The Effects of Advertising and Personalization on Search Costs: Evidence from an Online Field Experiment*

Cristiana Firullo† *Cornell*

Cristobal Cheyre *Cornell*

Alessandro Acquisti *MIT*

September 2026

Abstract

We study the impact of advertising and personalization on consumer search costs using data from an online field experiment where participants browse with either advertising blocked, or with behaviorally targeted advertising, or with reduced tracking that limits ad personalization. We introduce a method to identify and analyze *shopping intents* using browsing and platforms' metadata. Relative to participants who see no advertising, those exposed to targeted advertising spend 22 percent more time per product evaluated and 25 percent more time per shopping intent. Yet they evaluate no more products, visit no more merchants, and purchase products of similar price and quality. Participants exposed to less personalized advertising are statistically indistinguishable from those who see no ads on these outcomes. We then evaluate whether the additional time spent under targeting reflects more intensive product evaluation, using a duration-based search model to estimate how quickly time spent on a product improves the consumer's information about it. Under targeted advertising, this learning rate is about one-third lower than when advertising is blocked. Thus, the additional search time generated by targeting appears to be less productive rather than more informative. Valued at the average hourly wage, the additional time costs the consumer about 1.5 cents per personalized ad seen and $17 to $32 per participant over three months.

*This work was supported by the Alfred P. Sloan Foundation, the Carnegie Mellon University Block Center for Technology and Society, the Carnegie Mellon University CyLab Security and Privacy Institute, the Institute of Humane Studies and the Program on the Economics & Privacy at George Mason University, the University of Pennsylvania Center for Technology, Innovation and Competition and the Warren Center for Network & Data Sciences, the Pittsburgh Supercomputing Center, the NET Institute, the Cornell Center for Social Sciences, and the National Science Foundation (awards 2237327/2237328/2237329).

†Email: cf528@cornell.edu.

1. Introduction

The conventional wisdom in the debate around digital advertising is that online ads, and particularly behaviorally targeted ones, benefit consumers by lowering their search costs. A primer of the FTC Bureau of Economics states that "targeting benefits the consumer because it effectively reduces their search costs" (Lau, 2020, p. 5), and an FTC Commissioner argues that "a relevant ad lowers consumers' search costs, which is the entire economic justification for advertising in the first place" (Holyoak, 2024, p. 6). The advertising industry makes the same claim: the Network Advertising Initiative writes that tailored advertising "reduces shopping and product discovery costs" (Network Advertising Initiative, 2026), and Meta describes personalized ads as an important way for consumers to discover small businesses (Meta, 2021). However, an intermediary paid per impression monetizes attention itself, and therefore may have an incentive to divert rather than resolve consumer search, whereas an intermediary paid per transaction shares the seller's interest in shortening the path from goal formation to purchase (Hagiu & Jullien, 2011, 2014; de Cornière & Taylor, 2014, 2019). This potential misalignment is consistent with research on incentives to prolong engagement on advertising-funded platforms (Ichihashi & Kim, 2023; Hébert & Zhong, 2026; Kleinberg et al., 2022; Aridor et al., 2024), while studies of platforms that earn revenue from completed transactions document design changes that reduce consumer search costs (Dinerstein et al., 2018; Ursu, 2018). [1]

Empirical tests of the relationship between online ads and search costs face two challenges. First, many studies are not designed to separate the effect of advertising from the effect of targeting and personalization of the ads. Second, search costs are often estimated from searches that culminate in a purchase (Hong & Shum, 2006; Ursu, 2018; Ursu et al., 2020), or in studies that focus on the effectiveness of specific ad campaigns, rather than through designs that capture the impact of ecosystem-wide exposure to advertising and are able to record search paths that end without a transaction. We use data from the Internet Behavior Experiment ('IBE'), an experimental infrastructure designed to investigate the impact of online tracking, targeting, and advertising on consumer behavior and welfare (Cheyre et al., 2024), to examine the impact of advertising and personalization on consumer search costs. Our experiment separates these two margins by design, with

[1] The mapping from revenue model to search frictions is not mechanical: retailers themselves obfuscate when search compresses margins (Ellison & Ellison, 2009).

one arm that removes advertising, one that shows personalized advertising, and one that shows advertising with tracking curtailed, holding organic personalization fixed, while a browser extension records the complete search path of every episode, including those that end without a purchase. IBE participants install a research browser extension and are randomly assigned for three months to *Ad-Blocking* (advertising blocked), *Opt-In* (tracking is accepted and advertising is personalized as usual), or *Anti-tracking* (an extension blocks third-party trackers, so ads appear but are less likely to be personalized). The extension records every retail results page with its product tiles, organic and sponsored, every product page opened with the time spent on it, and every purchase. A language model classifies each product seen by the user into a category, and groups related viewing, clicking, and purchasing events that correspond to a common shopping goal. These grouped events define a "shopping intent", our unit of analysis. Viewing several running shoes across three days is one shopping intent, while an unrelated grocery search during the same period constitutes another. We keep the intents in which the participant opens at least one product page, or sees the products of the intent on at least two result pages. We allocate time spent on a results page across the product tiles appearing before the product that is opened, and count time spent on opened product pages in full. Only about 15 percent of intents end in a purchase, with no significant differences across arms. Hence, an analysis limited to completed purchases would miss most of the search behavior captured by the IBE. We restrict the analysis to the shopping that participants do while logged into the merchant. When a participant is logged into a merchant such as Amazon or Walmart, the merchant knows who is browsing and personalizes rankings, recommendations and its own sponsored listings with its data on that customer, and it does so in every arm, since neither the ad blocker nor the anti-tracking extension removes the merchant's own personalization. Among logged-in participants, therefore, the three arms share the same organic personalization and differ in the advertising they see, which allows us to separate the personalization delivered through advertising from the personalization that the merchant applies to every customer. The experiment is ongoing, and the data currently include 31,049 shopping intents from 699 participants.

We find that participants in the Opt-In arm spend 22 percent more time per product evaluated and 25 percent more per intent than participants in the Ad-Blocking arm. The number of products evaluated and of merchant sites visited does not differ across

conditions, and the Anti-tracking arm does not differ significantly from Ad-Blocking on these outcomes. One reading is that personalized content, tailored to the consumer's inferred tastes, takes longer to process even when it does not change what consumers buy. Prior studies generally cannot observe this margin because they focus either on a single platform or on searches that end in a purchase. The IBE instead records every shopping intent across merchants, whether or not it ends in a purchase. The longer dwell time we observe in the Opt-In arm is surprising, but it could reflect deeper product evaluation rather than higher search cost. Following the duration-augmented search model of Ursu et al. (2020), we estimate how quickly additional time spent on a product sharpens what the consumer knows about it. The time spent examining a product is informative in every arm, but under targeting the estimated learning rate is about a third lower than with no ads, and it is lower by a similar amount under untargeted advertising, so that the two advertising arms do not differ. Combining the two margins, the time needed to acquire the same information about a product is almost twice as high under targeting as with no ads, and 1.5 times as high under untargeted advertising. The additional time does not lead to a better purchase, since relative to the alternatives the participant opened, the product bought under targeted advertising is neither cheaper nor better rated, and the price, rating and review count of the chosen product are indistinguishable across arms.

Relative to the Anti-tracking arm, which shows the same amount of ads, an Opt-In intent lasts 19 percent longer, half a minute more for the median intent, and shows 15 ads on the merchant pages it visits. The additional time therefore amounts to 1.4 seconds per personalized ad seen, which at the average hourly wage of U.S. private employees (Bureau of Labor Statistics, 2026) is worth 1.5 cents, and $17 to $32 per participant over the three months, an upper bound that treats every extra minute as pure loss. In number, participants in the Opt-In and in the Anti-tracking arm see similar ads per merchant page. On the price margin, existing estimates of the value added to ad impressions by tracking range from 4 percent (Marotta et al., 2019) to 30 percent (Alcobendas et al., 2025; Aridor et al., 2025). At a conservative price of $8 per thousand impressions,[2] the advertiser pays 0.8 cents for the impression the consumer sees, half of what the additional time costs the consumer, and between 0.03 and 0.24 cents for its personalization.

[2]The 2025 industry estimates of the price per thousand impressions range from $5 to $6 for digital display and from $2 to $8 for social media (Solomon Partners, 2025); $8 is the top of these ranges.

Our findings speak to three groups of studies. First, in the economics of advertising and information, we provide a direct measurement of search that separates personalized from contextual ads. Second, in the empirical study of consumer search, we include the many unsuccessful attempts that purchase-based analyses leave out. Third, in the economics of online targeting and privacy, we provide direct behavioral evidence on how personalization changes the search process. We discuss in detail these streams of work in Section 1.1. The rest of the paper is organized as follows. Section 2 describes the experiment, the reconstruction of shopping intents by a language model, the data, and the empirical approach. Section 3 reports the effect of the experimental conditions on search effort and its robustness, decomposes the additional time across pages and products, examines the candidate channels, and turns to purchase outcomes. Section 4 estimates the informativeness of search time by arm and discusses welfare and limitations, with the data construction documented in the appendix, and Section 5 concludes.

1.1. Related literature

First, a large literature asks whether advertising informs consumers or diverts their attention (Bagwell, 2007). In the informative view, advertising conveys information that lowers search costs and improves the match between buyers and products (Stigler, 1961; Nelson, 1974), and higher-ranked options are searched and bought more often (Ursu, 2018), which gives prominent positions a value that platforms sell in auctions (Edelman et al., 2007; Athey & Ellison, 2011). More recent evidence from online marketplaces looks at what these positions contain: on Amazon, about a third of the three most clicked products in a search are sponsored, the position of a product predicts its clicks even when it is more expensive or worse rated, and advertised top products are about 17 percent more expensive (Rock et al., 2024), while in simulated searches products ranked beyond the hundredth organic position appear as sponsored results above the top organic ones (Dash et al., 2024). When a marketplace randomizes its sponsored listings, their effects vary by category, and they raise the platform's revenue without hurting consumers' overall response in the short run (Abhishek et al., 2025). In a laboratory study, participants who search with ads on the page take considerably longer to complete the same tasks and retrieve a similar number of relevant documents (Foulds et al., 2021), and a field experiment estimates the cost to users of annoying display ads (Goldstein et al., 2014). These studies measure what

the platform shows, or how people search in a controlled setting, while we observe the complete search of the same consumers across merchants over three months, so we can measure how much time advertising adds to each product evaluated, 22 percent under targeted advertising, and show that the additional time does not lead to more products evaluated or to better purchases. Our results are also close to the literature on obfuscation, in which firms raise the cost of acquiring or comparing information in ways that soften competition (Gabaix & Laibson, 2006; Ellison & Ellison, 2009), and to the trade-off online intermediaries face between helping consumers complete purchases efficiently (Bakos, 1997; Dinerstein et al., 2018; Ursu, 2018) and extracting value from the time consumers spend on the platform (Anderson & Coate, 2005; Allcott et al., 2022; Ichihashi & Kim, 2023).

Second, the empirical study of consumer search, building on Weitzman (1979), increasingly uses clickstream data to recover how consumers browse before buying (De los Santos et al., 2012; Ursu, 2018; Dinerstein et al., 2018; Hodgson & Lewis, 2025), and recent methods recover preferences and the value of information when consumers are imperfectly informed (Abaluck et al., 2026). Ursu et al. (2020) extend the search model to the time spent on each option, which we use to estimate how quickly time sharpens the consumer's information about a product. By design or by data, much of this work centers on the searches that end in a purchase (Hong & Shum, 2006), and the effect of advertising is easy to overstate when it is inferred from data selected on consumer behavior: in large eBay field experiments, the returns to paid search are a small fraction of observational estimates (Blake et al., 2015), and advertising returns are in general very hard to measure (Lewis & Rao, 2015). With URL-level clickstream data, ad-initiated arrivals are more than half of the arrivals on apparel websites, occur early in the search, and lead to less in-depth search and fewer transactions (Ursu et al., 2025). We keep every shopping intent, including the 85 percent that end without a purchase, and we build these intents with a language model, as recent work does to generate and validate intent taxonomies from search logs (Shah et al., 2025).

