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

Dheer Avashia

April 16, 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 general signal-set sizes, that better underlying reviews lead to *less* review-reading when AI summaries are present: the verification threshold diverges as raw-signal precision approaches complete certainty, and 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.

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.

That observation motivates the central question of this paper: how does a free AI summary change costly search and downstream choice? Because the AI summary compresses the same evidence that consumers could unpack at a cost, it does not merely add a signal; it shifts search between two distinct margins. Within an option, the summary can substitute for costly verification: a consumer who reads a well-compressed AI overview may rationally skip the underlying reviews. Across options, the attention freed by that substitution can expand consideration: the consumer enters more options because the effective per-option screening cost falls. How those margins move, and whether the net effect improves outcomes for consumers, depends on the precision and quantity of the underlying evidence being summarized.

The within-option margin is where the model's analytical contribution lies. I introduce a compression-consistent information structure—the H-function—that tracks the dependence between an AI summary and partial inspection of the same finite evidence pool, and use it to derive exact posteriors and a closed-form cost threshold for verification initiation. That threshold characterizes the maximum inspection cost at which a consumer who has already seen the summary still finds it worthwhile to open one more underlying review. It is a new object because the information structure is new: unlike standard Bayesian updating with independent signals, the summary and the inspected signals are deterministically linked through a finite evidence set, and the H-function enforces that consistency.

The verification threshold diverges in the limits, as signal precision or the number of summarized signals grows, but the more interesting economic results lie at intermediate values, where a stronger evidence base simultaneously improves the overview and raises the value of each individual review if inspected. Those competing forces make the threshold hump-shaped in signal precision, and the paper characterizes how that within-option tension interacts with across-option entry in a fully solved nested model. 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. These comparative statics generate testable predictions. Platforms introducing AI summaries should see the largest decline in within-option click-depth for products whose reviews are most individually informative. Consideration sets should expand most where the summary makes initial screening cheapest.

This paper contributes to the search literature by embedding a compression-based information structure into a dynamic search environment with both across-option consideration and within-option verification. Standard sequential-search models (McCall, 1970; Stigler, 1961; Weitzman, 1979) focus on which alternative to sample next and when to stop; entering an option reveals its value. Here, entering an option reveals only a free summary, after which the consumer may pay to inspect the underlying signals. The model therefore separates costly consideration from final choice in a way that fits the broader consideration-set literature (Hauser and Wernerfelt, 1990) and recent evidence that reviews matter at the consideration stage (Gavilan, Avello, and Martinez-Navarro, 2018; Hu and Yang, 2020).

The paper is also related to Bayesian learning models of costly product inspection, including Branco, Sun, and Villas-Boas (2012) and Ke, Shen, and Villas-Boas (2016). It is likewise related to two-stage information-acquisition environments (Anderson and Renault, 2006; Gibbard, 2022) and to work on algorithmic advice (Bundorf, Polyakova, and Tai-Seale, 2024; Dietvorst, Simmons, and Massey, 2015; Logg, Minson, and Moore, 2019). The no-summary benchmark shares structure with Branco, Sun, and Villas-Boas (2012), though the present paper works with finite binary signals rather than continuous-time Brownian learning. The overview-then-inspection architecture is related to Gibbard (2022), who studies two-stage search with additive learning, and to Ke, Shen, and Villas-Boas (2016), who studies multi-product learning with switching. My specification differs because the first signal is a compression of the same finite evidence set that later inspection unpacks, creating a distinct AI-versus-no-AI comparison. The adjacent question of how much information to reveal for free before costly search, studied by Anderson and Renault (2006) in an advertising context, is closely connected to the endogenous-summary extension discussed in the conclusion.

2 Model

2.1 Environment

Each option $j \in \{1, \dots, J\}$ 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]

[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

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

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

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. 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. 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.

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.

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

No-AI environment: observe $(\mathbf{x}_j, p_j) \rightarrow$ pay $c^A(1)$ to enter option $j \rightarrow$ no summary arrives $\rightarrow$ 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.

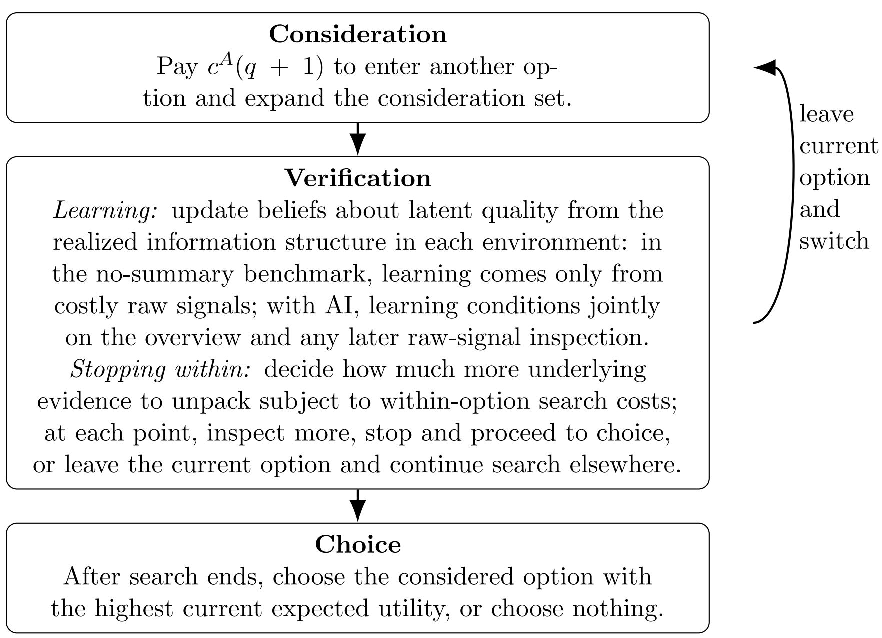

Figure 2: High-level structure of the model. Consideration determines which options enter the set, verification governs learning and within-option evidence unpacking, and choice occurs only after search stops.

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:

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

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 naively, 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,

[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.

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 B.1. 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 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), 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

[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.

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 probablities.

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:

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

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

Step-by-step derivations of (2.11) *and* (2.12) *appear in Appendix B.2 and Appendix B.4.* 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

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

[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.

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*

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

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

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

[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.

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 B.3.*

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 on the same reachable histories for which $H_{ja}(m, y; p) > 0$; 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 favorable is therefore

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

A step-by-step derivation of (2.17) and (2.18) appears in Appendix B.4.

