Google AI Overviews and Publisher Traffic: Evidence from a Field Experiment*

Saharsh Agarwal† Ananya Sen‡

April 2026

Preliminary – Feedback welcome. Draft will be updated with additional results.

Abstract

Online search engines depend on external content to respond to user queries, while content providers rely on search engines for visibility and user acquisition. The integration of GenAI-based summaries into search interfaces (eg, Google's AI Overviews (AIOs)) may reduce users' incentives to visit downstream publishers. Despite growing regulatory and industry attention, causal evidence on the effects of AI-generated summaries on search behavior remains limited. This paper provides causal evidence on the impact of AIOs on user behavior and downstream engagement using a field experiment that randomly assigns users to either the standard Google Search experience with AIOs or a version where AIOs are removed. We develop a custom Chrome browser extension that passively tracks detailed browsing behavior while also enabling real-time modifications to the search engine results page (SERP), allowing us to precisely manipulate AIO exposure in a natural browsing environment. When they appear, AIOs reduce outbound organic clicks by 38% and increase the probability of a zero-click search by 33%. We find no effects on sponsored clicks or overall search frequency, implying that AIOs reduce total external traffic by lowering organic clicks per search. Heterogeneity analyses show that effects are concentrated when AIOs appear at the top of the page and among the highest-ranked organic results. Using survey evidence, we find that despite these behavioral effects, removing AIOs does not affect users' overall search experience, including perceived ease of finding information and the quality of results. Exploratory evidence from a third AI Mode condition, where users are redirected to a Gemini-powered conversational AI interface, suggests even stronger reductions in engagement with downstream publishers. Overall, the findings suggest that AIOs divert traffic away from publishers without delivering measurable improvements in user experience, raising important implications for competition policy, platform power, and publisher sustainability.

Keywords: Generative AI, Online Search, AI Overviews.

*This study was pre-registered in the AEA RCT Registry (ID: AEARCTR-0017393). The study protocol was approved by the Institutional Review Board at the Indian School of Business (ISB). We are indebted to the Office of Advancement and the Office of Alumni Engagement at ISB for facilitating funding for this project through the Impact Innovation Grant as well as the Emerging Scholar Research Award. We are thankful to Rohit Akash Rajendran for excellent research assistance. We gratefully acknowledge the ISB Institute of Data Science for funding the research assistance.

†Indian School of Business. Email: saharsh_agarwal@isb.edu

‡Carnegie Mellon University. Email: ananyase@andrew.cmu.edu

1 Introduction

Google Search has long served as a primary gateway to information on the internet, accounting for a dominant share of global search activity.[1] For many downstream content providers—especially news publishers, review platforms, and informational websites (e.g., Wikipedia)—visibility on the search engine results page (SERP) is a critical driver of user acquisition and monetization. This relationship is fundamentally symbiotic: while publishers rely on search engines for traffic and discovery, search engines depend on a steady supply of high-quality content to deliver relevant results and sustain user engagement.

Over the past decade, this inter-dependence has intensified as publishers have increasingly shifted from direct traffic toward platform-mediated discovery through search engines, social media, and aggregators. At the same time, this growing reliance has given rise to tensions over value capture, with content producers arguing that platforms appropriate and monetize their content without providing adequate compensation. These concerns have fueled regulatory scrutiny and legal disputes across multiple jurisdictions (De Vynck, 2021).

The recent integration of generative artificial intelligence into search interfaces represents a potentially significant shift in this ecosystem. In particular, Google's introduction of AI Overviews (AIOs henceforth)—AI-generated summaries that appear prominently in response to certain queries—changes how users interact with search results (Google, 2024). Rather than directing users to external websites to synthesize information, AIOs consolidate content from underlying sources and present it directly on the search engine results page (SERP). While such summaries may improve convenience for user, publishers and other downstream stakeholders have raised concerns that AIOs may further reduce referral traffic without requiring clicks to source websites (Bellan, 2025). Moreover, such AIOs use underlying data from these content websites, which is necessary to produce the AI summary in the first place, raising concerns that search platforms may increasingly internalize content consumption while externalizing content production costs.

These concerns have quickly entered policy debates. Competition authorities in the United Kingdom have suggested that publishers should have the option to opt out of AI-generated summaries, and complaints regarding AIOs have been filed with European antitrust regulators (Muvija, 2026). In the United States, the introduction of AIOs has raised concerns that Google

[1] See https://knowledge.wharton.upenn.edu/article/why-google-dominates-the-search-engine-market/ for some details

may be leveraging its monopoly power in the search market in ways that violate the Sherman Act (Singh and Scott Morton, 2026). Moreover, shareholders of information platforms such as Reddit are suing the platform itself over concerns about Google AIO stealing traffic.[2] At the core of these discussions is a fundamental empirical question: do AI-generated search summaries materially reduce engagement with the very sources from which they draw information?

Despite the importance of this issue for platform design, publisher sustainability, and competition policy, systematic causal evidence remains limited.[3] Importantly, the direction and magnitude of these effects are theoretically ambiguous. On the one hand, AIOs may substitute for clicks by directly satisfying user information needs. On the other hand, they extensively include citations and links to underlying sources, which may serve as referrals and potentially increase traffic to certain publishers. Moreover, even if clicks per search decline, overall engagement need not fall if AIOs alter search behavior – for instance, by increasing the total number of searches. As such, the net effect on publisher traffic is ultimately an empirical question that depends on how these competing mechanisms interact in practice. The implications for user experience are also unclear. While removing AIOs may increase downstream engagement, it could come at the cost of reduced convenience or slower information access, making the welfare effects of such interventions similarly ambiguous.

This paper provides novel and rigorous empirical evidence on how AI Overviews affect user behavior and downstream engagement. We do so using a field experiment that directly manipulates the search experience of real users in a natural browsing environment. Our approach allows us to isolate the causal impact of AIO exposure on click behavior, scrolling patterns, and user-perceived utility, thereby contributing to an emerging yet nascent literature on generative AI in search environments.

To implement the experiment, we developed a custom Chrome extension that modifies users' Google Search experience in real time. We recruited over 1,000 participants who installed the extension and consented to passive measurement of their browsing behavior over a two-week period. Participants were randomly assigned to one of three conditions. In the control group, users experienced Google Search as usual, with AIOs appearing when available. In the primary treatment group ("Hide AIO", HAIO), the extension removed AIOs in real time

[2] See https://www.dandodiary.com/2025/06/articles/artificial-intelligence/reddit-downplayed-impact-of-googles-ai-related-changes-suit-alleges/ for some more details.

[3] There are some surveys and industry reports that suggest a negative impact of AIOs on downstream website traffic – see, for example, https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/ and https://ahrefs.com/blog/ai-overviews-reduce-clicks/.

whenever they would otherwise appear. Importantly, the remaining search results were automatically shifted upward to fill the space, ensuring that the interface remained seamless and that no blank or disrupted layout signaled the intervention to users. In a third condition ("AI Mode"), users were redirected to Google's AI-focused search interface for all queries. This design enables us to compare outcomes both when AI summaries are suppressed and when users are exposed to a more AI-centric conversational search environment. Because the AI Mode interface represents a substantial departure from traditional search and we anticipated substantial (selective) attrition in this arm, we treat results from this condition as exploratory.

Our results yield several insights. First, descriptively, we document that AIOs were triggered in approximately 42% of observed queries during our study period. This high prevalence indicates that these features are not marginal but constitute a substantial share of the search experience, with potentially large implications for downstream traffic. Second, for queries in which an AIO was (intended to be) triggered, removing AIOs increases outbound organic clicks from 0.38 to 0.61 per search, implying that AIOs reduce clicks by approximately 38%. In contrast, for queries where AIO was not triggered, clicking behavior is statistically indistinguishable between the Control and Hide AIO conditions, supporting the causal interpretation of our estimates. Similarly, we find that the probability of a zero-click session is 33% higher in the control group compared to HAIO users for AIO-triggering queries. Third, we find no meaningful change in clicks on sponsored links or other Google-owned page elements, suggesting that the primary behavioral substitution occurs between AI summaries and organic website visits rather than across the broader page ecosystem. Finally, we find no change in the number of searches that users make in the HAIO and control groups, suggesting that the reduction in clicks per search translates into a meaningful decline in overall outbound web traffic.

Heterogeneity analysis indicates that the treatment effect on clicks is concentrated in cases where the AIOs appears at the top of the SERP, aligning with prior evidence that top-of-page real estate captures disproportionate user attention (Agarwal et al., 2011). Moreover, looking at heterogeneity by search result rank, we find that search results in ranks one, two, and three had the largest increase in clicks in the HAIO group relative to the control. We also examine heterogeneity across user characteristics and baseline attitudes toward AI. Overall, the effects are broadly similar across demographic groups, including gender and race. We find some evidence that the substitution away from external clicks induced by AIOs is stronger among older users and among users who report higher comfort with AI tools, suggesting that these groups

may rely more heavily on AI-generated summaries. However, the reduction in outbound engagement in the presence of AIOs holds consistently across the sample.

Importantly, post-experiment survey responses indicate that users' perceived utility and satisfaction with search were unaffected by the removal of AIOs, suggesting that they divert traffic away from publishers without delivering measurable improvements in user experience.[4]

We complement these findings with descriptive evidence from the AI Mode condition. Adoption and retention in this arm were substantially lower, with notable attrition suggesting meaningful switching costs or user resistance to a fully AI-mediated search experience. We also observe attempts by numerous participants in this treatment arm to circumvent the extension using alternative extensions that bypass our intervention. Among the subset of users who remained compliant in the AI Mode condition, overall click-through rates for publishers were substantially lower than in the control group. While selection limits causal interpretation in this arm, the patterns provide suggestive evidence about potential equilibrium effects if AI-forward search modes become more prevalent.

This paper contributes to three related strands of literature. First, a growing body of work compares the use of large language models and AI assistants to traditional search (Agarwal and Ghosh, 2026; Aral et al., 2026; Gholami et al., n.d.; Padilla et al., 2025). These studies are complementary to ours, as they primarily rely on observational comparisons of search behavior across traditional platforms such as Google and AI assistants such as ChatGPT. Khosravi and Yoganarasimhan (2026) provides early evidence on this question by leveraging the rollout of AIOs in the US and comparing traffic to English Wikipedia articles with corresponding articles in other languages, which were presumably unaffected. While this approach is compelling, it relies on an observational difference-in-differences design and therefore cannot fully rule out contemporaneous shocks that may differentially affect traffic over the same period. By contrast, our study provides causal evidence from a randomized field experiment embedded directly in users' natural browsing environments. Moreover, our design enables a more granular decomposition of engagement, including differential effects on organic versus sponsored links, scroll behavior, and self-reported search experience. Instead of focusing on a single publisher (Wikipedia), our analysis captures effects across the broader web ecosystem, offering a more complete picture of the welfare implications of AI-generated search summaries.

[4]Participants also generally did not recognize the experimental manipulation, mitigating concerns about experimenter demand effects.

Second, our work relates to studies examining AI-generated summaries in other contexts, such as product reviews (Su et al., 2024; Zhai and Ching, 2026). We contribute to this strand by examining multiple behavioral margins—clicking, scrolling, and time-based engagement—and by exploring positional mechanisms specific to search environments. Importantly, the search engine setting features unique competitive dynamics among outbound links, making the welfare and traffic implications of AI summarization distinct from previously studied domains.

Third, we contribute to the literature on conflicts between news publishers and digital platforms. A substantial body of work examines how services such as Google News, Facebook, and other aggregators affect publisher traffic (Calzada and Gil, 2020; Song and Manchanda, 2025; Chiou and Tucker, 2017). Our study introduces generative AI as a new technological layer in this longstanding relationship, where data from such downstream websites is used as a key input into the GenAI product. Whereas earlier disputes focused on link aggregation and snippet display, AIOs represent a more intensive form of content synthesis that may further shift value capture toward the platform. By providing causal evidence on user behavior in this setting, we inform ongoing regulatory discussions in the United Kingdom, the European Union, and other jurisdictions considering the competitive implications of AI-mediated search.