Third, a literature studies the value of tracking and targeting. Opt-In impressions sell at a premium, and restricting tracking degrades ad performance: EU privacy regulation cut the effectiveness of display advertising by roughly two-thirds (Goldfarb & Tucker,

2011a), ad targeting and obtrusiveness each raise purchase intent on their own but not in combination (Goldfarb & Tucker, 2011b), third-party cookies raise winning bids in banner auctions by up to eighty percent (Alcobendas et al., 2025), and e-commerce advertisers lose performance and revenue when tracking is withdrawn (Aridor et al., 2025), although publishers earn only about four percent more per impression when a user's cookie is available (Marotta et al., 2019). Newer work on the GDPR traces its effects on advertising outcomes and market structure (Goldberg et al., 2024; Johnson et al., 2023), and the welfare effects of privacy protection depend on the context (Acquisti et al., 2016), including when sellers set prices on tracked purchase histories (Acquisti & Varian, 2005). On the consumer side, products shown in targeted ads are more relevant than random products but are more often sold by lower-quality vendors and at higher prices than competing products found through search (Mustri et al., 2023), and personalized rankings of organic results can raise consumer welfare (Donnelly et al., 2024). After the GDPR, which limits the data that every firm can collect, consumers in the EU submit more search terms for the same topic and spend 11 percent more time browsing products than consumers outside the EU (Zhao et al., 2023), and when a marketplace switches off the personalization of its recommendations, customers search more (Yuan et al., 2024). Evidence on ad blockers points in different directions: in panel data, adopting an ad blocker lowers online spending by 1.45 percent and reduces search activity (Todri, 2022), new users of an ad blocker spend 28 percent more active time in the browser (Miroglio et al., 2018), and in a laboratory experiment participants search and buy with and without an ad blocker (Frik et al., 2020). The largest experiments on the consumer side change the ads on a single platform: in a nine-year holdout in which a random half percent of Facebook users never see ads, users value the platform as much as the others (Brynjolfsson et al., 2025), and removing targeting on Facebook sharply reduces ad clicks and engagement on the platform (Beknazar-Yuzbashev et al., 2025). Closer to our design, Farronato et al. (2026) use a browser extension to switch the ad personalization settings of participants on Amazon, Google and Meta for two weeks and elicit what they would accept to be switched: most participants prefer personalized advertising, 27 percent of their platform settings do not match their stated preference, and two weeks under the non-preferred setting move the valuation of personalization by about two dollars, toward personalization on average. In theory, a consumer who receives a personalized recommendation searches less beyond it, and firms respond by raising prices (Rhodes & Zhou, 2025). This literature evaluates the impact of targeting through

firms' outcomes (cost per thousand impressions, ad effectiveness, revenue, concentration), through prices, or through stated preferences and valuations. In this work, we study search costs by isolating personalization from advertising, with a randomized experiment at the browser level, in which advertising is removed or de-personalized on every website the consumer visits, and we find that personalization delivered through advertising lengthens search by 16 percent relative to untargeted advertising, while the products consumers buy remain the same.

2. Data and Empirical Approach

2.1. The experiment

The Internet Behavior Experiment recruits U.S. Chrome users who install a browser extension, which records the pages participants visit, the rendered HTML of web pages, the products they see and open, and their purchases (Cheyre et al., 2024). Participants were randomly assigned to one of three experimental arms for three months. In the Ad-Blocking arm, the software installs an ad blocker that removes display advertising on every website and the sponsored product tiles inside the search results of Amazon and Walmart, the two merchants where most of the shopping in the sample takes place. In the Anti-tracking arm, a browser extension blocks third-party trackers and registers opt-outs, so that participants still see advertising, but advertising that is less likely to be personalized. In the Opt-In arm, the software accepts tracking on the consent banners that participants encounter, and they browse with the usual personalized advertising. The analysis sample includes the participants who completed the onboarding of the study and have at least 14 days with browsing activity. Random assignment makes the arms comparable in expectation, and the estimation sample is balanced on the seven pre-treatment covariates we use as controls (Appendix Table A4; the joint F-tests of the covariates predicting assignment give $p = 0.42$ for Opt-In and $p = 0.42$ for Anti-tracking against Ad-Blocking). On eBay, Etsy, Temu, Target, AliExpress and Home Depot, instead, promoted listings are served by the merchant itself together with the organic results and pass through the blocker, so participants in every arm see them. For instance, Ad-Blocking participants land on Etsy listings marked as ads at a rate close to the ad arms (13.3 percent of their Etsy listing visits, against 12.4 percent for Opt-In and 5.6 percent for Anti-tracking). For this reason the estimation sample excludes the intents whose main merchant is one of these six, and

the assignment does not change the share of a participant's intents that fall on them (12.1, 13.4 and 14.4 percent in the Ad-Blocking, Opt-In and Anti-tracking arms, not significantly different). Appendix E uses these merchants to compare the same Ad-Blocking participant with and without sponsored listings.

2.2. Shopping intents and search time

The extension records the page-level clickstream: each page view with a timestamp, URL, host, and the underlying HTML of the page. Each page view also records whether the page was active, and the results are unchanged when time is built from active pages only (Appendix D). The unit of analysis is the shopping *intent*: a shopping need pursued through an episode of activity directed at one goal, for example "running shoes". Every product interaction the extension records, a product listed on a results page, a product page opened, a purchase, is assigned to an intent by a large language model. The model reads each participant's chronological event stream, with product names, timestamps, and the page each product appeared on, groups the events that pursue the same goal, and keeps episodes open across days. Viewing several running shoes across three days is one intent; an unrelated search for grocery products the same week is a different one. An intent need not involve a search engine or a query: any product interaction opens one. A consumer who clicks an ad and buys directly, without typing a query, generates an intent containing that purchase. Episodes with no results-page activity are retained throughout, and their "search time" is the time spent on the product pages involved. These episodes are not rare: 27 percent of the intents in the estimation sample contain no results-page activity. The model does not see the participant's experimental arm, so any classification error is orthogonal to treatment. An intent can span several sessions and days, since shopping often takes more than one sitting, and we do not break it into single-sitting sessions.

For each intent we measure the number of distinct products evaluated, the total search time, and the time per product evaluated. A "product evaluated" is a distinct product-detail page, identified from URL patterns across the major merchants (Amazon, Walmart, eBay, Target, Etsy, and others). Search time follows the scan-until-click rule: the time spent on a results page is divided among the products listed before the one the participant opened, and the time spent on opened product pages is added in full. The dwell time on a page is

the difference between the timestamp of that navigation and the next one. This attributes to each product the time spent evaluating it, and it keeps interruptions that belong to other intents apart: if a participant clicks on an ad while searching for a product, the visit is counted as a different intent and the dwell times are computed separately. Product attributes (price, rating, review count) come from the rendered HTML of each page. Prices are taken only from each participant's own sessions, since prices are dynamic and personalized and sharing them across participants would contaminate the measure.

We keep the intents in which the participant opens at least one product page, or, if they open none, sees the products of the intent on at least two distinct result pages. About half of the intents contain no opened product, mostly in grocery shopping, where products are often added to the cart directly from the result list. Of the 43,501 intents pursued by our participants on logged-in participant-merchant pairs, 22,746 contain an opened product and 8,303 contain products seen on at least two result pages, and the share of intents we keep is similar across arms (72.6, 69.9 and 71.5 percent in the Ad-Blocking, Opt-In and Anti-tracking arms; the difference is not significant). The estimation sample keeps the 31,049 of these intents whose main merchant is one where the ad blocker removes sponsored listings (699 participants; Table 1); the intents on eBay, Etsy, Temu, Target, AliExpress and Home Depot, where promoted listings pass through the blocker, serve as a placebo sample (Table A1). We observe the search process at scale, and the large majority of it never ends in a transaction, so the analysis is not confined to the comparatively few episodes that end in a sale. Because the intent is itself constructed from behavior, we check how its formation responds to treatment. Advertising does open intents: 10.4 percent of intents in the Opt-In arm and 9.4 percent in the Anti-tracking arm begin within five minutes of a click on an ad, against 1.7 percent in the Ad-Blocking arm. The total number of intents, however, is no different from Ad-Blocking: the median participant pursues 20, 24 and 21 intents in the Ad-Blocking, Opt-In and Anti-tracking arms, and the arm effect on a participant's (log) intent count is not significantly different from zero ($+0.013$, $p = 0.92$ for Opt-In; -0.063, $p = 0.62$ for Anti-tracking, with the demographic controls), so ad-initiated intents displace others rather than add to them. Initiation is also nearly identical across the two ad arms, so it cannot explain the time effect specific to the Opt-In arm. The raw differences in mean intents per participant in Table 1 are driven by a few participants with very many intents in the tails of the distribution; Section 3.1 returns to this with unconditional participant-level

outcomes.

TABLE 1. Intents and purchases by arm, LLM-assigned intents with measurable search activity. The unit of analysis is the intent; most never end in an observed purchase. Products opened are distinct product-detail pages visited within the intent.

Arm	Participants	Intents	Per participant	Products opened per intent	Intents with purchase
Ad-Blocking	241	11,340	47.0	2.87	1,562
Opt-In	237	10,110	42.7	3.12	1,678
Anti-tracking	221	9,599	43.4	2.85	1,493

Figure 1 shows what people shop for. Grocery, home and kitchen, and apparel account for more than half of all intents (55 percent). The composition is the same across arms, both for all intents and for the ones that end in a purchase, so the comparisons across arms below are not driven by differences in what participants shop for, even before category fixed effects. Purchases within those intents are similar across arms: 14 to 17 percent of intents end in a linked purchase, and 19 to 21 percent of participants never complete an observed purchase, with no significant difference across arms on either margin. Median purchase prices are close, between $9 and $11. Nothing in the raw data suggests that the Opt-In arm buys more, or more expensive, items; Section 3.4 tests this formally.

2.3. Login status and empirical approach

We restrict the analysis to the shopping that participants do while logged into the merchant, because this allows us to separate the personalization delivered through advertising from the personalization that the merchant applies to every customer. When a participant is logged into a merchant such as Amazon or Walmart, the merchant knows who is browsing and personalizes rankings, recommendations and its own sponsored listings with its data on that customer, and it does so in every arm, since neither the ad blocker nor the anti-tracking extension removes the merchant's own personalization. Among logged-in participants, therefore, the three arms share the same organic personalization and differ in the advertising they see: none in the Ad-Blocking arm, ads selected with less information about the person in the Anti-tracking arm, and ads selected with the full set of information in the Opt-In arm.

We measure login status from the rendered HTML of every merchant page. For each page we read the header of the page and classify it as logged in when the header shows the

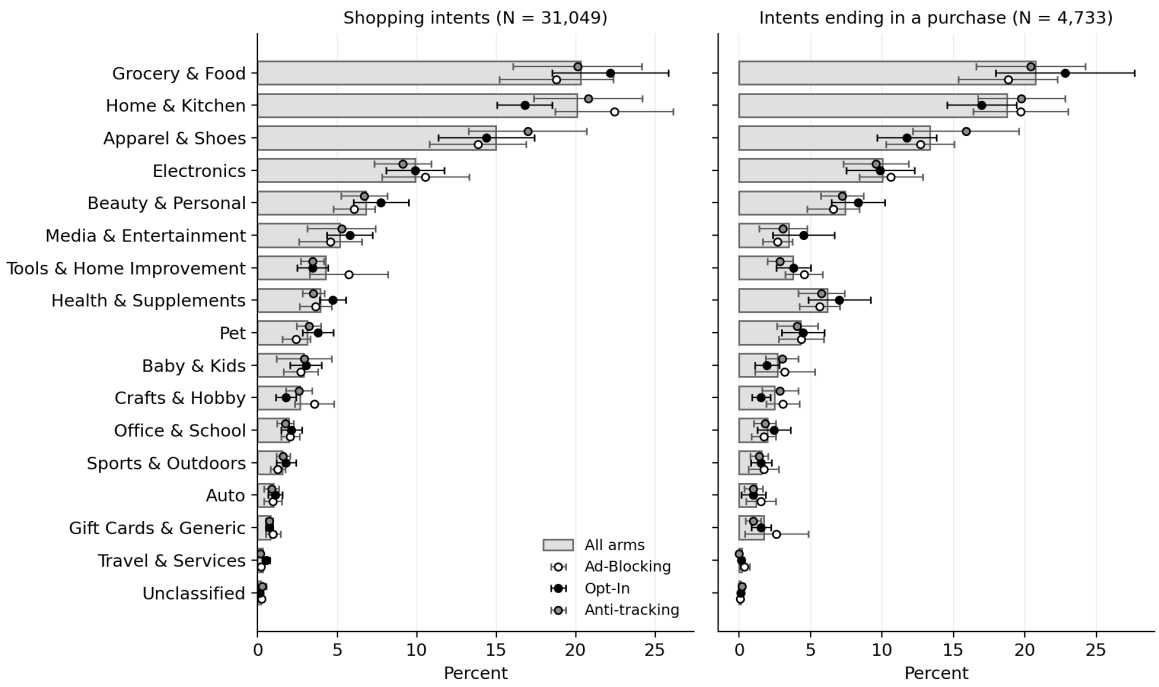


FIGURE 1. Composition of shopping intents by product category. Bars are all arms pooled; markers are the shares within each arm, with 95 percent confidence intervals from standard errors clustered by participant. Left: all intents with measurable search activity. Right: intents ending in a linked purchase. An intent's category is its goal category, the category of the fixed effects in equation (1); every category is shown.

account greeting with the participant's name (for example, "Hello, [name]" on Amazon) or a sign-out link, and as logged out when it shows a sign-in prompt without a greeting (for example, "Hello, sign in" on Amazon). For the merchants whose header differs we use merchant-specific rules: for instance, on Home Depot the account button reads "Log In" when the participant is logged out and their first name otherwise, and on Instacart the store pages cannot be browsed without an account. Since a participant who logs in usually stays logged in for the rest of a browsing session, every page of a merchant in a session counts as logged in when at least one page of that merchant in the session is. Of the 327,806 merchant pages with a rendered capture, 92 percent are logged in, 5 percent logged out, and 3 percent cannot be classified. We then classify a participant as logged into a merchant when more than half of their pages at that merchant with a status are logged in (at least three pages). Intents on logged-in participant-merchant pairs are 97 percent of the intents in every arm (97.5, 97.1 and 97.1 percent in the Ad-Blocking, Opt-In and Anti-tracking arms). Since participants choose whether to log in after assignment, we also report the results on

all intents (Table A1, columns 6 and 7), which are very similar. One possible objection to this choice is that in a logged-in, first-party session the anti-tracking extension has nothing left to remove. However, assignment to that experimental condition has several other channels through which it can influence user behavior. The sponsored placements a retailer shows an identified user are selected in auctions whose bids and audiences draw on cross-site identity and third-party behavioral segments: advertisers retarget users seen elsewhere, match them through identity graphs, and bid on segments built from off-site browsing. The extension curtails exactly these inputs under the Anti-tracking condition, by blocking third-party trackers and registering opt-outs, so even placements served first-party are selected with less information about the person under Anti-tracking than under Opt-In. Login, in turn, allows the retailer to act on that information on its own pages.