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 $(m+1)$-

[10] This ratio is well-defined on reachable histories for which $H_{ja}(m, y; p) > 0$ 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.

th human signal within option j is

$$c_j^W(m+1), \tag{2.19}$$

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

$$c_j^W(m+1) = \kappa_j^W + \lambda_j^W m, \qquad \kappa_j^W > 0,\ \lambda_j^W \geq 0. \tag{2.20}$$

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

$$c^A(q+1), \tag{2.21}$$

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

$$c^A(q+1) = \kappa^A + \lambda^A q, \qquad \kappa^A > 0,\ \lambda^A \geq 0. \tag{2.22}$$

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:

$$M_j^N(q) = \underbrace{\max_{\ell \neq j} V_\ell^N(q+1) - c^A(q+1)}_{\text{best continuation value from leaving option } j \text{ and continuing search without AI}}. \tag{2.23}$$

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 (in the reduced form incomplete way) isolates the within-option margin without restricting how the across-option problem is solved.

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 ultimate(terminal) condition is

$$\widetilde{W}_j(N_j, y; M_j^N(q)) = \max\{0, \delta_j(\widetilde{\mu}_j(N_j, y)), M_j^N(q)\}, \tag{2.24}$$

and for $m < N_j$,

$$\widetilde{W}_j(m, y; M_j^N(q)) = \max\Big\{0,\ \delta_j(\widetilde{\mu}_j(m, y)),\ M_j^N(q),\ \widetilde{\mathcal{K}}_j(m, y; M_j^N(q))\Big\}, \tag{2.25}$$

where

$$\begin{aligned} \widetilde{\mathcal{K}}_j(m, y; M_j^N(q)) = \quad & \underbrace{-c_j^W(m+1)}_{\text{pay current within-option cost}} \\ & + \underbrace{\widetilde{\pi}_j^+(m, y)\widetilde{W}_j(m+1, y+1; M_j^N(q))}_{\text{continuation value after a positive next signal}} \\ & + \underbrace{(1 - \widetilde{\pi}_j^+(m, y))\widetilde{W}_j(m+1, y; M_j^N(q))}_{\text{continuation value after a negative next signal}}. \end{aligned} \tag{2.26}$$

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

$$V_j^N(q) = \widetilde{W}_j(0, 0; M_j^N(q)). \tag{2.27}$$

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,

$$M_j^A(q) = \underbrace{\max_{\ell \neq j} V_\ell^A(q+1) - c^A(q+1)}_{\text{best continuation value from leaving option } j \text{ and continuing search with AI}}, \tag{2.28}$$

where $V_\ell^A(q+1)$ is the ex ante value of option 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 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

$$W_j(a, N_j, y; M_j^A(q)) = \max\{0, \delta_j(\widetilde{\mu}_j^{AI}(a, N_j, y)), M_j^A(q)\}. \tag{2.29}$$

For $m < N_j$,

$$W_j(a, m, y; M_j^A(q)) = \max\Big\{0,\ \delta_j(\widetilde{\mu}_j^{AI}(a, m, y)),\ M_j^A(q),\ \mathcal{K}_j(a, m, y; M_j^A(q))\Big\}, \tag{2.30}$$

where

$$\begin{aligned}\mathcal{K}_j(a, m, y; M_j^A(q)) = \quad & \underbrace{-c_j^W(m+1)}_{\text{pay current within-option cost}} \\ & + \underbrace{\pi_j^+(a, m, y) W_j(a, m+1, y+1; M_j^A(q))}_{\text{continuation value after a favorable next signal}} \\ & + \underbrace{(1 - \pi_j^+(a, m, y)) W_j(a, m+1, y; M_j^A(q))}_{\text{continuation value after an unfavorable next signal}}.\end{aligned} \tag{2.31}$$

Define the value of one additional human signal as

$$\begin{aligned}\Delta_j^A(a, m, y; q) = \quad & \underbrace{\mathcal{K}_j(a, m, y; M_j^A(q)) + c_j^W(m+1)}_{\text{expected continuation value before subtracting the current signal cost}} \\ & - \underbrace{\max\{0, \delta_j(\widetilde{\mu}_j^{AI}(a, m, y)), M_j^A(q)\}}_{\text{best stopping payoff at the current state}}.\end{aligned} \tag{2.32}$$

The consumer continues verifying within option j if and only if

$$\Delta_j^A(a, m, y; q) \geq c_j^W(m+1). \tag{2.33}$$

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

$$V_j^A(q) = \mathbb{E}_{a_j}\left[W_j(a_j, 0, 0; M_j^A(q))\right], \tag{2.34}$$

with the analogous no-AI object

$$V_j^N(q) = \widetilde{W}_j(0, 0; M_j^N(q)). \tag{2.35}$$

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

$$s_t = \left(q_t, \mathcal{V}_t, \{a_{jt}, m_{jt}, y_{jt}\}_{j\in\mathcal{V}_t}, \mathcal{U}_t\right), \tag{2.36}$$

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] 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

$$\mu_{jt} = \widetilde{\mu}_j^{AI}(a_{jt}, m_{jt}, y_{jt}), \tag{2.37}$$

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

$$S(s_t) = \max\left\{0, \max_{j\in\mathcal{V}_t} \delta_j(\mu_{jt})\right\}. \tag{2.38}$$

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

$$\begin{aligned} K_j^W(s_t) = & \underbrace{-c_j^W(m_{jt}+1)}_{\text{pay within-option inspection cost}} \\ & + \underbrace{\pi_j^+(a_{jt}, m_{jt}, y_{jt}) V(s_{t+1}^{W,j,+})}_{\text{continuation value after a favorable signal}} \\ & + \underbrace{\left(1-\pi_j^+(a_{jt}, m_{jt}, y_{jt})\right) V(s_{t+1}^{W,j,-})}_{\text{continuation value after an unfavorable signal}} \end{aligned} \tag{2.39}$$

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

$$K_k^A(s_t) = \underbrace{-c^A(q_t+1)}_{\text{pay across-option entry cost}} + \underbrace{\mathbb{E}_{a_k}\left[V(s_{t+1}^{A,k}(a_k))\right]}_{\text{expected continuation value after entering a new option}}, \tag{2.40}$$

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.

