ_j denote the probability that the majority-rule overview is correct. Then*

$$\zeta_j > \rho_j. \tag{A.12}$$

More generally, if

$$h_N(p) = \sum_{m=(N+1)/2}^{N} \binom{N}{m} p^m (1 - p)^{N - m}, \tag{A.13}$$

then for every odd $N > 1$, $h_N(p) > p$ for $p > 1/2$, $h_N(1/2) = 1/2$, and $h_N(p) < p$ for $p < 1/2$.

Proof. This is the standard Condorcet-jury-theorem result for majority aggregation of independent binary signals that are each more likely than not to be correct. In the notation of this paper, $p = \rho_j$, $N = N_j$, and $h_N(p) = \zeta_j$. Therefore odd-N_j majority aggregation is strictly more accurate than a single underlying signal whenever $\rho_j > 1/2$. See Austen-Smith and Banks (1996) and Boland (1989). □

A.5 No-AI posterior after m signals

This subsection derives equation (2.11). In the no-AI environment, the consumer learns only from costly raw-signal inspection. Fix an option j and suppress the option index when no confusion arises. 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

$$\mathbb{P}(\theta_j = 1 \mid r_1 = 1) = \frac{\mathbb{P}(r_1 = 1 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1)}{\mathbb{P}(r_1 = 1 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1) + \mathbb{P}(r_1 = 1 \mid \theta_j = 0)\mathbb{P}(\theta_j = 0)} \tag{A.14}$$

$$= \frac{\mu_{j0}\rho_j}{\mu_{j0}\rho_j + (1 - \mu_{j0})(1 - \rho_j)}. \tag{A.15}$$

After one negative signal,

$$\mathbb{P}(\theta_j = 1 \mid r_1 = 0) = \frac{\mathbb{P}(r_1 = 0 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1)}{\mathbb{P}(r_1 = 0 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1) + \mathbb{P}(r_1 = 0 \mid \theta_j = 0)\mathbb{P}(\theta_j = 0)} \tag{A.16}$$

$$= \frac{\mu_{j0}(1 - \rho_j)}{\mu_{j0}(1 - \rho_j) + (1 - \mu_{j0})\rho_j}. \tag{A.17}$$

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

$$\rho_j^y (1 - \rho_j)^{m-y}, \tag{A.18}$$

while conditional on $\theta_j = 0$ it is

$$(1 - \rho_j)^y \rho_j^{m-y}. \tag{A.19}$$

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

$$\mathbb{P}(Y = y \mid \theta_j = 1, m) = \binom{m}{y} \rho_j^y (1 - \rho_j)^{m-y}, \tag{A.20}$$

$$\mathbb{P}(Y = y \mid \theta_j = 0, m) = \binom{m}{y} (1 - \rho_j)^y \rho_j^{m-y}. \tag{A.21}$$

Applying Bayes' rule,

$$\mathbb{P}(\theta_j = 1 \mid m, y) = \frac{\mu_{j0}\binom{m}{y}\rho_j^y(1-\rho_j)^{m-y}}{\mu_{j0}\binom{m}{y}\rho_j^y(1-\rho_j)^{m-y} + (1-\mu_{j0})\binom{m}{y}(1-\rho_j)^y\rho_j^{m-y}} \tag{A.22}$$

$$= \frac{\mu_{j0}\rho_j^y(1-\rho_j)^{m-y}}{\mu_{j0}\rho_j^y(1-\rho_j)^{m-y} + (1-\mu_{j0})(1-\rho_j)^y\rho_j^{m-y}}. \tag{A.23}$$

This is equation (2.11). 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) rather than on the order in which the signals arrived.

A.6 AI posterior after an overview and partial inspection

This subsection derives equation (2.16). In the AI environment, the consumer first observes the overview realization $a_j \in \{0, 1\}$ 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 an option j and suppress the option index when no confusion arises. Let the majority threshold be

$$T_j = \frac{N_j + 1}{2}. \tag{A.24}$$

Given an overview outcome $a \in \{0, 1\}$ and an inspected history (m, y), define

$$H_{ja}(m, y; p) = \begin{cases} \mathbb{P}(\text{Bin}(N_j - m, p) \geq T_j - y), & a = 1, \\ \mathbb{P}(\text{Bin}(N_j - m, p) < T_j - y), & a = 0. \end{cases} \tag{A.25}$$

This is the probability that the remaining uninspected signals complete the finite signal set in a way that is consistent with the 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_j. Under $\theta_j = 1$, a positive signal is the correct realization, so

$$H_{j1}(m,y;\rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, \rho_j) \geq T_j - y)\,, \tag{A.26}$$

$$H_{j0}(m,y;\rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, \rho_j) < T_j - y)\,. \tag{A.27}$$

Under $\theta_j = 0$, a positive signal is an incorrect realization, so

$$H_{j1}(m,y;1-\rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, 1-\rho_j) \geq T_j - y)\,, \tag{A.28}$$

$$H_{j0}(m,y;1-\rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, 1-\rho_j) < T_j - y)\,. \tag{A.29}$$

So the H-function is just a compact way to write 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 $\theta_j = 1$, a favorable overview occurs with probability

$$L_{j1}(1,0,0) = H_{j1}(0,0;\rho_j) = \zeta_j, \tag{A.30}$$

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

$$L_{j0}(1,0,0) = H_{j1}(0,0;1-\rho_j) = 1 - \zeta_j. \tag{A.31}$$

Hence

$$\widetilde{\mu}_j^{AI}(1,0,0) = \frac{\mu_{j0}\zeta_j}{\mu_{j0}\zeta_j + (1-\mu_{j0})(1-\zeta_j)}, \tag{A.32}$$

and analogously for $a = 0$.