We compare the arms on three margins of search: how many products participants evaluate in an intent, how long they spend on each product, and what they buy. In principle, if advertising informs consumers, targeted ads should point them to good options sooner, so that they evaluate fewer products, or buy better ones, at no extra time per product. For instance, a participant looking for running shoes who is shown a sponsored listing for the model they would have chosen anyway could stop searching earlier. As a counterexample, if advertising is a friction, the sponsored and personalized content competes for the attention that each evaluation needs, and participants spend more time on each product without evaluating more products or buying better ones. A friction can take two forms: sponsored listings that participants inspect and discard, which add time on the result lists and extra inspections, or content that makes each inspection take longer, which adds time on the product pages. Table 2 summarizes the predictions of the two accounts and where we test them. The Anti-tracking arm separates advertising from its personalization: both ad arms see advertising, and the anti-tracking extension limits the information with which ads are selected, so an effect specific to the Opt-In arm points to personalization.

TABLE 2. What the two accounts of advertising predict, and where we test them.

Margin	Information	Friction	Tested in
Products evaluated per intent	Fewer or unchanged	Unchanged	Table 4
Time per product evaluated	Unchanged	Higher	Tables 4 and 5
Price and quality of the purchase	Better	Unchanged	Table 7
Information per minute of search	Unchanged or higher	Lower	Section 4

For intent i we estimate the following model:

$$y_i = \beta_1 \, \mathbf{1}\{\text{Opt-In}_i\} + \beta_2 \, \mathbf{1}\{\text{Anti-tracking}_i\} + X_i'\gamma + \delta_{c(i)} + \varepsilon_i, \tag{1}$$

where the omitted category is Ad-Blocking, X_i collects seven pre-treatment demographic controls (age, female, non-white, college, log income, employed, IT skill), $\delta_{c(i)}$ are fixed effects for the intent's goal category, and standard errors are clustered by participant. The goal category is assigned by the language model from the episode label (sixteen categories; where the label is too generic to classify, 41 percent of intents, we use the modal product category of the episode's members), so the comparison is between intents pursuing goals in the same category. Appendix Table A6 shows the estimate under four alternative constructions of the category control, including a k-means clustering of the raw labels; the coefficient moves between $+0.18$ and $+0.20$. Because treatment is randomized, β_1 and β_2 have a causal interpretation.[3] The primary outcome is log time per product evaluated. The estimation sample is the $N = 31{,}049$ intents with an opened product or with products seen on at least two result pages, on merchants where the blocker removes sponsored listings, pursued while logged into the merchant (Section 2.3), from the 699 participants with complete demographics.

2.4. Advertising exposure by arm

We now look at how the three arms change the advertising that participants see and click on. We count the ads on merchant pages from the browser's record of the ad elements rendered on each page, one per ad element, and classify them as sponsored results in the merchant's result lists, sponsored products in the carousels of the merchant's pages (mostly on product pages), and display banners. Participants also click on sponsored results in two places: on the results pages of a search engine such as Google or Bing, which we call sponsored results (search), and on the merchant's own pages, which we call sponsored results (merchant). The browser does not record the ads that the blocker removes, so for participants in the Ad-Blocking arm we observe the ads they click on but not the ads they are shown, and we compare the volume of ads between the two advertising arms only.

[3] The controls improve precision; they are not needed to identify the treatment effects.

TABLE 3. Advertising exposure by arm. Panel A: ads seen per 100 pages on the shopping merchants of the estimation sample, counting one ad per rendered ad element, in the Opt-In arm relative to the Anti-tracking arm; OLS at the participant × merchant level with merchant fixed effects, weighted by pages. The browser does not record the ads removed by the blocker, so the Ad-Blocking arm is not in Panel A. Panel B: clicks on sponsored results per shopping intent, relative to the Ad-Blocking arm, with demographic controls and goal-category fixed effects; under Ad-Blocking these are the sponsored results the blocker does not remove. Standard errors clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel A. Ads seen per 100 pages on shopping merchants

	All ads	Sponsored results (merchant)	Sponsored carousels (merchant)	Display
Anti-tracking, mean	128.7	51.7	59.0	18.0
Opt-In − Anti-tracking	$+3.6$	-4.9	$+9.6$	-1.1
	(12.3)	(5.6)	(7.2)	(1.6)
Observations (participant × merchant)		1,667		

Panel B. Clicks on sponsored results per shopping intent

	Sponsored results (search)	Sponsored results (merchant)
Ad-Blocking, mean	0.033	0.016
Opt-In − Ad-Blocking	$+0.180^{***}$	$+0.182^{***}$
	(0.022)	(0.021)
Anti-tracking − Ad-Blocking	$+0.095^{***}$	$+0.131^{***}$
	(0.014)	(0.016)
Opt-In − Anti-tracking	$+0.086^{***}$	$+0.051^{**}$
	(0.026)	(0.025)
Observations (intents)	31,049	

We find that participants in the Anti-tracking arm see about 130 ads per 100 merchant pages, and participants in the Opt-In arm a similar number ($+3.6$, $p = 0.77$), with no significant difference for the sponsored results in the result lists, for the sponsored carousels or for the display banners (Table 3, Panel A). Participants in the Ad-Blocking arm click on 0.03 sponsored results (search) and 0.02 sponsored results (merchant) per intent, against 0.21 and 0.20 in the Opt-In arm and 0.13 and 0.15 in the Anti-tracking arm (Panel B), which means that the ad blocker removes about 84 percent of the clicks on sponsored results (search) and 92 percent of the clicks on sponsored results (merchant) of the Opt-In arm. In addition, participants in the Opt-In arm click on more sponsored results (search) than participants in the Anti-tracking arm ($+0.086$ per intent, $p = 0.001$) and on more sponsored results (merchant) ($+0.051$, $p = 0.05$). Taken together, these findings suggest that the two advertising arms see a similar volume of ads in view and differ in how the ads they see are selected.

3. Results

We now study whether the experimental arms change the way participants search and what they buy.

3.1. Search effort

We look at whether participants who see targeted advertising search differently from those who see no advertising. In principle, if targeted ads point participants to products they value, they could evaluate fewer products or spend less time on each of them. We find that, instead, participants in the Opt-In arm spend 22 percent more time per product evaluated than participants in the Ad-Blocking arm ($+0.195$ log points, $p = 0.001$) and 25 percent more time per intent ($+0.226$, $p < 0.001$), while they evaluate as many products and visit as many merchants (Table 4, columns 1 to 4). Participants in the Anti-tracking arm, who see advertising that is less likely to be personalized, spend a time that is not significantly different from that of the Ad-Blocking arm ($+0.048$, $p = 0.49$), and the difference between the two advertising arms is significant both for time per product ($+0.147$, $p = 0.02$) and for time per intent ($+0.170$, $p = 0.01$). Since both advertising arms see advertising and differ in how it is selected, this difference suggests that the additional time comes with the personalization of advertising. The probability that an intent ends in a purchase is not significantly different across arms (column 5).

These results are robust to the construction of the goal-category control, to controls for the experience that participants accumulate during the study, to the inclusion of the intents pursued while logged out, of the participants who left the study and of every merchant, and to removing the time that pages take to load (Appendix Table A1). On the merchants where the ad blocker leaves sponsored listings, which the Ad-Blocking participants also see, the difference between the Opt-In and the Ad-Blocking arm is close to zero and not significant ($+0.032$, $p = 0.75$). The additional time is also present across the whole distribution of search time, from the shortest to the longest intents (Appendix Figure A1).

3.2. Where the additional time goes

We now ask in which part of the search participants spend the additional time. As discussed in Section 2.3, sponsored listings that participants inspect and discard would add time on the result lists, whereas content that makes each inspection take longer would

TABLE 4. Main results. OLS of each outcome on indicators for the Opt-In and Anti-tracking arms, relative to Ad-Blocking; outcomes (1) to (4) in logs, (5) a linear probability. The Ad-Blocking mean is in levels: minutes per product and per intent, number of products and merchants, share of intents with a purchase. Standard errors clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1) Time per product	(2) Time per intent	(3) Products evaluated	(4) Merchants visited	(5) P(purchase)
Opt-In	$+0.195^{***}$	$+0.226^{***}$	$+0.031$	$+0.003$	$+0.023$
	(0.061)	(0.062)	(0.027)	(0.012)	(0.016)
Anti-tracking	$+0.048$	$+0.056$	$+0.008$	$+0.009$	$+0.020$
	(0.069)	(0.073)	(0.028)	(0.014)	(0.015)
p-value, Opt-In = Anti-tracking	0.02	0.01	0.44	0.67	0.84
Ad-Blocking mean (levels)	2.94	6.25	2.87	1.21	0.14
Observations (intents)	31,049	31,049	31,049	31,049	31,049
Participants	699	699	699	699	699
Demographic controls	Yes	Yes	Yes	Yes	Yes
Goal-category FE	Yes	Yes	Yes	Yes	Yes

add time on the product pages. We find that most of the additional time is spent on the product pages. An intent in the Opt-In arm lasts about 1.4 minutes more than an intent in the Ad-Blocking arm, of which about 1.2 minutes are spent on product pages ($+0.32$ log points, $p = 0.007$) and 0.2 minutes on result lists, a difference that is not statistically significant ($p = 0.73$; Table 5, columns 1 and 2).

On each product page they open, participants in the Opt-In arm spend about 24 percent more time than Ad-Blocking participants ($+0.218$, $p = 0.01$), and 24 percent more when we compare products with a similar price, rating and number of reviews ($+0.218$, $p = 0.002$; columns 3 and 4). Participants in the Anti-tracking arm also spend more time on each product page ($+0.093$), although the difference from Ad-Blocking is not significant. In other words, a product with given attributes takes more time to evaluate when the advertising that participants see is personalized. Finally, the additional time is spent on the products that participants then discard: in the intents that end in a purchase, Opt-In participants spend 0.23 log points more time than Ad-Blocking participants on the products they do not buy ($p = 0.001$) and a similar time on the product they buy ($+0.13$, $p = 0.18$; columns 5 and 6). In addition, we distinguish the products that participants reach from the organic results, from the recommendation carousels and from the sponsored results (merchant). Participants

TABLE 5. Where the additional time goes. Columns (1) and (2): OLS of log(1+seconds) spent on result lists and on product pages in the intent. Columns (3) to (6): OLS of log dwell on each opened product page; column (4) restricts to products with a parsed price, rating and review count and controls for them; columns (5) and (6) use the intents that end in a purchase, for the product bought and for the products opened and not bought. Raw median dwell per product page: Ad-Blocking 22 seconds, Opt-In 29, Anti-tracking 26. Standard errors clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Log seconds per intent		Log dwell per product page			
	(1)	(2)	(3)	(4)	(5)	(6)
	Result lists	Product pages	Product page	Product page	Bought product	Products not bought
Opt-In	$+0.051$	$+0.315^{***}$	$+0.218^{**}$	$+0.218^{***}$	$+0.133$	$+0.233^{***}$
	(0.146)	(0.117)	(0.087)	(0.069)	(0.100)	(0.072)
Anti-tracking	$+0.000$	$+0.118$	$+0.093$	$+0.073$	-0.078	$+0.038$
	(0.153)	(0.127)	(0.086)	(0.068)	(0.104)	(0.076)
p-value, Opt-In = Anti-tracking	0.70	0.13	0.19	0.02	0.03	0.01
Observations	31,049	31,049	31,831	25,029	2,599	10,323
Participants	699	699	641	625	453	448
Unit	Intent	Intent	Product	Product	Product	Product
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Goal-category FE	Yes	Yes	No	No	No	No
Price, rating, reviews	No	No	No	Yes	No	No
Sample	All	All	All	Attributes	Purchase	Purchase

in the Opt-In arm spend 0.24 log points more time than Ad-Blocking participants on the products they reach from the organic results ($p = 0.002$), 0.17 more on the products they reach from the recommendation carousels ($p = 0.05$), and a similar time on the products they reach from the sponsored results (merchant) ($+0.07$, $p = 0.65$). In other words, the additional time is spent on the evaluation of the products that participants find in the organic results and in the recommendations.