Taken together, our findings highlight a central tension in the design of AI-enhanced search. While AIOs may reduce search friction, they come at a high cost in reduced traffic to external content providers. At the same time, user-reported utility appears largely unchanged when these summaries are removed, raising questions about the marginal consumer benefits relative to the ecosystem-level costs. As generative AI becomes more deeply integrated into information intermediaries, understanding these trade-offs will be increasingly important for platform strategy, publisher sustainability, and competition policy.

The remainder of the paper proceeds as follows. Section 2 describes the data and experimental design. Section 3 outlines the empirical strategy and the main results, and Section 4 concludes.

2 Data and Experimental Design

At the core of this study is a custom-built Chrome desktop browser extension. The underlying architecture of our extension builds on Webmunk (Farronato et al., 2024). The extension enables the collection of high-frequency behavioral data including URL visits, clicks, scroll depth,

and time spent, while simultaneously allowing us to modify the content of the Google search engine results page (SERP) in a controlled and unobtrusive manner.

This data allows us to identify all searches conducted on Google, along with detailed information on the corresponding search results page. In particular, we observe whether an AI Overview (AIO) was shown, its relative position on the SERP, the full set of search results including titles, links, and their positions on the page, data on every single scroll as well as all user clicks on both organic and sponsored results. For participants in the AI Mode condition, in addition to capturing all scroll and click activity, we also observe the full set of queries entered by users within the conversational interface.

2.1 Recruitment and Study Flow

Participants (US only) were recruited from Prolific in a staggered manner between 7 January and 27 Jaunary 2026. Recruitment through Prolific allowed access to a diverse pool of active internet users while also enabling prescreening based on behavioral criteria. Recruitment proceeded in multiple stages. First, participants completed a pre-screening survey. In this stage, we restricted the sample to individuals who (i) reported Google as their primary search engine, (ii) used Chrome as their primary (and only) browser, given that the extension was compatible only with Chrome, and (iii) reported using their current desktop for general browsing activities beyond simply survey platforms such as Prolific. Those who did not pass the pre-screening were screened out.

Participants who passed the pre-screening stage were asked for informed consent and then administered a baseline survey, which elicited information on their trust in AI-generated information online and their comfort with using AI tools for everyday information-seeking tasks. Following this, participants were directed to install our browser extension, which was listed on the Chrome Web Store and had undergone Google's review process for security and compliance.[5]

After installation, the extension conducted a browser history check based on the preceding 20 days. To qualify for the main study, participants were required to meet minimum activity thresholds based on their browser history in the last 20 days: (i) at least 10 visits to Google, (ii)

[5] As part of installation, participants were required to (i) enable incognito access for the extension and (ii) enter a study-specific passcode to prevent installations by people outside the study. Chrome extensions by default do not work in incognito mode. Rather than simply requesting users to enable this permission, the extension automatically directed them to the relevant settings page. The passcode prompt was triggered only after incognito access was granted.

visits to at least 25 unique domains, and (iii) browsing activity on at least 10 distinct days. These criteria were designed to exclude low-activity users and those using their devices primarily for survey-taking.

Participants who satisfied these conditions were invited to continue in the main study and were eligible for an additional bonus payment conditional on keeping the extension installed for two weeks. The bonus was paid on the 16th day following installation, after which participants were invited to complete an endline survey. The endline instrument measured self-reported compliance, overall satisfaction with the search process, perceived usefulness of search results, among others. These self-reported outcomes complement the behavioral data collected via the extension. Importantly, the endline survey also included items designed to assess whether participants correctly inferred the study's purpose. This allows us to evaluate the risk of experimenter demand effects. Participants who uninstalled the extension before completing the study period were not eligible for the bonus.

A total of 1,065 participants satisfied the browser history checks and took part in the main two-week study. Data collection began on January 7, 2026 (the first day of recruitment) and concluded on February 10, 2026, with up to two weeks of browsing data observed per participant.

2.2 Randomization and Treatment Arms

Eligible participants who passed all data checks were randomly assigned to one of three experimental conditions. Randomization occurred in real time at the individual level using a Bernoulli assignment procedure, as participants were enrolled sequentially. The two primary groups of interest were the control group (status quo) and the "Hide AI Overview" (HAIO) treatment group, each assigned with a probability of 36 percent. The remaining assignment probability of 28 percent corresponded to a third "AI Mode" group. This unequal allocation reflects the study's primary focus on estimating the causal effect of removing AIOs relative to the status quo, while still allowing exploratory analysis of the emerging conversational AI Mode interface.

The control group experienced no changes to their browsing environment (figure 2a). When these participants conducted searches on Google (or via their browser's search bar), they saw the standard search engine results page (SERP) exactly as Google presented it. If a given query triggered an AIO, the overview appeared normally. This group, therefore, represents the status

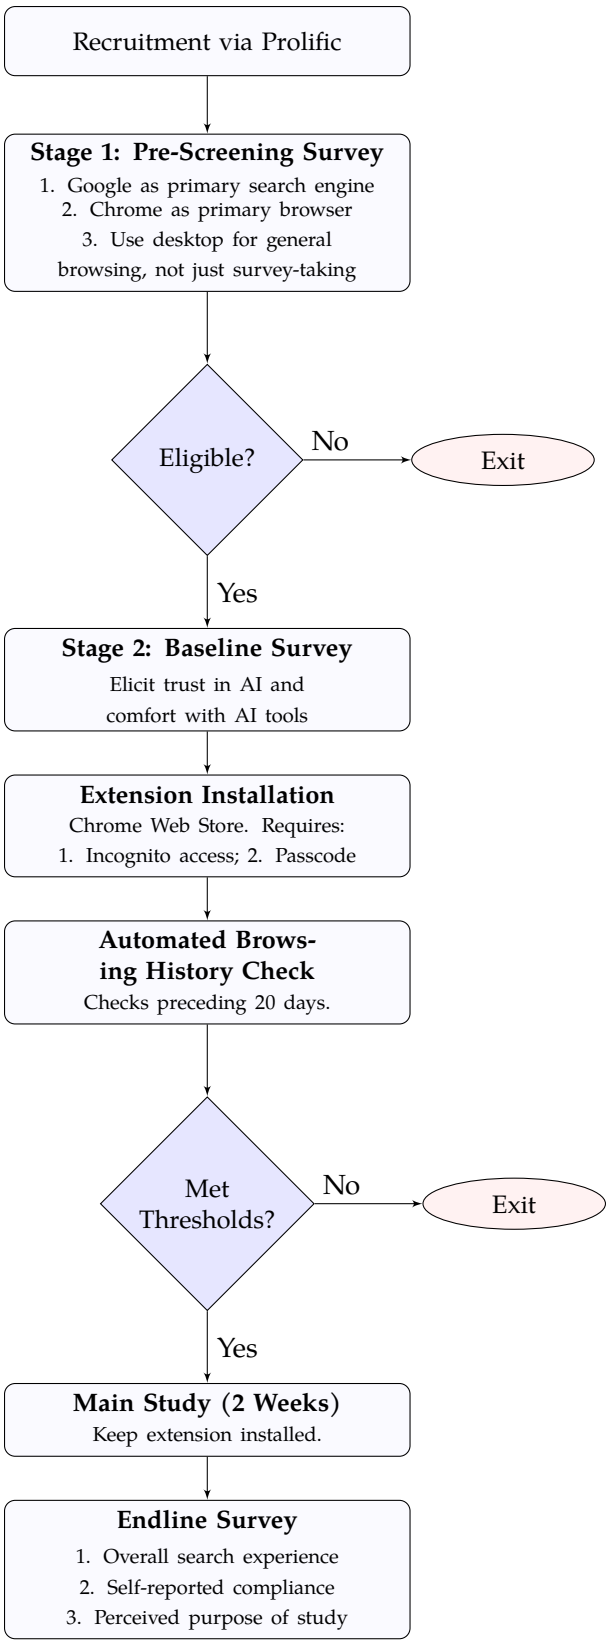

Figure 1: Participant Recruitment and Study Flow.

quo user experience and provides the primary counterfactual against which treatment effects are measured.

Participants assigned to the HAIO group experienced the study's main experimental manipulation. For these users, the Chrome extension automatically removed any AIOs that would otherwise have appeared on the SERP during the two-week experimental window (figure 2b). Importantly, the remaining search results were automatically shifted upward to fill the space previously occupied by the AIO, ensuring that the interface remained seamless and visually indistinguishable from a standard results page without AIOs. All other aspects of the search page remained unchanged. This design isolates the causal impact of AIO exposure by holding constant the broader search environment.

The third group—the AI Mode group—was exposed to Google's Gemini-powered conversational search interface whenever they conducted searches (figure 2c). This condition represents a more substantial departure from the traditional SERP experience. Because this interface is relatively new, we anticipated higher attrition and noncompliance in this arm. Accordingly, analyses involving the AI Mode group are conceptualized as primarily descriptive rather than strictly causal. Nevertheless, including this group provides useful contextual evidence about how users respond to a fully AI-mediated search experience.

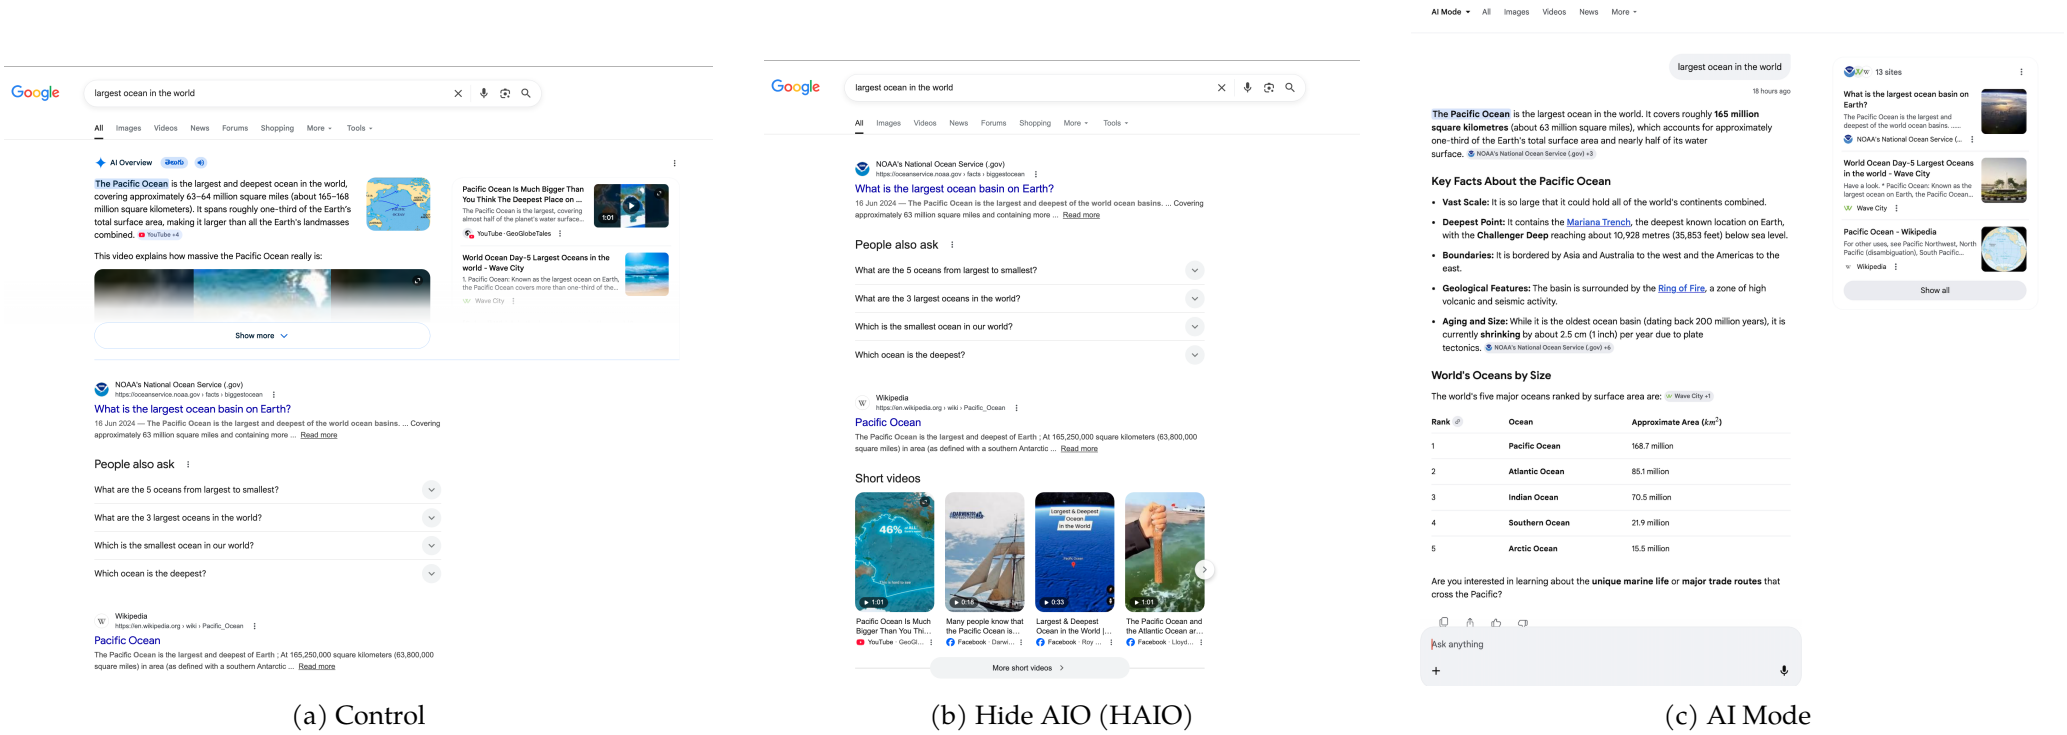

(a) Control (b) Hide AIO (HAIO) (c) AI Mode

Figure 2: Search interfaces across experimental conditions.