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

$$V(s_t) = \max\Big\{ \underbrace{S(s_t)}_{\text{stop and choose from the current consideration set}}, \quad \underbrace{\max_{j \in \mathcal{V}_t} K_j^W(s_t)}_{\text{inspect further within a visited option}}, \quad \underbrace{\max_{k \in \mathcal{U}_t} K_k^A(s_t)}_{\text{enter a new option and expand consideration}} \Big\}. \tag{2.41}$$

Equation (2.41) 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 even hope for analytic solutions, I must simplfy the scope of the problem. The most useful simplification is to collapse across-option search into these reduced-form outside values, which I do below. 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 though it was useful to study simpler special cases that preserve the core tradeoff between compression and verification. Using a natural starting point fixing a constant continuation value $\bar{M}$ from leaving the current option does away with the across-option choice problem, so that the problem becomes a single-option stopping problem after the overview arrives. This reduced-form object may summarize the value of alternative opportunities, including the option to stop with normalized no-purchase utility zero. In that environment, the central question is whether the summary raises or lowers the incentive to inspect additional underlying signals relative to the no-summary benchmark. Because every proposition below conditions on the fixed $\bar{M}$, the results hold regardless of how $\bar{M}$ is determined.

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 1 (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. 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*

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

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

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

b. *If*

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

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

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

Proof in Appendix C.

[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 1 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 1(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*

$$\zeta(\rho_j) = 3\rho_j^2 - 2\rho_j^3 \tag{3.1}$$

and define the posteriors after a favorable overview and, respectively, after a favorable or unfavorable inspected signal. These are the $N_j = 3$ specializations of the general notation μ_j^a, μ_j^{a+}, μ_j^{a-} from Proposition 1, evaluated at $a = 1$:

$$\mu_j^{+} = \frac{\mu_{j0}\zeta(\rho_j)}{\mu_{j0}\zeta(\rho_j) + (1-\mu_{j0})(1-\zeta(\rho_j))}, \tag{3.2}$$

$$\mu_j^{++} = \frac{\mu_{j0}\rho_j^2(2-\rho_j)}{\mu_{j0}\rho_j^2(2-\rho_j) + (1-\mu_{j0})(1-\rho_j)^2(1+\rho_j)}, \tag{3.3}$$

$$\mu_j^{+-} = \frac{\mu_{j0}\rho_j}{\mu_{j0}\rho_j + (1-\mu_{j0})(1-\rho_j)}. \tag{3.4}$$

Let

$$\pi_j^{+} = \frac{\mu_{j0}\rho_j^2(2-\rho_j) + (1-\mu_{j0})(1-\rho_j)^2(1+\rho_j)}{\mu_{j0}\zeta(\rho_j) + (1-\mu_{j0})(1-\zeta(\rho_j))} \tag{3.5}$$

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

$$\bar{M} = \mathbf{x}_j'\beta - \alpha p_j + \gamma\bar{\mu}_j \tag{3.6}$$

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

$$\mu_j^{+-} < \bar{\mu}_j \le \mu_j^{+} < \mu_j^{++}. \tag{3.7}$$

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

$$\kappa_j^W \le \gamma(1-\pi_j^{+})(\bar{\mu}_j - \mu_j^{+-}). \tag{3.8}$$

Equivalently, there is a threshold outside-option belief

$$\bar{\mu}_j^* = \mu_j^{+-} + \frac{\kappa_j^W}{\gamma(1-\pi_j^+)} \tag{3.9}$$

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

Derivation in Appendix C.

Condition (3.7) 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 C records the algebra. By symmetry, an analogous expression applies after an unfavorable overview.

Remark 1 (Outside the threshold region)*.* If $\bar{\mu}_j \leq \mu_j^{+-}$, 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 1. Finally, if $\bar{\mu}_j \geq \mu_j^{++}$, 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 (3.7), 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. Using the notation of Proposition 1, posterior beliefs satisfy*

$$\mu_j^a = \pi_j^a \mu_j^{a+} + (1-\pi_j^a)\mu_j^{a-}. \tag{3.10}$$

Proof in Appendix C.

Lemma 1 records the standard martingale property of Bayesian posteriors. For $N_j = 3$ with a favorable overview, equation (3.10) reduces to $\mu_j^+ = \pi_j^+ \mu_j^{++} + (1 - \pi_j^+)\mu_j^{+-}$ using the closed-form posteriors in Example 1.

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. Then for each overview realization $a \in \{0,1\}$,*

$$1 - \pi_j^a \to 0 \qquad \text{as } \rho_j \to 1, \tag{3.11}$$

and consequently the outside-option threshold $\bar{\mu}_j^*$ in Proposition 1(a) satisfies $\bar{\mu}_j^* \to \infty$ as $\rho_j \to 1$. 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.

Proof in Appendix C.

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. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$ as in Appendix A. Then for a favorable overview ($a = 1$),*

$$\frac{\partial \bar{\mu}_j^*}{\partial \rho_j} > 0 \tag{3.12}$$

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.

Proof in Appendix C.

Corollary 3 (High precision eliminates finite verification)**.** *Fix any finite $N_j > 1$ and any finite inspection horizon $\bar{m}_j \le N_j$. Suppose within-option marginal inspection costs satisfy $c_j^W(m) > 0$ for every $m \in \{1, \dots, \bar{m}_j\}$. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$ as in Appendix A. 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 C.

For even N_j, the proofs differ only 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 1–3; Appendix A gives the corresponding tie-breaking formula and Figure 7 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 7 shows that the same geometry is qualitatively similar for even N_j under neutral tie-breaking.