3.3. What explains the additional time

We now ask which feature of the advertising environment explains the additional time. Since the analysis is restricted to logged-in participants, both advertising arms receive the merchant's own personalization of organic content, and the comparison between them isolates the advertising and its personalization. We consider three explanations: participants in the Opt-In arm might see more ads or more products, might be pulled away from the merchant more often, or might open different products.

TABLE 6. Advertising volume and the time gap between the two advertising arms. OLS of log time per product on an indicator for the Opt-In arm, relative to the Anti-tracking arm, on the intents with a merchant page in the browser's navigation record. Column (2) adds the number of ads seen per page during the intent, as log(1 + ads), separately for the sponsored results in the result lists, the sponsored carousels on the other merchant pages and the display banners (the ad types of Table 3); an ad is counted once per rendered ad element when it comes into view. Standard errors clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Log time per product	
	(1)	(2)
Opt-In − Anti-tracking	$+0.143^{**}$	$+0.129^{**}$
	(0.060)	(0.059)
Ads seen per page during the intent, log(1 + ads):		
Sponsored results (merchant)		$+0.245^{***}$
		(0.032)
Sponsored carousels (merchant)		$+0.205^{***}$
		(0.030)
Display banners		$+0.255^{***}$
		(0.080)
Observations (intents)	19,203	19,203
Participants	456	456
Demographic controls	Yes	Yes
Goal-category fixed effects	Yes	Yes

First, the two advertising arms see a similar volume of advertising (Section 2.4), and participants in the Opt-In arm see as many products in an intent as participants in the Ad-Blocking arm (53.6 against 52.3 on average, $p = 0.92$); controlling for the number of products seen leaves the difference in time per product at $+0.185$ ($p = 0.001$). Moreover, holding fixed the ads per page seen during the intent, the difference in time per product between the two arms changes little, from $+0.143$ ($p = 0.02$) to $+0.129$ ($p = 0.03$; Table 6). Since the browser records an ad when it comes into view, a participant who stays longer on a page records more ads, so we repeat the exercise counting every ad element present on the page, whether or not it came into view: the pages seen by Anti-tracking participants contain 13 percent fewer ads than those seen by Opt-In participants ($p = 0.03$), mostly display banners and carousels, and the gap in time per product holding these ads fixed remains at $+0.123$ ($p = 0.03$). The gap is also similar when we account for the presence of a display banner on the merchant pages of the intent (Appendix Table A3).

Second, participants in the Opt-In arm are not pulled away from the merchant more often. The time measure already excludes the time spent away from the merchant,

since the time on a page ends at the next navigation. Counting any departure from a merchant page to a non-shopping page, the number of exits in the intents with at least one exit is not significantly higher in the Opt-In arm ($+0.15$ log points, $p = 0.21$), and the time effect is present in the intents without any exit ($+0.260$ against Ad-Blocking, $p = 0.002$).

Finally, participants in the Opt-In arm open the same kind of products. We compare the products participants click with the products they only see on the same result pages. With no ads, participants click products about 8 percent cheaper than the ones they see and with far more reviews, and in the Opt-In arm the differences between clicked and seen products in price, rating and number of reviews are not significantly different ($p = 0.64$, 0.16 and 0.23). The only difference we find is in the Anti-tracking arm, where participants click products 4.7 percent more expensive relative to the ones they see ($p = 0.10$). Taken together, these findings suggest that the additional time comes with the way ads are selected for each participant, since participants in the Opt-In arm see as many ads, leave the merchant as often, and open the same kind of products as participants in the Anti-tracking arm.

3.4. Purchase outcomes

We now ask whether the additional search effort leads to better purchases. We measure purchases by linking each intent to confirmed transactions parsed from cart and checkout pages, matched to the intent first by the purchased product's own page, then by a same-merchant page in the intent's session, and finally by product-text similarity. About 15 percent of intents end in a purchase we can link, and 19 to 21 percent of participants never complete an observed purchase, with no significant difference across arms. Conditional on an intent, the probability that it ends in a purchase is not significantly higher in the Opt-In arm ($+2.3$ percentage points, $p = 0.15$; Table 4), and the two advertising arms do not differ ($p = 0.84$).

Moreover, the products bought are similar across arms. We measure the quality of a purchase in six ways. The first two compare the price and the rating of the product bought with the mean price and rating of the other products the participant opened in the same intent. The third, which we call overpayment, is the difference between the price paid and the lowest price at which we observe the same product. The fourth and the fifth are the price of the product bought given its rating, number of reviews and category, and its

TABLE 7. Purchase outcomes. OLS of each outcome of the bought product on indicators for the Opt-In and Anti-tracking arms, relative to Ad-Blocking. Relative price and rating: log price and rating of the bought product minus the mean of the other products the participant opened in the intent. Overpayment: log price paid minus the log of the lowest price observed for the same product. Quality-adjusted price and price-adjusted rating: residuals of hedonic regressions of price on rating, reviews and category, and of rating on price, reviews and category. Value for money: standardized index of the two residuals. Standard errors clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1) Relative price	(2) Relative rating	(3) Over-payment	(4) Quality-adj. price	(5) Price-adj. rating	(6) Value for money
Opt-In	-0.195	$+0.015$	$+0.022$	$+0.033$	$+0.008$	-0.006
	(0.150)	(0.025)	(0.036)	(0.055)	(0.016)	(0.041)
Anti-tracking	-0.227	$+0.037$	$+0.003$	$+0.016$	-0.023	-0.035
	(0.168)	(0.038)	(0.033)	(0.055)	(0.016)	(0.040)
p-value, Opt-In = Anti-tracking	0.84	0.57	0.58	0.77	0.04	0.49
Observations (purchases)	5,485	4,961	6,501	9,131	7,367	9,131
Participants	498	490	546	552	546	552
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Goal-category FE	Yes	Yes	Yes	Yes	Yes	Yes

rating given its price, number of reviews and category, each measured as the residual of a regression on these characteristics. The sixth, which we call value for money, averages the last two after standardizing them, with the sign of the price residual reversed, so that a higher value means a lower price for a given quality and a higher rating for a given price. Compared with the other products the participant opened in the same intent, the product bought in the Opt-In arm is neither cheaper ($p = 0.19$) nor better rated ($p = 0.54$), and the overpayment, the price given quality, the rating given price and the value for money are not significantly different from the Ad-Blocking arm (Table 7). The product bought in the Anti-tracking arm is not significantly different from Ad-Blocking on these margins either. Therefore, the additional time in the Opt-In arm does not lead to a better or cheaper purchase. In addition, in the satisfaction surveys administered during the experiment, Opt-In participants do not report finding what they want more often ($+0.04$ points on the five-point scale, $p = 0.60$), and their satisfaction with the prices paid declines between the onboarding and the last survey relative to Ad-Blocking (-0.20 points, $p = 0.04$).

4. The Informational Return to Search Time and Welfare

4.1. The learning rate by arm

We now ask whether the additional time produces additional information. In principle, participants in the Opt-In arm could spend more time on each product because they evaluate it more carefully, and in that case the additional time would be a benefit rather than a cost. For instance, a participant who reads the reviews of a product more carefully might end up with a product that better fits their needs. We study this question with a sequential-search model in the tradition of Weitzman (1979). A consumer with a shopping need faces a set of products and learns each product's match value only by inspecting it; inspection costs time, and the search cost is proportional to it, $c = \chi \cdot$time, where χ is the cost per unit of inspection time. The consumer inspects products in order of expected value, stops when the option in hand beats the value of continuing, and buys the best inspected product or takes the outside option. Following Ursu et al. (2020), the consumer's signal about a product has noise $\sigma_\varepsilon(t) = \sigma_0\, t^{-\kappa}$, so that κ, which we call the *learning rate*, measures how quickly the time spent on a product sharpens what the consumer knows about it.

We estimate κ arm by arm over all the intents with at least two products with measured time, including those that end without a purchase. Each inspected product contributes a three-stage likelihood,

$$\mathcal{L} \;=\; \underbrace{\Phi(z_j/\sigma_z)}_{\text{click}} \;\times\; \underbrace{f(\log t_j \mid X)}_{\text{duration}} \;\times\; \underbrace{\frac{\exp\big(\mu_j/\sigma_\varepsilon(t_j)\big)}{\sum_k \exp\big(\mu_k/\sigma_\varepsilon(t_k)\big) + 1}}_{\text{purchase}}, \qquad \sigma_\varepsilon(t) = \sigma_0\, t^{-\kappa}, \tag{2}$$

where t_j is the time spent on product j; z_j is the index of the click stage, with scale σ_z, which captures which listed products are opened; the duration stage $f(\log t_j \mid X)$ is lognormal in the demographic controls and absorbs differences in the level of time; and μ_j is the product's value in the purchase stage. The learning rate enters the purchase stage only: within an intent, how strongly the relative time spent on each opened product predicts which one is bought identifies κ, and the intents without a purchase discipline the outside option. We therefore estimate that stage directly, as a conditional logit inside each intent: the participant chooses among the products they opened, or chooses not to buy, and the utility of an opened product increases with the log time spent on it, net of its mean

within the intent, so that differences in the level of time across people and categories drop out. We estimate the model jointly on the three arms, with a learning rate and a value of the outside option specific to each arm, and, in our preferred specification, with the price, rating and number of reviews of each product, whose coefficients are common to the three arms. The learning rate is not constrained to be positive. We test whether the learning rate differs between two arms with a Wald test whose standard errors are clustered by participant, since the same participant contributes many intents; a likelihood ratio test of the same restriction would treat the intents as independent.

We find that the time spent on a product is informative in every arm, since the learning rate is positive and precisely estimated in each of them (Table 8). However, the learning rate is about a third lower under targeted advertising than with no ads ($\hat{\kappa} = 0.24$ against 0.36, $p = 0.02$), and it is lower by a similar amount under untargeted advertising (0.26, $p = 0.05$), while the two advertising arms do not differ ($p = 0.76$). Without the product attributes the estimates are similar (0.24 against 0.35, $p = 0.03$). In 1,453 intents the product bought was seen on a result list but never opened, so it is not in the choice set and the estimation treats the intent as ending without a purchase; dropping these intents gives 0.24 against 0.36 ($p = 0.02$), and 0.24 under Anti-tracking. In other words, an additional minute spent on a product provides less information about it when participants see advertising, whether or not the advertising is personalized, which is consistent with a higher cost of attention rather than with a more careful evaluation.

The lower learning rate and the longer time on each product are closely related. As shown in Section 3.2, in the intents that end in a purchase Opt-In participants spend 0.23 log points more time than Ad-Blocking participants on the products they do not buy and a similar time on the product they buy, and a longer time on the products that are then discarded weakens the link between time and choice, which is what a lower learning rate measures. Combining the two, the time needed to reach a given precision about a product, which is proportional to the time on the product divided by the learning rate, is 1.84 times higher under Opt-In than under Ad-Blocking ($p < 0.01$) and 1.48 times higher under Anti-tracking ($p = 0.06$; Figure 2b). Since the two margins are related, we interpret this ratio as the time result expressed in the units of the model. We also estimated an aggregate stopping model, which recovers the search cost from the number of products evaluated

TABLE 8. Learning rate by arm. Conditional logit inside each intent among the opened products with measured time and the outside option of not buying, estimated jointly on the three arms with an arm-specific learning rate $\hat{\kappa}$ and outside-option value; column (2) adds the log price, rating and log number of reviews of each product, with coefficients common to the arms, on the intents in which every opened product has the three attributes. The choice is the purchase within the intent window. Standard errors clustered by participant in parentheses; p-values of Wald tests with the same standard errors.