Notes: The Control group sees the standard SERP with AIOs when available. The HAIO group sees the SERP with AIOs removed and results shifted upward. The AI Mode group is exposed to a conversational AI interface.

2.3 Internal Validity

The extension was designed to operate seamlessly in the background, without introducing any noticeable changes to page load times or overall browsing performance. All modifications to the search interface were executed client-side and in real time, ensuring that users did not experience any visible lag, flicker, or disruption when interacting with the SERP. This was critical not only for preserving a naturalistic browsing environment and minimizing the likelihood that participants would detect the intervention, but also for maintaining an "all else equal" comparison across conditions, thereby strengthening internal validity.[6]

To ensure comparability across experimental conditions, we conducted balance tests on a range of pre-treatment characteristics, including demographic attributes, baseline attitudes toward AI, and prior browsing behavior. Demographic variables were obtained from Prolific's participant records. Measures of trust in AI-generated information and comfort with using AI tools were collected through the baseline survey administered prior to treatment assignment. Finally, browsing history measures were constructed using data collected by the extension during the eligibility screening stage. Table 1 reports the results of the balance tests. Across all observable characteristics, we find no statistically significant differences between the experimental groups, consistent with successful randomization.

Table 1: Balance in Means Across Experimental Groups

Variable	N	Control (C) (N = 396)	AI Mode (A) (N = 295)	HAIO (H) (N = 374)	p(C - A)	p(C - H)	p(A - H)
Age	1,063	38.22	38.08	39.03	0.87	0.37	0.32
Male	1,063	0.43	0.40	0.41	0.50	0.55	0.91
White	1,060	0.70	0.75	0.71	0.21	0.87	0.27
Latitude	1,065	37.36	37.31	37.71	0.90	0.33	0.32
Longitude	1,065	-91.04	-91.80	-89.85	0.53	0.28	0.12
Trust in AI info	1,065	3.37	3.39	3.34	0.71	0.75	0.51
Comfort with AI tools	1,065	3.59	3.52	3.64	0.38	0.46	0.12
Prior Google visits	1,045	242.51	241.54	219.02	0.97	0.29	0.40
Prior active browsing days	1,045	18.78	18.55	18.88	0.35	0.67	0.19
Prior unique domains visited	1,045	203.10	201.72	214.01	0.89	0.24	0.22

Notes: Trust in AI information and comfort with AI tools are measured on a Likert scale from 1 to 5. All prior browsing variables are computed over the 20 days preceding enrollment. Prior Google visits refer to the total number of unique URL visits to Google (excluding refreshes). Prior active browsing days measure the number of distinct days with any recorded browsing activity. Prior unique domains visited refers to the number of distinct domains visited.

While randomization ensures balance at baseline, differential attrition across experimental

[6]To further illustrate the user experience under each condition, we provide short video demonstrations of the extension in operation. Video demonstrations of the browsing experience under each experimental condition are available at https://youtu.be/-toedlcDP3s (Control), https://youtu.be/O5LeerRMDgs (HAIO), and https://youtu.be/Qz2D6TZwkdQ (AI Mode).

groups can undermine this comparability. We therefore examine participant retention across treatment arms.

We measure attrition using two complementary indicators. First, we track whether participants uninstall the extension, which is directly observed through the extension itself. Second, we capture attempts to circumvent the intervention. This latter channel is primarily a concern in the AI Mode condition, where users are exposed to a substantially different interface; in contrast, participants in the Control and HAIO groups experience minimal or no perceptible changes (as confirmed by our endline survey as well), and thus have little incentive to bypass the intervention. In particular, some users install alternative browser extensions that bypass our modifications. These extensions work by appending "udm=14" to Google search queries, which forces the search interface into a plain web-only mode. We treat the first occurrence of such a query as an attrition event, as it reflects a deliberate attempt to bypass the assigned condition.

Figure A.1 plots survival rates—defined as the proportion of participants who remain active (i.e., have neither uninstalled the extension nor engaged in bypass behavior)—over time since installation.

As shown in the figure, attrition patterns for the Control and HAIO groups are broadly similar over the study period, suggesting that removing AI Overviews does not meaningfully affect participant retention or compliance. In contrast, the AI Mode group exhibits substantially higher attrition, with a sharper decline in the share of active users over time. This pattern is consistent with the greater deviation of AI Mode from the standard search experience and suggests that users are less willing to remain in a fully AI-mediated interface. We also observe a discrete drop in survival around day 16, which coincides with the timing of the bonus payment and the administration of the endline survey.

Given this differential attrition, we treat the results from the AI Mode arm as primarily exploratory rather than causal, as selection into continued participation may bias estimates in this group.

We complement these behavioral checks with evidence from the endline survey, which provides additional insights into compliance and potential behavioral responses to the intervention. Importantly, the endline survey was administered after bonus payments had been disbursed, and participants were explicitly informed that their compensation would not depend on their responses. This design helps mitigate concerns about strategic or socially desirable

reporting.

We first examine whether participants adjusted their browsing behavior in ways that could undermine the intervention, such as shifting activity to a different browser or device. Such behavior could arise either as a direct response to the treatment or due to privacy concerns associated with installing the extension. Figures A.2 and A.3 in the Appendix show no meaningful differences between the Control and HAIO groups along these dimensions, suggesting that participants did not differentially substitute away from the observed browsing environment.

We next assess whether participants detected changes to the search interface. Over 95% of users in both the Control and HAIO groups report either not noticing any changes or being unsure (Figure A.4). For those who reported being unsure, we manually reviewed open-ended responses and found no evidence that participants correctly inferred the nature of the intervention. This provides strong evidence that the manipulation was largely imperceptible to users, reinforcing the credibility of our design and strengthening internal validity.

Finally, we examine self-reported compliance with the study protocol. Consistent with the behavioral measures discussed above, we find no meaningful differences in reported compliance between the Control and HAIO groups (Figure A.5). Taken together, these results suggest that the intervention did not induce systematic changes in user behavior beyond the intended treatment and that our estimates are unlikely to be confounded by differential compliance or awareness across groups.

3 Empirical Framework and Results

We begin by focusing on the HAIO treatment arm, which provides the cleanest setting for causal inference. We return to the AI Mode condition later, where we present results as exploratory.

We first document the prevalence of AIOs in our sample. Figure 3 shows that AIOs are triggered in approximately 42% of observed queries. This estimate is somewhat lower than that reported in Aral et al. (2026) (67%), likely reflecting differences in query composition: while Aral et al. (2026) rely on a curated set of queries, our measure is based on organically generated user queries in a natural browsing environment. Importantly, this prevalence is similar across the Control and HAIO groups, suggesting that removing AIOs does not meaningfully

alter the types of queries users submit[7]. This provides further reassurance that the intervention does not induce changes in search behavior at the query generation stage, and that our estimates primarily capture differences in responses to search results rather than differences in underlying search intent.

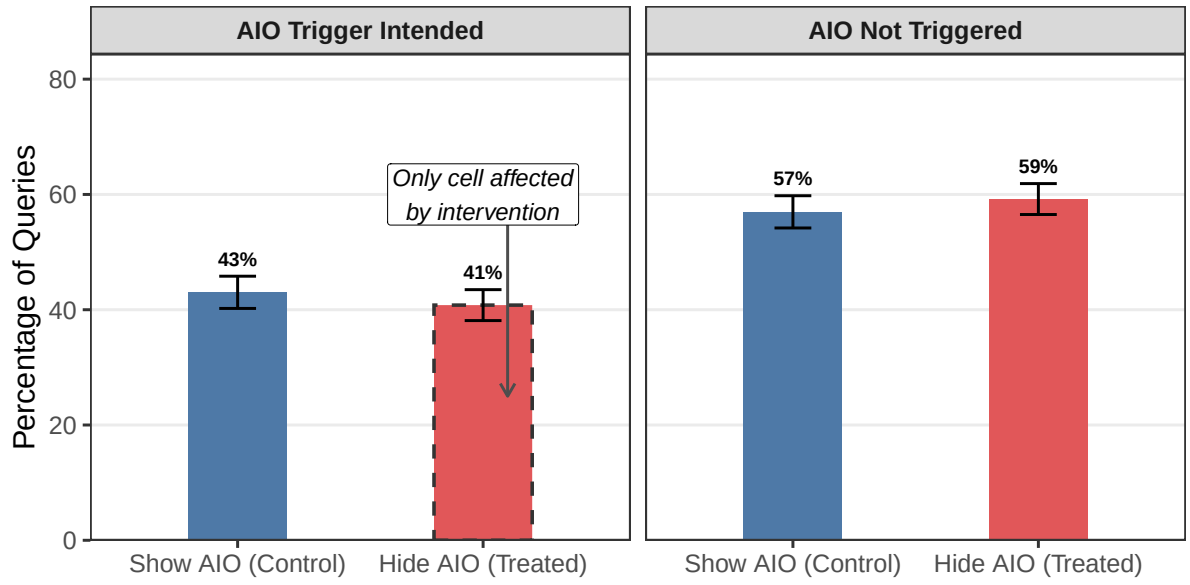

Figure 3: AIO Prevalence

Notes: The figure reports the share of queries in each condition for which an AI Overview (AIO) was (intended to be) triggered (left panel) and not triggered (right panel). The estimates are computed using search-level data across the Control and HAIO groups over the two-week period following enrollment. For queries in the HAIO group, AIO presence is inferred based on whether an overview would have been shown in the absence of the intervention. Error bars denote 95% confidence intervals clustered at the individual level.

A central question in the ongoing debate around AI-generated search summaries concerns their impact on outbound traffic to external websites or publishers. At a conceptual level, total clicks from Google can be decomposed as the product of the number of searches and the number of clicks conditional on a search. AIOs may affect both margins: they may reduce (or increase) the number of clicks on a given query (intensive margin), and they may also alter users' overall search frequency (extensive margin). Accordingly, we structure our empirical analysis to separately examine these two channels.

Our empirical strategy exploits the fact that the HAIO intervention only affects queries for which an AIO would have been shown (figure 3). We therefore structure all subsequent analyses by partitioning queries into two groups: those for which an AIO was intended to be triggered, and those for which it was not[8]. Queries in the latter category serve as a natural placebo group, as we do not expect the intervention to affect behavior in these cases. Accordingly, any

[7] A more detailed analysis of query types will be added in subsequent versions of this paper.
[8] We use the wording 'intended to be triggered' because in the HAIO group, the AIO was not actually shown.

treatment effects should be concentrated among queries where AIOs would have appeared in the absence of the intervention.