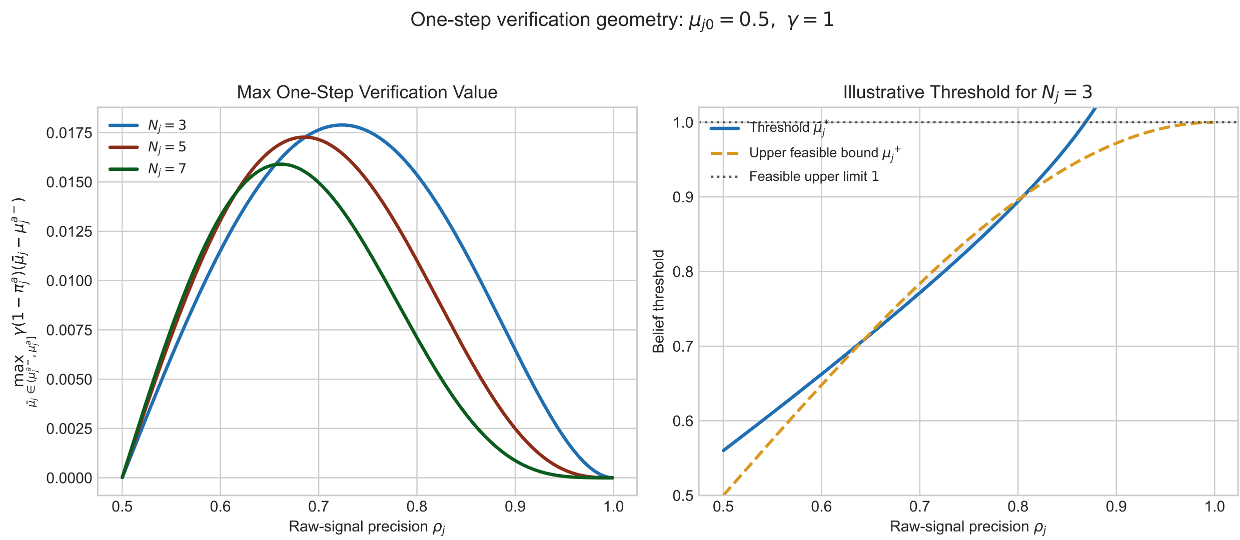

Figure 3: One-step verification geometry under $\mu_{j0} = 0.5$ and $\gamma = 1$. The left panel plots the maximal one-step verification value as a function of raw-signal precision for illustrative signal-set sizes $N_j = 3$, $N_j = 5$, and $N_j = 7$. The right panel specializes to $N_j = 3$ and translates 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 7 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. Let $V_j^A(q)$ denote the ex ante value of entering option j when the consumer observes the AI summary before choosing whether to inspect raw signals, and let $V_j^N(q)$ denote the ex ante value when no AI summary is shown. If the summary is costless and the consumer updates correctly, then

$$V_j^A(q) \geq V_j^N(q) \tag{3.13}$$

for every option j. 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 2 (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*

$$\bar{c}_j(a, m, y; q) = \Delta_j^A(a, m, y; q). \tag{3.14}$$

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

$$c_j^W(m + 1) \leq \bar{c}_j(a, m, y; q). \tag{3.15}$$

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

Proof in Appendix C. 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

$$\begin{aligned} &\mathbb{E}[\text{options entered}], \qquad \mathbb{E}[\text{raw signals opened}], \\ &\mathbb{P}(\text{stop immediately}), \qquad \mathbb{P}(\text{choose high-quality option}), \end{aligned} \tag{3.16}$$

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

$$V = \mathbb{E}[\text{actual payoff}] - \mathbb{E}[\text{within-option search cost}] - \mathbb{E}[\text{across-option search cost}]. \tag{3.17}$$

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

$$\begin{aligned} &\Delta E[\text{options}], \qquad \Delta E[\text{signals}], \\ &\Delta \mathbb{P}(\text{high quality}), \qquad \Delta\text{payoff}, \qquad \Delta V. \end{aligned} \tag{3.18}$$

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, strong signal environments should exhibit the largest 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.

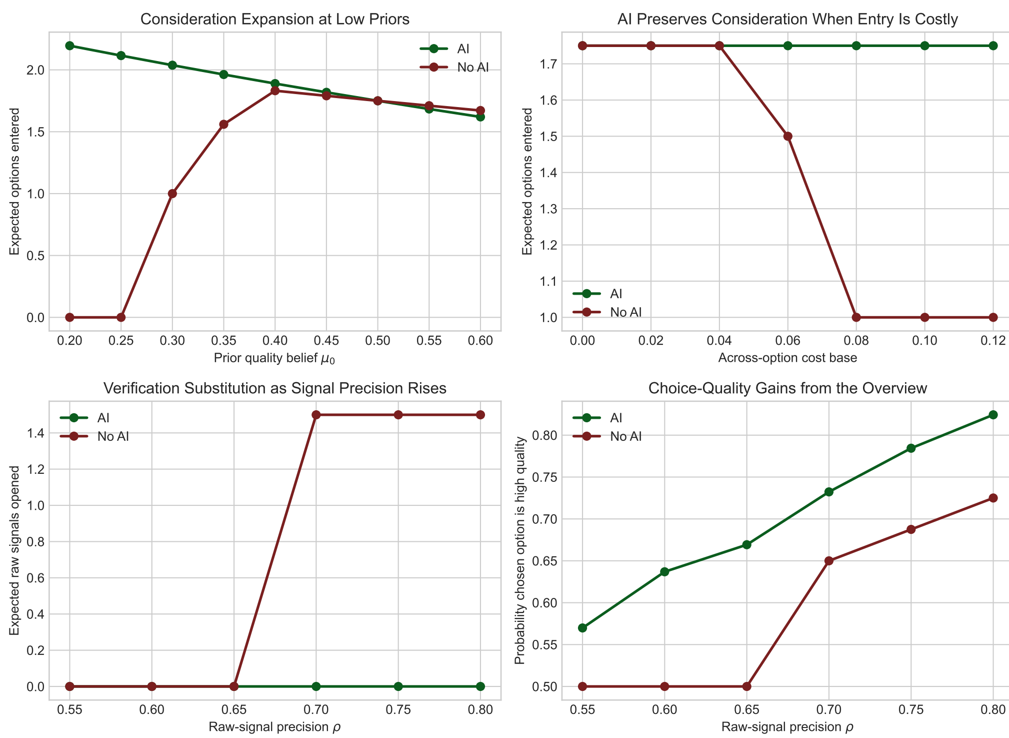

Figure 4: Preliminary numerical illustrations from the nested $J = 3$ model. The top-left panel shows that AI expands consideration when prior quality beliefs are weak. The top-right panel shows that AI preserves consideration when across-option entry becomes costly. The bottom-left panel shows that 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.

