

The Summary Comes First: How AI Overviews Reshape Consumer Product Search

Pre-print version of the following publication: | Versione pre-print della seguente pubblicazione:

Original Citation/Citazione:

De Angelis, Matteo; Mazzù, Marco Francesco; Sabatini, Nicola. (2026). The Summary Comes First: How AI Overviews Reshape Consumer Product Search. <https://ssrn.com/abstract=7397080>

Availability/Disponibilità:

This version is available at: [11385/265198](https://iris.luiss.it/11385/265198) since: 2026-09-03T11:54:05Z - Questa versione è disponibile alla pagina: [11385/265198](https://iris.luiss.it/11385/265198) dal: 2026-09-03T11:54:05Z

Publisher/Casa editrice:

Published version/Pubblicato:

DOI: <https://dx.doi.org/10.2139/ssrn.7397080>

License/Licenza:

Tutti i diritti riservati

Availability/Termini d'uso:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. Works made available under a Creative Commons license can be used according to the terms and conditions of said license. For all terms of use and more information see the publisher's website. | I termini e le condizioni relativi al riutilizzo della presente versione della pubblicazione sono disciplinati dalla politica editoriale. Le opere messe a disposizione con licenze Creative Commons possono essere utilizzate conformemente ai termini e alle condizioni previste da tali licenze. Per l'insieme delle condizioni di utilizzo e per ulteriori informazioni si rinvia al sito web dell'editore.

This item was downloaded from IRIS Luiss (<https://iris.luiss.it/>). When citing, please refer to the published version. | Questo documento è stato scaricato da IRIS Luiss (<https://iris.luiss.it/>). Per la citazione, fare riferimento alla versione pubblicata sul sito dell'editore.

(Article begins on next page | Il contributo inizia nella pagina successiva)

The Summary Comes First: How AI Overviews Reshape Consumer Product Search

Abstract

Search engines increasingly place AI Overviews—generative-AI syntheses that answer a query within the results page—above traditional links. This changes product search from a source-led process, in which consumers navigate and integrate information themselves, toward a summary-led process, in which the interface performs part of that work. Across three online experiments (total N = 584), participants used realistic, clickable search environments recreated in Qualtrics, allowing us to observe source clicks and time spent on sources as search unfolded. Study 1 shows that an AI Overview, compared with a traditional search engine results page, reduces both source clicks and source-viewing time during initial product search. Study 2 replicates this effect across search, experience, and credence product categories. Product type alone does not moderate the effect, but interface type, product type, and positive attitudes toward AI jointly shape exposure clicks. Study 3 shows that AI Overviews reduce epistemic vigilance—the critical appraisal of information and its sources—which indirectly increases purchase intention. The findings identify summary-led search as a phase-specific search architecture that reallocates epistemic labor from consumers to the interface, compressing initial source navigation while separately reducing critical scrutiny of product information.

Keywords

AI Overview; online search; consumer information search; source navigation; search goods; experience goods; credence goods; generative AI.

1. Introduction

Search engines have traditionally helped consumers find information without resolving it for them. A conventional search engine results page (SERP) presents ranked links, short snippets, and source cues. Consumers must decide which results to open and in what order, compare claims across pages, and integrate what they find. Consumer search is therefore commonly understood as a costly, consumer-directed process of information acquisition (Bettman et al., 1998; Krefeld-Schwalb et al., 2026; Moorthy et al., 1997; Stigler, 1961).

AI Overviews change this familiar arrangement. Consistent with Google's own terminology for the feature, we define an AI Overview as a generative-AI synthesis displayed within a search-results page, typically above the organic links, that produces a direct response to the user's query by combining information drawn from multiple sources. Google's system triggers this synthesis selectively, typically for queries judged complex or informational. Unlike a traditional snippet, which previews one particular webpage, an AI Overview integrates claims across sources into a single, apparently self-contained answer while retaining links that users may consult for support. For example, a consumer searching "How can I tell whether a cabin trolley is durable?" may first see a generated paragraph explaining that shell material, wheel construction, handles, zippers, and warranties are key indicators of durability. The consumer can still open luggage guides, retailers, or reviews, but she may no longer need to visit those sources to satisfy her needs.

This apparently modest interface change alters who performs the work of search. In a traditional SERP, the platform retrieves and ranks information, whereas the consumer selects sources and constructs an answer. With an AI Overview, the platform also performs part of the selection, organization, and synthesis before the consumer visits a source. With an AI Overview, source-led search becomes summary-led: consumers begin with an integrated interpretation, while individual sources serve primarily as supporting material.

Consumer search models show that consumers balance the value of additional information against the time and effort required to acquire it (Diehl, 2005; Krefeld-Schwalb et al., 2026; Moorthy

et al., 1997). Interactive decision aids reduce these costs by filtering and organizing marketplace information, while contemporary recommendation agents increasingly tailor results to consumers' explicit and inferred preferences (Baccelloni et al., 2026; Häubl & Trifts, 2000). For example, a recommendation agent may rank laptops according to a consumer's preferred price, battery life, and portability. Such personalization can facilitate choice but also shape—and potentially narrow—the attributes and alternatives consumers consider (Baccelloni et al., 2026; Häubl & Murray, 2003). Recent evidence similarly indicates that learning from LLM syntheses requires less effort than learning through web links (Melumad & Yun, 2025). AI Overviews extend these developments by inserting synthesis directly into the conventional search environment. Indeed, unlike the previously explored comparison between LLM-based learning and web search, this paper considers a setting that embeds generative synthesis within a conventional SERP while preserving the underlying source environment, allowing to examine how the addition of an AI-generated overview changes navigation of the sources that remain available. Yet it remains unclear whether, compared to traditional SERP, AI Overviews reduce navigation of underlying sources, whether this reduction extends to later verification, whether it varies across product contexts and consumers and why it may happen.

Beyond reducing search effort, summary-led search may change how consumers evaluate information. Epistemic vigilance refers to the critical appraisal of communicated content and of the competence and trustworthiness of its sources (Sperber et al., 2010). Whereas a conventional SERP presents sources separately, an AI Overview integrates their claims into a single fluent response. This synthesis may make information easier to use while rendering its evidentiary basis less salient. Research on generative systems suggests that users may consequently treat cited sources as implicitly endorsed without inspecting them directly (Dabran-Zivan et al., 2026). We therefore investigate whether AI Overviews reduce epistemic vigilance and, in turn, influence downstream purchase responses.

Across three controlled online experiments featuring realistic, clickable search environments recreated in Qualtrics, we examine three linked questions: whether AI Overviews reduce initial

77 engagement with underlying sources; how that effect varies across product categories and consumers’
78 attitudes toward AI; and whether AI Overviews influence purchase intention through epistemic
79 vigilance. Study 1 compares a traditional SERP with a SERP containing an AI Overview in a cabin-
80 trolley search task. Study 2 crosses the interface manipulation with search, experience, and credence
81 product categories (notebook, VR headset, and omega-3 supplements) and examines positive
82 attitudes toward AI as an individual boundary condition. Study 3 returns to a health-related
83 smartwatch context and examines an indirect pathway linking AI Overview exposure, epistemic
84 vigilance, and purchase intention. In all studies, participants interacted with realistic, clickable search
85 interfaces, allowing us to observe source clicks and time rather than infer search from self-reports.

86 The studies reveal a consistent behavioral main effect and a more nuanced psychological
87 consequence. Across contexts, AI Overviews reduce both the number of source-opening clicks and
88 the time spent on them during initial exposure. This effect does not extend uniformly to a later
89 verification phase, indicating that summary-led search compresses the beginning of the search process
90 rather than eliminating checking altogether. Product type alone does not moderate the effect. Instead,
91 Study 2 identifies a three-way interaction among interface, product type, and positive attitudes toward
92 AI for exposure clicks. Study 3 further shows that AI Overview exposure lowers epistemic vigilance
93 and thereby indirectly increases purchase intention.

94 These findings make three contributions. First, we extend consumer information-search theory
95 by showing that AI-generated synthesis changes not merely the amount or order of information
96 available, but the allocation of epistemic labor between the consumer and the interface (Diehl, 2005;
97 Krefeld-Schwalb et al., 2026; Moorthy et al., 1997; Stigler, 1961). Whereas established models treat
98 consumers as the agents who acquire and integrate source-level information, our behavioral measures
99 show that an integrated answer systematically reduces initial source clicks and time. We thereby
100 identify summary-led search as a distinct search interface architecture for product search and
101 complement evidence that LLM synthesis reduces learning effort and depth (Melumad & Yun, 2025)
102 by documenting how it restructures source navigation in product search. The separation between

103 initial exposure and later verification further shows that the effect is phase-specific rather than a
104 uniform withdrawal from search. By observing source clicks and viewing time as search unfolds, the
105 studies examine the search process directly rather than inferring it from self-reported effort or final
106 choice alone.

107 Second, we extend research on interactive decision aids and consumer responses to AI
108 (Baccelloni et al., 2026; Castelo et al., 2019; Dorotic et al., 2024; Häubl & Murray, 2003; Häubl &
109 Trifts, 2000; Puntoni et al., 2021) by locating AI influence upstream of recommendation and choice.
110 AI Overviews preserve the links and visual conventions of traditional search while adding a
111 synthesized interpretation before them. They therefore constitute a hybrid retrieval-and-synthesis
112 interface rather than a conventional recommendation agent or stand-alone chatbot. Our findings show
113 that this upstream assistance compresses search across product contexts, while the three-way
114 interaction among interface, product type, and positive AI attitudes indicates that the effect of the
115 interface can vary jointly with product context and consumers' orientation toward AI.

116 Third, we introduce epistemic vigilance as a consumer-relevant pathway through which
117 generative search can shape marketplace responses (Dabran-Zivan et al., 2026; Sperber et al., 2010;
118 Tseng et al., 2021). Prior research shows that people assess communicated claims through both
119 content plausibility and source-credibility cues (Metzger et al., 2010; Sperber et al., 2010), while
120 emerging evidence indicates that generative systems reshape how users evaluate sources and
121 corroborate information (Dabran-Zivan et al., 2026). We show experimentally that an embedded AI
122 Overview can lower such critical appraisal and indirectly increase purchase intention. Importantly,
123 epistemic vigilance does not mediate the reduction in source navigation. This result distinguishes the
124 behavioral efficiency of summary-led search from its epistemic consequence: fewer source visits and
125 lower scrutiny are related effects of the same interface, but they are not a single process.

126 The shift from traditional SERP to SERP containing an AI Overview has important implications
127 for firms and policymakers. Brands and publishers no longer compete only for visibility and clicks
128 within ranked search results; they also compete to have their information selected and accurately

represented in AI-generated summaries. Because consumers may rely on these summaries without visiting the underlying webpages, AI Overviews can influence both referral traffic and the product information consumers encounter. For platforms and regulators, ensuring accurate content and disclosing AI involvement may therefore be insufficient. Interface design must also make sources visible, traceable, and easy to verify so that consumers can evaluate the evidence underlying generated answers.

2. Theoretical Background

2.1 From source-led to summary-led consumer search

Consumer search theory treats information acquisition as a purposeful but costly activity. Consumers search because they are uncertain about product attributes, expected performance, personal fit, quality, risk, or value, but each additional source requires time, attention, and cognitive effort (Bettman et al., 1998; Moorthy et al., 1997; Ratchford, 2001; Stigler, 1961). Search therefore continues only while the expected value of additional information exceeds its perceived cost. Recent work integrates several models of this process and emphasizes that search is fundamentally an allocation of limited attention in the service of choice (Krefeld-Schwalb et al., 2026).