Across the 396 users in the control group and 374 in HAIO, we observe a total of 70,132 unique searches over the two-week period following each participant's enrollment. We focus on the search-level data to estimate the following regression at the individual i search s level.

$$Y_{is} = \gamma + \beta T_i + \epsilon_{is} \tag{1}$$

where T_i is an indicator for assignment to the HAIO treatment and β captures the causal effect of the HAIO on the outcome of interest. Because treatment is randomly assigned at the individual level, we estimate this specification without additional control variables in our baseline. Standard errors are clustered at the individual level to account for within-user correlation in outcomes across searches.

3.1 Search-level Analysis: Engagement with SERP Content

We examine four primary outcomes at the search level[9]. First, we measure the number of outbound organic clicks, defined as clicks to external websites from anywhere on the results page. Second, we consider the likelihood of a zero-click search, defined as an indicator for searches with no outbound organic clicks. Third, we analyze the number of internal clicks, which capture interactions within Google's ecosystem, such as navigating to images, news, or other result tabs, as well as engaging with query refinements or suggested searches. Finally, we examine the number of sponsored clicks.

Figure 4 presents the main results. We begin by focusing on queries for which an AIO was (intended to be) triggered (top row). In these queries, removing AIOs leads to a large and statistically significant increase in outbound organic clicks. The average number of external clicks per search increases from 0.38 in the control group to 0.61 in the HAIO group. These estimates suggest that AIOs reduce external clicks per search by approximately 38% when they appear. Consistent with this, the likelihood of a zero-click search declines sharply from 0.72 to 0.54 in the HAIO group, suggesting that AIOs increase the likelihood of a zero-click search by 33%. In contrast, for queries in which an AIO was not intended to be triggered (bottom row), the intervention does not alter the search interface. Consistent with this, we find no mean-

[9]More detailed analyses related to scroll depth, time spent on the clicked websites and bounce back rates will be added in subsequent versions.

ingful differences in outcomes between the Control and HAIO groups for these queries. The absence of effects in this placebo group strengthens the causal interpretation of our estimates and provides additional reassurance that the observed changes are driven by the removal of AIOs rather than by other aspects of the experimental setup.

We find no meaningful changes in other forms of engagement. The number of internal clicks remains largely unchanged, suggesting that the increase in outbound traffic is not driven by reduced interaction with other Google features. Similarly, sponsored clicks are unaffected, indicating that the primary behavioral substitution occurs between AI-generated summaries and organic website visits, rather than across the broader page ecosystem. Effects on internal and sponsored clicks, however, are less precisely measured, and detecting smaller effects on these margins may require greater statistical power.

These findings have important implications for the ongoing debate on how AI-generated summaries affect publisher traffic. Our results provide direct causal evidence that AIOs substitute for outbound clicks to external websites on a per-search basis. This is consistent with concerns raised by publishers and industry reports that AI summaries may reduce referral traffic by satisfying user intent directly on the search page.

Our estimates also highlight that the aggregate impact of AIOs depends critically on their prevalence. Currently, AIOs are triggered in roughly 42% of queries in our sample, implying that a majority of searches remain unaffected. However, our intensive margin estimates suggest that the effect per affected query is large. As the share of queries displaying AIOs increases over time (Aral et al., 2026), the overall reduction in outbound traffic could become substantially larger. In other words, even if the per-query effect remains constant, increasing AIO prevalence will mechanically amplify its aggregate impact on publisher engagement.

These findings speak to broader concerns about the evolution of search toward "answer engines" that directly resolve queries rather than directing users to external sources. By reducing the need to click through to underlying content, AIOs shift the role of search engines from intermediaries that facilitate discovery to destinations that aggregate and deliver information. This shift has important implications for content producers. Much of the online content ecosystem is sustained by referral traffic from search, which underpins both advertising revenue and subscription models. To the extent that AIOs reduce outbound clicks, they may weaken these incentives by lowering the returns to producing high-quality, discoverable content, potentially affecting both the quantity and quality of content over time.

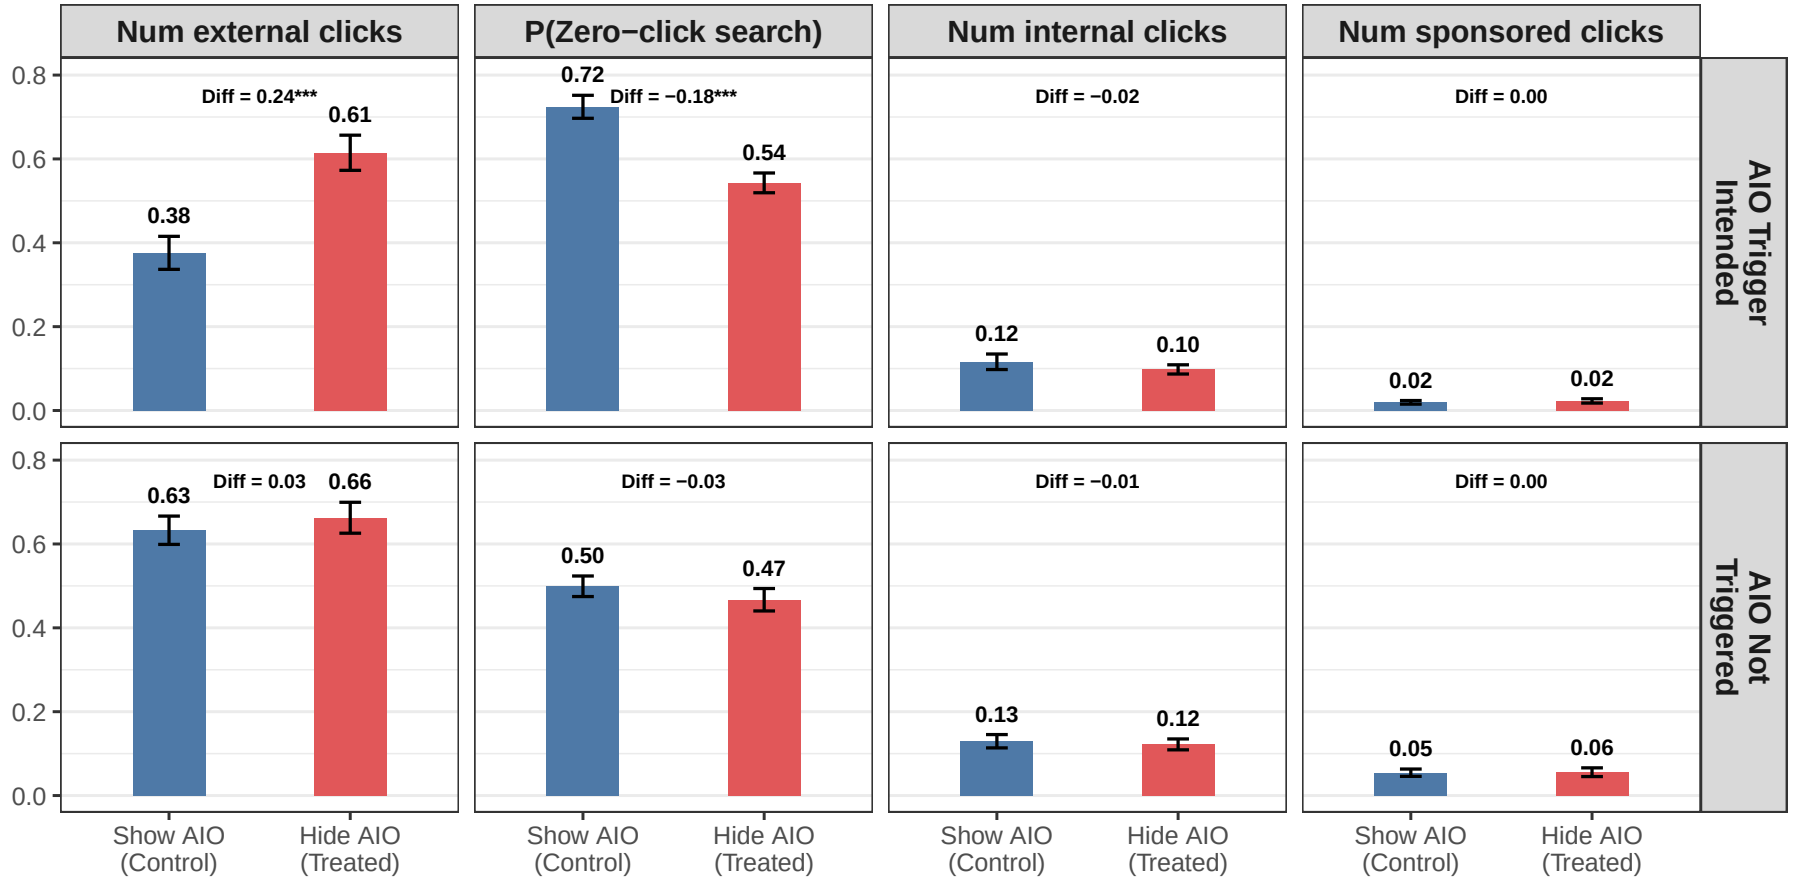

Figure 4: Search-level Treatment Effects

Notes: The top row includes queries for which an AIO was (intended to be) triggered, while the bottom row includes queries for which an AIO was not triggered. Each bar represents the corresponding average of search-level outcomes for the Control (Show AIO) and HAIO (Hide AIO) groups. Error bars denote 95% confidence intervals clustered at the individual level. The difference (HAIO − Control) shown between the bars corresponds to the estimate from Equation 1, with standard errors clustered at the individual level.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

3.2 Heterogeneous Treatment Effects

3.2.1 AIO Position

We next examine heterogeneity in treatment effects based on the position of the AI Overview (AIO) on the search results page[10]. This analysis was pre-registered prior to data collection. This dimension is particularly relevant given the strong role of position in shaping user attention and click behavior on search engine results pages (Agarwal et al., 2011). Prior work has consistently documented that content appearing at the top of the page captures a disproportionate share of user engagement.

In our setting, this consideration is especially salient because AIO placement is highly skewed. As shown in Figure A.6 in the Appendix, more than 80% of queries for which an AIO was (intended to be) triggered display the overview in the top position (position 0), above all organic results. This motivates a natural partition of queries into two groups: those where the AIO appears at the top of the page and those where it appears in a lower position.

We find that the effects documented in the previous section are almost entirely driven by queries in which the AIO appears at the top of the page (Figure 5). For these queries, removing the AIO leads to large and statistically significant increases in outbound organic clicks (88% increase relative to the control) and corresponding reductions in zero-click searches. In contrast, for queries where the AIO appears in a lower position, we observe no effect on user behavior.

These results suggest that the impact of AIOs is closely tied to their prominence on the page, with top-of-page placement playing a central role in driving substitution away from organic results. At the same time, it is important to note that this analysis captures causal heterogeneity with respect to observed AIO position, rather than the causal effect of position itself. Since AIO placement is not randomly assigned, differences across positions may reflect underlying differences in query types or user intent. Experimental variation in AIO placement would be required to cleanly isolate the causal effect of position from other confounding factors. Understanding this distinction is important for interpreting the mechanism behind our observed effects.

[10] We will also examine heterogeneity by query type in subsequent versions of the paper.