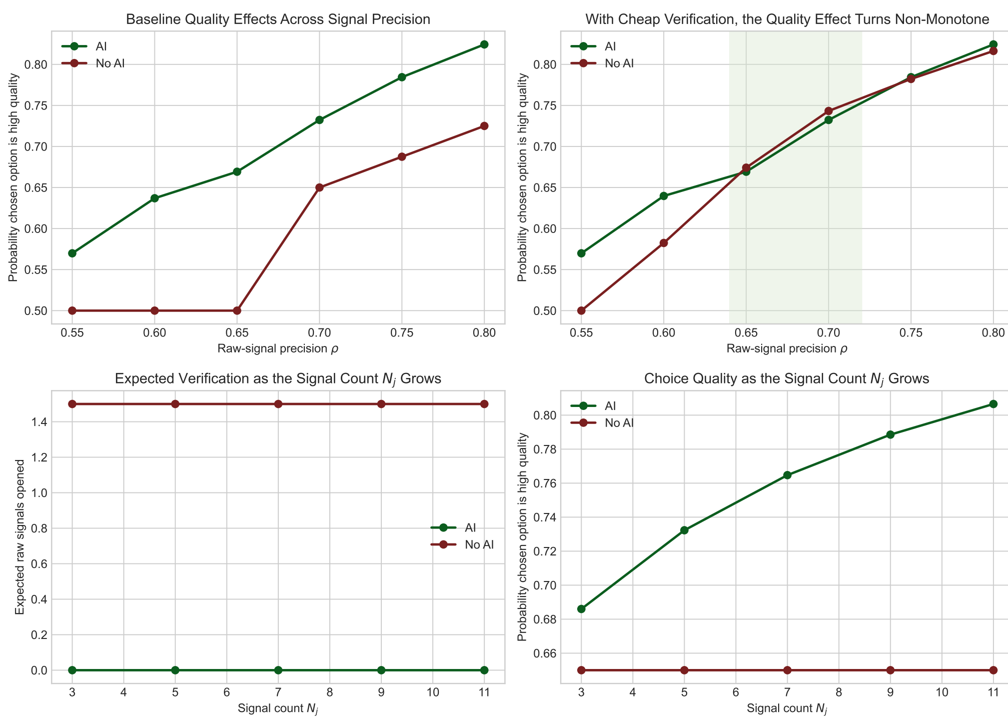

Figure 5: Additional numerical checks. The top-left panel reports the baseline quality sweep across raw-signal precision. The top-right panel repeats that sweep under a lower within-option cost schedule, $c_j^W(m+1) = 0.01 + 0.005m$, and shows that the AI effect on choice quality can become non-monotone 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.

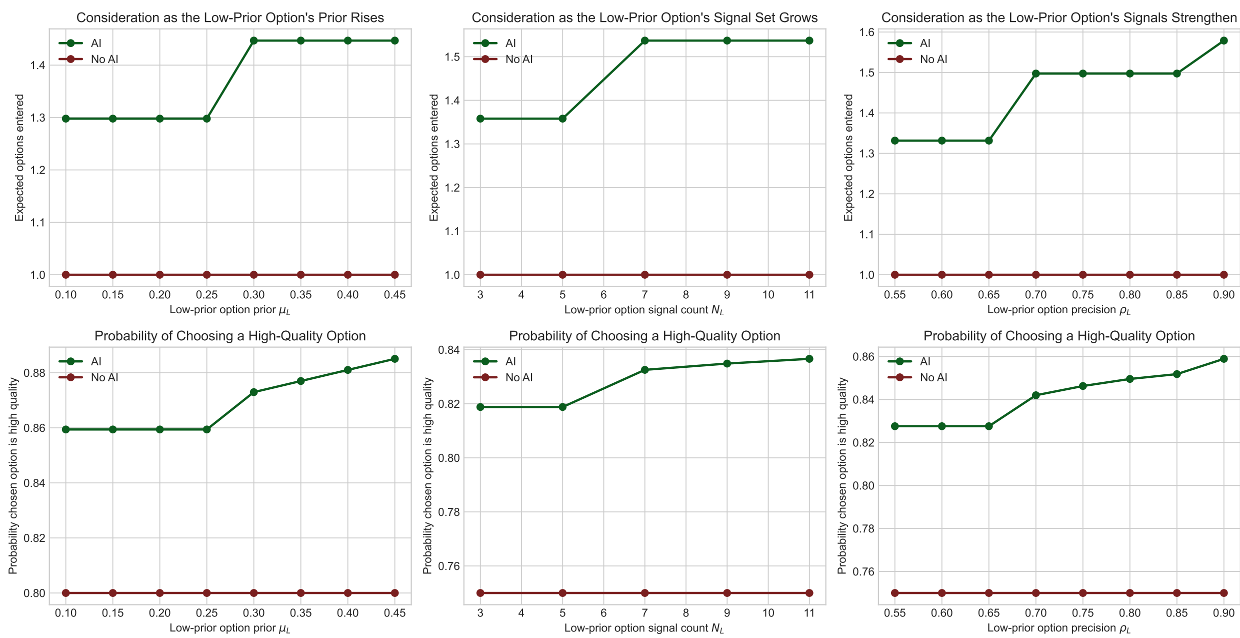

Figure 6: Heterogeneous comparative statics. The left column varies the low-prior option's prior while holding the other options fixed. The middle column varies the low-prior option's signal-set size. The right column varies the low-prior option's raw-signal precision. The top row reports expected options 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.

4 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 analytical results concern the within-option verification margin. For any finite signal-set size $N_j > 1$ and one-step post-summary inspection, the verification threshold diverges as raw-signal precision rises (Corollary 1) and the set of states in which verification is optimal shrinks monotonically near perfect precision (Corollary 2), with even N_j handled by a tie-breaking rule. 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 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.

Those within-option results interact with the across-option margin in the numerical analysis. When verification costs fall because the summary substitutes for inspection, the effective cost of screening new options drops. The numerical model shows that AI can expand consideration sets—especially for lower-prior options that the no-AI consumer would never have entered—while simultaneously reducing total within-option search. Whether this reallocation improves or reduces choice quality depends on the balance between the information lost from reduced verification and the information gained from broader consideration. The numerical comparative statics characterize that balance across the parameter space.

Several testable predictions follow directly from the model. First, platforms that introduce AI summaries should observe the largest decline in within-option click-depth for products whose underlying reviews are most individually informative. Second, consideration sets should expand most where AI summaries make initial screening cheapest—low-prior or unfamiliar product categories. Third, when within-option verification is sufficiently cheap, the net effect of AI on choice quality can be non-monotone in signal precision: at intermediate precision, the reallocation from depth to breadth may reduce the probability of choosing the best option before summary accuracy dominates again at higher precision. These predictions differ from those of standard sequential-search models and suggest empirical tests using platform experiments that randomize the availability of AI overviews, though implementation would need to account for consumer beliefs, trust, and understanding of the aggregation process. More broadly, the consideration-expansion predictions are qualitatively consistent with the evidence in Fang et al. (2024) that richer platform-provided search information can reshape downstream search and purchasing behavior over time.