In a traditional SERP, the platform retrieves and ranks potentially relevant information, but it does not resolve the consumer's information problem. Consumers must still decide which sources merit attention, in what sequence to open and inspect them, determine which claims are relevant and credible, compare information presented in different formats, and integrate sometimes fragmented or conflicting evidence into a coherent understanding of the product category (Diehl, 2005; Ghose et al., 2014). Thus, although ranking reduces the cost of locating information, the cognitive work of constructing an answer remains largely with the consumer.

Interactive decision aids and recommendation agents redistribute part of this work to the interface. Decision aids allow consumers to filter, sort, and compare alternatives using selected attributes, whereas recommendation agents can infer preferences and prioritize options that appear

155 most compatible with them (Baccelloni et al., 2026; Häubl & Trifts, 2000). For example, a laptop
156 recommendation agent may ask consumers to indicate their preferred price, battery life, portability,
157 and intended use and then rank the available models accordingly. These tools reduce the size and
158 complexity of the choice set, but consumers ordinarily remain responsible for examining the
159 recommended alternatives and interpreting the attribute information on which the ranking is based.
160 Moreover, because the interface determines which options and attributes become prominent, such
161 assistance can direct—and potentially narrow—consumers’ attention and consideration sets
162 (Baccelloni et al., 2026; Häubl & Murray, 2003). This distinction provides an important baseline for
163 AI Overviews: conventional decision aids support consumers in navigating and evaluating
164 marketplace information, whereas generative synthesis can perform part of the integrative work
165 through which consumers construct the information they will subsequently use.

166 AI Overviews relocate this assistance to an earlier point in the search process. They do not
167 necessarily recommend a product; instead, they synthesize category-level information before
168 consumers visit individual sources. In this sense, they function as upstream shopping support,
169 structuring the informational basis on which later product evaluation occurs (Liu et al., 2024).
170 Whereas source-led search requires consumers to construct an answer from ranked links, summary-
171 led search begins with an integrated answer and relegates sources to optional supporting material.
172 Recent evidence that learning from LLM syntheses requires less effort than learning through web
173 links — even when the underlying information is comparable — underscores the importance of this
174 distinction (Melumad & Yun, 2025).

175 A synthesized answer should reduce the perceived incremental value of source navigation.
176 When the overview appears sufficiently complete, opening another source offers less immediate
177 benefit because the interface has already selected, organized, and integrated relevant information.
178 Consumers should therefore open fewer sources and spend less time consulting them during initial
179 product search. The prediction concerns observable search behavior rather than stated trust or
180 acceptance. Formally, we hypothesize the following:

- H1a. AI Overview reduces the number of source clicks during initial product search relative to a traditional SERP.
- H1b. AI Overview reduces time spent on sources during initial product search relative to a traditional SERP.

These predictions should apply most directly to consumers' initial acquisition of product information. During this stage, consumers are constructing an informational basis for judgment: identifying relevant attributes, learning how the product works, and determining which criteria should guide evaluation. By selecting, organizing, and integrating information before consumers visit individual sources, an AI Overview performs part of this constructive work. It therefore reduces the incremental value of opening sources at precisely the point at which consumers would otherwise use them to assemble an initial understanding of the product category.

The same effect may not necessarily extend to subsequent verification. Once consumers have formed an initial understanding, source consultation serves a different purpose: rather than discovering and integrating information, consumers may use it to confirm, qualify, or challenge the answer already encountered. An AI Overview could discourage this checking by creating a sense that sufficient information has already been acquired. At the same time, awareness that the summary is AI-generated could prompt caution and motivate consumers to inspect the underlying evidence. Accordingly, we do not make a directional prediction about verification behavior. Instead, we examine whether the reduction in source engagement extends beyond initial information acquisition to a later, optional verification stage. This distinction allows us to establish whether AI Overviews cause a general withdrawal from source engagement or more specifically compress the initial construction of product knowledge. Formally:

RQ1: Does exposure to an AI Overview affect consumers' post-exposure verification behavior?

207 2.2 Product diagnosticity and attitudes toward AI as joint boundary conditions

208 The value consumers assign to an AI-generated synthesis should depend partly on the type of
209 product attributes to be evaluated. Marketing research distinguishes among search, experience, and
210 credence attributes (Darby & Karni, 1973; Nelson, 1970). Search attributes can be evaluated before
211 purchase through information acquisition. Experience attributes are difficult to assess without use,
212 whereas credence attributes may remain difficult to evaluate even after consumption because their
213 assessment requires specialized knowledge or trusted expertise (Ford et al., 1990; Klein, 1998; Mitra
214 et al., 1999).

215 This classification shapes the extent to which an AI Overview can inform the decision. For a
216 search good such as a notebook computer, specifications including processor, memory, storage,
217 display, and price are comparatively codifiable and verifiable. For an experience good such as a VR
218 headset, comfort, immersion, usability, and motion-sickness risk are harder to infer from a textual
219 synthesis. For a credence good such as an omega-3 supplement, consumers must interpret technical
220 or health-related signals such as concentration, purity, chemical form, certification, and dosage.
221 Automated product information may therefore perform differently across categories (Roelen-
222 Blasberg et al., 2023). Yet the direction of a standalone product-type effect is not clear: difficult-to-
223 evaluate information may make synthesis more valuable, but it may also increase the need for external
224 checking. We consequently examine product type alone as a secondary boundary element rather than
225 propose a directional hypothesis.

226 How consumers respond to the same informational challenge should also depend on their
227 orientation toward AI. Consumer responses to AI are context dependent and reflect beliefs about its
228 usefulness, competence, appropriateness, risks, and implications for autonomy (Castelo et al., 2019;
229 Dorotic et al., 2024; Puntoni et al., 2021). Consumers with more positive attitudes toward AI should
230 be more willing to treat an AI Overview as useful search support, whereas less AI-positive consumers
231 may see the same summary as an additional layer that requires scrutiny. Recommendation-agent

232 research similarly shows that personalization can facilitate choice while perceptions of narrow
233 knowledge or reduced control can undermine acceptance (Baccelloni et al., 2026).

234 Product diagnosticity and AI attitudes should therefore be considered jointly. When attributes
235 are concrete and readily verifiable, consumers may use the overview regardless of their general AI
236 orientation. When attributes are experiential or credence-based, willingness to rely on a synthesized
237 interpretation should depend more strongly on whether consumers view AI as capable and useful in
238 that context. Thus, product type alone need not produce a uniform moderation pattern; category
239 characteristics should become consequential in combination with consumers' positive attitudes
240 toward AI.

- 241 • H2. Product type and positive attitudes toward AI jointly moderate the effect of AI
242 Overview on source navigation.

243

244 *2.3 Epistemic vigilance and downstream purchase responses*

245 Summary-led search may affect both information acquisition and the scrutiny consumers apply
246 to the information they encounter. Epistemic vigilance refers to the cognitive and social mechanisms
247 through which individuals protect themselves from being misinformed by communicated claims
248 (Sperber et al., 2010). It involves evaluating both the content of a message—whether its claims are
249 plausible, internally consistent, and supported by evidence—and its source—whether the
250 communicator is competent, knowledgeable, and trustworthy. In online environments, these
251 judgments commonly rely on cues such as source identity, expertise, supporting evidence, agreement
252 across sources, and corroboration with prior knowledge (Metzger et al., 2010). Epistemic vigilance
253 is particularly relevant to product search because consumers must judge both what information is
254 available and whether it provides a credible basis for evaluation and choice.

255 The design of the search interface may influence how readily these safeguards are activated. A
256 traditional SERP presents sources as distinct information objects. Consumers encounter separate
257 titles, publishers, and excerpts and must decide which pages deserve attention, whether individual

258 claims are credible, and how evidence from different sources should be reconciled. Although this
259 process requires effort, it keeps source identity, variation across accounts, and the need for
260 comparison relatively salient. The consumer remains visibly responsible for constructing,
261 sequencing, and evaluating the informational basis of the decision.

262 An AI Overview changes this evaluative environment by presenting claims from multiple
263 sources as a single, coherent response. In addition to reducing the effort required to integrate
264 information, synthesis can obscure the boundaries among individual claims and the sources from
265 which they originate. Citations may remain technically available, but they become supporting
266 elements of an already-constructed answer rather than separate information objects that consumers
267 must actively select and assess. Fluency and coherence may further give the response the appearance
268 of completeness, even when the underlying sources vary in expertise, evidentiary quality, or
269 interpretation. Consequently, provenance, disagreement, omissions, and the possibility that particular
270 claims require corroboration may become less behaviorally salient.

271 Emerging research on generative systems supports this possibility. Users may shift from
272 evaluating individual sources to evaluating whether the system appears to have synthesized them
273 appropriately, and they may interpret the inclusion of citations as an implicit endorsement of those
274 sources without opening or inspecting them directly (Dabran-Zivan et al., 2026). In effect, part of the
275 responsibility for source selection and evidentiary integration is transferred from the consumer to the
276 interface. Because the answer has already been organized and presented as a unified account,
277 consumers may perceive less need to question its claims, examine the competence of the underlying
278 sources, or seek corroborating evidence. We therefore predict:

- 279 • H3. AI Overview reduces epistemic vigilance toward product-search information relative
280 to a traditional SERP.

281

282 Reduced epistemic vigilance may, in turn, shape how readily consumers translate product
283 information into purchase responses. When epistemic vigilance is high, consumers subject

284 communicated claims to greater scrutiny: they question whether the information is sufficiently
285 supported, consider the competence and motives of its sources, and remain alert to omissions or
286 alternative interpretations. Such scrutiny introduces uncertainty into the movement from information
287 exposure to product evaluation. Even when the information is favorable, consumers may hesitate to
288 rely on it until its evidentiary basis has been examined or corroborated. When epistemic vigilance is
289 lower, by contrast, consumers are less likely to challenge the claims or investigate the sources behind
290 them. Favorable product information can therefore be incorporated more readily into the evaluation
291 of the product, increasing purchase intention.

292 Importantly, this proposed pathway is distinct from the behavioral reduction in source
293 navigation. Consumers may open fewer sources because an AI Overview makes additional
294 information acquisition appear less valuable: the interface has already selected, organized, and
295 synthesized an answer. Epistemic vigilance instead concerns how critically consumers evaluate the
296 information already available to them. A consumer could therefore perceive little need to acquire
297 further information while still evaluating the overview critically; conversely, a consumer could
298 consult several sources without carefully assessing their credibility. Search compression and reduced
299 epistemic vigilance are thus conceptually related consequences of the same interface, but they need
300 not constitute a single psychological process. We consequently propose that AI Overview influences
301 purchase intention indirectly by lowering epistemic vigilance, even if the interface has no direct effect
302 on purchase intention and even if epistemic vigilance does not account for the reduction in clicks or
303 source-viewing time.

- 304 • H4. AI Overview indirectly increases purchase intention by reducing epistemic vigilance.

305

306 *2.4 Overview of the studies*

307 The empirical sequence follows this theoretical progression. Study 1 tests the predicted effect
308 of AI Overview on initial source clicks and source-viewing time and examines whether it extends to
309 optional verification. Study 2 tests the generalizability of the behavioral effect across search,

310 experience, and credence product categories and evaluates product type and positive AI attitudes as
311 joint moderators. Study 3 replicates the behavioral effect in a health-related smartwatch context and
312 examines whether the data are consistent with a statistical indirect pathway linking AI Overview
313 exposure, epistemic vigilance, and purchase intention.

314 Across all our experiments, we employed a realistic, interactive product-search scenario
315 designed to closely reproduce how consumers ordinarily use a search engine when gathering
316 information before a purchase. Rather than merely viewing static stimuli or reporting how they
317 believed they would search, participants could navigate the simulated results page, open underlying
318 sources, return to the search interface, and decide when they had acquired sufficient information. This
319 design also allowed us to capture actual search behavior—including the number of sources opened
320 and the time spent consulting them—rather than relying solely on self-reported engagement or
321 intended behavior.