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

$$\binom{m}{y} \rho_j^y (1-\rho_j)^{m-y}. \tag{A.33}$$

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

$$H_{ja}(m,y;\rho_j). \tag{A.34}$$

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

$$L_{j1}(a,m,y) = \binom{m}{y} \rho_j^y (1-\rho_j)^{m-y} H_{ja}(m,y;\rho_j). \tag{A.35}$$

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

$$\binom{m}{y}(1-\rho_j)^y \rho_j^{m-y}, \tag{A.36}$$

and the consistency of the remaining uninspected signals contributes

$$H_{ja}(m, y; 1-\rho_j), \tag{A.37}$$

so

$$L_{j0}(a, m, y) = \binom{m}{y}(1-\rho_j)^y \rho_j^{m-y} H_{ja}(m, y; 1-\rho_j). \tag{A.38}$$

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

$$\widetilde{\mu}_j^{AI}(a, m, y) = \mathbb{P}(\theta_j = 1 \mid a, m, y) \tag{A.39}$$

$$= \frac{\mu_{j0} L_{j1}(a, m, y)}{\mu_{j0} L_{j1}(a, m, y) + (1-\mu_{j0}) L_{j0}(a, m, y)}. \tag{A.40}$$

This is equation (2.16). Relative to the no-AI posterior, the new object is the H-function: it enforces joint consistency between the inspected signals and the overview.

A.7 Next-signal probabilities in the no-AI and AI environments

This subsection derives equations (2.12) and (2.18).

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

$$\widetilde{\pi}_j^+(m, y) = \mathbb{P}(r_{m+1} = 1 \mid m, y) \tag{A.41}$$

$$= \mathbb{P}(r_{m+1} = 1 \mid \theta_j = 1, m, y)\mathbb{P}(\theta_j = 1 \mid m, y) + \mathbb{P}(r_{m+1} = 1 \mid \theta_j = 0, m, y)\mathbb{P}(\theta_j = 0 \mid m, y) \tag{A.42}$$

$$= \widetilde{\mu}_j(m, y)\rho_j + (1 - \widetilde{\mu}_j(m, y))(1-\rho_j), \tag{A.43}$$

which is equation (2.12).

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 $N_j - m$ 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 $N_j - m - 1$ 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

$$H_{ja}(m+1, y+1; p). \tag{A.44}$$

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

$$H_{ja}(m, y; p). \tag{A.45}$$

Hence

$$\psi_{ja}(m, y; p) = \mathbb{P}(r_{m+1} = 1 \mid a, m, y;\ p) = p\frac{H_{ja}(m+1, y+1; p)}{H_{ja}(m, y; p)}, \tag{A.46}$$

which is equation (2.17).

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

$$\begin{aligned} \pi_j^+(a, m, y) &= \mathbb{P}(r_{m+1} = 1 \mid a, m, y) && \text{(A.47)} \\ &= \mathbb{P}(r_{m+1} = 1 \mid a, m, y, \theta_j = 1)\mathbb{P}(\theta_j = 1 \mid a, m, y) \\ &\quad + \mathbb{P}(r_{m+1} = 1 \mid a, m, y, \theta_j = 0)\mathbb{P}(\theta_j = 0 \mid a, m, y) && \text{(A.48)} \\ &= \widetilde{\mu}_j^{AI}(a, m, y)\psi_{ja}(m, y; \rho_j) + \left(1 - \widetilde{\mu}_j^{AI}(a, m, y)\right)\psi_{ja}(m, y; 1 - \rho_j), && \text{(A.49)} \end{aligned}$$

which is equation (2.18).

The likelihood-normalized version (2.19) follows by substituting the posterior formula into (2.18) 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:

$$
\begin{aligned}
\pi_j^+(a,m,y) &= \frac{A_{j1}(a,m,y)+A_{j0}(a,m,y)}{D_{j1}(a,m,y)+D_{j0}(a,m,y)},\\
A_{j1}(a,m,y) &= \mu_{j0}\rho_j^y(1-\rho_j)^{m-y}\rho_j H_{ja}(m+1,y+1;\rho_j),\\
A_{j0}(a,m,y) &= (1-\mu_{j0})(1-\rho_j)^y\rho_j^{m-y}(1-\rho_j)H_{ja}(m+1,y+1;1-\rho_j),\\
D_{j1}(a,m,y) &= \mu_{j0}\rho_j^y(1-\rho_j)^{m-y}H_{ja}(m,y;\rho_j),\\
D_{j0}(a,m,y) &= (1-\mu_{j0})(1-\rho_j)^y\rho_j^{m-y}H_{ja}(m,y;1-\rho_j).
\end{aligned}
\tag{A.50}
$$

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 $\psi_{ja}(m,y;p)$ is needed at latent states that have zero posterior probability.

A.8 Proofs

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

$$
H^3_{ja}(m,y;p) = \mathbb{P}(y+B \in \mathcal{C}_{ja}), \qquad B \sim \mathrm{Bin}(N_j - m, p),
$$

which gives (2.21)–(2.22).