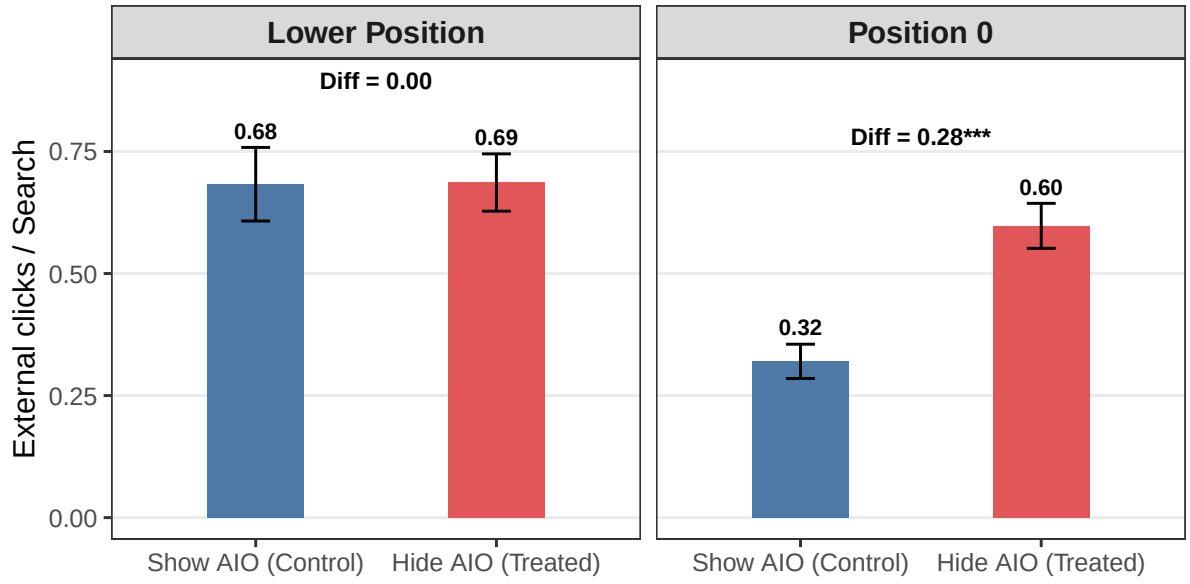

Figure 5: Position of AIOs on the SERP

Notes: The figure reports treatment effects separately by the position of the AIO on the search results page. Queries are split into those where the AIO appears at the top of the page (position 0) and those where it appears in a lower position. Queries where AIO were not triggered have been excluded. Each bar represents the corresponding search-level average number of organic external clicks for the Control (Show AIO) and HAIO (Hide AIO) groups. Error bars denote 95% confidence intervals clustered at the individual level. The difference (HAIO − Control) shown between the bars corresponds to the estimate from Equation 1, with standard errors clustered at the individual level.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

3.2.2 Search Position

We next examine how treatment effects vary across the position of organic search results on the page. In our data, most queries display between 7 and 10 organic results, and this distribution is balanced across the Control and HAIO groups, ensuring that differences in click behavior are not driven by differences in the composition of the results page (Figure A.8).

We therefore estimate treatment effects separately by the position of the clicked result on the SERP. Figure 6 presents results for queries in which the AIO appears at the top of the page (position 0), while Figure A.7 presents results for queries in which the AIO appears in a lower position.

Focusing first on cases where the AIO is shown at the top (Figure 6), we find that the increase in outbound clicks is concentrated among the highest-ranked results. The largest absolute increase occurs at position 1, followed by positions 2 and 3, with progressively smaller effects at lower ranks. This pattern suggests that AIOs primarily displace attention from the most prominent organic results, rather than from deeper links on the page.

In contrast, when the AIO appears below the top of the page (Figure A.7), we observe no

meaningful changes in click behavior at any position. Clicks across all ranks remain nearly identical between the Control and HAIO groups, reinforcing the earlier finding that AIOs only affect behavior when they occupy the most salient position on the page.

Taken together, these results provide further evidence that the effects of AIOs are driven by their prominence. By occupying the top position, AIOs divert attention away from the highest-ranked organic results that traditionally receive the majority of user engagement. When AIOs are placed lower on the page, they appear to have little to no impact on user behavior.

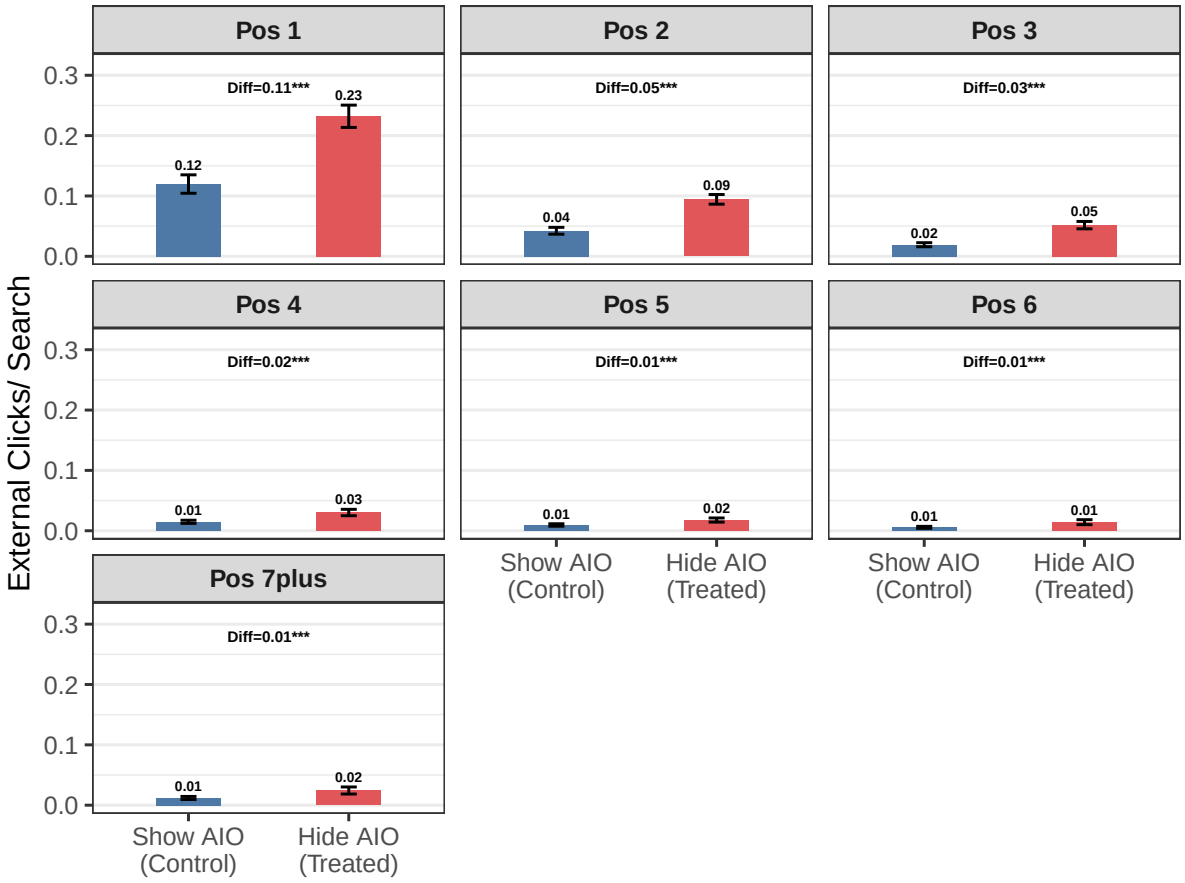

Figure 6: Organic Clicks by Search Position: AIO Shown at Top Position

Notes: The figure reports treatment effects on external organic clicks by the position of the clicked result on the search engine results page (SERP). Panels correspond to result positions (1 through 6 and 7+). The sample is restricted to queries for which an AIO was (intended to be) triggered at the top position. Each bar represents the average number of external clicks per search for the Control (Show AIO) and HAIO (Hide AIO) groups at the corresponding position. Error bars denote 95% confidence intervals clustered at the individual level. The reported differences (HAIO − Control) correspond to estimates from Equation 1, with standard errors clustered at the individual level. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

3.2.3 Demographics and Baseline AI Attitudes

Next, we examine whether treatment effects vary across demographic characteristics and baseline attitudes toward AI. This analysis was pre-registered prior to data collection.

We estimate the following specification at the search level:

$$Y_{is} = \alpha + \beta T_i + \gamma M_i + \delta(T_i \times M_i) + \epsilon_{is} \tag{2}$$

where T_i is an indicator for assignment to the HAIO treatment, and M_i is a moderator capturing either demographic characteristics (age, gender, race) or baseline attitudes toward AI (trust and comfort). For age, trust, and comfort, we construct binary indicators based on median splits: age is split at 37 years, while trust and comfort are split at a value of 4 or higher on the Likert scale. The coefficient β captures the treatment effect for the reference group, while δ captures the differential treatment effect for the group defined by the moderator. Standard errors are clustered at the individual level.

Table 2 reports the results. Across all specifications, we continue to find a positive and statistically significant baseline treatment effect, consistent with our earlier findings that removing AIOs increases outbound clicks.

Table 2: Heterogeneous Treatment Effects

	(1) Age	(2) Male	(3) White	(4) Trust	(5) Comfort
Treatment	0.183***	0.247***	0.264***	0.191***	0.147***
	(0.037)	(0.042)	(0.064)	(0.042)	(0.039)
Moderator	-0.023	0.078*	0.050	-0.098**	-0.141***
	(0.038)	(0.038)	(0.046)	(0.036)	(0.032)
Treatment × Moderator	0.132*	-0.018	-0.039	0.076	0.128*
	(0.057)	(0.055)	(0.070)	(0.054)	(0.054)
Observations	29,423	29,379	29,384	29,423	29,423

Notes: Each column reports results from a separate regression interacting the treatment indicator with a moderator. DV is number of organic external clicks at the search level. Moderators correspond to indicators for being above the median age (37), male, white, above-median trust in AI (4+), and above-median comfort with AI tools (4+). Standard errors are clustered at the individual level.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

For gender and race, the interaction terms are small in magnitude and statistically insignificant, suggesting that the effects of AIOs on click behavior are broadly similar across these groups. For age, we find some evidence of heterogeneity. The interaction term is positive and statistically significant, indicating that the increase in outbound clicks from removing AIOs is larger for individuals above the median age. This implies that AIOs reduce outbound clicks more strongly for older users, suggesting that they are relatively more likely to rely on the overview rather than clicking through to external websites.

Turning to baseline attitudes toward AI, we find modest heterogeneity along some dimensions. The interaction term for trust in AI is positive but not statistically significant, suggest-

ing that treatment effects are broadly similar across users with different levels of trust in AI-generated information. In contrast, we find some evidence of heterogeneity by comfort with AI tools. The interaction term is positive and statistically significant, indicating that the increase in outbound clicks from removing AIOs is larger for individuals who report higher comfort with AI. This suggests that more AI-comfortable users are better able to extract value directly from AIOs, reducing their need to click through to external websites when the overview is present.

Overall, these results suggest that the effects of AIOs are relatively broad-based, with only limited variation across observable user characteristics. While some heterogeneity emerges along age and comfort with AI, the core finding that AIOs substitute for external website visits holds consistently across groups.

3.3 Additional Results: AI Mode Analysis

We now turn to the AI Mode condition. As discussed earlier, participants were assigned to this arm with a lower probability than to the Control and HAIO groups. This design choice, combined with the higher anticipated attrition in this condition, reflects the fact that our analysis of AI Mode is intended to be exploratory rather than strictly causal.

Analyzing behavior in AI Mode presents additional measurement challenges relative to the standard search interface. In the traditional SERP setting, each search instance can be cleanly defined by a query. In contrast, AI Mode involves a conversational interface in which users may issue multiple follow-up inputs within the same interaction. These follow-ups can include genuine new queries, refinements of the original query, or non-informational inputs (e.g., acknowledgments or clarifications). As a result, defining a "search" is less straightforward in this setting.

To address this, we adopt a simplified and conservative approach. We define a search instance in AI Mode as the initial query that initiates the conversational window, and we aggregate all subsequent interactions and clicks within that session into a single observation. Thus, even if a user issues multiple follow-up prompts, all resulting clicks are attributed to the same search instance. Overall, we observer 12,507 such unique search instances among the AI mode users in the two-week observation period. This approach may overstate clicks per search if users effectively conduct multiple distinct searches within a single session, since all the clicks are attributed to the original query.

Despite this potential upward bias, we find that overall engagement with external content

is substantially lower in AI Mode relative to both the Control and HAIO groups. As shown in Figure 7, the average number of external clicks per search is markedly lower in AI Mode relative to even the default with AIO, while the likelihood of a zero-click search is correspondingly higher. For comparability, we do not partition Control and HAIO queries by AIO trigger status in this analysis, as AI Mode does not have a direct analogue to this distinction.

These patterns suggest that a fully conversational, AI-mediated interface may lead to even stronger substitution away from external websites than the AIOs embedded within the traditional SERP. While these results should be interpreted with caution due to the non-causal nature of this comparison and the higher attrition in this arm, they provide suggestive evidence that as large language model interfaces become more prevalent, the overall impact on outbound traffic could be substantially larger than what we estimate in the HAIO setting.