322

323 **3. Study 1: Direct effect of AI Overview on source navigation**

324 *3.1 Design and sample*

325 Aiming to test H1a and H1b, Study 1 examined the direct effect of AI Overview on consumer
326 source navigation in a single product-search context. Participants were asked to imagine that they
327 were looking for a cabin trolley and wanted to understand which characteristics make a trolley
328 resistant, safe, and suitable for air travel. The study used a between-subjects design with two interface
329 conditions: a traditional SERP and a SERP with an AI Overview. A total of 205 responses were
330 collected. Following pre-specified quality checks and data-cleaning procedures, 46 responses were
331 excluded: 4 were identified as bot-like responses, 4 were duplicates, 2 were incomplete, and 36 were
332 excluded due to inconsistencies in the recording of embedded behavioral data. The final sample
333 included 159 valid Italian respondents: 75 in the SERP only condition and 84 in the AI Overview
334 condition. Product expertise was included as a covariate.


335 *3.2 Procedure and stimuli*

336 Participants first read a short scenario in Italian and a knowledge primer about cabin trolleys.
337 They were then shown a simulated Google Search results page modeled on an actual Google search
338 for the query “How to tell if a trolley is sturdy?”, reproducing a real screenshot, not an artificially
339 constructed layout. The page was presented as a static, simulated interface rather than a live search,
340 since navigating to Google directly would have returned a different, non-controlled results page for
341 the same query at the time of participation. In the SERP-only condition, participants saw a traditional
342 list of clickable search results, drawn from the same underlying screenshot. In the AI Overview
343 condition, participants saw the same search context with an AI-generated summary positioned above
344 the organic results (see Figure 1; the complete set of stimuli is available in Web Appendix).

345

346

Figure 1. AI Overview interface



347

348

349 Participants could open sources, return to the results page, and continue once they felt they had
350 enough information. The interfaces and accessible sources were recreated from screenshots of actual

351 Google results and webpages. HTML and CSS reproduced the layouts. JavaScript recorded
352 navigation behavior within Qualtrics, including clicks, viewing time, and source navigation. Both
353 interface conditions displayed the same underlying results and sources. The AI Overview condition
354 additionally included an expandable AI-generated summary. Participants could freely explore
355 clickable sources during the exposure phase. Opened sources were shown as partial webpage
356 screenshots. Immediately after the exposure phase, participants completed the purchase-intention
357 measure.

358 Afterwards, participants entered a verification phase (see Figure 2) in which they were shown
359 three additional trolley-related sources and could consult any of them or choose to skip the
360 verification step altogether.

361

362 Figure 2. Verification phase, source selection

**Seleziona una fonte da consultare, oppure continua con il
questionario.**

1. Trolley resistenti e sicuri: i materiali chiave

2. Ruote e manico: come influenzano la tenuta di un trolley?

3. L'importanza di zip e chiusure nei trolley da cabina

Continua con il questionario

363

364 Note. The heading shown to participants reads: ‘Select a source to consult or continue with the questionnaire’.

365

366 Specifically, participants were given the following instructions: “Now, if you think it’s needed,
367 you can check what you’ve read about trolleys by reading one or more in-depth articles. Once you’ve
368 opened an article, you can return to this page to open another one or carry on filling in the
369 questionnaire”. Verification sources provided targeted trolley-related information from different,

370 non-overlapping angles, with an approximately equal word count (150 words). This phase was
371 included to observe whether the initial interface type changed subsequent verification behavior.

372 *3.3 Measures*

373 The main dependent variables were two different realistic behaviors, namely exposure clicks
374 count and exposure time on sources. Exposure clicks count captured the total number of source-
375 opening events during the initial search phase; repeated openings of the same source were counted
376 separately. Exposure time on sources corresponded to the total time of exposure to all opened sources
377 combined – thus excluding the time spent on the SERP interface before or between clicks – and was
378 log-transformed as $\ln(\text{time} + 1)$. We also measured verification clicks count and verification time on
379 sources during the post-exposure verification phase, informing respondents that, if they wished, they
380 could verify the information they had viewed by visiting up to three relevant sources. Count outcomes
381 were analyzed using Poisson generalized linear models with a log link and heteroskedasticity-robust
382 standard errors. Log-transformed time outcomes were analyzed using ordinary least squares models.
383 Interface type was effect-coded, and product expertise was included as a covariate in all models.
384 Purchase intention was measured with three items adapted from Dodds, Monroe, and Grewal (1991),
385 capturing the perceived likelihood, consideration, and propensity to purchase a trolley. Product
386 expertise was measured with three trolley-specific items adapted from Trzebinski et al. (2023) and
387 used as a covariate: (1) I keep up to date with the latest developments in the cabin trolley market; (2)
388 I could give excellent advice on cabin trolleys if asked; (3) I know a great deal about cabin trolleys.
389 The purchase-intention scale showed high internal consistency (Cronbach's $\alpha = .940$). The three-item
390 product-expertise measure, though, showed lower internal consistency (Cronbach's $\alpha = .658$).

391 *3.4 Results*

392 Study 1 supports the effect predicted in H1a and H1b: when looking at exposure clicks count, the
393 effect of interface type was significant (Wald $\chi^2 = 10.414$, $p = .001$). In particular, participants in the
394 AI Overview condition made fewer source-opening clicks during the initial exposure phase than

395 participants in the traditional SERP condition. The effect of product expertise was not significant.
396 The effect remained unchanged when product expertise was omitted from the models.

397 For exposure time on sources, the effect of interface type was also significant ($F = 9.769$, $p =$
398 $.002$, log-transformed SERP condition mean = 4.32, AI Overview = 3.81). Participants in the AI
399 Overview condition spent less time consulting sources than participants in the SERP condition. The
400 effect of product expertise was not significant. Also in this case, the effect remained unchanged when
401 product expertise was omitted from the models.

402 The interface-type difference observed during initial exposure was no longer present in the
403 verification phase. Participants previously exposed to an AI Overview did not differ significantly
404 from those exposed to the traditional SERP in either verification click count (Wald $\chi^2 = .427$, $p =$
405 $.514$) or verification time on sources ($F = .053$, $p = .818$). The effect on purchase intention was also
406 not significant ($F = .362$, $p = .548$). These findings indicate that AI Overview exposure reduces initial
407 source navigation, whereas subsequent verification behavior no longer differs between the two
408 interface conditions.

409 *3.5 Discussion*

410 Using a realistic, navigable product-search environment, Study 1 provides initial evidence that
411 placing an AI-generated synthesis above conventional search results compresses consumers'
412 engagement with the underlying sources. Participants exposed to an AI Overview opened fewer
413 sources and spent less time consulting them than participants exposed to a traditional SERP. This
414 difference reflects a change in how that information was organized and presented. By selecting and
415 integrating relevant content before consumers navigate individual webpages, the AI Overview
416 appears to reduce the perceived incremental value of consulting those pages directly.

417 The convergence between clicks and viewing time is important, with a reduction in clicks
418 indicating that consumers were less likely to move from the search interface to the underlying sources,
419 and reduced source-viewing time reflecting lower overall engagement with source content. These
420 measures show that the AI overview affected more than a superficial interface action: it reduced both

421 the frequency with which consumers accessed sources and the attention allocated to those sources.
422 This supports the conceptualization of AI Overviews as a shift from source-led to summary-led
423 search. In a traditional SERP, consumers must construct an answer by selecting sources, comparing
424 their claims, and integrating their content. When an AI Overview is present, part of this constructive
425 work is performed by the interface, allowing consumers to obtain an initial understanding of the
426 product while undertaking less independent navigation.

427 This pattern also clarifies how AI Overviews differ from conventional search aids. Ranking
428 already lowers the cost of locating potentially relevant information, but it leaves consumers
429 responsible for opening and interpreting the results. An AI Overview goes further by providing a
430 synthesized answer before source consultation occurs. The observed reduction in navigation therefore
431 suggests that synthesis does not merely supplement the ranked links; it can partially substitute for
432 them. Even though the underlying sources remain visible and accessible, their behavioral role changes
433 once an apparently complete answer precedes them.

434 This conclusion should be read alongside the absence of significant differences during the
435 subsequent verification phase. Study 1 does not indicate that exposure to an AI Overview produces a
436 uniform withdrawal from source engagement across the entire search process. Instead, its immediate
437 effect appears concentrated in initial information acquisition, when consumers are attempting to
438 identify relevant product attributes and construct a preliminary understanding of the category. Once
439 participants were explicitly offered an opportunity to verify what they had encountered, those
440 previously exposed to an AI Overview did not consult fewer sources or spend less time on them than
441 participants who had viewed a traditional SERP. Thus, the AI overview appears to compress the
442 initial construction of product knowledge without necessarily eliminating consumers' willingness to
443 check information later.

444 This distinction suggests that initial acquisition and verification serve different informational
445 purposes. During initial search, source consultation helps consumers discover and integrate
446 information; a synthesized answer can perform much of that function and thereby reduce the value of

opening additional pages. During verification, however, sources are consulted to confirm, qualify, or challenge an answer that has already been formed. The AI Overview may therefore reduce the need to construct an answer without equivalently reducing the potential value of checking it. Because the verification effects were nonsignificant, this interpretation should not be taken as evidence that AI Overviews actively stimulate verification. Rather, the findings show that the reduction observed during initial exposure did not reliably carry over to the later checking opportunity.

Study 1 also found no significant effect on purchase intention. This result indicates that reduced source navigation does not automatically translate into a more favorable product response. Consumers can acquire less information from underlying sources without necessarily changing their willingness to purchase the focal product. The behavioral effect should therefore not be interpreted as direct persuasion or as evidence that consumers uncritically accepted the overview. Instead, Study 1 establishes a more specific consequence: AI-generated synthesis changes the amount of independent source engagement consumers undertake when forming an initial understanding of a product. This finding provides the behavioral foundation for Studies 2 and 3, which examine when this compression varies and whether summary-led search has a separate effect on how critically consumers evaluate the information they encounter.

463

4. Study 2: Product-type and AI-attitude boundary conditions

4.1 Design and sample

Aiming to test H2, Study 2 examined whether the direct effect observed in Study 1 generalizes across product categories and is affected by consumers' positive attitudes toward AI. The study used a 3×2 between-subjects design, whereby participants were randomly assigned to one of three product categories and one of two interface conditions. The product categories were notebook (search product, Mazzù et al., 2025), VR headset (experience product, Mazzù et al., 2025), and omega-3 supplements (credence product, Nagler et al., 2011). The interface conditions were the same as those in Study 1, namely a traditional SERP and one with the AI Overview interface.

473 A total of 334 responses were collected. Following pre-specified quality checks and data-cleaning
474 procedures, 69 responses were excluded: 24 were identified as bot-like responses, 19 were duplicates,
475 3 did not provide consent, 16 were incomplete, 2 failed the attention check, and 5 were completed on
476 tablet or smartphone. The final dataset contained 265 valid responses.

477 Cell sizes were balanced enough for the planned experimental comparisons. In the notebook
478 condition, 37 participants were assigned to the SERP condition and 45 to the AI Overview condition.
479 In the VR condition, 45 were assigned to SERP and 45 to AI Overview. In the omega-3 condition, 46
480 were assigned to SERP and 47 to AI Overview.

481 *4.2 Procedure and stimuli*

482 Participants first completed measures of positive and negative attitudes toward AI. Product
483 expertise was assessed later in the questionnaire and used as a covariate in the analyses. They then
484 received a product-specific scenario asking them to imagine they wanted to understand how to
485 recognize a good product in the assigned category.