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

$$
H^3_{ja}(m,y;p) = pH^3_{ja}(m+1,y+1;p) + (1-p)H^3_{ja}(m+1,y;p).
$$

This is the only probabilistic identity used to construct ψ_{ja} and the unconditional transition probability. Therefore all posterior and transition formulas from (2.16)–(2.19) apply after replacing H_{ja} by H^3_{ja} and allowing a to range over the three cells. The martingale property follows from the law of iterated expectations exactly as in Lemma 1, and the one-step cutoff in Proposition 2 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 $\mathcal{C}_{j\mathsf{m}} = \{T_j^{lo}+1, \ldots, T_j^{hi}-1\}$. If $T_j^{lo} + T_j^{hi} = N_j$, then $y \in \mathcal{C}_{j\mathsf{m}}$ if and only if $N_j - y \in \mathcal{C}_{j\mathsf{m}}$. If $Y \sim \text{Bin}(N_j, \rho_j)$, then $N_j - Y \sim \text{Bin}(N_j, 1-\rho_j)$. Hence

$$H^3_{j\mathsf{m}}(0,0;\rho_j) = \mathbb{P}(Y \in \mathcal{C}_{j\mathsf{m}}) = \mathbb{P}(N_j - Y \in \mathcal{C}_{j\mathsf{m}}) = H^3_{j\mathsf{m}}(0,0;1-\rho_j).$$

Substituting this equality into Bayes' rule gives $\widetilde{\mu}_j^{AI}(\mathsf{m},0,0) = \mu_{j0}$ for any interior prior. Finally, as $\rho_j \to 1$, the high-quality distribution of Y_j converges to a point mass at N_j and the low-quality distribution converges to a point mass at zero. Since the mixed interval excludes both endpoints, $\mathbb{P}(a_j = \mathsf{m} \mid \theta_j) \to 0$. The positive cell contains N_j and the negative cell contains zero, so Bayes' rule gives $\widetilde{\mu}_j^{AI}(+,0,0) \to 1$ and $\widetilde{\mu}_j^{AI}(-,0,0) \to 0$ for any $\mu_{j0} \in (0,1)$. $\square$

Proof of Proposition 2. Fix overview realization a. With at most one post-summary inspection, the consumer either stops immediately and receives

$$S_j(\mu_j^a; \bar{M}) = \max\{0, \delta_j(\mu_j^a), \bar{M}\},$$

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

$$-\kappa_j^W + \pi_j^a S_j(\mu_j^{a+}; \bar{M}) + (1-\pi_j^a) S_j(\mu_j^{a-}; \bar{M}).$$

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

$$\kappa_j^W \le \pi_j^a S_j(\mu_j^{a+}; \bar{M}) + (1-\pi_j^a) S_j(\mu_j^{a-}; \bar{M}) - S_j(\mu_j^a; \bar{M}) = \mathcal{I}_j^a(\bar{M}),$$

which proves the threshold characterization. $\square$

Proof of Proposition 3. By Proposition 2, verification is optimal if and only if $\kappa_j^W \le \mathcal{I}_j^a(\bar{M}_j)$. If this inequality fails, the consumer chooses the best stopping action immediately. The stopping payoff is

$$S_j^a = \max\{0, \delta_j(\mu_j^a), \bar{M}_j\}.$$

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 $\delta_j(\mu_j^a)$, and leaving the current product for the best remaining product is optimal whenever it is attained by $\bar{M}_j$. A fixed tie-breaking rule selects among multiple maximizers. If $\kappa_j^W \leq \mathcal{I}_j^a(\bar{M}_j)$, 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. $\square$

Proof of Proposition 4. Fix overview realization $a \in \{0, 1\}$. With at most one post-summary inspection, the consumer faces three terminal actions: keep the current option with expected payoff $\delta_j(\mu)$ at whatever posterior μ she holds, take the outside option with payoff $\bar{M}$, or choose nothing with payoff 0. Since the normalized no-purchase option is weakly dominated by assumption, the relevant comparison at each terminal node is between $\delta_j(\mu)$ and $\bar{M}$.

Closed-form posterior and transition objects. Equation (2.16) implies that the posterior after the overview alone is

$$\mu_j^a = \frac{\mu_{j0} H_a(0, 0; \rho_j)}{\mu_{j0} H_a(0, 0; \rho_j) + (1 - \mu_{j0}) H_a(0, 0; 1 - \rho_j)}.$$

After one additional favorable inspected signal,

$$\mu_j^{a+} = \frac{\mu_{j0} \rho_j H_a(1, 1; \rho_j)}{\mu_{j0} \rho_j H_a(1, 1; \rho_j) + (1 - \mu_{j0})(1 - \rho_j) H_a(1, 1; 1 - \rho_j)}.$$

After one additional unfavorable inspected signal,

$$\mu_j^{a-} = \frac{\mu_{j0}(1 - \rho_j) H_a(1, 0; \rho_j)}{\mu_{j0}(1 - \rho_j) H_a(1, 0; \rho_j) + (1 - \mu_{j0}) \rho_j H_a(1, 0; 1 - \rho_j)}.$$

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

$$\pi_j^a = \frac{\mu_{j0} \rho_j H_a(1, 1; \rho_j) + (1 - \mu_{j0})(1 - \rho_j) H_a(1, 1; 1 - \rho_j)}{\mu_{j0} H_a(0, 0; \rho_j) + (1 - \mu_{j0}) H_a(0, 0; 1 - \rho_j)}.$$

Thus the posteriors and transition probabilities are available in closed form for every finite $N_j > 1$, with even N_j using the tie-adjusted analogue from Appendix A.1; 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_j^W to inspect one signal, the next signal is favorable with probability π_j^a and unfavorable with probability $1 - \pi_j^a$. After a favorable signal the posterior rises to μ_j^{a+}; after an unfavorable signal it falls to μ_j^{a-}.