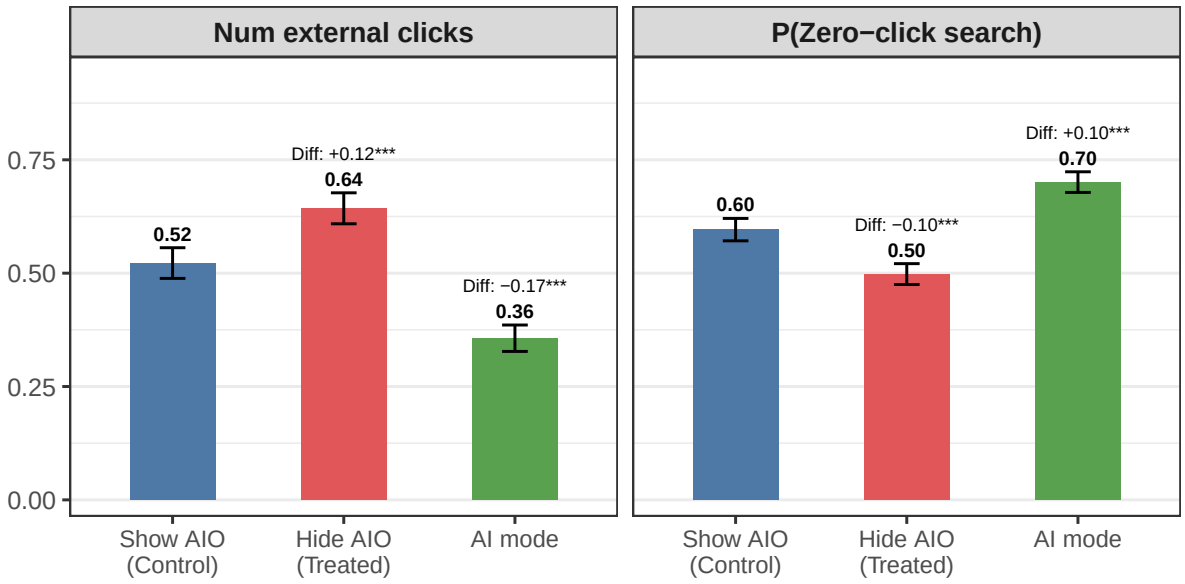

Figure 7: Treatment Effects for AI Mode Group

Notes: The figure compares engagement outcomes across the Control (Show AIO), HAIO (Hide AIO), and AI Mode groups. Outcomes are aggregated at the level of a search instance. In the AI Mode condition, a search instance is defined as the initial query that initiates a conversational session, with all subsequent interactions and clicks within that session aggregated into a single observation. For comparability, queries in the Control and HAIO groups are not partitioned by AIO trigger status. Each bar represents the corresponding average outcome at the search level. Error bars denote 95% confidence intervals clustered at the individual level. The differences across groups are based on regressions analogous to Equation 1, with standard errors clustered at the individual level. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

3.4 Extensive Margin Effects

The search-level results establish that AIOs reduce clicks conditional on a search. A natural next question is whether they also affect the overall volume of search activity. If AIOs make search more convenient, they may increase the total number of searches users conduct. To

examine this extensive-margin channel, we aggregate behavior to the user level and focus on participants who remained active for the full two-week study period.

Table 3 reports the total number of searches conducted by users (averaged across all users) in the Control and HAIO groups over the study period. We find no meaningful difference between the two groups. The point estimate is small in magnitude and statistically insignificant. This result suggests that removing AIOs does not materially affect how often users search. In other words, the treatment does not appear to change overall search frequency.

Table 3 reports the corresponding user-level totals for outbound organic clicks. Here, we find a strong pattern. For queries in which an AIO was (intended to be) triggered, users in the HAIO group generate significantly more external clicks over the two-week period than users in the Control group. In contrast, for queries where AIOs were not triggered, total external clicks are statistically similar across the two groups. This mirrors the search-level results and indicates that the increase in clicks per search translates into a higher total volume of outbound traffic at the user level.

Table 3: Extensive Margin Effects

	(1) Total Searches	(2) Total Clicks (AIO Triggered)	(3) Total Clicks (AIO Not Triggered)
Hide AIO (Treated)	-5.05	7.56***	1.18
	(8.57)	(2.03)	(3.92)
Observations	713	713	713
Control Mean	98.06	15.95	35.42

Notes: Each column reports the coefficient from a user-level regression of the indicated outcome on an indicator for assignment to the HAIO treatment. Column (1) reports effects on the total number of searches conducted over the study period. Columns (2) and (3) report effects on the total number of outbound organic clicks generated from queries for which an AIO was (intended to be) triggered and not triggered, respectively. The analysis is restricted to Control and HAIO users who remained active for the full two-week study period. Because the data are collapsed to the user level, standard errors are not clustered. Standard errors are reported in parentheses.
$^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Taken together, these findings suggest that the effects of AIOs operate primarily through the intensive margin rather than the extensive margin. AIOs do not meaningfully change how often users search, but they do reduce the number of external clicks generated from those searches. This implies that the decline in publisher traffic associated with AIOs reflects lower engagement per search rather than a change in overall search activity.

3.5 Endline Survey: User Experience on Google Search

While our results show that AI Overviews (AIOs) reduce outbound clicks to external websites, an important complementary question is whether they improve user experience. If AIOs deliver substantial benefits to users by making information easier to access or improving satisfaction, then reductions in publisher traffic may reflect an efficiency gain rather than a welfare loss.

To examine this, we analyze responses from the endline survey, which elicited participants' perceptions of their search experience. We had more than 90% take-up in the endline survey. We restrict this analysis to participants who remained active for the full two-week study period (i.e., did not uninstall the extension or engage in bypass behavior), ensuring that responses reflect sustained exposure to the assigned condition. Figure 8 reports average responses across several dimensions, including trust in AI-generated information, comfort with AI tools, overall satisfaction with Google Search, perceived quality of information, and ease of finding information.

Across all these measures, we find no meaningful differences between the Control and HAIO groups. Self-reported satisfaction, perceived quality, and ease of use are virtually identical across the two conditions (precisely estimated). This suggests that removing AIOs does not materially affect users' perceived search experience. In fact, as discussed earlier, the vast majority of participants in the HAIO group did not even detect any changes to the interface during the two-week study period (figure A.4).

In contrast, we observe a markedly different pattern for participants assigned to the AI Mode condition. Across all dimensions—overall satisfaction, perceived quality of information, and ease of finding information—responses are substantially lower relative to both the Control and HAIO groups. This suggests that users find the fully conversational, AI-mediated search experience less satisfactory than the traditional SERP.

It is important to note that these estimates are likely upward biased due to selection. As documented in the attrition analysis, the AI Mode group experienced significantly higher dropout rates over the study period. The endline responses therefore reflect the subset of users who chose to remain in the study despite being exposed to this interface. If anything, the average user experience in AI Mode is likely worse than what is captured in these responses.

Taken together, these findings suggest that AIOs and conversational AI searches divert traf-

fic away from external websites without delivering measurable improvements in user experience, at least along the dimensions captured in our survey. This has important implications for the broader debate on the role of generative AI in search. If generative AI in search reduce referral traffic to publishers without providing commensurate benefits to users, this raises concerns about the sustainability of the content ecosystem that underpins search, as well as about platform incentives and market power. As generative AI becomes more deeply integrated into search interfaces, understanding these trade-offs will be critical for informing both platform design and competition policy.

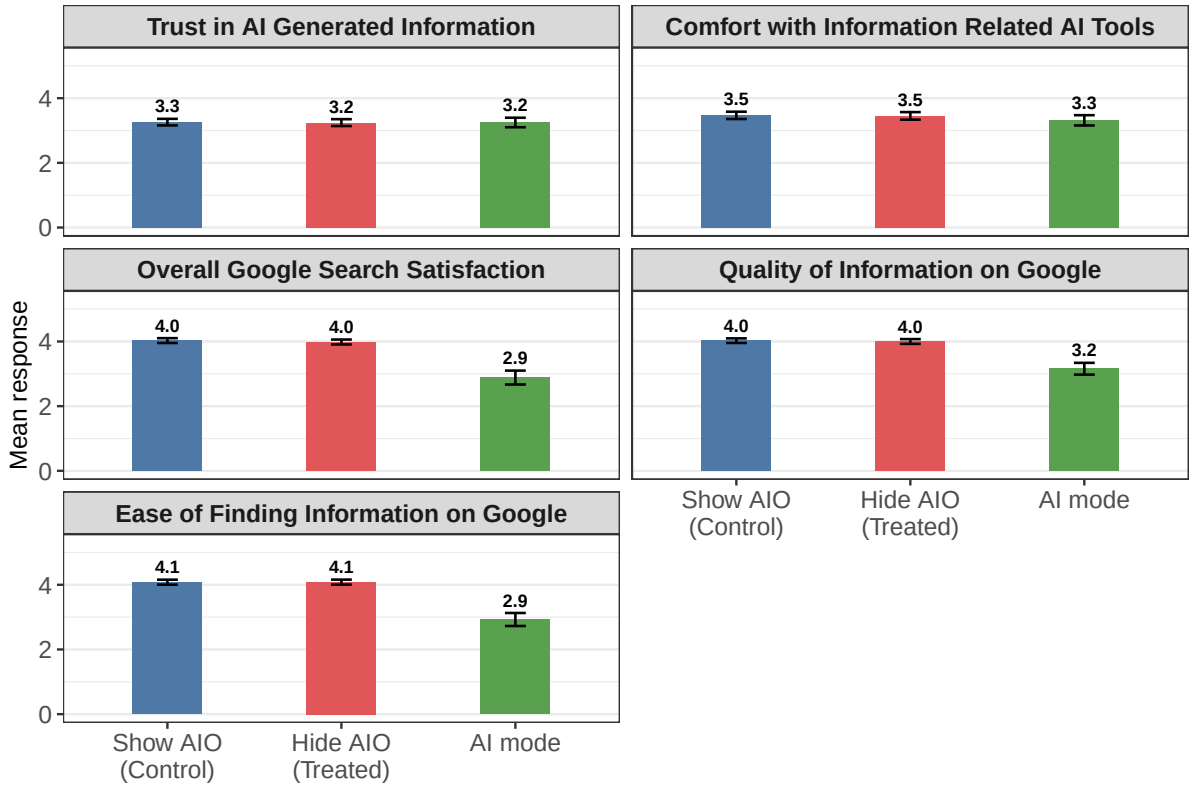

Figure 8: Endline Survey Responses

Notes: The figure reports average responses from the endline survey across treatment groups. Outcomes include trust in AI-generated information, comfort with AI tools, overall satisfaction with Google Search, perceived quality of information, and ease of finding information. All measures are collected on Likert scales ranging from 1 to 5. The survey was administered after bonus payments were disbursed, and participants were informed that their compensation would not depend on their responses. Each bar represents the mean response for the corresponding group. Error bars denote 95% confidence intervals.

Finally, we examine responses related to trust in AI-generated information and comfort with AI-based tools. Across all three groups—Control, HAIO, and AI Mode—we find no meaningful differences in these measures at endline. This suggests that short-term exposure to different search interfaces does not significantly shift users' underlying attitudes toward AI. In other words, while AI-driven features can influence behavior on the margin, they do not appear to

alter deeper perceptions of trust or comfort over the time horizon of our study.

4 Conclusion

This paper aims to provide timely evidence to platforms, publishers, and policymakers on how integrating generative AI into search interfaces reshapes user behavior and the broader digital information ecosystem. Leveraging a field experiment with real users, we isolate the causal impact of AI Overviews (AIOs) on engagement patterns, offering insights that move beyond speculative or correlational claims. Our findings demonstrate that the prevalence of AIOs substantially reduces clicks to organic search results, increasing the prevalence of zero-click searches while leaving interactions with sponsored content largely unchanged. These results suggest that AI-generated summaries primarily substitute for visits to external websites rather than redistributing attention across the search page.