	(1) No controls	(2) Price, rating, reviews
$\hat{\kappa}$, Ad-Blocking	0.350 (0.034)	0.361 (0.036)
$\hat{\kappa}$, Opt-In	0.242 (0.038)	0.238 (0.038)
$\hat{\kappa}$, Anti-tracking	0.269 (0.045)	0.256 (0.040)
Opt-In − Ad-Blocking	-0.108 ($p = 0.03$)	-0.123 ($p = 0.02$)
Anti-tracking − Ad-Blocking	-0.081 ($p = 0.15$)	-0.106 ($p = 0.05$)
Opt-In − Anti-tracking	-0.027 ($p = 0.65$)	-0.017 ($p = 0.76$)
Intents (with a purchase)	12,442 (2,040)	11,047 (1,898)
Participants	673	662

and from conversion alone, without using time; on this construction the differences between arms are not significant and imprecisely estimated, as one would expect for an aggregate cost that averages over the many inspections the friction does not touch.

We also ask whether a longer intent leads to a choice that is better on the attributes we observe. Among the intents in which the product bought is one of the opened ones, we build for each opened product an index of its rating, number of reviews and price relative to the other products opened in the same intent, and we let the weight of the index in the choice vary with the time of the intent, within and between participants. The weight does not increase with time in any arm, and the differences between arms are small and not significant. The bought product is the best of the opened ones on the index in 39 percent of the intents with no ads, 38 percent under Opt-In and 35 percent under Anti-tracking, differences that are not significant ($p = 0.34$ and $p = 0.21$); the products reached through an ad, which are 7 percent of the opened products under advertising, rank like the others on the index. Whether and how a purchase happens does not differ across arms either: 84 percent of the intents end without a purchase, 11 percent with a purchase from a result list without opening the product and 6 percent with a purchase among the opened products with no ads, against 81, 12 and 7 percent under Opt-In.

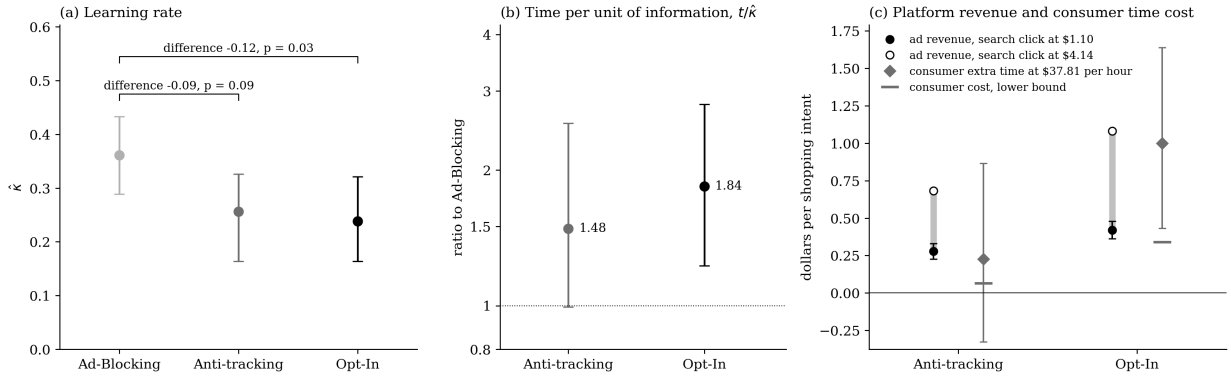


FIGURE 2. Learning rate, time per unit of information, and dollars per shopping intent by arm. (a) $\hat{\kappa}$ by experimental arm, from the conditional logit with price, rating and reviews (Table 8, column 2); brackets give the differences from Ad-Blocking with the participant-clustered Wald p-values of Table 8. (b) Time per unit of information, $t/\hat{\kappa}$, as a ratio to Ad-Blocking on a logarithmic scale; dwell per product (Table 4) and $\hat{\kappa}$ are bootstrapped on the same participant draws. (c) Platform advertising revenue inside the intent, from the clicks on sponsored results (merchant) and (search) assigned to the nearest intent, valued at benchmark prices (sponsored results (merchant) at the Skai cost per click by quarter, weighted by the Ad Badger category index; sponsored results (search) at $1.10, filled circles, or $4.14, hollow circles; the shaded segment joins the two), against the value of the consumer's extra search time relative to Ad-Blocking at $37.81 per hour (diamonds), with the lower bound scaled by $1 - \hat{\kappa}_a/\hat{\kappa}_{\text{Ad-Blocking}}$ (horizontal ticks). Whiskers are 95% confidence intervals from a participant-clustered bootstrap (a, b) or participant-clustered standard errors (c).

4.2. Costs and benefits

We now ask how large the cost of the additional time is for consumers, and how it compares with the revenue that personalized advertising brings to the platforms.

4.2.1. Consumers

The cost for consumers depends on the additional time per intent, on the number of intents, and on the value of time. Relative to the Anti-tracking arm, which shows the same ads without personalization, an intent in the Opt-In arm lasts 19 percent longer: half a minute more for the median intent and 1.2 minutes more on average. Over the three months of the study, the median Opt-In participant spends 28 additional minutes and the mean participant 51. We value time at the average hourly earnings of employees on U.S. private nonfarm payrolls, $37.81 in September 2026 (Bureau of Labor Statistics, 2026): the additional time is worth $17 for the median participant and $32 on average, about $6 to $11 per month. This is an upper bound, since it treats every additional minute as a loss. A lower bound counts as lost only the information that the additional time no longer

provides: scaling by one minus the ratio of the learning rates, $1 - \hat{\kappa}_{\text{Opt-In}}/\hat{\kappa}_{\text{Ad-Blocking}} \approx 0.34$, leaves \$6 to \$11 per participant. This linear scaling is a heuristic between the two bounds, and under the power form of equation (2) the time equivalent in information is larger, so the lower bound is conservative. Since the product chosen and the price paid do not change, this cost is not compensated by a better match, and it falls mostly on the searches that never end in a purchase. Expressed per ad, the median Opt-In intent shows 15 ads on the merchant pages it visits (52 on average), so the additional time amounts to 1.4 seconds, or 1.5 cents, per personalized ad seen.

4.2.2. Platforms

Whether the platforms gain from personalization enough to offset this cost depends on how much advertising revenue personalization adds. We do not observe the prices of the ads, so we compare the cost of the time with the price of the impression at a conservative benchmark of \$8 per thousand impressions (Solomon Partners, 2025), 0.8 cents per ad seen. Per ad, the additional time costs the consumer about twice what the advertiser pays for the whole impression; per month, the 730 ads that the mean Opt-In participant sees on merchant pages are worth about \$6 at that price, against \$11 of additional time. The premium that advertisers pay for the personalization of an impression is a fraction of its price, between 4 percent (Marotta et al., 2019) and 30 percent (Alcobendas et al., 2025; Aridor et al., 2025), that is, between 0.03 and 0.24 cents per ad seen. Therefore, the time cost of personalization for the consumer is at least six times the premium that advertisers pay for it. Valuing instead the clicks on sponsored results that we observe inside the shopping intents at public benchmark prices per click gives the same order of magnitude (Figure 2c): the revenue that personalization adds inside the intent is between 14 and 40 cents, against 34 cents to one dollar of additional time.[4] Outside the shopping intents, Opt-In participants also click more often on sponsored search results over the three months (25 more clicks per participant than under Anti-tracking), a revenue earned outside the searches that generate the cost, which we do not price.

[4]Time is valued at \$37.81 per hour, the average hourly earnings of employees on U.S. private nonfarm payrolls in September 2026 (Bureau of Labor Statistics, 2026).

5. Conclusion

Using a randomized field experiment with a reduced-tracking arm that isolates personalization, and a full record of search behavior that keeps the unsuccessful attempts, we ask how advertising changes the way consumers shop online. We find that personalization lengthens search, with 22% more time per product and 25% more per intent, without leading consumers to evaluate more products, find better ones, or report higher satisfaction; a product with given attributes takes longer to evaluate under personalized advertising, and the time needed to acquire the same information about a product is almost twice as high. Among logged-in participants, who receive the same organic personalization in every arm, the extra time is spent on products that are then discarded, while the number of ads per merchant page is similar under targeted and untargeted advertising. Valued at the average hourly wage, it amounts to $17 to $32 per participant per quarter, or 1.5 cents per personalized ad seen, about twice what the advertiser pays for the impression at a conservative price and at least six times the premium paid for its personalization. In other words, personalization changes the way people search while leaving what they buy unchanged, and on that margin it raises the cost of search. Because most search ends without a purchase, with about one consumer in five never buying, the friction is largely missed by analyses restricted to the searches that end in a sale.

These findings are relevant for the goals of online retailers, which earn revenue both from the transactions they complete and from the advertising they show. A retailer paid per transaction shares the seller's interest in shortening the path from a shopping need to a purchase, and it would design its pages to raise the probability that a search ends in a sale (Dinerstein et al., 2018; Ursu, 2018). A retailer that also sells advertising monetizes the attention of its customers, and it may benefit from longer searches that generate more impressions and more clicks on sponsored listings, even when they do not raise the probability of a purchase (Hagiu & Jullien, 2011; de Cornière & Taylor, 2019; Ichihashi & Kim, 2023). In our data, personalized advertising lengthens search by about a quarter to a third without changing the probability that an intent ends in a purchase or the products bought, while the advertiser pays for the personalization of an impression a small fraction of what the consumer's additional time costs. We do not observe how retailers weigh the two goals, but our results are consistent with an environment in which advertising

revenue rewards the time consumers spend searching more than the purchases they complete, and in which the cost of that time falls on consumers.

The results we presented in this work have some limitations. The clickstream captures only within-browser desktop behavior, so search and purchases off the browser or on mobile are unobserved. Observed never-buyers therefore include consumers who bought off-platform. The cross-arm comparison is unaffected, since randomization balances it, but the level of the never-buyer share represents an upper bound. Product attributes come from rendered pages with high but incomplete coverage (price, rating and number of reviews for 96 percent of opened products), using only each participant's own observed prices; we flag missing values rather than impute them across participants, since prices are personalized. The rank a product held on the result page is recovered only when the product was reached through a results page we parsed; the realized order in which products were examined, which the analysis relies on, is observed throughout from timestamps. The intent assignment is produced by a language model and, though blind to the experimental arm, is measured with error; the main estimate is similar when intents are built from query clusters instead, so the result does not depend on the language-model assignment. Finally, login status is chosen by the participant after assignment, so the logged-in sample conditions on a choice made after treatment; the share of intents on logged-in participant-merchant pairs is the same in the three arms (97 percent), and the estimates on all intents are very similar (Table A1, columns 6 and 7).

References

Abaluck, J., Compiani, G., & Zhang, F. (2026). A Method to Estimate Discrete Choice Models That Is Robust to Consumer Search. *Journal of Political Economy*, forthcoming.

Acquisti, A., Taylor, C., & Wagman, L. (2016). The Economics of Privacy. *Journal of Economic Literature*, 54(2), 442–492.

Acquisti, A., & Varian, H. R. (2005). Conditioning Prices on Purchase History. *Marketing Science*, 24(3), 367–381.

U.S. Bureau of Labor Statistics (2026). Average Hourly Earnings of All Employees, Total Private (CES0500000003), September 2026. Retrieved from FRED, Federal Reserve Bank of St. Louis.

Alcobendas, M., Kobayashi, S. J., Shi, K., & Shum, M. (2025). The Impact of Privacy Protection on Online Advertising Markets. Working paper.

Aridor, G., Che, Y.-K., Hollenbeck, B., Kaiser, M., & McCarthy, D. (2025). Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency. *Management Science*, forthcoming.

Allcott, H., Gentzkow, M., & Song, L. (2022). Digital Addiction. *American Economic Review*, 112(7), 2424–2463.

Anderson, S. P., & Coate, S. (2005). Market Provision of Broadcasting: A Welfare Analysis. *Review of Economic Studies*, 72(4), 947–972.

Athey, S., & Ellison, G. (2011). Position Auctions with Consumer Search. *Quarterly Journal of Economics*, 126(3), 1213–1270.

Bagwell, K. (2007). The Economic Analysis of Advertising. In M. Armstrong & R. Porter (Eds.), *Handbook of Industrial Organization* (Vol. 3, pp. 1701–1844). Elsevier.

Bakos, J. Y. (1997). Reducing Buyer Search Costs: Implications for Electronic Marketplaces. *Management Science*, 43(12), 1676–1692.

Beknazar-Yuzbashev, G., Jiménez-Durán, R., Simonov, A., & Stalinski, M. (2025). Social Media Advertising Loads as Prices. Working paper, Columbia Business School.

Blake, T., Nosko, C., & Tadelis, S. (2015). Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment. *Econometrica*, 83(1), 155–174.

Brynjolfsson, E., Collis, A., Liaqat, A., Kutzman, D., Garro, H., Deisenroth, D., & Wernerfelt, N. (2025). The Consumer Welfare Effects of Online Ads: Evidence from a Nine-Year Experiment. *American Economic Review: Insights*, forthcoming.