Therefore the expected payoff from starting verification is

$$\mathcal{V} = -\kappa_j^W + \pi_j^a \max\{\delta_j(\mu_j^{a+}), \bar{M}\} + (1-\pi_j^a)\max\{\delta_j(\mu_j^{a-}), \bar{M}\}. \tag{V}$$

Part (a). Suppose the ordering

$$\mu_j^{a-} < \bar{\mu}_j \le \mu_j^a < \mu_j^{a+}.$$

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

$$\delta_j(\mu_j^{a-}) < \bar{M} \le \delta_j(\mu_j^a) < \delta_j(\mu_j^{a+}).$$

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

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

$$\mathcal{V} = -\kappa_j^W + \pi_j^a\delta_j(\mu_j^{a+}) + (1-\pi_j^a)\bar{M}.$$

Verification is optimal if and only if $\mathcal{V} \ge \delta_j(\mu_j^a)$, i.e.,

$$-\kappa_j^W + \pi_j^a\delta_j(\mu_j^{a+}) + (1-\pi_j^a)\bar{M} \ge \delta_j(\mu_j^a).$$

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

$$-\kappa_j^W + \pi_j^a\gamma\mu_j^{a+} + (1-\pi_j^a)\gamma\bar{\mu}_j \ge \gamma\mu_j^a,$$

which rearranges to

$$\kappa_j^W \le \gamma\left[\pi_j^a\mu_j^{a+} + (1-\pi_j^a)\bar{\mu}_j - \mu_j^a\right].$$

By Lemma 1, Bayesian posteriors satisfy the martingale property

$$\mu_j^a = \pi_j^a\mu_j^{a+} + (1-\pi_j^a)\mu_j^{a-}.$$

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

$$\begin{aligned} \gamma\left[\pi_j^a \mu_j^{a+} + (1-\pi_j^a)\bar{\mu}_j - \mu_j^a\right] &= \gamma\left[\mu_j^a - (1-\pi_j^a)\mu_j^{a-} + (1-\pi_j^a)\bar{\mu}_j - \mu_j^a\right] \\ &= \gamma(1-\pi_j^a)(\bar{\mu}_j - \mu_j^{a-}). \end{aligned} \tag{A.51}$$

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

$$\kappa_j^W \le \gamma(1-\pi_j^a)(\bar{\mu}_j - \mu_j^{a-}),$$

which proves part (a).

Part (b). Suppose instead the ordering

$$\mu_j^{a-} < \mu_j^a \le \bar{\mu}_j < \mu_j^{a+}.$$

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

Resolving the max operators in display (V):

$$\mathcal{V} = -\kappa_j^W + \pi_j^a \delta_j(\mu_j^{a+}) + (1-\pi_j^a)\bar{M}.$$

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

$$-\kappa_j^W + \pi_j^a \delta_j(\mu_j^{a+}) + (1-\pi_j^a)\bar{M} \ge \bar{M}.$$

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

$$-\kappa_j^W + \pi_j^a\left[\delta_j(\mu_j^{a+}) - \bar{M}\right] \ge 0.$$

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

$$\kappa_j^W \le \gamma\pi_j^a(\mu_j^{a+} - \bar{\mu}_j),$$

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 $W_j^1(a, m, y; M)$ denote the value of the problem in which at most one additional inspection is allowed, and let $W_j^\infty(a, m, y; M)$ 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

$$W_j^{\infty}(a, m, y; M) \geq W_j^1(a, m, y; M)$$

at every reachable history. In particular, at the post-overview entry state, the one-step verification payoff characterized in Proposition 4 is a lower bound on the unrestricted verification value. The cutoff inequalities in Proposition 4 therefore give sufficient conditions for verification to start in the unrestricted problem as well. Corollary 3 closes the other side at high precision by showing that, for sufficiently large ρ_j, even the unrestricted continuation value is too small to justify any further inspection.

Derivation for Example 1. In the notation of Proposition 4, the worked example sets $a = 1$ and uses the shorthand

$$\mu_j^+ \equiv \mu_j^1, \qquad \mu_j^{++} \equiv \mu_j^{1+}, \qquad \mu_j^{+-} \equiv \mu_j^{1-}, \qquad \pi_j^+ \equiv \pi_j^1.$$

Write $\zeta(\rho_j) = \mathbb{P}(a_j = 1 \mid \theta_j = 1)$. Under majority aggregation with $N_j = 3$,

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

Bayes' rule immediately yields

$$\mu_j^+ = \frac{\mu_{j0}\zeta(\rho_j)}{\mu_{j0}\zeta(\rho_j) + (1 - \mu_{j0})(1 - \zeta(\rho_j))}.$$

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

$$L_{j1}(1, 1, 1) = \rho_j\left[1 - (1 - \rho_j)^2\right] = \rho_j^2(2 - \rho_j),$$

$$L_{j0}(1, 1, 1) = (1 - \rho_j)\left[1 - \rho_j^2\right] = (1 - \rho_j)^2(1 + \rho_j),$$

which gives

$$\mu_j^{++} = \frac{\mu_{j0}\rho_j^2(2 - \rho_j)}{\mu_{j0}\rho_j^2(2 - \rho_j) + (1 - \mu_{j0})(1 - \rho_j)^2(1 + \rho_j)}.$$

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

$$L_{j1}(1, 1, 0) = \rho_j^2(1 - \rho_j), \qquad L_{j0}(1, 1, 0) = \rho_j(1 - \rho_j)^2,$$

so

$$\mu_j^{+-} = \frac{\mu_{j0}\rho_j}{\mu_{j0}\rho_j + (1-\mu_{j0})(1-\rho_j)}.$$

Remark 3 (A notable cancellation in μ_j^{+-}). After a favorable overview ($a = 1$) and an unfavorable inspected signal, the posterior μ_j^{+-} 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 $\rho_j/(1-\rho_j)$, exactly the likelihood ratio of one positive raw signal. Formally, the posterior odds after the favorable overview and unfavorable inspection are

$$\frac{\mu_{j0}L_{j1}(1,1,0)}{(1-\mu_{j0})L_{j0}(1,1,0)} = \frac{\mu_{j0}\rho_j^2(1-\rho_j)}{(1-\mu_{j0})\rho_j(1-\rho_j)^2} = \frac{\mu_{j0}\rho_j}{(1-\mu_{j0})(1-\rho_j)},$$

which is the posterior odds from exactly one positive signal. The cancellation is specific to the $N_j = 3$ 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

$$\pi_j^+ = \frac{\mu_{j0}\rho_j^2(2-\rho_j) + (1-\mu_{j0})(1-\rho_j)^2(1+\rho_j)}{\mu_{j0}\zeta(\rho_j) + (1-\mu_{j0})(1-\zeta(\rho_j))}.$$