Given that AIOs appear in a significant share of queries, a share that is likely to rise in the coming months and years, their aggregate impact on publisher traffic is likely to be substantial. At the same time, the absence of measurable improvements in user satisfaction or perceived utility challenges the notion that these features deliver clear consumer welfare gains. Instead, our evidence points to a redistribution of value within the ecosystem, where platforms may increasingly capture user attention. The exploratory results from the AI Mode condition further highlight the complexities of transitioning toward more AI-mediated search environments. More immersive AI interfaces appear to further drive substitution away from external websites.

Taken together, our findings contribute to ongoing debates in competition policy, platform regulation, and digital market design. They highlight a fundamental tension: while AI integration can streamline access to information, it may simultaneously undermine the economic incentives that sustain high-quality content production. Ultimately, as generative AI becomes more deeply embedded in core digital infrastructures, understanding its behavioral and economic consequences will be critical for shaping a more innovative and equitable online information environment.

References

Agarwal, Ashish, Kartik Hosanagar, and Michael D Smith, "Location, location, location: An analysis of profitability of position in online advertising markets," *Journal of marketing research*, 2011, *48* (6), 1057–1073.

Agarwal, Saharsh and Saptashya Ghosh, "AI at the Fingertips: How do LLM Apps Change Smartphone Usage," 2026.

Aral, Sinan, Haiwen Li, and Rui Zuo, "The Rise of AI Search: Implications for Information Markets and Human Judgement at Scale," *arXiv preprint arXiv:2602.13415*, 2026.

Bellan, Rebecca, "Google's AI Overviews are killing traffic for publishers," https://techcrunch.com/2025/06/10/googles-ai-overviews-are-killing-traffic-for-publishers/ June 2025. Accessed: 2026-02-24.

Calzada, Joan and Ricard Gil, "What do news aggregators do? Evidence from Google News in Spain and Germany," *Marketing Science*, 2020, *39* (1), 134–167.

Chiou, Lesley and Catherine Tucker, "Content aggregation by platforms: The case of the news media," *Journal of Economics & Management Strategy*, 2017, *26* (4), 782–805.

Farronato, Chiara, Andrey Fradkin, and Chris Karr, "Webmunk: A new tool for studying online behavior and digital platforms," Technical Report, National Bureau of Economic Research 2024.

Gholami, Samira, Cristiana Firullo, Cristobal Cheyre, and Alessandro Acquisti, "Beyond Search: LLM Adoption and the Concentration of Web Traffic."

Google, "AI Overviews in Search," https://blog.google/products-and-platforms/products/search/generative-ai-google-search-may-2024/ May 2024. Accessed: 2026-02-24.

Khosravi, Mehrzad and Hema Yoganarasimhan, "Impact of AI Search Summaries on Website Traffic: Evidence from Google AI Overviews and Wikipedia," *Available at SSRN 6164926*, 2026.

Muvija, M., "UK regulator proposes changes to Google search for publishers," https://www.reuters.com/legal/litigation/uk-regulator-proposes-changes-google-search-publishers-2026-01-28/ January 2026. Accessed: 2026-02-24.

Padilla, Nicolas, H Tai Lam, Anja Lambrecht, and Brett Hollenbeck, "The impact of llm adoption on online user behavior," *Available at SSRN*, 2025.

Singh, Madhavi and Fiona M. Scott Morton, "A Roadmap for a Monopolization Case Against Google: Monopsony Power and AI Overviews," Technical Report, SSRN January 2026. Forthcoming in University of Pennsylvania Law Review, Vol. 175.

Song, Yu and Puneet Manchanda, "Frontiers: Does Carrying News Increase Engagement with Non-News Content on Social Media Platforms?," *Marketing Science*, 2025, *44* (4), 748–759.

Su, Yi, Qili Wang, Liangfei Qiu, and Runyu Chen, "Navigating the sea of reviews: Unveiling the effects of introducing AI-generated summaries in E-commerce," *Available at SSRN 4872205*, 2024.

Vynck, Gerrit De, "In the fight for tech giants to pay for news, Google relented. Facebook didn't.," *The Washington Post*, 2021, pp. NA–NA.

Zhai, Sihan and Andrew T Ching, "AI Summary Overrepresents Fake Reviews: Evidence from Amazon," 2026.