The model is deliberately silent on two fronts. First, it takes the aggregation rule as exogenous. Making the summary rule endogenous would open a Bayesian persuasion problem (Kamenica and Gentzkow, 2011) in which the platform designs the compression to maximize its own objective, which need not align with consumer welfare. The wedge between a consumer-optimal and a profit-maximizing summary is where strategic summary design and self-preferencing would arise, and it is best treated as a separate paper. Second, the model 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 aggregation rule itself, which would introduce a joint learning problem over both quality and summary precision. Under the lossy compression provision, both extensions would naturally preserve the paper's core reallocation mechanism while enriching the comparative statics and findings about search behavior.

A 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 comparative statics emphasized in Corollaries 1–3, however, even-N_j tie-breaking does not introduce a qualitatively new force. The exact-tie term changes the constant on the leading asymptotic term but not the fact that favorable overviews become nearly sufficient statistics as $\rho_j \to 1$. Figure 7 illustrates this 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.

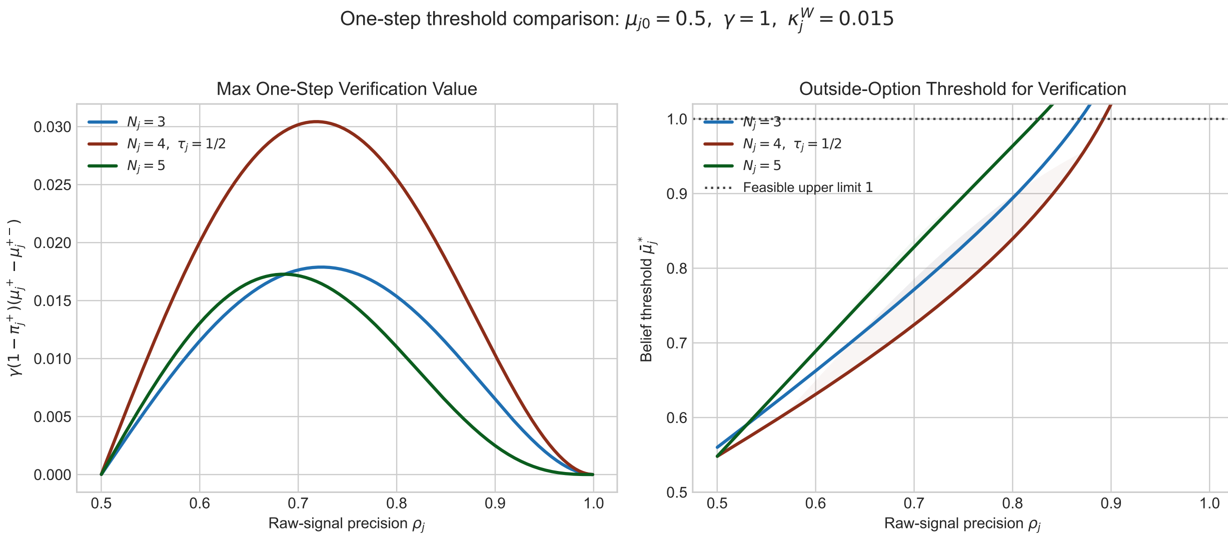

Figure 7: One-step verification thresholds with odd and even signal counts. The left panel plots the maximal one-step verification value after a favorable overview for $N_j = 3$, $N_j = 4$ with neutral tie-breaking, and $N_j = 5$ under $\mu_{j0} = 0.5$ and $\gamma = 1$. The right panel translates those value objects into outside-option thresholds $\bar{\mu}_j^*$ for illustrative inspection cost $\kappa_j^W = 0.015$. Verification is possible only where $\bar{\mu}_j^* \leq \mu_j^+$, 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 the threshold diverges as ρ_j approaches one.

A.1 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 7 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 7: 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.

B Derivations

B.1 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:

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

By assumption,

$$\mathbb{P}(X_{jk} = 1 \mid \theta_j) = \rho_j, \tag{B.2}$$

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

$$C_j = \sum_{k=1}^{N_j} X_{jk}, \tag{B.3}$$

is binomial:

$$C_j \mid \theta_j \sim \text{Binomial}(N_j, \rho_j). \tag{B.4}$$

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

$$C_j \geq \frac{N_j + 1}{2}. \tag{B.5}$$

Therefore

$$\zeta_j = \mathbb{P}(a_j = \theta_j \mid \theta_j) = \mathbb{P}\left(C_j \geq \frac{N_j + 1}{2} \mid \theta_j\right). \tag{B.6}$$

Using the binomial probability mass function,

$$\mathbb{P}(C_j = m \mid \theta_j) = \binom{N_j}{m} \rho_j^m (1 - \rho_j)^{N_j - m}, \tag{B.7}$$

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

$$\zeta_j = \sum_{m=(N_j+1)/2}^{N_j} \binom{N_j}{m} \rho_j^m (1 - \rho_j)^{N_j - m}, \tag{B.8}$$

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*

$$\zeta_j > \rho_j. \tag{B.9}$$

More generally, if

$$h_N(p) = \sum_{m=(N+1)/2}^{N} \binom{N}{m} p^m (1 - p)^{N - m}, \tag{B.10}$$

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). □

B.2 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

$$\mathbb{P}(\theta_j = 1 \mid r_1 = 1) = \frac{\mathbb{P}(r_1 = 1 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1)}{\mathbb{P}(r_1 = 1 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1) + \mathbb{P}(r_1 = 1 \mid \theta_j = 0)\mathbb{P}(\theta_j = 0)} \tag{B.11}$$
$$= \frac{\mu_{j0}\rho_j}{\mu_{j0}\rho_j + (1 - \mu_{j0})(1 - \rho_j)}. \tag{B.12}$$

After one negative signal,

$$\mathbb{P}(\theta_j = 1 \mid r_1 = 0) = \frac{\mathbb{P}(r_1 = 0 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1)}{\mathbb{P}(r_1 = 0 \mid \theta_j = 1)\mathbb{P}(\theta_j = 1) + \mathbb{P}(r_1 = 0 \mid \theta_j = 0)\mathbb{P}(\theta_j = 0)} \tag{B.13}$$
$$= \frac{\mu_{j0}(1 - \rho_j)}{\mu_{j0}(1 - \rho_j) + (1 - \mu_{j0})\rho_j}. \tag{B.14}$$