Under (3.11), 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

$$-\kappa_j^W + \pi_j^+\delta_j(\mu_j^{++}) + (1-\pi_j^+)\bar{M},$$

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

$$-\kappa_j^W + \pi_j^+\delta_j(\mu_j^{++}) + (1-\pi_j^+)\bar{M} \geq \delta_j(\mu_j^+).$$

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

$$\kappa_j^W \leq \gamma\left[\pi_j^+\mu_j^{++} + (1-\pi_j^+)\bar{\mu}_j - \mu_j^+\right].$$

By Lemma 1,

$$\mu_j^+ = \pi_j^+\mu_j^{++} + (1-\pi_j^+)\mu_j^{+-},$$

so the right-hand side simplifies to

$$\gamma(1-\pi_j^+)(\bar{\mu}_j-\mu_j^{+-}).$$

This proves (3.12). Rearranging yields

$$\bar{\mu}_j^* = \mu_j^{+-} + \frac{\kappa_j^W}{\gamma(1-\pi_j^+)},$$

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

Proof of Lemma 1. Fix any odd $N_j > 1$ and any overview realization $a \in \{0,1\}$. Let $\mu_j^a = \widetilde{\mu}_j^{AI}(a,0,0)$ denote the posterior after the overview alone, and let $\mu_j^{a+} = \widetilde{\mu}_j^{AI}(a,1,1)$ and $\mu_j^{a-} = \widetilde{\mu}_j^{AI}(a,1,0)$ 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,

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

which is exactly (3.14). This is the standard martingale property of Bayesian posteriors: it holds for any finite N_j and any aggregation rule, because the law of iterated expectations does not depend on the signal structure. □

Proof of Lemma 2. Let $N_j = 2s+1$ with $s \geq 1$ and let A denote the event that the majority overview is favorable. Conditional on the high-quality state and on the event that the next inspected signal is negative, the remaining $2s$ signals must contain at least $s+1$ positives for the final overview to remain favorable. Hence

$$\mathbb{P}(r_{j,next} = 0, A \mid \theta_j = 1) = (1-\rho_j)\mathbb{P}(\text{Bin}(2s,\rho_j) \geq s+1) \sim 1-\rho_j,$$

while

$$\mathbb{P}(A \mid \theta_j = 1) = \mathbb{P}(\text{Bin}(2s+1,\rho_j) \geq s+1) \to 1.$$

Therefore

$$\mathbb{P}(r_{j,next} = 0 \mid A, \theta_j = 1) \sim 1-\rho_j.$$

Conditional on the low-quality state, a favorable majority requires at least $s+1$ erroneous positive signals, so $\mathbb{P}(A \mid \theta_j = 0) = O((1-\rho_j)^{s+1})$. Bayes' rule then gives $\mathbb{P}(\theta_j = 0 \mid A) = O((1-\rho_j)^{s+1})$, which is $o(1-\rho_j)$ because $s \geq 1$. Combining the two latent states,

$$1-\pi_j^+(1,0,0) = \mathbb{P}(r_{j,next} = 0 \mid A) \sim 1-\rho_j.$$

This proves (3.15). The unfavorable-overview statement follows by replacing positives with negatives and $\theta_j = 1$ with $\theta_j = 0$. $\square$

Proof of Corollary 1. Fix any finite $N_j > 1$. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$ and replace the odd-N_j consistency terms by their tie-adjusted analogues. Consider first $a = 1$. Conditional on $\theta_j = 1$, a favorable overview has probability tending to one and a negative next inspected signal has probability $1-\rho_j \to 0$. Conditional on $\theta_j = 0$, a favorable overview has probability tending to zero. Since the posterior after a favorable overview satisfies $\widetilde{\mu}_j^{AI}(1,0,0) \to 1$, the unconditional probability of a negative next signal after $a = 1$ satisfies

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

Under the additional simple-zero condition used in Corollary 2, this convergence has the first-order expansion

$$1 - \pi_j^+(1,0,0) \sim C_j^\star(1-\rho_j) \qquad \text{as } \rho_j \to 1. \tag{A.52}$$

for some $C_j^\star > 0$ that depends on the signal count and, when relevant, the tie-breaking rule. 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

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

Thus $d_j^a(\rho_j) \to 0$ for both summary realizations.

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

$$S_j(\mu; \bar{M}) = \max\{0, \delta_j(\mu), \bar{M}\}$$

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 $\mu, \bar{\mu}_j \in [0,1]$; let $\bar{B}_j$ be such a bound. For $a = 1$, $\mu_j^a \to 1$ and $\mu_j^{a+} \to 1$, while the possibly nonconvergent payoff after the disconfirming signal is multiplied by $1 - \pi_j^a \to 0$. Hence