486 The query was held constant across categories except for the target product: “How to recognise a
487 good [target product]?”. Participants then saw either a simulated traditional Google SERP or a
488 simulated Google SERP with an AI Overview. Both conditions were designed to resemble a realistic
489 product-search environment. The same procedures employed in Study 1 were used to design the
490 SERP and AI Overview interfaces during the exposure phase, as well as the subsequent verification
491 phase.

492 The exposure and verification procedures followed Study 1. Participants interacted with the same
493 type of simulated Google SERP and could open and revisit the underlying sources. A procedural
494 difference was present, as in Study 2 opened sources were displayed as full-page webpage
495 screenshots. After exposure, participants could consult up to three additional product-specific sources
496 or skip the verification phase.

497 *4.3 Measures*

498 Study 2 used the same real behavioral measures as Study 1. The four focal outcomes were
499 exposure clicks count, exposure time on sources, verification click count, and verification time on
500 sources. Exposure clicks count captured the total number of source-opening events during the
501 exposure phase; repeated openings of the same source were counted separately. Exposure time
502 captured the total time spent consulting opened sources and was log-transformed as $\ln(\text{time} + 1)$.
503 Verification clicks count captured the total number of source-opening events during the verification
504 phase, including repeated openings, and verification time was again log-transformed as $\ln(\text{time} + 1)$.
505 Positive attitudes toward AI were measured with the 12-item positive GAAIS subscale (Sacco et al.,
506 2025; Schepman & Rodway, 2023), which showed high internal consistency (Cronbach's $\alpha = .909$).
507 The negative GAAIS subscale was also administered; because H2 focused on positive attitudes
508 toward AI, supplementary models using the negative subscale are reported in Web Appendix. Product
509 expertise was measured with three product-specific items adapted from Trzebinski et al. (2023),
510 showed acceptable reliability (Cronbach's $\alpha = .855$), and was included as a covariate in our analysis.

511 We estimated separate models for exposure clicks count, exposure time, verification click count,
512 and verification time. Count outcomes were analyzed using Poisson generalized linear models with a
513 log link and heteroskedasticity-robust standard errors. Log-transformed time outcomes were analyzed
514 using ordinary least squares models. Product expertise was included as a covariate throughout.
515 Interface type and product type were effect-coded, and positive GAAIS was mean-centered.

516 The analyses proceeded in three steps. Step 1 tested the main effect of interface type, providing a
517 replication test of Study 1. Step 2 added product type and the interface type $\times$ product type interaction.
518 Step 3 added positive GAAIS and included all constituent two-way interactions and the interface type
519 $\times$ product type $\times$ positive GAAIS three-way interaction. Multi-parameter effects were evaluated using
520 Type III Wald tests for count outcomes and Type III F tests for time outcomes. Conditional effects
521 were probed at low, mean, and high positive GAAIS, defined as one standard deviation below the
522 mean, the mean, and one standard deviation above the mean, respectively (3.316, 4.426, and 5.537).

523 H2 was evaluated through the three-way interaction for the exposure outcomes, whereas the
524 verification analyses addressed RQ1.

525 *4.4. Results*

526 *4.4.1 Main effects of interface type*

527 Step 1 tested whether the main effect observed in Study 1 replicated in a multi-category design.
528 Specifically, we examined whether interface type affected search behavior during exposure and
529 verification while controlling for product expertise. For exposure clicks count, the effect of interface
530 type was significant (Wald $\chi^2 = 21.641$, $p < .001$): participants exposed to the AI Overview made
531 fewer source-opening clicks during initial exposure than participants exposed to the traditional SERP.
532 Product expertise was not significant.

533 For exposure time on sources, the effect of interface type was also significant ($F = 14.972$, $p <$
534 $.001$): participants in the AI Overview condition spent less time consulting sources during exposure.
535 Product expertise was significant in this model ($F = 4.038$, $p = .046$), while the interface-type effect
536 remained significant after controlling for expertise.

537 The interface type effect was not significant in the verification phase. Interface type did not
538 significantly affect verification click count (Wald $\chi^2 = 1.579$, $p = .209$). It also did not significantly
539 affect verification time on sources ($F = 2.250$, $p = .135$). Product expertise was not significant in
540 either verification model.

541 Overall, analyses performed in step 1 support H1a and H1b. AI Overview reduces source
542 navigation during initial exposure. In the verification stage, instead, there was no evidence supporting
543 a general carryover to verification behavior.

544 *4.4.2 Moderation by product type*

545 Step 2 tested whether product type moderated the interface type effect. For exposure clicks count,
546 the interface type effect remained significant (Wald $\chi^2 = 17.665$, $p < .001$), but neither the main effect
547 of product type (Wald $\chi^2 = .160$, $p = .923$) nor the effect of the interface type $\times$ product type interaction
548 (Wald $\chi^2 = 1.946$, $p = .378$) was significant.

549 Similar pattern was observed when looking at exposure time on sources. The interface type effect
550 remained significant ($F = 15.421, p < .001$), while neither the effect of product type ($F = 2.078, p =$
551 $.127$) or the effect of the interface type $\times$ product type interaction ($F = 1.891, p = .153$) was significant.
552 Product expertise resulted significant ($F = 5.256, p = .023$).

553 For verification click count, neither interface type (Wald $\chi^2 = 2.046, p = .153$) nor product type
554 (Wald $\chi^2 = 4.699, p = .095$) nor the interface type $\times$ product type interaction (Wald $\chi^2 = .908, p =$
555 $.635$) was significant. For verification time, the interface type $\times$ product type interaction was again
556 not significant ($F = .169, p = .844$).

557 Overall, analyses performed in step 2 show that product type alone does not moderate the effect
558 of AI Overview on our focal outcomes.


559 *4.4.3 Moderation by product type and positive attitudes toward AI*

560 In step 3 we introduced positive attitudes toward AI and tested whether the interface type effect
561 depended jointly on product type and GAAIS positive.

562 For exposure clicks count, interface type remained significant (Wald $\chi^2 = 25.851, p < .001$), while
563 neither the effect of interface type $\times$ product type interaction (Wald $\chi^2 = 1.981, p = .371$) nor the effect
564 of interface type $\times$ positive GAAIS interaction (Wald $\chi^2 = .011, p = .918$) were significant. However,
565 the three-way interaction among interface type, product type, and positive GAAIS was found to be
566 significant (Wald $\chi^2 = 6.303, p = .043$).

567

568 Figure 3. Effect of interface type across low, medium, and high values of positive GAAIS



569

570

571

572

573

574

575

576

577

578

579

580

581

582

To clarify how the effect of interface type varied across search (notebook), experience (VR headset), and credence (omega-3 supplements) goods, conditional effects were tested at low (-1 SD), mean, and high (+1 SD) levels of positive GAAIS, as shown in figure 3. For the Poisson models, conditional effects are reported as incidence rate ratios (IRRs), obtained by exponentiating the regression coefficients; values below 1 indicate fewer expected clicks in the AI Overview condition relative to the traditional SERP. At low positive GAAIS, AI overview significantly reduced exposure clicks relative to SERP only for the experience good, VR headset (IRR = .470, $p < .001$; predicted means SERP = 2.74, AI Overview = 1.29), while the contrast was not significant either for the search good, notebook (IRR = .771, $p = .209$; predicted means SERP = 2.45, AI Overview = 1.89) or for the credence good, omega-3 supplements (IRR = .821, $p = .275$; predicted means SERP = 2.46, AI Overview = 2.02). At mean positive GAAIS, the reduction reached the conventional significance threshold for both the experience good, VR headset (IRR = .590, $p < .001$, predicted means SERP =

2.54, AI Overview = 1.50), and the credence good, omega-3 supplements (IRR = .639, $p = .002$, predicted means SERP = 2.39, AI Overview = 1.53), whereas the contrast for the search good, notebook, was marginally significant (IRR = .772, $p = .054$, predicted means SERP = 2.22, AI Overview = 1.72). At high positive GAAIS, the pattern shifted again, as AI overview significantly reduced exposure clicks for the credence good, omega-3 supplements (IRR = .497, $p < .001$, predicted means SERP = 2.33 versus, AI Overview = 1.16), but not for the search good, notebook (IRR = .773, $p = .109$, predicted means SERP = 2.02, AI Overview = 1.56), or the experience good, VR headset (IRR = .739, $p = .193$, predicted means SERP = 2.36, AI Overview = 1.74). Overall, the reduction in exposure clicks associated with AI overview was concentrated on the experience good at low positive GAAIS, extended to both the experience and credence good at mean positive GAAIS, and influenced significantly only the credence good at high positive GAAIS. The three-way interaction therefore shows that the joint effect of interface type and product category changes in magnitude and significance across GAAIS levels.

For exposure time on sources, the interface type effect remained significant ($F = 16.118$, $p < .001$), although neither the effect of the interface type $\times$ product type interaction ($F = 1.202$, $p = .302$) nor the effect of the interface type $\times$ positive GAAIS interaction ($F = .087$, $p = .768$) proved significant. The three-way interaction was not significant either ($F = 2.123$, $p = .122$); thus, the reduction in exposure time was more general than the product- and attitude-dependent pattern observed for exposure clicks.

For verification click count, the interface type effect was not significant (Wald $\chi^2 = 2.393$, $p = .122$) and neither the effect of interface type $\times$ product type interaction (Wald $\chi^2 = 1.045$, $p = .593$) nor the three-way interaction (Wald $\chi^2 = 1.149$, $p = .563$) was significant, while the effect of interface type $\times$ positive GAAIS interaction was significant (Wald $\chi^2 = 7.225$, $p = .007$). At low positive GAAIS, exposure to an AI Overview increased verification clicks relative to a traditional SERP ($b = .620$, $SE = .246$, $z = 2.516$, $p = .012$, IRR = 1.858, 95% CI for the IRR [1.147, 3.011]). The interface type difference was not significant at mean positive GAAIS ($b = .239$, $p = .122$) or high positive

609 GAAIS ($b = -.142$, $p = .389$). Because product type did not further qualify the interaction, this
610 compensatory-verification pattern was not category-specific.

611 For verification time, the average interface type effect was not significant ($F(1, 252) = 2.748$, $p =$
612 $.099$). Neither the interface type $\times$ product type interaction ($F(2, 252) = .130$, $p = .878$) nor the three-
613 way interaction ($F(2, 252) = .510$, $p = .601$) was significant. Interestingly, the effect of the interface
614 type $\times$ positive GAAIS interaction was significant ($F(1, 252) = 5.665$, $p = .018$). At low positive
615 GAAIS, participants spent more time on verification sources after exposure to an AI Overview than
616 after exposure to a traditional SERP ($b = .824$, $SE = .294$, $t = 2.802$, $p = .005$, 95% CI [.245, 1.402]).
617 The conditional effect was not significant at mean positive GAAIS ($b = .335$, $p = .099$), or high
618 positive GAAIS ($b = -.153$, $p = .588$). The verification-time results therefore reproduce the pattern
619 observed for verification clicks. The significant three-way interaction for exposure clicks counts,
620 combined with the nonsignificant interaction for exposure time, provides partial support for H2.
621 Product type and positive AI attitudes jointly qualified source-opening clicks, but not the time spent
622 consulting sources. The conditional pattern did not follow a monotonic search–experience–credence
623 ordering. Instead, the significant interface effect shifted from the experience good at low positive
624 GAAIS to the experience and credence goods at mean positive GAAIS, and to the credence good at
625 high positive GAAIS. The results were robust when click counts were replaced by the number of
626 distinct sources opened. The interface type main effect remained significant (Wald $\chi^2 = 23.995$, $p <$
627 $.001$) as did the interface type $\times$ product type $\times$ positive GAAIS interaction effect (Wald $\chi^2 = 6.873$,
628 $p = .032$). The effect of the interface type $\times$ positive GAAIS interaction also remained significant for
629 the number of distinct verification sources consulted (Wald $\chi^2 = 7.577$, $p = .006$).