Cheyre, C., Jiang, L., Schaub, F., Ding, Z., Firullo, C., Li, Y., & Acquisti, A. (2024). An Experimental Infrastructure to Investigate the Impact of Online Tracking, Targeting, and Advertising on Consumer Behavior and Consumer Welfare. Working paper.

De los Santos, B., Hortaçsu, A., & Wildenbeest, M. R. (2012). Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior. *American Economic Review*, 102(6), 2955–2980.

Dinerstein, M., Einav, L., Levin, J., & Sundaresan, N. (2018). Consumer Price Search and Platform Design in Internet Commerce. *American Economic Review*, 108(7), 1820–1859.

Donnelly, R., Kanodia, A., & Morozov, I. (2024). Welfare Effects of Personalized Rankings. *Marketing Science*, 43(1), 92–113.

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). Internet Advertising and the Generalized Second-Price Auction. *American Economic Review*, 97(1), 242–259.

Ellison, G., & Ellison, S. F. (2009). Search, Obfuscation, and Price Elasticities on the Internet. *Econometrica*, 77(2), 427–452.

Frik, A., Haviland, A., & Acquisti, A. (2020). The Impact of Ad-Blockers on Product Search and Purchase Behavior: A Lab Experiment. *Proceedings of the 29th USENIX Security Symposium*.

Farronato, C., Fradkin, A., Lin, T., & Yu, C. (2026). Experience and the Demand for Ad Personalization: Evidence from a Field Experiment. Working paper, August 2026.

Gabaix, X., & Laibson, D. (2006). Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quarterly Journal of Economics*, 121(2), 505–540.

Goldberg, S. G., Johnson, G. A., & Shriver, S. K. (2024). Regulating Privacy Online: An Economic Evaluation of the GDPR. *American Economic Journal: Economic Policy*, 16(1), 325–358.

Goldfarb, A., & Tucker, C. (2011). Privacy Regulation and Online Advertising. *Management Science*, 57(1), 57–71.

Goldfarb, A., & Tucker, C. (2011). Online Display Advertising: Targeting and Obtrusiveness. *Marketing Science*, 30(3), 389–404.

Hodgson, C., & Lewis, G. (2025). You Can Lead a Horse to Water: Spatial Learning and Path Dependence in Consumer Search. *Econometrica*, 93(4).

Johnson, G. A., Shriver, S. K., & Goldberg, S. G. (2023). Privacy and Market Concentration: Intended and Unintended Consequences of the GDPR. *Management Science*, 69(10), 5695–5721.

Lewis, R. A., & Rao, J. M. (2015). The Unfavorable Economics of Measuring the Returns to Advertising. *Quarterly Journal of Economics*, 130(4), 1941–1973.

Marotta, V., Abhishek, V., & Acquisti, A. (2019). Online Tracking and Publishers' Revenues: An Empirical Analysis. *Workshop on the Economics of Information Security (WEIS)*.

Ichihashi, S., & Kim, B.-C. (2023). Addictive Platforms. *Management Science*, 69(2), 1127–1145.

Iyengar, S. S., & Lepper, M. R. (2000). When Choice is Demotivating: Can One Desire Too Much of a Good Thing? *Journal of Personality and Social Psychology*, 79(6), 995–1006.

Kuksov, D., & Villas-Boas, J. M. (2010). When More Alternatives Lead to Less Choice. *Marketing Science*, 29(3), 507–524.

Nelson, P. (1974). Advertising as Information. *Journal of Political Economy*, 82(4), 729–754.

Rhodes, A., & Zhou, J. (2025). Personalization and Privacy Choice. Working paper, Toulouse School of Economics and Yale University.

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

Ursu, R. M. (2018). The Power of Rankings: Quantifying the Effect of Rankings on Online Consumer Search and Purchase Decisions. *Marketing Science*, 37(4), 530–552.

Ursu, R. M., Wang, Q., & Chintagunta, P. K. (2020). Search Duration. *Marketing Science*, 39(5), 849–871.

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

Aridor, G., Jiménez-Durán, R., Levy, R., & Song, L. (2024). The Economics of Social Media. *Journal of Economic Literature*, 62(4), 1422–1474.

Yuan, Z., Chen, Y., Wang, Y., & Sun, T. (2024). How Recommendation Affects Customer Search: A Field Experiment. *Information Systems Research*, forthcoming.

Zhao, Y., Yildirim, P., & Chintagunta, P. K. (2023). Privacy Regulations and Online Search Friction: Evidence from GDPR. Marketing Science Institute Working Paper 21-110.

de Cornière, A., & Taylor, G. (2014). Integration and Search Engine Bias. *RAND Journal of Economics*, 45(3), 576–597.

de Cornière, A., & Taylor, G. (2019). A Model of Biased Intermediation. *RAND Journal of Economics*, 50(4), 854–882.

Hagiu, A., & Jullien, B. (2011). Why Do Intermediaries Divert Search? *RAND Journal of Economics*, 42(2), 337–362.

Hagiu, A., & Jullien, B. (2014). Search Diversion and Platform Competition. *International Journal of Industrial Organization*, 33, 48–60.

Hébert, B., & Zhong, W. (2026). Engagement Maximization. Working paper, Stanford University.

Hong, H., & Shum, M. (2006). Using Price Distributions to Estimate Search Costs. *RAND Journal of Economics*, 37(2), 257–275.

Kleinberg, J., Mullainathan, S., & Raghavan, M. (2022). The Challenge of Understanding What Users Want: Inconsistent Preferences and Engagement Optimization. *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*, 29.

Abhishek, V., Jerath, K., & Sharma, S. (2025). The Impact of "Retail Media" on Online Marketplaces: Insights from a Field Experiment. *Information Systems Research*, 36(1), 456–473.

Dash, A., Ghosh, S., Mukherjee, A., Chakraborty, A., & Gummadi, K. P. (2024). Sponsored is the New Organic: Implications of Sponsored Results on Quality of Search Results in the Amazon Marketplace. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 7.

Foulds, O., Azzopardi, L., & Halvey, M. (2021). Investigating the Influence of Ads on User Search Performance, Behaviour, and Experience during Information Seeking. In *Proceedings of the 2021 Conference on Human Information Interaction and Retrieval (CHIIR '21)*, 107–117.

Goldstein, D. G., Suri, S., McAfee, R. P., Ekstrand-Abueg, M., & Diaz, F. (2014). The Economic and Cognitive Costs of Annoying Display Advertisements. *Journal of Marketing Research*, 51(6), 742–752.

Miroglio, B., Zeber, D., Kaye, J., & Weiss, R. (2018). The Effect of Ad Blocking on User Engagement with the Web. In *Proceedings of the 2018 World Wide Web Conference (WWW '18)*.

Mustri, E. S., Adjerid, I., & Acquisti, A. (2023). Behavioral Advertising and Consumer Welfare: An Empirical Investigation. SSRN Working Paper No. 4398428.

Rock, R., Strauss, I., O'Reilly, T., & Mazzucato, M. (2024). Behind the Clicks: Can Amazon Allocate User Attention as It Pleases? *Information Economics and Policy*, 69, 101115.

Shah, C., White, R. W., Andersen, R., Buscher, G., Counts, S., et al. (2025). Using Large Language Models to Generate, Validate, and Apply User Intent Taxonomies. *ACM Transactions on the Web*, 19(3), Article 34, 1–29.

Todri, V. (2022). Frontiers: The Impact of Ad-Blockers on Online Consumer Behavior. *Marketing Science*, 41(1), 7–18.

Ursu, R. M., Simonov, A., & An, E. (2025). Online Advertising as Passive Search. *Management Science*, 71(2), 1050–1073.

Ad Badger (2026). Amazon Advertising Benchmarks 2026. Online report (data January–August 2026).

Skai (2026a). Q4 2025 Quarterly Trends Report. https://skai.io/wp-content/uploads/2026/01/Q4_2025_Quarterly-Trends-Report.pdf.

Skai (2026b). Q1 2026 Quarterly Trends Report. https://skai.io/wp-content/uploads/2026/04/Skai-Q1-2026-Quarterly-Trends-Report.pdf.

WordStream by LocaliQ (2026). Google Ads Benchmarks 2026. Online report (median cost per click, U.S. campaigns, April 2025–March 2026).

Holyoak, M. (2024). A Path Forward on Privacy, Advertising, and AI. Remarks at the National Advertising Division Keynote 2024, Federal Trade Commission, September 17, 2024. https://www.ftc.gov/system/files/ftc_gov/pdf/Holyoak-NAD-Speech-09-17-2024.pdf.

Lau, Y. (2020). A Brief Primer on the Economics of Targeted Advertising. Economic Issues Paper, Bureau of Economics, Federal Trade Commission, January 2020. https://www.ftc.gov/reports/brief-primer-economics-targeted-advertising.

Meta (2021). Let's Talk About Why Personalized Ads Matter. February 2021. https://about.fb.com/news/2021/02/lets-talk-about-why-personalized-ads-matter/.

Network Advertising Initiative (2026). Benefits of Tailored Advertising. Last updated January 2026. https://thenai.org/about-online-advertising/benefits-of-tailored-advertising/.

Solomon Partners (2025). Solomon's 2025 Major Media CPM Comparison. Solomon Partners Securities, June 2025. https://oaaa.org/wp-content/uploads/2025/06/2025-Solomon-Partners-US-Core-Media-CPM-Comparison.pdf.

A. Robustness

The main result holds across the specifications that address the obvious confounds, summarized in Table A1. The k-means clustering of the raw labels, an alternative category control that involves no taxonomy judgment, gives $+0.182$ ($p = 0.003$); the modal product category of the episode's members gives $+0.198$ ($p = 0.001$). Adding learning controls leaves the coefficient significant (Section A.1). The earlier search-seeded construction of intents, built from query clusters rather than from the language-model assignment, gives

$+0.063$ ($p = 0.07$) on its own sample, so the result does not depend on how intents are delimited. On all intents, including those pursued while logged out, the estimate is $+0.186$ ($p = 0.002$), and on the participants of the live roster, who were still enrolled when the data were extracted, it is $+0.212$ ($p = 0.002$; Table A1, columns 6 and 7). On the merchants where we could not block advertising, participants in the Ad-Blocking arm also see sponsored listings, which allows us to observe the same Ad-Blocking participant with and without them. Once the difference between the two types of merchant is measured on the Opt-In participants, who see sponsored listings on both, an Ad-Blocking participant spends about 0.19 log points more time per product where they see sponsored listings ($p = 0.03$), an effect similar in size to the main estimate (Appendix E).

The intent is formed after treatment, so per-intent outcomes condition on which episodes exist. The margins this could distort do not move. Treatment does not change how many intents a participant forms (Section 2), and fully unconditional participant-level aggregates point the same way: total search time over the quarter is $+0.108$ ($p = 0.46$) for Opt-In, time per intent aggregated at the participant is $+0.094$ ($p = 0.17$), consistent with the intent-level estimate though less precise, which is why the analysis is run at the intent level with goal-category fixed effects, and total products evaluated per participant does not change (-0.005, $p = 0.97$).

Dropping short intents attenuates the estimate (Table A2). Conditioning on a minimum duration truncates the outcome from below, and the treatment moves the whole duration distribution, so the cut removes part of the margin on which the effect operates. The attenuation below each threshold is what a truncation from below produces when the treatment shifts the distribution. A fixed rendering delay would predict the opposite pattern, a difference that is largest in the shortest episodes and vanishes in the longest, and it is ruled out also by the direct measurement of page load times (Appendix C). Figure A1 shows the shift across the distribution: the Opt-In minus Ad-Blocking difference in the quantiles of log time per product is between 0.17 and 0.24 log points from the tenth to the eightieth percentile and 0.10 at the ninetieth, the difference in log time per intent is between 0.20 and 0.26 at every decile, and neither falls at the top once the ten-minute dwell cap is lifted. Table A2 places next to each observed coefficient the one obtained by shifting the Ad-Blocking distribution by the full-sample estimate and applying the same cut. The two paths agree at the first two cuts, and the observed share of intents crossing each threshold rises with the threshold where the simulated share is flat, so the treatment

TABLE A1. Robustness of the effect on time per product. OLS of log time per product on indicators for the Opt-In and Anti-tracking arms, relative to Ad-Blocking. Column (2) is the baseline of Table 4. Category FE: the goal category of the intent (Goal), a k-means clustering of the intent labels with $K = 60$, or the modal product category of the intent (Product). Learning controls: logs of the participant's prior intents and prior intents in the same category. Login status: intents on logged-in participant-merchant pairs, or all intents. Roster: participants with at least 14 active days, or participants still enrolled. Merchants: the merchants where the blocker removes sponsored listings (Blocked), all merchants, or only eBay, Etsy, Temu, Target, AliExpress and Home Depot (Unblocked, placebo). Standard errors clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Opt-In	$+0.199^{***}$	$+0.195^{***}$	$+0.182^{***}$	$+0.198^{***}$	$+0.170^{***}$	$+0.186^{***}$	$+0.212^{***}$	$+0.174^{***}$	$+0.032$
	(0.072)	(0.061)	(0.060)	(0.061)	(0.055)	(0.061)	(0.067)	(0.059)	(0.101)
Anti-tracking	$+0.051$	$+0.048$	$+0.031$	$+0.050$	$+0.036$	$+0.048$	$+0.084$	-0.030	-0.131
	(0.089)	(0.069)	(0.071)	(0.069)	(0.057)	(0.069)	(0.074)	(0.080)	(0.103)
p-value, Opt-In = Anti-tracking	0.06	0.02	0.02	0.01	0.01	0.02	0.04	0.01	0.11
Observations (intents)	31,049	31,049	31,049	31,049	31,049	31,923	26,630	38,519	7,470
Participants	699	699	699	699	699	786	616	710	354
Demographic controls	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Category FE	None	Goal	k-means	Product	Goal	Goal	Goal	Goal	Goal
Learning controls	No	No	No	No	Yes	No	No	No	No
Login status	Logged in	Logged in	Logged in	Logged in	Logged in	All	Logged in	Logged in	Logged in
Roster	$\geq$14 days	$\geq$14 days	$\geq$14 days	$\geq$14 days	$\geq$14 days	$\geq$14 days	Enrolled	$\geq$14 days	$\geq$14 days
Merchants	Blocked	Blocked	Blocked	Blocked	Blocked	Blocked	Blocked	All	Unblocked

moves somewhat more mass across the lowest thresholds than a constant proportional shift implies.