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

$$\rho_j^y (1 - \rho_j)^{m-y}, \tag{B.15}$$

while conditional on $\theta_j = 0$ it is

$$(1 - \rho_j)^y \rho_j^{m-y}. \tag{B.16}$$

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

$$\mathbb{P}(Y = y \mid \theta_j = 1, m) = \binom{m}{y} \rho_j^y (1 - \rho_j)^{m-y}, \tag{B.17}$$
$$\mathbb{P}(Y = y \mid \theta_j = 0, m) = \binom{m}{y} (1 - \rho_j)^y \rho_j^{m-y}. \tag{B.18}$$

Applying Bayes' rule,

$$\mathbb{P}(\theta_j = 1 \mid m, y) = \frac{\mu_{j0}\binom{m}{y}\rho_j^y(1-\rho_j)^{m-y}}{\mu_{j0}\binom{m}{y}\rho_j^y(1-\rho_j)^{m-y} + (1-\mu_{j0})\binom{m}{y}(1-\rho_j)^y\rho_j^{m-y}} \tag{B.19}$$

$$= \frac{\mu_{j0}\rho_j^y(1-\rho_j)^{m-y}}{\mu_{j0}\rho_j^y(1-\rho_j)^{m-y} + (1-\mu_{j0})(1-\rho_j)^y\rho_j^{m-y}}. \tag{B.20}$$

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.

B.3 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

$$T_j = \frac{N_j + 1}{2}. \tag{B.21}$$

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

$$H_{ja}(m, y; p) = \begin{cases} \mathbb{P}(\mathrm{Bin}(N_j - m, p) \geq T_j - y)\,, & a = 1, \\ \mathbb{P}(\mathrm{Bin}(N_j - m, p) < T_j - y)\,, & a = 0. \end{cases} \tag{B.22}$$

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

$$H_{j1}(m, y; \rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, \rho_j) \geq T_j - y)\,, \tag{B.23}$$

$$H_{j0}(m, y; \rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, \rho_j) < T_j - y)\,. \tag{B.24}$$

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

$$H_{j1}(m, y; 1-\rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, 1-\rho_j) \geq T_j - y)\,, \tag{B.25}$$

$$H_{j0}(m, y; 1-\rho_j) = \mathbb{P}(\mathrm{Bin}(N_j - m, 1-\rho_j) < T_j - y)\,. \tag{B.26}$$

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

$$L_{j1}(1, 0, 0) = H_{j1}(0, 0; \rho_j) = \zeta_j, \tag{B.27}$$

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

$$L_{j0}(1, 0, 0) = H_{j1}(0, 0; 1-\rho_j) = 1 - \zeta_j. \tag{B.28}$$

Hence

$$\widetilde{\mu}_j^{AI}(1, 0, 0) = \frac{\mu_{j0}\zeta_j}{\mu_{j0}\zeta_j + (1-\mu_{j0})(1-\zeta_j)}, \tag{B.29}$$

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

$$\binom{m}{y}\rho_j^y(1-\rho_j)^{m-y}. \tag{B.30}$$

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

$$H_{ja}(m, y; \rho_j). \tag{B.31}$$

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

$$L_{j1}(a, m, y) = \binom{m}{y}\rho_j^y(1-\rho_j)^{m-y}H_{ja}(m, y; \rho_j). \tag{B.32}$$

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

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

and the consistency of the remaining uninspected signals contributes

$$H_{ja}(m, y; 1 - \rho_j), \tag{B.34}$$

so

$$L_{j0}(a, m, y) = \binom{m}{y}(1 - \rho_j)^y \rho_j^{m-y} H_{ja}(m, y; 1 - \rho_j). \tag{B.35}$$

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

$$\widetilde{\mu}_j^{AI}(a, m, y) = \mathbb{P}(\theta_j = 1 \mid a, m, y) \tag{B.36}$$

$$= \frac{\mu_{j0} L_{j1}(a, m, y)}{\mu_{j0} L_{j1}(a, m, y) + (1 - \mu_{j0}) L_{j0}(a, m, y)}. \tag{B.37}$$

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.

B.4 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

$$\widetilde{\pi}_j^+(m, y) = \mathbb{P}(r_{m+1} = 1 \mid m, y) \tag{B.38}$$

$$\begin{aligned} &= \mathbb{P}(r_{m+1} = 1 \mid \theta_j = 1, m, y)\mathbb{P}(\theta_j = 1 \mid m, y) \\ &\quad + \mathbb{P}(r_{m+1} = 1 \mid \theta_j = 0, m, y)\mathbb{P}(\theta_j = 0 \mid m, y) \end{aligned} \tag{B.39}$$

$$= \widetilde{\mu}_j(m, y)\rho_j + (1 - \widetilde{\mu}_j(m, y))(1 - \rho_j), \tag{B.40}$$

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

$$H_{ja}(m + 1, y + 1; p). \tag{B.41}$$

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

$$H_{ja}(m, y; p). \tag{B.42}$$

Hence

$$\psi_{ja}(m, y; p) = \mathbb{P}(r_{m+1} = 1 \mid a, m, y;\ p) = p\frac{H_{ja}(m + 1, y + 1; p)}{H_{ja}(m, y; p)}, \tag{B.43}$$

which is equation (2.17).

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

$$\begin{aligned}
\pi_j^+(a, m, y) &= \mathbb{P}(r_{m+1} = 1 \mid a, m, y) && \text{(B.44)}\\
&= \mathbb{P}(r_{m+1} = 1 \mid a, m, y, \theta_j = 1)\mathbb{P}(\theta_j = 1 \mid a, m, y)\\
&\quad + \mathbb{P}(r_{m+1} = 1 \mid a, m, y, \theta_j = 0)\mathbb{P}(\theta_j = 0 \mid a, m, y) && \text{(B.45)}\\
&= \widetilde{\mu}_j^{AI}(a, m, y)\psi_{ja}(m, y; \rho_j) + \left(1 - \widetilde{\mu}_j^{AI}(a, m, y)\right)\psi_{ja}(m, y; 1 - \rho_j), && \text{(B.46)}
\end{aligned}$$

which is equation (2.18).

C Proofs

Proof of Proposition 1. 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; 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 $\mathcal{V} \geq \delta_j(\mu_j^a)$, i.e.,

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

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

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

which rearranges to

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

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:

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

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

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

which proves part (a).