630 4.4.4 Discussion

631 Study 2 replicates and extends the behavioral effect observed in Study 1. Across product
632 categories, exposure to an AI Overview reduced both the number of source-opening clicks and the
633 time spent consulting underlying sources during initial search. This replication suggests that
634 summary-led search consistently reduces consumers' direct engagement with source material.

635 However, the significant three-way interaction for exposure clicks shows that the magnitude of this
636 reduction cannot be understood from either product type or attitudes toward AI alone. Instead,
637 consumers' willingness to let the overview substitute for source navigation depends on the
638 combination of what must be learned about the product and how positively they view AI as a source
639 of assistance.

640 This conditional pattern clarifies the joint dependence, showing that for the experience good,
641 the AI Overview reduced source opening even among consumers with low positive attitudes toward
642 AI. Experience attributes such as comfort, immersion, and usability cannot be established
643 conclusively through webpage consultation because they ultimately require product use. Opening
644 additional sources may therefore offer only limited resolution of the underlying uncertainty. In this
645 context, even consumers who are not especially favorable toward AI may accept the overview as an
646 efficient provisional synthesis, because conventional source navigation cannot fully substitute for
647 direct experience. As positive attitudes toward AI increased, however, the relative reduction in clicks
648 for the experience good became less pronounced. One possible interpretation is that AI-positive
649 consumers regard both the overview and its underlying digital sources as useful inputs when forming
650 expectations about an unfamiliar experience, making the overview less likely to replace source
651 consultation completely.

652 The credence good followed a different pattern, with the reduction in source opening for omega-
653 3 supplements growing progressively stronger as positive attitudes toward AI increased. Evaluating
654 a supplement requires consumers to interpret specialized claims concerning concentration, purity,
655 chemical form, certification, and dosage—claims that may remain difficult to assess even after
656 consumption. Under these conditions, source navigation is potentially valuable but cognitively
657 demanding. Consumers who view AI positively may therefore be especially willing to delegate the
658 organization and interpretation of this technical information to the overview, treating its synthesis as
659 an adequate substitute for consulting the underlying sources themselves. Consumers who are less
660 positively disposed toward AI appear less willing to make this substitution, presumably because the

661 same difficulty that increases the value of synthesis also raises the perceived cost of relying on an
662 automated interpretation. Thus, the results suggest that positive AI attitudes become particularly
663 consequential when consumers must decide whether to delegate the interpretation of information they
664 cannot readily evaluate on their own.

665 For the search good, by contrast, the reduction in clicks was comparatively stable across levels
666 of positive AI attitudes, although the conditional contrasts did not consistently reach statistical
667 significance. Notebook attributes are relatively codifiable and can be checked directly across sources.
668 Consumers therefore retain a comparatively clear route for independently evaluating the overview,
669 which may limit the extent to which their general orientation toward AI determines whether they
670 consult additional sources. More broadly, the three-way interaction does not support a simple
671 progression in which AI Overviews become uniformly more influential as products move from search
672 to experience and credence qualities. It instead indicates that different forms of informational
673 uncertainty interact differently with consumers' readiness to rely on AI. Experience-related
674 uncertainty can make synthesis useful even among consumers who are not strongly AI-positive,
675 whereas credence-related uncertainty appears to produce substantial delegation mainly when
676 consumers already regard AI favorably.

677 This joint boundary condition emerged for source-opening clicks but not for source-viewing
678 time. The distinction suggests that product type and AI attitudes primarily influence the decision to
679 cross from the synthesized answer to an underlying source—the extensive margin of source
680 engagement. Once consumers decided to open a source, however, the amount of time devoted to
681 inspecting it was reduced relatively consistently by the AI Overview across categories and attitude
682 levels. In other words, the three-way interaction concerns whether consumers seek additional source
683 input, rather than how deeply they process a source after deciding to consult it. H2 therefore receives
684 partial support: product type and positive AI attitudes jointly qualified source opening, but not the
685 time spent on opened sources.

686 The verification findings reinforce this interpretation while revealing a distinct role for AI
687 attitudes. Although AI Overview exposure did not reduce verification on average, consumers low in
688 positive AI attitudes subsequently opened more verification sources and spent more time consulting
689 them. This compensatory response was not product-specific and disappeared at mean and high levels
690 of positive AI attitudes. Less AI-positive consumers may therefore use the overview during initial
691 search—particularly when ordinary source navigation cannot fully resolve uncertainty—while later
692 attempting to validate or challenge its synthesis. More AI-positive consumers, by contrast, appear
693 less likely to compensate through additional checking. Taken together, the results suggest that
694 attitudes toward AI influence not only whether consumers delegate initial information integration to
695 the interface, but also whether they subsequently reclaim evaluative responsibility through
696 verification.

697

698 **5. Study 3: Epistemic vigilance and purchase intention**

699 *5.1 Design and sample*

700 Study 3 examined epistemic vigilance as a psychological consequence of AI Overview exposure
701 (H3) and tested the hypothesized indirect effect of interface type on purchase intention through
702 epistemic vigilance (H4). As in Studies 1 and 2, we used a two-condition between-subjects design
703 comparing a traditional SERP with a SERP containing an AI Overview. The product context was in
704 this case smartwatch use for health monitoring. Participants imagined that they were looking for a
705 smartwatch to monitor their health and wanted to understand how smartwatches work, which features
706 matter, and whether the product is reliable. A total of 211 responses were collected following pre-
707 specified quality checks and data-cleaning procedures, 51 responses were excluded: 7 were identified
708 as bot-like responses, 6 were duplicates, 3 were incomplete, 1 failed the attention check, and 34 were
709 excluded due to inconsistencies in the recording of embedded behavioral data. The final sample
710 included 160 valid respondents: 90 in the SERP only condition and 70 in the AI Overview condition.
711 Again, product expertise was included as a covariate.

712 5.2 Procedure and stimuli

713 Similarly to the previous studies, participants first read a short scenario and knowledge primer in
714 Italian, about smartwatches. They were then shown a simulated Google Search page based on the
715 query “How does a smartwatch for health work?”. In the SERP only condition, participants saw a
716 traditional list of clickable results. In the AI Overview condition, participants saw an AI-generated
717 summary positioned above the same search results. Participants could open sources, return to the
718 results page, and continue once they felt they had enough information. Opened sources were shown
719 as partial webpage screenshots. Immediately after the exposure phase, participants completed the
720 purchase-intention measure. Following the same procedure used in the previous studies, participants
721 were then offered up to three additional smartwatch-related sources for optional verification and
722 subsequently completed the epistemic-vigilance and other post-search measures, including
723 processing fluency, message credibility, knowledge insufficiency, source transparency, and product
724 expertise.

725 5.3 Measures

726 Study 3 measured the same behavioral outcomes as the previous studies: exposure clicks count,
727 exposure time on sources, verification click count, and verification time on sources. Purchase
728 intention was measured using the same scale as in Study 1 (Cronbach’s $\alpha = .940$). Epistemic vigilance
729 was operationalized using nine reverse-coded items adapted from Tseng, Bonilla, and MacPherson
730 (2021), assessing reliance on the information, evidentiary support, evidence quality, source
731 competence, persuasiveness, reliability, recommendation, and information sharing. Higher scores
732 indicated greater vigilance, and the nine-item measure showed high internal consistency (Cronbach’s
733 $\alpha = .915$). The three-item product-expertise scale also showed satisfactory reliability (Cronbach’s α
734 $= .801$).

735 Because the final three vigilance items referred to recommendation and sharing rather than
736 directly to evidentiary appraisal, we repeated the focal analysis using the first six items only. These
737 robustness results are reported below and in Web Appendix.

738 Other measures operationalized included processing fluency, measured using 1 item adapted from
739 Graf et al. (2018); message credibility (Cronbach's $\alpha = .819$.), measured through 3 items adapted
740 from Appelman & Sundar (2016); source transparency (Cronbach's $\alpha = .864$.), composed of 3 items
741 adapted from the sponsor clarity dimension of Wojdyski et al. (2018). Additionally, knowledge
742 insufficiency was operationalized based on Kahlor (2010) and subsequent related literature (Guo et
743 al., 2025; Kim et al., 2020); participants rated their current knowledge and perceived need for
744 additional information on 7-point scales, and a difference score was computed such that higher values
745 represented greater information insufficiency.

746 *5.4 Results*

747 Study 3 replicated the direct behavioral effect of AI Overview on initial source navigation. For
748 exposure clicks count, the effect of interface type was significant (Wald $\chi^2 = 6.576$, $p = .010$).
749 Participants in the AI Overview condition made fewer source-opening clicks during initial exposure
750 than participants in the traditional SERP condition. Product expertise was not significant (Wald $\chi^2 =$
751 $.026$, $p = .873$).

752 For exposure time on sources, the effect of interface type was also significant ($F = 12.696$, $p <$
753 $.001$). Participants in the AI Overview condition (mean log-transformed exposure time = 3.40) spent
754 less time consulting sources than participants in the SERP condition (mean log-transformed exposure
755 time = 4.12). Product expertise was not significant ($F = .003$, $p < .958$).

756 As in the previous studies, the effect did not extend to post-exposure verification. Interface type
757 did not significantly affect verification click count (Wald $\chi^2 = .192$, $p = .661$) or verification time on
758 sources ($F = .042$, $p = .837$). The total effect of interface type on purchase intention was not significant
759 as well ($F = 2.617$, $p = .108$), while product expertise significantly predicted purchase intention ($F =$
760 58.884 , $p < .001$).

761 We tested the proposed indirect effect using PROCESS Model 4 (Hayes, 2022) with HC3 standard
762 errors and 10,000 bootstrap samples. First of all, AI Overview exposure reduced epistemic vigilance
763 ($b = -.321$, $SE = .131$, $p = .015$, 95% CI $[-.580, -.062]$), supporting H3. Exploratory models found

764 no significant indirect effects through epistemic vigilance for exposure clicks count, exposure time,
765 verification click count, or verification time.

766 For purchase intention, the effect of interface type significantly affected epistemic vigilance as
767 exposure to AI overview reduced epistemic vigilance with respect to SERP ($b = -.321$, $SE = .131$, t
768 $= -2.452$, $p = .015$, 95% CI $[-.580, -.062]$); epistemic vigilance, in turn, was negatively associated
769 with purchase intention ($b = -.440$, $SE = .117$, $p < .001$, 95% CI $[-.671, -.209]$). Thus, while the
770 direct effect of interface type remained nonsignificant ($b = .174$, $SE = .196$, $p = .375$), the mediation
771 analysis showed a significant positive indirect effect of interface type on purchase intention via
772 epistemic vigilance (indirect effect $= .141$, BootSE $= .066$, 95% bootstrap CI $[.026, .282]$): exposure
773 to AI overview was associated with lower epistemic vigilance which, in turn, was associated with
774 higher purchase intention. The corresponding analysis based on the first six epistemic-vigilance items
775 produced the same conclusion (indirect effect $= .126$, 95% bootstrap CI $[.014, .266]$).