A Appendix: Additional Results

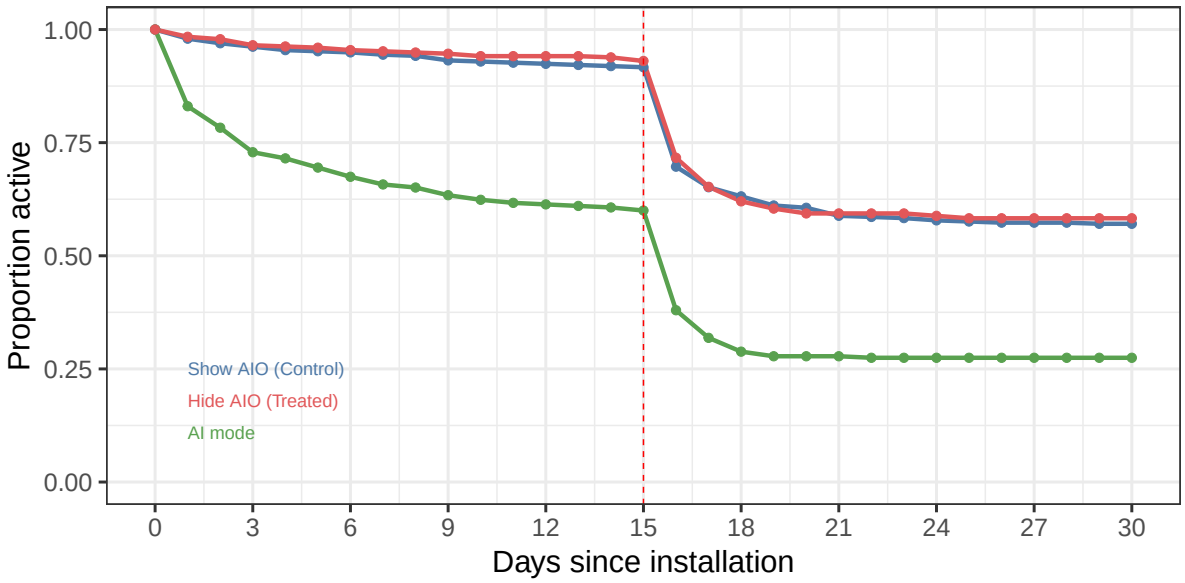

Figure A.1: Participant Survival Over Time by Treatment Group

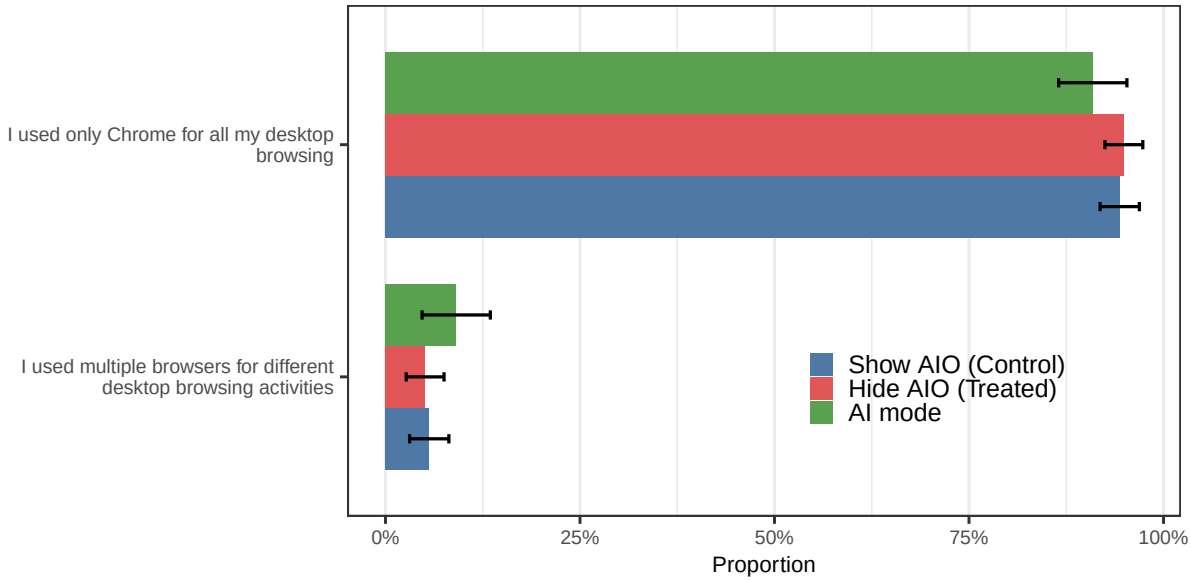

Figure A.2: Endline: Exclusive Use of Chrome Browser

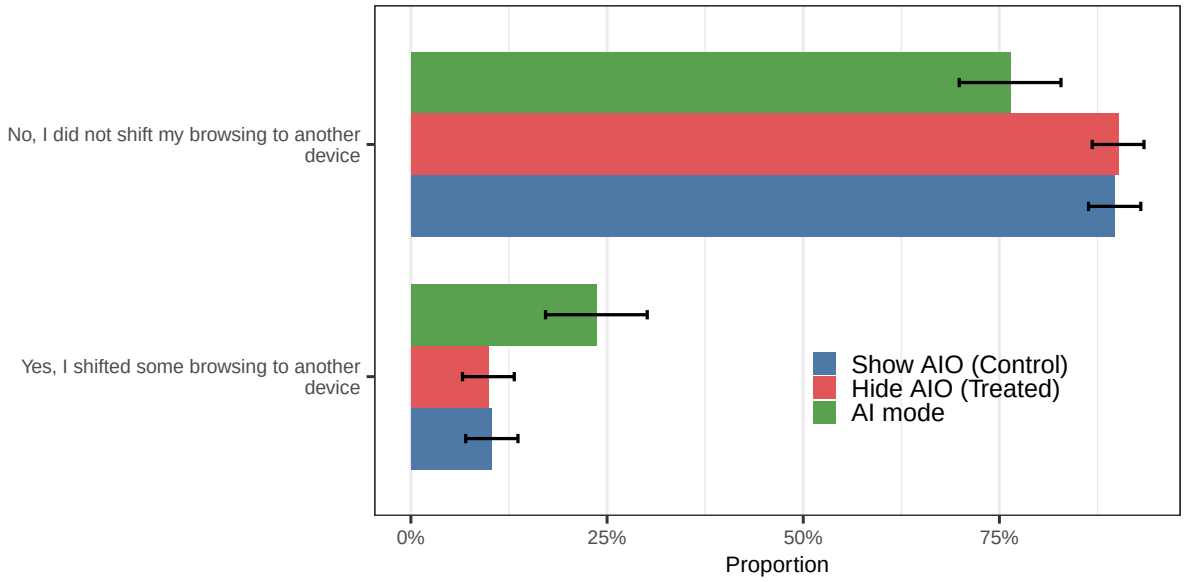

Figure A.3: Endline: Shift to Another Device

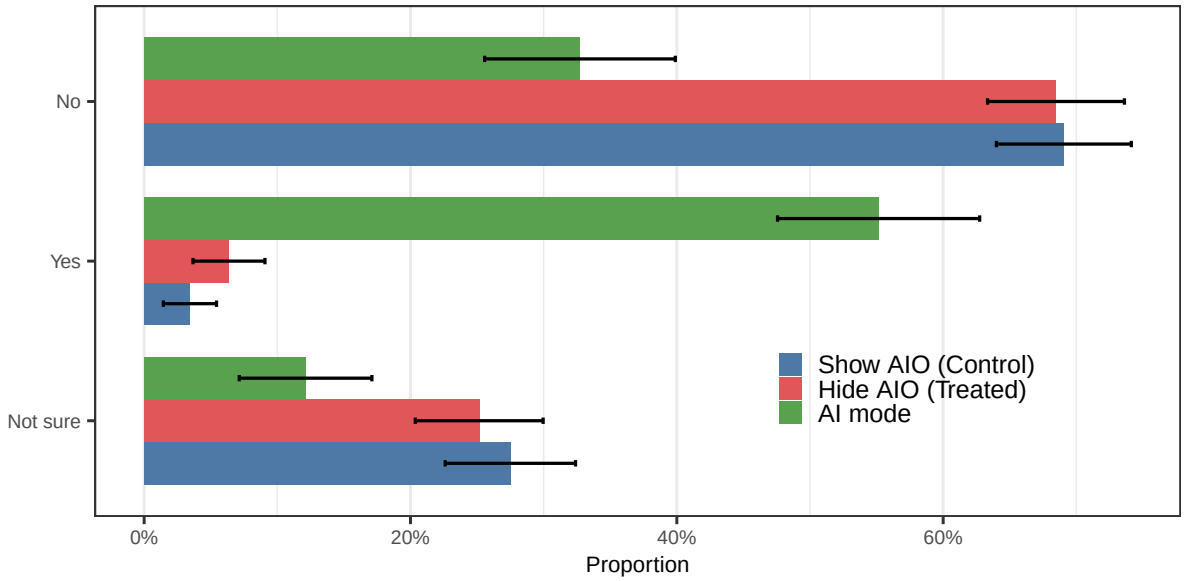

Figure A.4: Endline: Noticed Changes

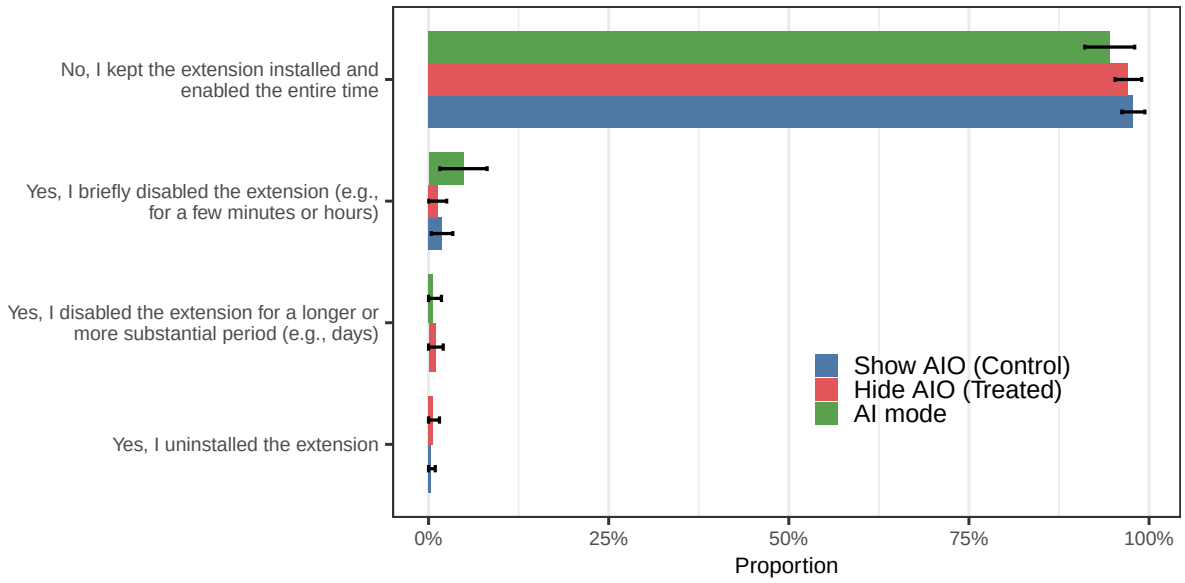

Figure A.5: Endline: Self-reported Compliance

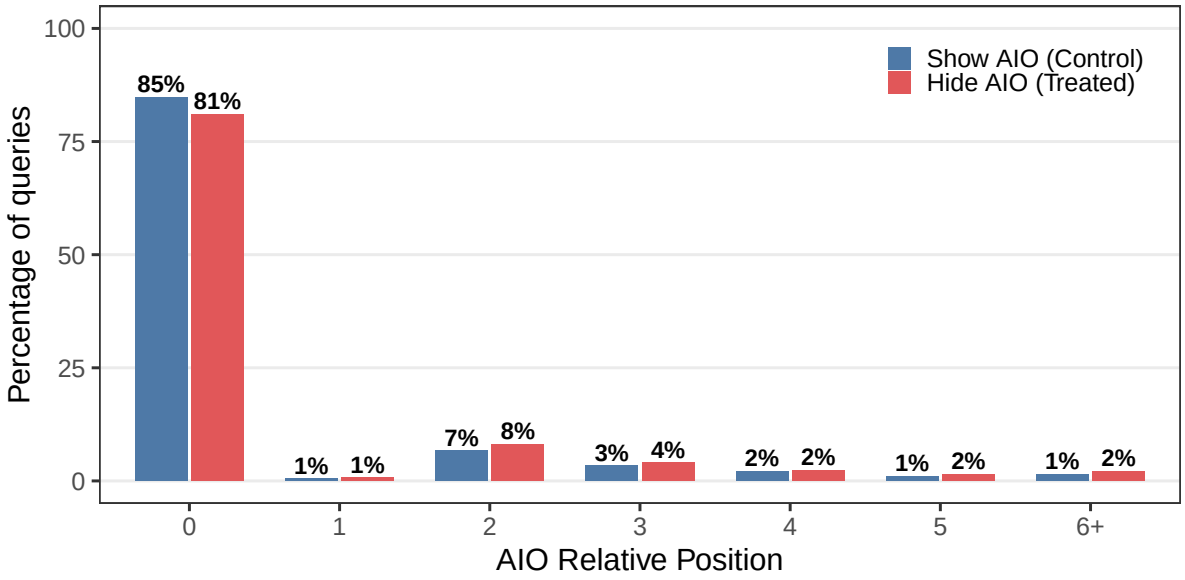

Figure A.6: AIO Position Distribution

Notes: The figure reports the distribution of AI Overview (AIO) positions on the search engine results page (SERP) across queries for which an AIO was (intended to be) triggered. Relative position 0 corresponds to cases where the AIO appears at the very top of the SERP, before any other content is shown. A relative position of x indicates that there are $x - 1$ elements above the AIO (e.g., organic search results, video snippets, or other SERP features).

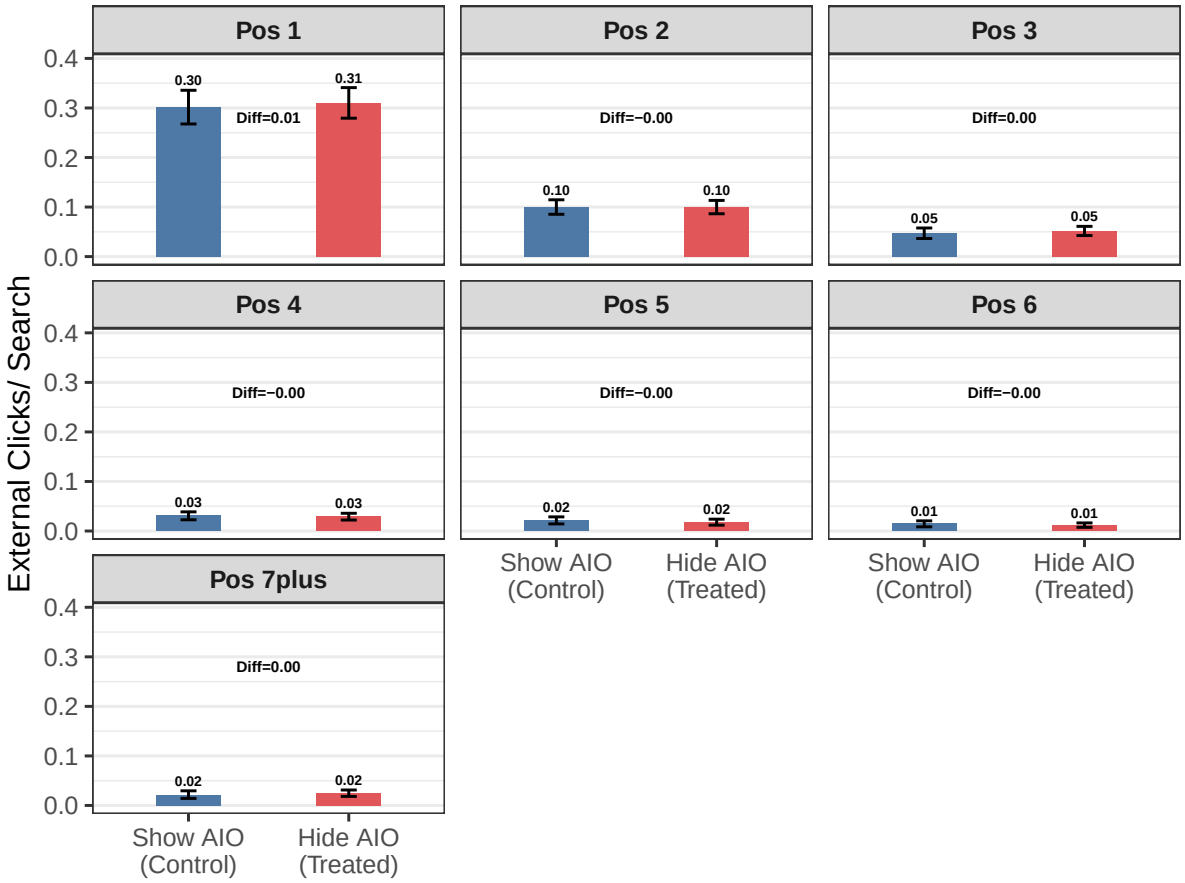

Figure A.7: Organic Clicks by Search Position: Overview Not Shown at Top position

Notes: The figure reports treatment effects on external organic clicks by the position of the clicked result on the search engine results page (SERP). Panels correspond to result positions (1 through 6 and 7+). The sample is restricted to queries for which an AIO was (intended to be) triggered, but not at the top position. Each bar represents the average number of external clicks per search for the Control (Show AIO) and HAIO (Hide AIO) groups at the corresponding position. Error bars denote 95% confidence intervals clustered at the individual level. The reported differences (HAIO − Control) correspond to estimates from Equation 1, with standard errors clustered at the individual level. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

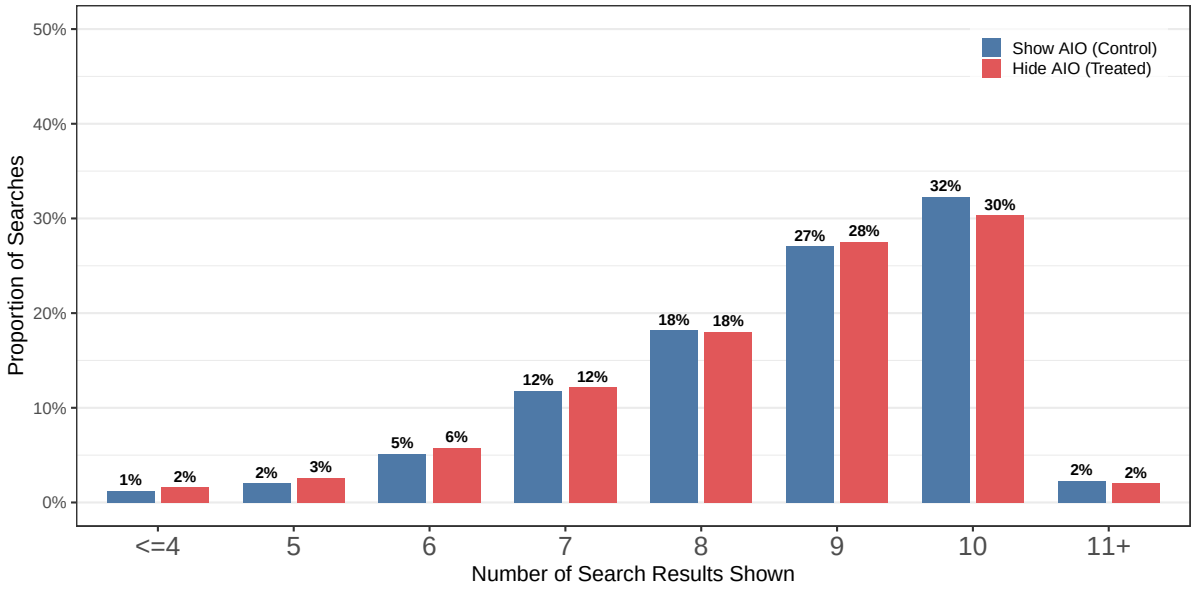

Figure A.8: Number of Search Results distribution