Part (b). Suppose instead the ordering

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

Now $\delta_j(\mu_j^a) \leq \bar{M}$, so stopping immediately yields $\bar{M}$ (the consumer prefers the outside option at the current posterior). After a favorable signal, $\delta_j(\mu_j^{a+}) > \bar{M}$, 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 $\mathcal{V} \geq \bar{M}$, i.e.,

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

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

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

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

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

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 1 is a lower bound on the unrestricted verification value. The cutoff inequalities in Proposition 1 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. 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

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

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

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 (3.7), 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 (3.8). 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,

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

which is exactly (3.10). 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. $\square$

Proof of Corollary 1. This proof is a work in progress and may be refined in a later draft.

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 overview-consistency term by its τ_j-adjusted analogue from Appendix A. Without loss of generality consider a favorable overview ($a = 1$); the argument for $a = 0$ is symmetric. Write $\mu_j^{a=1} \equiv \widetilde{\mu}_j^{AI}(1, 0, 0)$ and $\pi_j^{a=1} \equiv \pi_j^{+}(1, 0, 0)$ for short. From Proposition 1, the threshold in case (a) involves $1-\pi_j^a$, the probability that the next inspected signal is unfavorable given the overview. Using the transition probability $\psi_{j1}(0, 0; p)$ defined in (2.17) with $a = 1$, the unconditional probability of an unfavorable next signal given $a = 1$ can be decomposed as

$$1 - \pi_j^{a=1} = \mu_j^{a=1}\big(1 - \psi_{j1}(0, 0; \rho_j)\big) + (1 - \mu_j^{a=1})\big(1 - \psi_{j1}(0, 0; 1 - \rho_j)\big),$$

where $\psi_{j1}(0, 0; p) = p\, H_{j1}(1, 1; p)/H_{j1}(0, 0; p)$. Consider each factor as $\rho_j \to 1$:

Under $\theta_j = 1$: the overview is correct with probability tending to one, and after a favorable overview the next inspected signal is unfavorable only if a raw-signal mistake occurs. Hence $\psi_{j1}(0, 0; \rho_j) \to 1$ and $1 - \psi_{j1}(0, 0; \rho_j)$ is first order in $(1 - \rho_j)$.

Under $\theta_j = 0$: the probability of observing a favorable overview is polynomially small in $(1 - \rho_j)$. For odd N_j this comes from strict-majority events; for even N_j the same expression picks up an additional τ_j-weighted exact-tie term. In either case, $\psi_{j1}(0, 0; 1 - \rho_j)$ and $1 - \psi_{j1}(0, 0; 1 - \rho_j)$ remain bounded as $\rho_j \to 1$.

Because $\zeta_j \to 1$ as $\rho_j \to 1$ for odd N_j and, by Appendix A, also for even N_j under any fixed tie-breaking rule, the post-overview posterior converges to certainty: $\mu_j^{a=1} \to 1$ and $1 - \mu_j^{a=1} \to 0$. The dominant contribution to $1 - \pi_j^{a=1}$ therefore comes from the high-state probability of seeing one unfavorable inspected signal after a favorable overview, which is first order in $(1 - \rho_j)$. In the even-N_j case the τ_j-weighted tie term changes the leading constant but not the rate. Hence there exists a finite constant $C_{N_j,\tau_j} > 0$ such that

$$1 - \pi_j^{a=1} \sim C_{N_j,\tau_j}(1 - \rho_j) \qquad \text{as } \rho_j \to 1, \tag{C.2}$$

where for odd N_j one may take $C_{N_j,\tau_j} = 1$.

Consequently $\kappa_j^W/[\gamma(1 - \pi_j^a)] \to +\infty$ since $\kappa_j^W, \gamma > 0$ are fixed. The threshold from Proposition 1(a) satisfies

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

so the threshold diverges regardless of the bounded behavior of μ_j^{a-}. Because feasible outside-option beliefs satisfy $\bar{\mu}_j \in [0, 1]$, there exists $\rho_j^* < 1$ such that $\bar{\mu}_j^* > 1$ for all $\rho_j > \rho_j^*$. For those signal-precision values, the continuation condition $\bar{\mu}_j \geq \bar{\mu}_j^*$ cannot be satisfied by any feasible $\bar{\mu}_j$, so continuation after the overview is not optimal. □

Proof of Corollary 2. This proof is a work in progress and may be refined in a later draft.

Fix any finite $N_j > 1$. If N_j is even, fix a tie-breaking rule $\tau_j \in [0,1]$. From Proposition 1(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 smooth function of ρ_j mapping into $[0,1]$.

Define $g(\rho_j) = 1 - \pi_j^a$. From (C.2) (established in the proof of Corollary 1),

$$g(\rho_j) \sim C_{N_j,\tau_j}(1-\rho_j) \qquad \text{as } \rho_j \to 1.$$

Therefore $1/g(\rho_j) \sim 1/[C_{N_j,\tau_j}(1-\rho_j)]$, which gives

$$\frac{g(\rho_j)}{1-\rho_j} \to C_{N_j,\tau_j} \qquad \text{as } \rho_j \to 1,$$

and consequently

$$(1-\rho_j)^2 \frac{\partial}{\partial \rho_j}\left(\frac{1}{g(\rho_j)}\right) \to 1/C_{N_j,\tau_j}.$$

Therefore

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

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

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

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 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. Specifically, let

$$C_j = \sup_{\mu \in [0,1]} |\delta_j(\mu)| < \infty$$

(which is finite since $\delta_j(\mu)$ is affine in μ on a compact domain). Then for any feasible continuation plan σ after overview realization a_j,

$$\text{gross gain from } \sigma \leq C_j \mathbb{P}(E_{a_j,\sigma,\rho_j}),$$

where E_{a_j,σ,ρ_j} is the event that following σ changes the final action relative to immediate stopping after the overview. 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 feasible continuation plans after the overview is itself finite. Pointwise convergence of $\mathbb{P}(E_{a_j,\sigma,\rho_j})$ to zero for each 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 each additional inspected signal is favorable with probability converging to one. Since only finitely many signals can be inspected, the probability that any feasible continuation history drives the posterior below the stopping threshold converges to zero. Conditional on an unfavorable overview, the symmetric argument implies that the probability of crossing the relevant lower stopping boundary also 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

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

to a probability distribution over the three action classes:

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

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

$$V_j^A(q) \geq V_j^N(q).$$

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 2. 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. □

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.