776 When exploring for alternative explanations, namely processing fluency (indirect effect $= .039$,
777 BootSE $= .042$, 95% bootstrap CI $[-.029, .138]$), knowledge insufficiency (indirect effect $= .013$,
778 BootSE $= .045$, 95% bootstrap CI $[-.068, .119]$), message credibility (indirect effect $= .073$, BootSE
779 $= .059$, 95% bootstrap CI $[-.024, .205]$), and source transparency (indirect effect $= .008$, BootSE $=$
780 $.026$, 95% bootstrap CI $[-.042, .069]$), none of these resulted to act as significant mediator of the
781 effect of interface type on purchase intention. The observed indirect effect is consistent with H4,
782 subject to the measurement-order limitation discussed below.

783 *5.5 Discussion of Study 3*

784 Study 3 both replicates the behavioral effect found in the preceding studies and identifies a
785 distinct psychological consequence of summary-led search. In a health-related smartwatch context,
786 participants exposed to an AI Overview again opened fewer underlying sources and spent less time
787 consulting them during initial search. This replication strengthens the evidence that search
788 compression is not confined to the products examined in Studies 1 and 2. Across the three realistic,

789 navigable experiments, AI-generated synthesis consistently reduced consumers' direct engagement
790 with source content during initial information acquisition.

791 As in Study 1, this reduction did not extend to the optional verification phase. Participants
792 exposed to an AI Overview did not subsequently differ in either the number of verification sources
793 opened or the time spent consulting them. The repeated distinction between exposure and verification
794 suggests that AI Overviews primarily alter how consumers construct their initial informational basis
795 for judgment. Once consumers enter an explicit checking stage, the earlier interface does not appear
796 to determine their behavior in the same consistent way. Summary-led search therefore compresses
797 the front end of information acquisition, but the present evidence does not support the stronger
798 conclusion that it suppresses all later efforts to consult or verify sources.

799 Study 3 extends this behavioral account by showing that AI Overviews may change not only
800 how much consumers search, but also how critically they appraise the information made available to
801 them. Participants exposed to an AI Overview reported lower epistemic vigilance than those exposed
802 to a traditional SERP. In a conventional results page, individual sources remain visible and separate,
803 making differences in source identity and evidentiary support comparatively salient. Consumers must
804 decide which sources deserve attention and how their claims should be reconciled. An AI Overview
805 instead presents claims drawn from multiple sources as a single, fluent, and coherent response.
806 Although citations remain available, they become supporting elements of an already-integrated
807 answer rather than independent information objects that consumers must actively inspect and
808 compare.

809 This interface architecture may transfer part of the evaluative responsibility from the consumer
810 to the system. When the platform appears to have selected appropriate sources, reconciled their
811 content, and produced a unified answer, consumers may perceive less need to question whether
812 individual claims are adequately supported or whether the underlying sources are competent and
813 reliable. The significant reduction in epistemic vigilance is therefore consistent with the idea that
814 synthesis makes provenance and evidentiary quality less behaviorally salient. Importantly, this

815 consequence is not equivalent to an increase in perceived credibility. The alternative indirect effect
816 through message credibility was not significant. The findings instead point to reduced critical
817 scrutiny—a conceptually different process from actively judging the message to be more credible.

818 Lower epistemic vigilance was, in turn, associated with higher purchase intention, producing a
819 significant positive indirect effect of AI Overview exposure on purchase intention. When consumers
820 critically examine product claims, consider the quality of their evidence, and question source
821 competence, favorable information may be incorporated into product evaluations more cautiously.
822 Reduced vigilance can make that transition more direct: consumers may rely on favorable product
823 information without introducing the same degree of evidentiary qualification or uncertainty. AI
824 Overviews may therefore influence purchase responses not by directly making the product more
825 attractive, but by reducing the critical scrutiny applied to the informational basis of product
826 evaluation.

827 The direct effect on purchase intention was not significant: the AI Overview did not uniformly
828 increase participants' willingness to purchase the product. Instead, the results are consistent with a
829 more specific indirect association: the interface reduced epistemic vigilance, and lower vigilance was
830 associated with higher purchase intention. This distinction prevents the finding from being
831 characterized as straightforward AI persuasion. The overview did not simply make participants want
832 the product more; rather, its synthesized presentation was associated with a change in how critically
833 the information supporting that response was evaluated.

834 The behavioral and psychological findings should also be kept conceptually separate. The
835 exploratory analyses did not support epistemic vigilance as a mediator of exposure clicks, source-
836 viewing time, verification clicks, or verification time. Consumers may open fewer sources because
837 the overview makes additional information acquisition appear less valuable, regardless of how
838 critically they assess the answer itself. Epistemic vigilance instead concerns the scrutiny applied to
839 the information already encountered. A consumer can regard an overview as sufficient while still
840 evaluating it critically, just as consulting multiple sources does not guarantee careful assessment of

841 their quality. Study 3 therefore identifies two related but distinct consequences of summary-led
842 search: reduced initial source acquisition and lower critical appraisal of the information presented.
843 The lack of significant indirect effects through processing fluency, knowledge insufficiency, message
844 credibility, and source transparency further helps delimit this account. Within Study 3, the indirect
845 association with purchase intention was specifically observed through epistemic vigilance rather than
846 through a general sense that the information was easier to process, more credible, more transparent,
847 or sufficient to close a perceived knowledge gap. Although these null alternatives should not be
848 treated as evidence that such processes can never operate, their inclusion supports the conceptual
849 distinctiveness of epistemic vigilance as an explanation of how summary-led search may affect
850 downstream product responses.

851 Overall, Study 3 completes the empirical sequence by moving from the amount of search to its
852 evaluative quality. Studies 1 and 2 establish that AI Overviews reduce direct source engagement and
853 identify conditions under which the reduction in source opening varies. Study 3 suggests that the
854 same interface may also reduce scrutiny of the evidentiary basis of product information. Thus,
855 keeping source links technically available may not be sufficient to preserve source-oriented search:
856 when synthesis comes first, consumers may both consult underlying evidence less and evaluate the
857 synthesized answer less critically.

858

859 **6. General Discussion**

860 AI Overviews are changing the architecture of consumer search. Rather than presenting only a
861 ranked set of sources from which consumers must construct an answer, they place a synthesized
862 response before those sources and make further navigation optional. Across three realistic, navigable
863 experiments and multiple product contexts, this shift consistently reduced engagement with
864 underlying sources during initial information acquisition. Participants exposed to an AI Overview
865 opened fewer sources and spent less time consulting them than participants who encountered a
866 traditional SERP. Because the studies recorded clicks and source-viewing time as search unfolded,

867 the findings capture actual changes in search behavior rather than consumers' stated expectations
868 about how they would search. Collectively, they indicate that synthesis does not merely supplement
869 conventional results: it partially substitutes for the source-selection and information-integration
870 activities traditionally performed by consumers.

871 The studies also clarify the boundaries of this behavioral effect. Studies 1 and 3 found no
872 corresponding reduction during optional verification, suggesting that AI Overviews primarily
873 compress the initial construction of product knowledge rather than producing a uniform withdrawal
874 from source engagement. Study 2 further showed that product type does not operate as an independent
875 boundary condition. Instead, its relevance depends on consumers' attitudes toward AI. The significant
876 three-way interaction suggests that different forms of product uncertainty create different conditions
877 for reliance on synthesis. For experience goods, AI Overviews reduced source opening even among
878 consumers who were not especially favorable toward AI, possibly because conventional source
879 consultation cannot fully resolve qualities that ultimately require product use. For credence goods, by
880 contrast, reliance on synthesis became more pronounced as positive AI attitudes increased, consistent
881 with greater willingness to delegate the interpretation of technical information that consumers may
882 struggle to evaluate independently. Product characteristics and AI attitudes therefore jointly shape
883 whether synthesis substitutes for source navigation.

884 Study 3 reveals a separate epistemic consequence. Exposure to an AI Overview reduced
885 epistemic vigilance, and lower vigilance was associated with higher purchase intention, producing a
886 significant statistical indirect effect despite the absence of a significant direct effect on purchase
887 intention. Epistemic vigilance did not explain the reduction in clicks or source-viewing time,
888 indicating that search compression and critical appraisal should not be treated as a single process. The
889 former concerns whether consumers perceive value in acquiring additional information; the latter
890 concerns how critically they evaluate the information already available. The studies show that
891 summary-led search can affect both the amount of independent source engagement and the scrutiny
892 applied to synthesized information, but through conceptually distinct pathways.

893 6.1. Theoretical contributions

894 The first contribution is to consumer information-search theory. Established models
895 conceptualize search as a consumer-directed allocation of limited time, attention, and effort:
896 consumers continue acquiring information while its expected value exceeds its cost (Diehl, 2005;
897 Krefeld-Schwalb et al., 2026; Moorthy et al., 1997; Stigler, 1961). A traditional SERP reduces the
898 cost of locating potentially relevant information, but it leaves consumers responsible for choosing
899 sources, comparing claims, and integrating evidence. AI Overviews alter this allocation of epistemic
900 labor. By selecting, organizing, and synthesizing information before a source is opened, the interface
901 performs part of the work that consumer-search models ordinarily assign to the consumer. The lower
902 number of clicks and shorter source-viewing time observed across the studies are behavioral
903 manifestations of this reallocation.

904 This argument identifies summary-led search as a search architecture that is theoretically
905 different from a better-ranked or more extensive set of results. Summary-led search differs from a
906 better-ranked or more extensive result set because consumers begin with an integrated interpretation
907 that can satisfy part of the query before source-level search starts. They begin with an integrated
908 interpretation that can satisfy part of the query before source-level search starts. Consequently, the
909 expected incremental value of opening another source falls. The findings complement evidence that
910 learning from LLM syntheses can require less effort than learning through web links (Melumad &
911 Yun, 2025) by showing how embedded synthesis restructures observable navigation within a
912 conventional product-search environment.

913 The distinction between initial acquisition and later verification further refines this contribution.
914 Source consultation serves different functions across these phases. During initial acquisition,
915 consumers use sources to identify relevant attributes and construct an understanding of the product
916 problem. During verification, they use sources to confirm, qualify, or challenge an answer already
917 encountered. AI Overviews consistently compressed the first function but not the second. Summary-
918 led search should therefore not be equated with a general cessation of search. Its most reliable effect

919 is phase-specific: it substitutes for part of the consumer's initial construction of product knowledge
920 while leaving open the possibility of later checking. This phase distinction also answers calls to
921 connect consumer-search models more closely to process-tracing evidence (Krefeld-Schwalb et al.,
922 2026). Clicks and source-viewing time reveal when in the process search effort changes, rather than
923 inferring search only from a final judgment or choice.

924 The second contribution concerns the role of AI-enabled assistance in the consumer journey,
925 where new technologies can create simultaneous efficiency gains and new decision-making tensions
926 (Inman et al., 2024). Interactive decision aids and recommendation agents commonly help consumers
927 filter alternatives, compare attributes, or prioritize products according to their preferences (Baccelloni
928 et al., 2026; Häubl & Murray, 2003; Häubl & Trifts, 2000). AI Overviews intervene further upstream.
929 They need not recommend a particular product or execute a transaction; instead, they structure the
930 informational basis on which later product evaluation occurs. They are therefore best understood as
931 hybrid retrieval-and-synthesis interfaces: they retain the links and source cues of conventional search
932 while placing an AI-generated interpretation before them. This extends research on shopping-support
933 technologies (Liu et al., 2024) by demonstrating that upstream assistance changes whether consumers
934 engage with the source material from which product understanding is ordinarily constructed.

935 The moderation results qualify explanations that predict a simple match between AI assistance
936 and product diagnosticity. Search, experience, and credence products differ in how readily their
937 attributes can be assessed before or after consumption. Yet product type alone did not systematically
938 alter the interface type effect. Its role emerged only in combination with positive attitudes toward AI,
939 and only for exposure clicks. In the product contexts examined here, product category did not provide
940 sufficient boundary conditions by themselves. Their relevance depends partly on whether consumers
941 regard AI as an appropriate and useful interpreter of the information problem. This pattern is
942 consistent with research showing that responses to AI depend on the task, context, and perceived
943 trade-offs involved (Castelo et al., 2019; Dorotic et al., 2024; Puntoni et al., 2021), while cautioning
944 against a broad claim that AI summaries are inherently more consequential for any one product class.