A.1. Experience and learning effects

A natural alternative explanation is experience. Consumers who have searched a category before may have learned something about it and search faster the next time; if the arms differed in accumulated experience, that difference could look like a treatment effect. Time per product evaluated does fall with a consumer's cumulative search history, chiefly through the running count of prior intents (-0.099 in log time per product; prior intents in the same category add -0.016), which is evidence of learning within each experimental condition. Adding both learning controls leaves the targeted coefficient at $+0.170$ ($p = 0.002$).

TABLE A2. Duration cuts: observed and simulated coefficients (time per product, Opt-In vs Ad-Blocking). Observed: baseline specification on the intents above the cut, SE clustered by participant. Simulated: the Ad-Blocking intents shifted by the full-sample estimate ($+0.195$), pooled with the Ad-Blocking intents, same cut and specification. Extensive margin: share of intents above the cut, Opt-In minus Ad-Blocking, observed and simulated. Estimation sample, 31,049 intents.

	$\hat{\beta}_1$		Extensive margin (pp)	
Minimum duration	Observed	Simulated	Observed	Simulated
None	$+0.195^{***}$	$+0.195$		
30 seconds	$+0.140^{***}$	$+0.147$	$+3.2$	$+2.8$
1 minute	$+0.112^{**}$	$+0.122$	$+5.0$	$+4.8$
2 minutes	$+0.077$	$+0.120$	$+6.0$	$+5.1$

B. Sample Balance

C. Page Load Time

Dwell on a page runs from one navigation to the next, so it includes the time the page takes to load. The three arms could differ on this margin: the Ad-Blocking extension strips advertising and the Anti-tracking extension strips trackers, and both make pages lighter than the pages an Opt-In participant loads. The browser's navigation log records, for every full navigation, the moment the navigation starts, the moment the document is usable, and the moment the page has finished loading, with the browser's own clock. Table A5 reports these load times on the pages of shopping merchants. Opt-In and Anti-tracking pages load alike (a median of 2.96 and 3.00 seconds to fully loaded; the participant-level difference is -0.01 log points, $p = 0.89$), and Ad-Blocking pages load faster (-0.17 log points, $p = 0.003$). Loading accounts for 5 to 6 percent of the time spent in an intent in every arm. These events are recorded for full navigations and not for transitions inside a site, so the load time is measured on 32 percent of product-page views and 54 percent of result-list views and imputed elsewhere with the participant's median on that host.

Subtracting the load time of every page from the time variable before it enters the regression leaves the estimate where it is: $+0.203$ ($p = 0.002$) with the measured loads and $+0.179$ ($p = 0.014$) when the participant-host median is also charged to the pages without a measured load, against $+0.195$ in the same sample and specification. The second variant removes 11.2 percent of all measured time and is an upper bound, since it charges a full page load to transitions that do not reload the page.

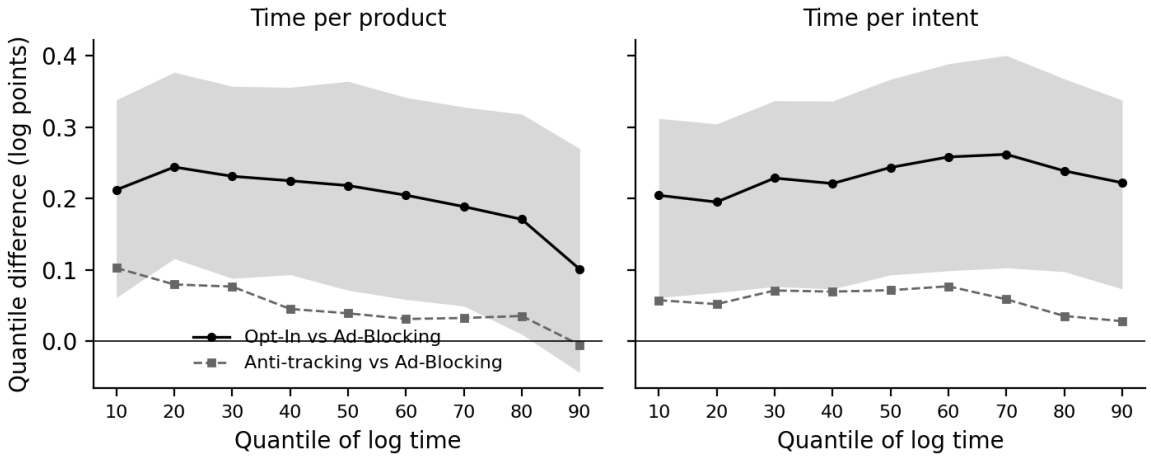


FIGURE A1. Quantile differences in log time, Opt-In and Anti-tracking minus Ad-Blocking. Each point is the difference between the arm's quantile and the Ad-Blocking quantile of log time per product (left) and log time per intent (right), after partialling out the seven controls and the goal-category effects. Bars are 95 percent intervals from 500 bootstrap replications that resample participants; the dotted line is the OLS coefficient of Table 4. Dwell capped at ten minutes; Figure A2 repeats the exercise with a sixty-minute cap. Estimation sample, 31,049 intents.

D. Data Construction

This appendix documents how the analysis variables are built from the raw clickstream.

The browser extension records one record per page view with a timestamp, URL, host, title, and rendered HTML, together with whether the page was active when recorded. Dwell on a page ends at the participant's next navigation, capped at ten minutes. Measured time reflects attention rather than open tabs: pages open in a background tab are excluded from every time measure (when the measures were first rebuilt from active pages only, the main estimate moved from $+0.185$ to $+0.181$).

An intent is an episode of shopping activity directed at one goal. Every product interaction the extension records (a product listed on a results page with its position, a product page opened, a purchase) enters a chronological event stream per participant, with the product's name, timestamp, and the page it appeared on. A large language model reads each participant's stream and groups events that pursue the same goal into intents, labeling each intent with the goal it infers and carrying open episodes across days; the model never sees the participant's experimental arm, so classification error is orthogonal to treatment. The assignment covers 99.9% of product events. The goal-category control is the product category the language model assigns to each intent's label (the paper's

TABLE A3. Display banners and the time gap between the two advertising arms. OLS of log time per product on an indicator for the Opt-In arm, relative to Anti-tracking, on the intents of the days on which the browser records ads, with an indicator for a display banner on the merchant pages of the intent (column 1) and its interaction with the Opt-In arm (column 2). The presence of a display banner depends on the pages visited, so it is not a pre-treatment control; the volume of display banners does not differ between the two arms (Table 3, Panel A). Standard errors clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Log time per product	
	(1)	(2)
Opt-In − Anti-tracking	$+0.123^{**}$	$+0.145^{*}$
	(0.060)	(0.076)
Display banner on the merchant pages	$+0.343^{***}$	$+0.364^{***}$
	(0.027)	(0.041)
Opt-In × display banner		-0.040
		(0.053)
Observations (intents)	19,188	19,188
Participants	457	457
Demographic controls	Yes	Yes
Goal-category fixed effects	Yes	Yes

baseline; Table A6 reports the alternatives, including the k-means clustering of the labels used in the companion extended abstract). An earlier construction built intents from query clusters and TF-IDF page tagging; the main estimate is similar on it (Table A1), and it is retained as a robustness sample.

Merchant classification is cleaned to drop false positives (survey-for-cash sites, streaming, gaming, marketplace-seller dashboards).

Table A7 shows what the assignment produces: for each product category, two episodes drawn at random with the label the model wrote for them, the merchants involved, and their size and span. The labels are short goal descriptions ("Wide Width Shoes", "185/65R15 Tires"), episodes routinely span more than one merchant, and many run over several days.

Table A8 shows how the assignment reads one participant's stream. Within a day the participant pursues two goals at once, a watch battery and a power tool, and the model keeps them apart: the results pages and product pages that belong to each goal are attributed to it as they arrive, the second episode moves across merchants, and one episode ends in a purchase while the other does not. This is the unit the analysis works with: the episode, with its time, its products, and its outcome, whether or not a purchase concludes it.

TABLE A4. Covariate balance across arms, estimation sample (699 participants). Means by arm; p-values from arm-wise comparisons to Ad-Blocking with robust standard errors. Joint F-tests of all seven covariates predicting assignment: $p = 0.42$ (Opt-In vs Ad-Blocking), $p = 0.42$ (Anti-tracking vs Ad-Blocking). Employed full-time is shown for description; the control is employment full or part time (60, 51 and 57 percent; $p = 0.05$ and 0.43).

	Ad-Blocking	Opt-In	p	Anti-tracking	p
Age	50.8	51.8	0.51	50.4	0.77
Female	68%	61%	0.14	61%	0.12
Non-white	24%	22%	0.74	29%	0.24
College	65%	66%	0.73	60%	0.31
Log income	10.87	10.81	0.47	10.79	0.32
Employed full-time	42%	42%	0.95	39%	0.64
IT skill (1–5)	3.39	3.44	0.51	3.49	0.19
Participants	241	237		221	

TABLE A5. Page load time on shopping-merchant pages, by arm. Seconds per page from the browser navigation log; "usable" is the time from the start of the navigation to the document being ready, "fully loaded" the time to the last resource. 525 participants, 267,672 pages.

	Usable, median	Fully loaded, median	Fully loaded, mean	Share above 5 s
Ad-Blocking	1.60	2.44	3.49	16%
Opt-In	1.69	2.96	4.47	25%
Anti-tracking	1.78	3.00	4.36	23%

A product evaluated is a distinct product-detail page, identified by merchant-specific URL patterns (e.g. Amazon `/dp/`, Walmart `/ip/`, eBay `/itm/`, and analogues). Time on a page is the difference to the next page view by the same participant, capped at ten minutes; the stored session identifier is not used for this, as it is a per-record counter. Time per product is total merchant time on tagged pages divided by the number of distinct products.

Price, rating, and review count are parsed from each rendered product page. Because prices are dynamic and personalized, a participant's price is taken only from that participant's own sessions; we never share a price across participants. Missing attributes are backfilled within participant (the participant's own observations of the same product or page) and flagged rather than imputed across participants.