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

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

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

Thus the primitive one-step inspection value converges to zero uniformly over feasible outside-option beliefs. Since $\kappa_j^W > 0$, there exists $\rho_j^* < 1$ such that $\mathcal{I}_j^a(\bar{M}) < \kappa_j^W$ for every $\rho_j > \rho_j^*$, every feasible $\bar{\mu}_j$, and both $a \in \{0,1\}$. Proposition 2 then implies that one-step verification is not optimal. Finally, in the favorable-overview case and in the case-(a) region of Proposition 4, the affine cutoff representation is

$$\bar{\mu}_j^* = \mu_j^{a-} + \frac{\kappa_j^W}{\gamma(1-\pi_j^a)}.$$

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

Proof of Corollary 2. Work under the simple-zero condition in the statement. From Proposition 4(a), the threshold after a favorable overview is

$$\bar{\mu}_j^* = \mu_j^{a-} + \frac{\kappa_j^W}{\gamma(1-\pi_j^a)}.$$

Differentiating gives

$$\frac{\partial \bar{\mu}_j^*}{\partial \rho_j} = \frac{\partial \mu_j^{a-}}{\partial \rho_j} - \frac{\kappa_j^W}{\gamma} \frac{\partial(1-\pi_j^a)/\partial \rho_j}{(1-\pi_j^a)^2}.$$

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

Define $g(\rho_j) = 1 - \pi_j^a$. By the simple-zero condition,

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

Since g is a ratio of finite polynomials with a simple zero at $\rho_j = 1$, this expansion can be differentiated in a neighborhood of one. Therefore $1/g(\rho_j) \sim 1/[C_j^\star(1-\rho_j)]$, which gives

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

and consequently

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

Therefore

$$(1-\rho_j)^2 \frac{\partial}{\partial \rho_j}\left(\frac{\kappa_j^W}{\gamma(1-\pi_j^a)}\right) \to \frac{\kappa_j^W}{\gamma C_j^\star} > 0.$$

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

$$(1-\rho_j)^2\frac{\partial\bar{\mu}_j^*}{\partial\rho_j} \to \frac{\kappa_j^W}{\gamma C_j^\star} > 0,$$

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

Proof of Corollary 3. Fix an overview realization $a_j \in \{0,1\}$. For odd N_j, Theorem 1 implies $\zeta_j \to 1$ as $\rho_j \to 1$; for even N_j, the tie-breaking formula in Appendix A.1 implies the same limit under any fixed $\tau_j \in [0,1]$. Hence the post-overview posterior $\mu_j^a \equiv \widetilde{\mu}_j^{AI}(a,0,0)$ converges to certainty:

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

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_j be a uniform bound on the stopping payoff $S_j(\mu;\bar{M}) = \max\{0, \delta_j(\mu), \bar{M}\}$ over $\mu, \bar{\mu}_j \in [0,1]$. The function S_j is also γ-Lipschitz in μ uniformly in $\bar{M}$.

For any feasible continuation plan σ after overview realization a_j, let E_{a_j,σ,ρ_j} be the event that at least one inspected raw signal contradicts the overview: after $a_j = 1$, at least one inspected signal is zero; after $a_j = 0$, 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 $E^c_{a_j,\sigma,\rho_j}$ is uniformly $o(1)$ over all feasible outside beliefs. On E_{a_j,σ,ρ_j}, the payoff difference is bounded by $2B_j$. Hence the gross gain from σ satisfies

$$\text{gross gain from } \sigma \leq o(1) + 2B_j\mathbb{P}(E_{a_j,\sigma,\rho_j}),$$

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

That uniformity follows from finiteness. With finite inspection horizon $\bar{m}_j$ 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 $\mathbb{P}(E_{a_j,\sigma,\rho_j})$ to zero for each pure feasible plan therefore implies

$$\sup_\sigma \mathbb{P}(E_{a_j,\sigma,\rho_j}) \to 0$$

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_j 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

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

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

The first additional inspection nevertheless costs at least $c_j^W(1) > 0$. Hence there exists $\rho_j^* < 1$ such that for every $\rho_j > \rho_j^*$, every feasible outside-option belief $\bar{\mu}_j \in [0,1]$, 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 $\bar{\mu}_j^*$ after a favorable overview, the same conclusion is equivalently stated as $\bar{\mu}_j^* > 1$ for all $\rho_j > \rho_j^*$. By symmetry, after an unfavorable overview the analogous lower continuation threshold lies below zero. □

Proof of Benchmark Observation. A strategy for the consumer is a mapping σ from each history

$$\big(q, \mathcal{V}_t, \{a_{jt}, m_{jt}, y_{jt}\}_{j\in\mathcal{V}_t}, \mathcal{U}_t\big)$$

to a probability distribution over the three action classes:

$$\{\text{stop and choose, inspect within a visited option, enter an unvisited option}\}.$$

Call a strategy *admissible* if it is measurable with respect to the history and the resulting expected total search cost is finite.

Fix any admissible strategy σ^N in the environment without an AI summary. In the environment with an AI summary, the consumer can adopt a strategy $\tilde{\sigma}$ that ignores the realization of the summary and replicates exactly the continuation, stopping, and raw-signal-acquisition choices prescribed by σ^N. Since the summary is free, $\tilde{\sigma}$ is feasible and yields exactly the same expected payoff as σ^N. Therefore

$$V^A(s) \geq V^N(s).$$

If there exists a set of histories with positive probability on which conditioning on the summary changes the consumer's optimal action, then the inequality is strict. □

Proof of Proposition 5. Let

$$\mathcal{S}_j(a,m,y;q) = \max\{0, \delta_j(\widetilde{\mu}_j^{AI}(a,m,y)), M_j^A(q)\}$$

denote the best stopping payoff at state (a,m,y). By the Bellman equation, continuation is optimal if and only if

$$\mathcal{K}_j(a,m,y;M_j^A(q)) \geq \mathcal{S}_j(a,m,y;q).$$

Adding $c_j^W(m+1)$ to both sides gives

$$c_j^W(m+1) \leq \mathcal{K}_j(a,m,y;M_j^A(q)) + c_j^W(m+1) - \mathcal{S}_j(a,m,y;q)$$

and the right-hand side is exactly $\Delta_j^A(a,m,y;q) = \bar{c}_j(a,m,y;q)$. Hence continuation is optimal if and only if

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

Holding the continuation value function fixed, a higher current marginal cost moves the left-hand side upward without changing the cutoff, so it weakly lowers the incentive to continue at that state. □

B Appendix to Part II: Multi-Dimensional Extension

B.1 Derivations for the multi-dimensional extension

This subsection expands the updating, continuation, and taste-region objects from Section 4 step by step. The goal is to show exactly where the new finite-pool topic-arrival friction enters relative to Part I and to record the closed-form objects available at a fixed history.