945 Study 2 also shows that verification depends on consumers' prior attitudes toward AI.
946 Consumers who were less positively disposed toward AI appeared to compensate for the summary
947 by consulting more sources later, whereas more AI-positive consumers did not. Attitudes toward AI
948 may influence both acceptance of AI assistance and the point at which consumers deploy additional
949 scrutiny. Skepticism can preserve verification, but it does not follow that the interface type itself
950 reliably encourages checking. This distinction shifts attention from general acceptance of AI to the
951 temporal organization of consumers' response: an overview may compress initial search for many
952 consumers while eliciting compensatory verification only among those already inclined to question
953 AI.

954 The third contribution is to research on epistemic vigilance in AI-mediated environments.
955 Epistemic vigilance protects individuals from misinformation by directing attention to the plausibility
956 of communicated content and to the competence and trustworthiness of its sources (Metzger et al.,
957 2010; Sperber et al., 2010). Traditional SERPs keep sources visibly separate and make the consumer's
958 responsibility for selecting and reconciling them salient. AI Overviews instead combine claims into
959 a fluent, apparently unified answer. The reduction in epistemic vigilance observed in Study 3 is
960 consistent with the idea that synthesis can make provenance, disagreement, and evidentiary variation
961 less salient, even when citations remain visible (Dabran-Zivan et al., 2026).

962 The results also separate this epistemic pathway from behavioral efficiency. Consumers may
963 stop opening sources because the overview makes further acquisition seem less useful yet still
964 evaluate the overview critically. Conversely, they may visit several sources without carefully
965 appraising their credibility. The nonsignificant indirect effects of epistemic vigilance on clicks and
966 viewing time are therefore theoretically informative rather than merely null results: they discourage
967 a single-process account in which reduced scrutiny automatically explains reduced navigation.
968 Subject to the temporal limitation of Study 3, the significant indirect effect on purchase intention
969 instead suggests that epistemic vigilance is more relevant to how readily favorable information is
970 incorporated into product evaluation than to how much additional information is acquired.

971

972 *6.2. Managerial and public-policy implications*

973 For search platforms, the findings demonstrate that placing synthesis before conventional
974 results is not a neutral interface decision. Across the three experiments, AI Overviews reduced both
975 the likelihood that consumers opened underlying sources and the total time devoted to those sources
976 during initial search. Platforms should therefore evaluate AI-assisted search not only in terms of
977 answer quality, user satisfaction, or task completion, but also in terms of how it redistributes attention
978 between the generated response and its evidentiary basis. Efficiency gains may come at the cost of
979 reduced source engagement, particularly when the overview appears sufficiently complete to
980 substitute for further navigation.

981 Making citations technically available may not be enough to address this concern. When
982 sources are presented as supporting elements of an already-integrated answer, users may have little
983 reason to inspect them or compare their claims. Platforms could make the relationship between
984 particular claims and their sources more explicit, clearly distinguish independent evidence from
985 repeated claims, signal disagreement or uncertainty, and provide low-friction ways to compare the
986 synthesis with the underlying material. Such design choices would preserve the convenience of AI-
987 generated answers while keeping provenance and opportunities for scrutiny behaviorally salient.

988 Study 2 further indicates that the consequences of interface design may vary across users and
989 product contexts. Less AI-positive consumers increased their subsequent verification after
990 encountering an AI Overview, whereas consumers with mean or high positive attitudes toward AI did
991 not display the same compensatory response. Platforms should therefore not assume that users'
992 natural skepticism will provide a reliable safeguard. The consumers most comfortable relying on AI
993 may be less inclined to verify its synthesis independently, especially when the information is technical
994 and difficult to evaluate. Context-sensitive interventions—such as more prominent evidence cues or
995 verification prompts for health, financial, safety, and other credence-related queries—may
996 consequently be more useful than uniform disclosures.

For brands and retailers, competition in search increasingly occurs at two levels. Firms still compete for ranking and referral traffic, but they also compete to have their information selected, interpreted, and accurately represented within an AI-generated answer that consumers may use without visiting the originating webpage. Search visibility can therefore no longer be assessed through rankings and click-through rates alone. Firms should also monitor how their products, attributes, and claims are represented within generated summaries, whether important qualifications are retained, and which sources the interface uses to construct those representations.

This development increases the strategic importance of clear, consistent, and well-substantiated product information. Structured specifications, unambiguous attribute definitions, expert evidence, and explicit support for performance claims may make it easier for AI systems to represent a product accurately. At the same time, firms should not treat optimization for AI Overviews as merely a new form of search-engine optimization. Claims designed to maximize inclusion in generated answers may be detached from their original context or presented without relevant qualifications. Brands therefore have an interest in making not only their claims but also their evidentiary status and limitations easy for both consumers and automated systems to identify.

These responsibilities are particularly important for products involving health, safety, sustainability, reliability, or technical performance. Study 3 suggests that synthesized presentation may be associated with lower scrutiny of the evidence and sources underlying product information. Brands operating in such contexts should clearly distinguish objective specifications, independently validated evidence, expert interpretation, and promotional claims. They should also make supporting evidence easy to access from the claim itself rather than requiring consumers to navigate through general corporate or product pages. This can reduce the risk that qualified evidence is transformed into an apparently unconditional claim when incorporated into an AI-generated synthesis.

Publishers, review sites, and comparison services face a related challenge. If consumers obtain an initial understanding directly from the overview, content that merely restates easily synthesized facts may attract fewer visits. The studies nevertheless show that source engagement remains possible

1023 during later verification. Publishers can strengthen their role at this stage by offering value that a brief
1024 synthesis cannot easily reproduce: transparent testing methods, primary data, expert credentials,
1025 interactive comparisons, personalized assessment tools, and explicit discussion of conflicting
1026 evidence or uncertainty. Their strategic value may increasingly lie not in providing another summary,
1027 but in enabling consumers to evaluate, contextualize, and verify the summary already encountered.

1028 Finally, policymakers should consider both the accuracy of AI-generated information and the
1029 search behavior produced by its presentation. Disclosure rules focused only on identifying AI
1030 involvement or listing sources may be insufficient if users rarely inspect those sources or treat their
1031 inclusion as evidence that they have already been validated. Effective oversight should therefore
1032 consider whether interfaces make provenance, evidentiary strength, disagreement, and opportunities
1033 for verification genuinely usable and salient. In high-stakes product domains, regulators could also
1034 examine whether generated claims preserve the qualifications and substantiation standards required
1035 of the original advertiser or publisher. The relevant policy question is not only whether the answer is
1036 accurate, but whether the interface supports an appropriately informed and critical basis for consumer
1037 judgment.

1038 *6.3. Limitations and future research*

1039 Several limitations define the scope of the conclusions and suggest priorities for further
1040 research. First, the studies used controlled simulations of a familiar search interface. This design is a
1041 strength because it enabled realistic navigation while holding the available sources constant and
1042 recording actual clicks and viewing time. Nevertheless, live search is more dynamic: consumers
1043 reformulate queries, encounter changing rankings and summaries, open multiple tabs, move across
1044 devices, and return to searches over time. Field studies and browser-based experiments should test
1045 whether the phase-specific compression observed here persists under those conditions and whether it
1046 affects eventual choice, satisfaction, returns, or post-purchase learning.

1047 Second, the studies focused on diagnostic product-information queries that invite consumers to
1048 learn which attributes matter. Navigational, comparative, review-seeking, risk-oriented,

1049 troubleshooting, and explicitly transactional queries may create different expectations about what
1050 constitutes a sufficient answer. Future research should vary query goals and examine whether AI
1051 Overview substitutes most strongly for exploratory learning, or whether their effects become larger
1052 when users seek a concise recommendation or factual resolution.

1053 Third, Study 2 represented each search, experience, and credence category with one product.
1054 Although Studies 1 and 3 broaden the set of contexts, product-specific characteristics remain
1055 confounded with the category manipulation. Replications should sample multiple products within
1056 each category and directly manipulate the type of attribute under evaluation. Such designs could
1057 distinguish whether the joint effect of product type and AI attitudes reflects verifiability, perceived
1058 risk, technical complexity, personal fit, or the availability of trusted external standards.

1059 Fourth, the interpretation of Study 3's indirect effect requires caution because purchase
1060 intention was measured before epistemic vigilance. Although the ordering of the statistical
1061 associations is consistent with the proposed pathway—AI Overview exposure was associated with
1062 lower epistemic vigilance, which was associated with higher purchase intention—the measurement
1063 sequence does not establish that reduced vigilance temporally preceded the purchase response. The
1064 indirect effect should therefore be interpreted as evidence compatible with the proposed process rather
1065 than as a definitive demonstration of causal mediation. Reverse ordering remains possible:
1066 participants' initial product evaluations may have influenced their subsequent reports of epistemic
1067 vigilance. Both variables may also have been shaped by an unmeasured response to the interface.
1068 Future research should test the causal role of epistemic vigilance using designs that establish temporal
1069 and experimental precedence. At a minimum, epistemic vigilance should be measured before
1070 purchase intention, although merely reversing the measurement order would not rule out common-
1071 method or demand effects. Stronger evidence would come from a moderation-of-process design that
1072 directly changes the opportunity or motivation for critical appraisal while holding the informational
1073 content of the AI Overview constant. For example, researchers could manipulate the visibility of
1074 source provenance, present claims and citations as clearly separated units, disclose disagreement

1075 among the underlying sources, prompt users to assess the evidentiary support for individual claims,
1076 or require them to inspect a source before making a product judgment. If the effect of AI Overview
1077 exposure on purchase intention is attenuated when epistemic vigilance is restored, this would provide
1078 more direct evidence for the proposed mechanism. Also, the nonsignificant alternative indirect effects
1079 should be interpreted cautiously. Processing fluency, message credibility, knowledge insufficiency,
1080 and source transparency were not supported as mediators in Study 3, but a single study cannot
1081 establish that these mechanisms are irrelevant across products, queries, or interface designs. Future
1082 research could manipulate these features independently to identify when they operate alongside
1083 epistemic vigilance, precede it, or compete with it. Such work would clarify whether reduced
1084 vigilance is a distinctive consequence of generative synthesis or one component of a broader process
1085 through which AI-mediated search changes consumers' responses to marketplace information.

1086 Fifth, the AI Overview citations in the present interfaces were not independently clickable,
1087 whereas the organic results were equally available across conditions. This design isolates engagement
1088 with the common source set but cannot distinguish navigation initiated through an overview citation
1089 from navigation initiated through a conventional result. Future interfaces should make both routes
1090 available and track them separately. Doing so would clarify whether AI Overview eliminates source
1091 visits, redirects them toward a narrower set of cited sources, or changes the kinds of sources
1092 consumers choose to inspect. The manipulation also captures AI Overview as an integrated interface
1093 intervention rather than isolating generative synthesis from other design features. Adding an overview
1094 changes both the information immediately available on the results page and the visual prominence of
1095 the organic results. The present studies therefore estimate the effect of the AI Overview interface as
1096 a whole; future research should disentangle informational synthesis from spatial displacement, visual
1097 salience, and AI labeling.

1098 Sixth, the studies identify the behavioral effect more clearly than its psychological mechanism.
1099 Epistemic vigilance did not explain the reduction in clicks or viewing time. Perceived informational
1100 sufficiency, expected value of additional search, cognitive effort, answer completeness, perceived

1101 autonomy (Ahmad Husairi & Rossi, 2024), and confidence are plausible mechanisms that should be
1102 compared directly. Process measures could also test whether consumers read the overview in place
1103 of sources or simply terminate the task earlier. Establishing these mechanisms would clarify when
1104 efficiency represents useful substitution and when it reflects premature search termination.

1105 Finally, the studies examine immediate responses to a single search episode. Repeated use
1106 may produce learning, calibration, or dependence. Consumers may become more skilled at checking
1107 summaries, or they may increasingly treat synthesized answers as the default endpoint of search.
1108 Longitudinal research should examine how experience, observed errors, and changes in AI literacy
1109 affect navigation and epistemic vigilance over time. It should also evaluate heterogeneous effects for
1110 consumers with different levels of domain expertise, digital literacy, need for cognition, and
1111 vulnerability to misleading claims. These questions are central to determining whether summary-led
1112 search produces durable efficiency gains, durable reductions in scrutiny, or both.

1113

1114 **References**

- 1115 • Ahmad Husairi, M., & Rossi, P. (2024). Delegation of purchasing tasks to AI: The role of
1116 perceived choice and decision autonomy. *Decision Support Systems*, 179, 114166.
1117 <https://doi.org/10.1016/j.dss.2023.114166>
- 1118 • Appelman, A., & Sundar, S. S. (2016). Measuring message credibility: Construction and
1119 validation of an exclusive scale. *Journalism & Mass Communication Quarterly*, 93(1), 59–
1120 79. <https://doi.org/10.1177/1077699015606057>
- 1121 • Baccelloni, A., De Angelis, M., Ricotta, F., & Mazzù, M. F. (2026). Too narrow to help?
1122 Unveiling how recommendation agents' specialization impacts user choices. *Journal of*
1123 *Interactive Marketing*, 61(2), 137–154. <https://doi.org/10.1177/10949968251358181>
- 1124 • Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive consumer choice
1125 processes. *Journal of Consumer Research*, 25(3), 187–217. <https://doi.org/10.1086/209535>

- 1126 • Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion.
1127 *Journal of Marketing Research*, 56(5), 809–825.
1128 <https://doi.org/10.1177/0022243719851788>
- 1129 • Dabran-Zivan, S., Klein-Avraham, I., & Baram-Tsabari, A. (2026). When the source is a
1130 bot: How people adapt their evaluation strategies to assess AI-generated content. *PLOS*
1131 *ONE*, 21(3), e0345300. <https://doi.org/10.1371/journal.pone.0345300>
- 1132 • Darby, M. R., & Karni, E. (1973). Free competition and the optimal amount of fraud. *The*
1133 *Journal of Law and Economics*, 16(1), 67–88. <https://doi.org/10.1086/466756>
- 1134 • Diehl, K. (2005). When two rights make a wrong: Searching too much in ordered
1135 environments. *Journal of Marketing Research*, 42(3), 313–322.
1136 <https://doi.org/10.1509/jmkr.2005.42.3.313>
- 1137 • Dodds, W. B., Monroe, K. B., & Grewal, D. (1991). Effects of price, brand, and store
1138 information on buyers' product evaluations. *Journal of Marketing Research*, 28(3), 307–
1139 319. <https://doi.org/10.1177/002224379102800305>
- 1140 • Dorotic, M., Stagno, E., & Warlop, L. (2024). AI on the street: Context-dependent responses
1141 to artificial intelligence. *International Journal of Research in Marketing*, 41(1), 113–137.
1142 <https://doi.org/10.1016/j.ijresmar.2023.08.010>
- 1143 • Ford, G. T., Smith, D. B., & Swasy, J. L. (1990). Consumer skepticism of advertising
1144 claims: Testing hypotheses from economics of information. *Journal of Consumer Research*,
1145 16(4), 433–441. <https://doi.org/10.1086/209228>
- 1146 • Ghose, A., Ipeirotis, P. G., & Li, B. (2014). Examining the impact of ranking on consumer
1147 behavior and search engine revenue. *Management Science*, 60(7), 1632–1654.
1148 <https://doi.org/10.1287/mnsc.2013.1828>
- 1149 • Graf, L. K. M., Mayer, S., & Landwehr, J. R. (2018). Measuring processing fluency: One
1150 versus five items. *Journal of Consumer Psychology*, 28(3), 393–411.
1151 <https://doi.org/10.1002/jcpy.1021>

- 1152 • Guo, Y., Liu, S., & He, Y. (2025). Examining the consequences of government information
1153 insufficiency during the COVID-19 public health emergency: A cross-sectional study in
1154 China. *Critical Public Health*, 35(1), Article 2520405.
1155 <https://doi.org/10.1080/09581596.2025.2520405>
- 1156 • Hayes, A. F. (2022). *Introduction to mediation, moderation, and conditional process*
1157 *analysis: A regression-based approach* (3rd ed.). The Guilford Press
- 1158 • Häubl, G., & Murray, K. B. (2003). Preference construction and persistence in digital
1159 marketplaces: The role of electronic recommendation agents. *Journal of Consumer*
1160 *Psychology*, 13(1–2), 75–91. https://doi.org/10.1207/S15327663JCP13-1&2_07
- 1161 • Häubl, G., & Trifts, V. (2000). Consumer decision making in online shopping
1162 environments: The effects of interactive decision aids. *Marketing Science*, 19(1), 4–21.
1163 <https://doi.org/10.1287/mksc.19.1.4.15178>
- 1164 • Inman, J. J., Meyer, R. J., Schweidel, D. A., & Srinivasan, R. (2024). Do great powers come
1165 with great responsibility? Opportunities and tensions of new technologies in marketing.
1166 *International Journal of Research in Marketing*, 41(1), 18–23.
1167 <https://doi.org/10.1016/j.ijresmar.2024.01.006>
- 1168 • Kahlor, L. (2010). PRISM: A planned risk information seeking model. *Health*
1169 *Communication*, 25(4), 345–356. <https://doi.org/10.1080/10410231003775172>
- 1170 • Kim, H. K., & Lai, C.-H. (2020). Seeking and sharing information about transboundary air
1171 pollution in Singapore: Effects of own and others' information insufficiency. *Environmental*
1172 *Communication*, 14(1), 68–81. <https://doi.org/10.1080/17524032.2019.1597751>
- 1173 • Klein, L. R. (1998). Evaluating the potential of interactive media through a new lens: Search
1174 versus experience goods. *Journal of Business Research*, 41(3), 195–203.
1175 [https://doi.org/10.1016/S0148-2963\(97\)00062-3](https://doi.org/10.1016/S0148-2963(97)00062-3)
- 1176 • Krefeld-Schwalb, A., Martinovici, A., Orquin, J. L., Aribarg, A., Erdem, T., Johnson, E. J.,
1177 Reeck, C., Schley, D. R., Smith, S. M., Speekenbrink, M., Stewart, N., Weber, E. U., &

- 1178 Yang, C. L. (2026). Search in service of choice: An agenda for comparing, extending, and
1179 integrating consumer search models. *International Journal of Research in Marketing*, 43(2),
1180 485–499. <https://doi.org/10.1016/j.ijresmar.2025.09.001>
- 1181 • Liu, Y., Yildirim, P., & Zhang, Z. J. (2024). Consumer preferences and firm technology
1182 choice. *International Journal of Research in Marketing*, 41(1), 41–55.
1183 <https://doi.org/10.1016/j.ijresmar.2023.06.008>
- 1184 • Melumad, S., & Yun, J. H. (2025). Experimental evidence of the effects of large language
1185 models versus web search on depth of learning. *PNAS Nexus*, 4(10), pgaf316.
1186 <https://doi.org/10.1093/pnasnexus/pgaf316>
- 1187 • Metzger, M. J., Flanagin, A. J., & Medders, R. B. (2010). Social and heuristic approaches
1188 to credibility evaluation online. *Journal of Communication*, 60(3), 413–439.
1189 <https://doi.org/10.1111/j.1460-2466.2010.01488.x>
- 1190 • Mitra, K., Reiss, M. C., & Capella, L. M. (1999). An examination of perceived risk,
1191 information search and behavioral intentions in search, experience and credence services.
1192 *Journal of Services Marketing*, 13(3), 208–228.
1193 <https://doi.org/10.1108/08876049910273763>
- 1194 • Moorthy, S., Ratchford, B. T., & Talukdar, D. (1997). Consumer information search
1195 revisited: Theory and empirical analysis. *Journal of Consumer Research*, 23(4), 263–277.
1196 <https://doi.org/10.1086/209482>
- 1197 • Nagler, M. G., Kronenberg, F., Kennelly, E. J., & Jiang, B. (2011). Pricing for a credence
1198 good: An exploratory analysis. *Journal of Product & Brand Management*, 20(3), 238–249.
1199 <https://doi.org/10.1108/10610421111134969>
- 1200 • Nelson, P. (1970). Information and consumer behavior. *Journal of Political Economy*,
1201 78(2), 311–329. <https://doi.org/10.1086/259630>

- 1202 • Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial
1203 intelligence: An experiential perspective. *Journal of Marketing*, 85(1), 131–151.
1204 <https://doi.org/10.1177/0022242920953847>
- 1205 • Ratchford, B. T. (2001). The economics of consumer knowledge. *Journal of Consumer*
1206 *Research*, 27(4), 397–411. <https://doi.org/10.1086/319617>
- 1207 • Roelen-Blasberg, T., Habel, J., & Klarmann, M. (2023). Automated inference of product
1208 attributes and their importance from user-generated content: Can we replace traditional
1209 market research? *International Journal of Research in Marketing*, 40(1), 164–188.
1210 <https://doi.org/10.1016/j.ijresmar.2022.04.004>
- 1211 • Sacco, F., Manzi, F., Di Dio, C., Schepman, A., Rodway, P., Massaro, D., & Marchetti, A.
1212 (2025). Validation of the General Attitudes towards Artificial Intelligence Scale in the
1213 Italian context and the role of trust. *Computers in Human Behavior Reports*, 19, 100751.
1214 <https://doi.org/10.1016/j.chbr.2025.100751>
- 1215 • Schepman, A., & Rodway, P. (2023). The General Attitudes towards Artificial Intelligence
1216 Scale (GAAIS): Confirmatory validation and associations with personality, corporate
1217 distrust, and general trust. *International Journal of Human–Computer Interaction*, 39(13),
1218 2724–2741. <https://doi.org/10.1080/10447318.2022.2085400>
- 1219 • Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origgi, G., & Wilson, D.
1220 (2010). Epistemic vigilance. *Mind & Language*, 25(4), 359–393.
1221 <https://doi.org/10.1111/j.1468-0017.2010.01394.x>
- 1222 • Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3),
1223 213–225. <https://doi.org/10.1086/258464>.
- 1224 • Trzebinski, W., Gaczek, P., & Marciniak, B. (2023). Is it better to communicate product
1225 information abstractly or concretely? The role of consumer product expertise and shopping-
1226 stage mindset. *Journal of Product & Brand Management*, 32(2), 273–285.
1227 <https://doi.org/10.1108/JPBM-05-2021-3470>

- Tseng, A. S., Bonilla, S., & MacPherson, A. (2021). Fighting “bad science” in the information age: The effects of an intervention to stimulate evaluation and critique of false scientific claims. *Journal of Research in Science Teaching*, 58(8), 1152–1178. <https://doi.org/10.1002/tea.21696>
- Wojdyski, B. W., Evans, N. J., & Hoy, M. G. (2018). Measuring sponsorship transparency in the age of native advertising. *Journal of Consumer Affairs*, 52(1), 115–137. <https://doi.org/10.1111/joca.12144>