Proof of the exact dynamic initiation cutoff

Suppress the option index. 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 6. 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 an option j and suppress the option index when no confusion arises. Let the post-summary history after t opened reviews be (a, h_t), where:

- $a = (a_d)_{d\in\mathcal{S}^K}$ is the realized top-K summary 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

$$\mathbb{P}(\theta, M \mid a, h_t) = \frac{\mathbb{P}(a, h_t \mid \theta, M)\mathbb{P}(\theta, M)}{\sum_{\theta'}\sum_{M'}\mathbb{P}(a, h_t \mid \theta', M')\mathbb{P}(\theta', M')}. \tag{B.1}$$

By prior independence between latent quality and coverage,

$$\mathbb{P}(\theta, M) = \mathbb{P}(\theta)\, G(M). \tag{B.2}$$

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

$$\mathcal{L}(a, h_t \mid \theta, M) \equiv \mathbb{P}(a, h_t \mid \theta, M) = \mathbf{1}\{\mathcal{S}^K(M) = \mathcal{S}^K\} \sum_{R \in \mathcal{R}(M; h_t)} \mathbb{P}(R \mid \theta, M)\mathbb{P}(a \mid R, M). \tag{B.3}$$

The probability $\mathbb{P}(a \mid R, M)$ is one or zero under deterministic tie-breaking and includes the appropriate τ_d factors under probabilistic tie-breaking. 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 summary and the opened reviews.

Substituting (B.2) and (B.3) into (B.1) yields

$$\mathbb{P}(\theta, M \mid a, h_t) \propto \mathbb{P}(\theta)\, G(M)\, \mathbf{1}\{\mathcal{S}^K(M) = \mathcal{S}^K\}\mathcal{L}(a, h_t \mid \theta, M), \tag{B.4}$$

which is equation (4.12) in the main text.

The marginal posterior on quality dimension d is obtained by summing over all latent states and coverage matrices consistent with $\theta_d = 1$:

$$\mu_{d,t}(a, h_t) = \mathbb{P}(\theta_d = 1 \mid a, h_t) = \sum_M \sum_{\theta_{-d}} \mathbb{P}(\theta_d = 1, \theta_{-d}, M \mid a, h_t). \tag{B.5}$$

This is the dimension-level posterior appearing in (4.11). 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.

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

$$C_d(M) = \sum_{n=1}^{N} m_{nd} \qquad \text{and} \qquad c_{d,t} = \sum_{n \in \text{opened at } t} m_{nd}. \tag{B.6}$$

Since there are $N - t$ unopened reviews,

$$\mathbb{P}(m_{n_{t+1},d} = 1 \mid a, h_t, M) = \frac{C_d(M) - c_{d,t}}{N - t}. \tag{B.7}$$

Applying the law of iterated expectations over the posterior distribution of M gives

$$\eta_{d,t}(a, h_t) = \mathbb{P}(m_{n_t+1,d} = 1 \mid a, h_t) \tag{B.8}$$

$$= \sum_M \mathbb{P}(m_{n_t+1,d} = 1 \mid a, h_t, M)\mathbb{P}(M \mid a, h_t) \tag{B.9}$$

$$= \sum_M \frac{C_d(M) - c_{d,t}}{N - t}\mathbb{P}(M \mid a, h_t) \tag{B.10}$$

$$= \mathbb{E}\left[\frac{C_d(M) - c_{d,t}}{N - t} \,\middle|\, a, h_t\right], \tag{B.11}$$

which is equation (4.16).

Predictive sign probability and unconditional relevance

Now condition on the event that the next unopened review covers dimension d. Under the maintained separability condition

$$\mathbb{P}(\theta_d = 1 \mid m_{n_t+1,d} = 1, a, h_t) = \mathbb{P}(\theta_d = 1 \mid a, h_t) = \mu_{d,t}(a, h_t), \tag{B.12}$$

conditioning on coverage of the next review does not further shift beliefs about latent quality on that same dimension. Therefore

$$\pi^+_{d,t}(a, h_t) = \mathbb{P}(r_{n_t+1,d} = 1 \mid m_{n_t+1,d} = 1, a, h_t) \tag{B.13}$$

$$= \rho_d\, \mu_{d,t}(a, h_t) + (1 - \rho_d)\big(1 - \mu_{d,t}(a, h_t)\big), \tag{B.14}$$

which is equation (4.18). Finally, the factorization below uses the same maintained restriction. Without (B.12), the event $m_{n_t+1,d} = 1$ would itself shift beliefs about θ_d, so $\eta_{d,t}(a, h_t)$ would no longer be a pure coverage object and the predictive term would need to be written directly as $\mathbb{P}(r_{n_t+1,d} = 1 \mid a, h_t)$ rather than as a product of separate coverage and sign components:

$$\psi^+_{d,t}(a, h_t) = \mathbb{P}(m_{n_t+1,d} = 1, r_{n_t+1,d} = 1 \mid a, h_t) \tag{B.15}$$

$$= \mathbb{P}(m_{n_t+1,d} = 1 \mid a, h_t)\mathbb{P}(r_{n_t+1,d} = 1 \mid m_{n_t+1,d} = 1, a, h_t) \tag{B.16}$$

$$= \eta_{d,t}(a, h_t)\, \pi^+_{d,t}(a, h_t), \tag{B.17}$$

which is equation (4.19). At a fixed history, both $\eta_{d,t}$ and $\pi^+_{d,t}$ are finite sums, so $\psi^+_{d,t}$ is also available in closed form.

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 7. 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

$$S(s_t; \gamma) = \max\{0, \delta(a, h_t), \bar{M}\} \tag{B.18}$$

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, so the lower-bound policy ignores any other realized contents of the review and stops at $S(s_t; \gamma)$;
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

$$
\begin{aligned}
\underline{\mathcal{K}}_d(s_t) = &-c^W(t+1) + (1-\eta_{d,t}(s_t))S(s_t;\gamma) \\
&+ \eta_{d,t}(s_t)\pi^+_{d,t}(s_t)S^+_d(s_t) + \eta_{d,t}(s_t)(1-\pi^+_{d,t}(s_t))S^-_d(s_t),
\end{aligned}
\tag{B.19}
$$

which is equation (4.26). Subtracting $S(s_t;\gamma)$ from both sides gives

$$
\begin{aligned}
\underline{\mathcal{K}}_d(s_t) - S(s_t;\gamma) = &-c^W(t+1) + \eta_{d,t}(s_t)\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big] \\
&+ \eta_{d,t}(s_t)(1-\pi^+_{d,t}(s_t))\big[S^-_d(s_t) - S(s_t;\gamma)\big],
\end{aligned}
\tag{B.20}
$$

so a sufficient condition for continuation is

$$
c^W(t+1) \le \eta_{d,t}(s_t)\left(\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big] + (1-\pi^+_{d,t}(s_t))\big[S^-_d(s_t) - S(s_t;\gamma)\big]\right), \tag{B.21}
$$

which is exactly (4.27). 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. This proves Proposition 8.

Analytical taste-region characterization

The focal one-step sufficient taste region in (5.2) is just a restatement of (B.21). Define

$$
\mathcal{B}_d(s_t;\gamma) = \eta_{d,t}(s_t)\left(\pi^+_{d,t}(s_t)\big[S^+_d(s_t) - S(s_t;\gamma)\big] + (1-\pi^+_{d,t}(s_t))\big[S^-_d(s_t) - S(s_t;\gamma)\big]\right). \tag{B.22}
$$

Then the inspection region is

$$
\mathcal{T}^I_d(s_t) = \left\{\gamma \in \mathbb{R}^D_+ : c^W(t+1) \le \mathcal{B}_d(s_t;\gamma)\right\}, \tag{B.23}
$$

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

$$
S^+_d(s_t) - S(s_t;\gamma) = \gamma_d\Delta^+_{d,t}(s_t), \qquad S^-_d(s_t) - S(s_t;\gamma) = \gamma_d\Delta^-_{d,t}(s_t), \tag{B.24}
$$

substituting (B.24) into (B.22) yields

$$
\mathcal{B}_d(s_t;\gamma) = \gamma_d\eta_{d,t}(s_t)\left(\pi^+_{d,t}(s_t)\Delta^+_{d,t}(s_t) + (1-\pi^+_{d,t}(s_t))\Delta^-_{d,t}(s_t)\right). \tag{B.25}
$$

Hence, whenever the bracketed term is positive,

$$\gamma_d \geq \gamma_d^*(s_t) = \frac{c^W(t+1)}{\eta_{d,t}(s_t)\left(\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t)\right)}. \tag{B.26}$$

This is the closed-form local taste threshold reported in (5.5).

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

$$\Omega_{d,t}(s_t) = \eta_{d,t}(s_t)\left[\pi_{d,t}^+(s_t)\Delta_{d,t}^+(s_t) + (1-\pi_{d,t}^+(s_t))\Delta_{d,t}^-(s_t)\right].$$

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 (5.7).

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.

Coverage-residual ordering under homogeneous omitted dimensions

The qualified ordering in Proposition 11 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

$$\mathcal{B}^U(s_t;\gamma) = \sum_{d \notin \mathcal{S}^K} \gamma_d \eta_{d,t}(s_t) B_t(s_t) \tag{B.27}$$

$$= B_t(s_t)\mathrm{RL}_t(\gamma; a, h_t). \tag{B.28}$$

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 largest cost at which one-step verification is optimal is weakly lower for γ. 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 restriction, the same scalar RL need not be sufficient because dimensions may differ in precision, posterior uncertainty, and decision-changing payoff wedges.

Proof of the endpoint propositions

Proof of Proposition 9. 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 10. 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 Part I 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. □

References

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.

Anderson, Simon P., and Régis Renault. 2006. "Advertising Content." *American Economic Review* 96(1): 93–113.

Branco, Fernando, Monic Sun, and J. Miguel Villas-Boas. 2012. "Optimal Search for Product Information." *Management Science* 58(11): 2037–2056.

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.

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.

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.

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.

McCall, J. J. 1970. “Economics of Information and Job Search.” *Quarterly Journal of Economics* 84(1): 113–126.

Nitzan, Shmuel, and Jacob Paroush. 1982. “Optimal Decision Rules in Uncertain Dichotomous Choice Situations.” *International Economic Review* 23(2): 289–297.

Stigler, George J. 1961. “The Economics of Information.” *Journal of Political Economy* 69(3): 213–225.

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.

Weitzman, Martin L. 1979. “Optimal Search for the Best Alternative.” *Econometrica* 47(3): 641–654.