Purchases are parsed from cart, checkout, and confirmation HTML. For Amazon, the parser reads only the order region of the confirmation page (the shipment-summary block), not the whole document, which otherwise picks up recommendation carousels; validated against the order-history database, order-region parsing raises the overlap from 3% to

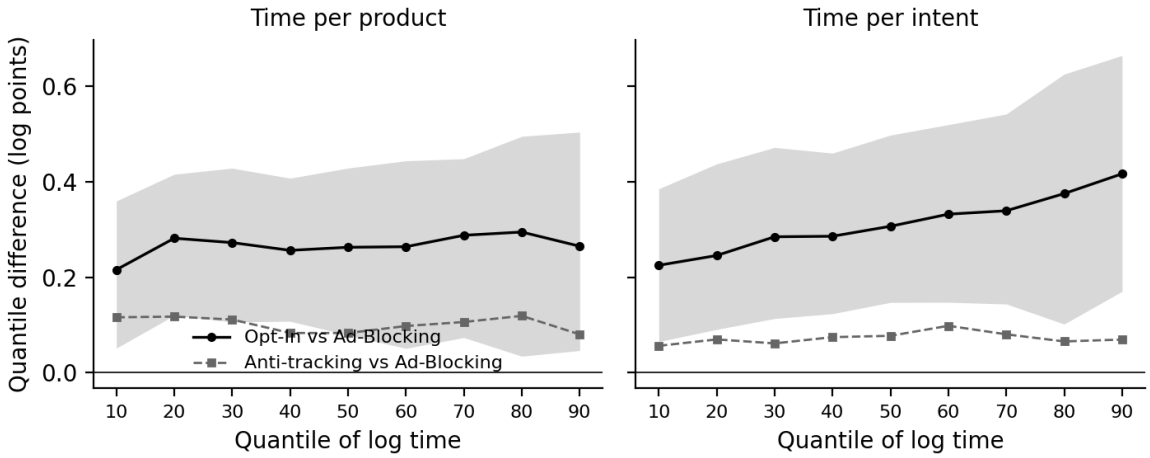


FIGURE A2. Quantile differences in log time with the dwell cap raised from ten to sixty minutes. Same construction as Figure A1; 200 bootstrap replications. With the ten-minute cap both arms are censored at the same absolute value, which shrinks a proportional shift at the top of the distribution; without it the difference at the ninetieth percentile is 0.26 log points for time per product and 0.42 for time per intent.

TABLE A6. The headline estimate under four constructions of the category control. Outcome is log search time per product; coefficient on Opt-In relative to Ad-Blocking, seven demographic controls, SE clustered by participant, $N = 31{,}049$. The LLM goal category is the paper's baseline; the k-means specification is the one used in the companion extended abstract.

Category control	$\hat{\beta}_1$	p
LLM goal category (16 categories, baseline)	$+0.195$	0.001
k-means clustering of labels ($K = 60$)	$+0.182$	0.003
Modal product category of members (15)	$+0.200$	0.001
Per-episode LLM classification with product names	$+0.199$	0.001

about 80%. An intent is matched to a purchase by a cascade, in order of reliability: (i) the purchased product's own page is tagged to the intent; (ii) a same-merchant page in the intent's session immediately precedes the purchase; (iii) product-text similarity. Text similarity alone is unreliable for this (consumers search one thing and buy another), so it is the last resort.

Visible ads come from the browser's record of in-view ad elements, each with its DOM container, the page it was on, its geometry, and a timestamp. Ads are classified by position into in-results versus out-of-results from the container and the page type. The same ad is re-observed on every scroll, so before counting we keep one impression per (participant, ad, page, position) per minute; this collapses the raw records to about an eighth of their

TABLE A7. Examples of LLM-assigned intents, two per product category (the first with a purchase where one exists). The label is the goal the assigner wrote for the episode; merchants are the two most frequent in the episode; listed and opened are the products seen on results pages and the product pages opened; the example product is one product from the episode, the purchased one when available; span is the number of calendar days the episode covers; time is measured search minutes. The category is the one the language model assigns to the label (the paper's baseline control). Episodes drawn at random among those with two to eight opened products and a classifiable label.

Category	Assigner's label	Example product	Merchants	Listed	Opened	Bought	Span	Min.
Home & Kitchen	5x7 Picture Frames	AVEAX Double 5x7 Picture Frame Wood Grain…	amazon	5	5	yes	1	1.1
	Paper Towels	Bounty Essentials Select-A-Size Paper…	walmart	60	7	no	1	3.2
Apparel & Shoes	Aviator Sunglasses	Sunglasses Men Polarized Aviator Sunglasses…	amazon	53	3	yes	1	4.5
	Womens Lounge Set	Womens Lounge Set 2025 Fall 2 Piece Outfits…	amazon	8	2	no	2	1.4
Grocery & Food	Coffee Pods	Happy Belly an Amazon Brand -Medium Roast…	amazon, walmart	36	6	yes	1	29.7
	Water Bottles	AUTOSEAL Spill-Proof Water Bottles, 24 Oz…	amazon, ebay	56	2	no	4	8.8
Electronics	iPad+Case	Apple iPad Pro (1st Generation, 9.7-inch,…	amazon, temu	46	5	yes	2	6.8
	Flip Phones	TCL Flip3 - Verizon - 16GB - Stone Gray -…	amazon	55	2	no	9	12.8
Media & Entertainment	Christian Books	Balaam: The Donkey Who Spoke Up (Children of…	amazon	12	2	yes	2	11.7
	Video+Games	Far Cry 6 Game Of The Year Edition | PC Code…	amazon	21	8	no	2	20.5
Beauty & Personal	Anti-Aging Serum	The Ordinary Niacinamide 10% + Zinc 1%,…	amazon	54	7	yes	4	10.7
	Hair Dye	Arctic Fox Semi-Permanent Hair Color Dye - 4…	target	25	6	no	1	4.6
Tools & Home Improvement	Toilet Supplies	Clorox Original Toilet Cleaning System,…	amazon	53	2	yes	1	1.4
	Bread Proofing Box	Bread Proofing Box Large Capacity Dough…	target, walmart	18	6	no	2	14.7
Health & Supplements	Incontinence Underwear	LivDry Adult Incontinence Underwear,…	amazon, walmart	19	2	yes	1	1.0
	Shoulder Brace	FIGHTECH Shoulder Brace for Torn Rotator…	amazon	23	6	no	1	18.6
Crafts & Hobby	Sewing Scissors	SINGER 00557 4-1/2-Inch ProSeries Detail…	amazon	49	2	yes	1	1.5
	Quilting Fabric	Kona Cotton - Snow 10 yard Bolt - Bulk…	etsy	58	2	no	1	2.6
Pet	Cat Food - Walmart	Purina Friskies Seafood Sensations Indoor…	walmart	6	2	yes	1	0.8
	Pet Health Care	2 Boxes Of Dog Probiotics, Containing 60…	temu	12	4	no	1	3.7
Sports & Outdoors	Rain Gear	Frogg Toggs Ultra Lite Rain Suit (Men's)	amazon, walmart	33	5	yes	1	3.8
	Swimsuit Accessories	size 8 girls swimsuit	walmart	21	5	no	1	1.4
Office & School	Printers	Canon imageCLASS LBP6230dw - Compact,…	amazon	22	2	yes	1	3.4
	Mechanical Pencils	uniball Kuru Toga Mechanical Pencil Starter…	amazon	13	2	no	1	3.1
Baby & Kids	Baby Diapers	Aquaphor Baby Healing Ointment Advanced…	target	23	6	yes	4	6.1
	Baby Beds & Accessories	Baby Relax Ivey 3 in 1 Convertible Crib White	walmart, temu	51	4	no	1	11.4
Auto	Carburetor	Carburetor For RuiXing Briggs & Stratton…	amazon, ebay	45	7	yes	2	9.3
	Jeep Parts	Right Rear Door Body Mount Weatherstrip…	walmart	40	5	no	38	3.6
Travel & Services	Travel+Adapters	US to Argentina Travel Adapter Plug for…	amazon	46	6	no	1	7.3

count. Display exposure per intent counts the out-of-results impressions whose timestamp falls inside the intent's event window.

Login status is read from the header of each rendered merchant page, as described in Section 2.3: an account greeting showing a name (the anonymizer's name marker, or a first name) or a sign-out link marks the page logged in, and a sign-in prompt without a greeting ("Hello, sign in", "Sign In", "Log In", "Sign in / Register") marks it logged out, with merchant-specific rules where the header differs. Every page of a merchant in a browsing session counts as logged in when at least one page of that merchant in the session is. Of the 327,806 merchant pages with a rendered capture, 92% are logged in, 5% logged out and 3% cannot be classified. A participant-merchant pair is classified as logged in when more than half of its pages with a status are (at least three pages).

Regressions use the seven pre-treatment controls (age, female, non-white, college, log income, employed, IT skill) and cluster standard errors by participant. The treatment is assigned one condition per participant, so participant fixed effects would absorb the treatment and are not used; the panel design instead uses condition (interacted with post, for the spending difference-in-differences) with month and product-category fixed effects. Price, spending, and cost regressions are reported with and without category fixed effects.

E. Within-Subject Analysis

On some merchants we were not able to block advertising for the Ad-Blocking arm. On eBay, Etsy, Temu, Target, AliExpress and Home Depot, sponsored listings are served by the merchant itself together with the organic results, and the ad blocker does not remove them, so participants in every arm see them. For this reason we exclude these merchants from the estimation sample. However, many participants shop on both types of merchant, and this allows us to observe the same Ad-Blocking participant without sponsored listings, on the merchants where the ad blocker removes them, and with sponsored listings, on the merchants where it does not.

In principle, if sponsored listings increase search time, a participant in the Ad-Blocking arm should spend more time per product on the merchants where they see sponsored listings than on the merchants where they do not. However, the two types of merchant also differ in ways that are unrelated to advertising: for instance, a product on eBay or Temu might take less time to evaluate than a product on Amazon, whatever the advertising. We measure this difference on the participants in the Opt-In arm, who see sponsored

listings on both types of merchant, so that their difference in time between the two types of merchant reflects only the merchants. The effect of sponsored listings on an Ad-Blocking participant is then the difference between their own gap across the two types of merchant and the gap of the Opt-In participants.

Figure A3 shows, for each arm, the difference in log search time between the intents on the merchants where we could not block advertising and the intents on the other merchants of the same participant, for the 343 participants with intents on both types of merchant (29,532 intents, 7,374 of them on the merchants where we could not block advertising). Participants in the Opt-In arm spend about 31 percent less time per product on the merchants where we could not block advertising (-0.366, $p < 0.001$), which is the difference between the two types of merchant. Participants in the Ad-Blocking arm also spend less time per product on these merchants, but by less (-0.162, $p = 0.01$). Therefore, once the difference between the merchants is taken into account, the same Ad-Blocking participant spends about 0.20 log points more time per product where they see sponsored listings ($p = 0.06$), an effect similar in size to the main estimate of Table 4.

Table A9 reports the estimates from a regression of log time on an indicator for the merchants where we could not block advertising and on its interactions with the Ad-Blocking and the Anti-tracking arm, with participant and goal-category fixed effects and the Opt-In arm as the reference. The participant fixed effects absorb the arm and everything that is constant about the person, so the interactions compare each participant with themselves across the two types of merchant. The effect for Ad-Blocking participants is similar when we add merchant fixed effects ($+0.193$, $p = 0.03$) and when we measure time per intent ($+0.213$, $p = 0.07$). The Anti-tracking arm lies in between ($+0.117$, $p = 0.26$), which is consistent with the smaller difference between the Anti-tracking and the Ad-Blocking arm on the merchants where the blocker removes sponsored listings (Table 4).

F. Professional Survey Takers

The panel recruits from a population that overlaps with paid online task work, so some participants spend most of their browsing on survey and get-paid-to platforms. If their shopping were part of that work rather than ordinary consumer search, the main estimate would measure a different behavior. We check this in two ways.

First, we ask whether the shopping of these participants originates from survey work.

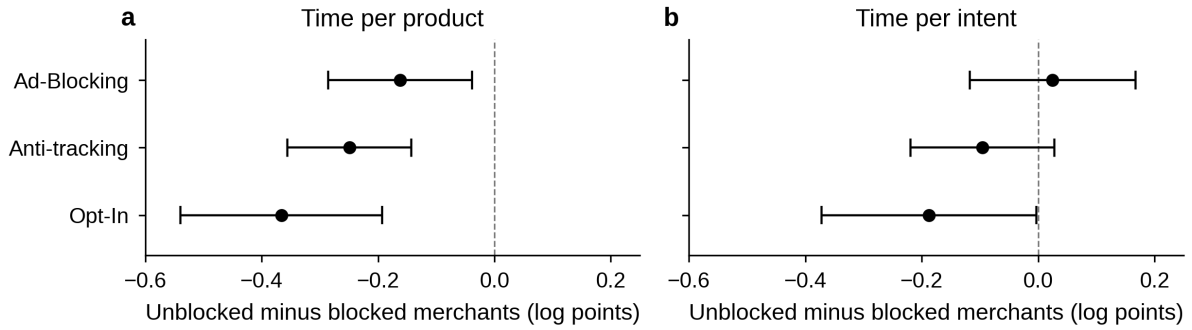


FIGURE A3. Within-subject analysis: difference in log search time between merchants where advertising could not be blocked (eBay, Etsy, Temu, Target, AliExpress, Home Depot) and the other merchants, same participant, by arm. Panel (a) time per product, panel (b) time per intent. Participant and goal-category fixed effects; 343 participants with intents on both types of merchant, 29,532 intents. Bars are 95% confidence intervals, standard errors clustered by participant.

We flag a participant by the share of their thirty-minute browsing sessions in which at least half the pageviews fall on a survey or get-paid-to website, using a list of 3,473 such websites. We exclude the study's own Qualtrics pages, which host the onboarding and the three monthly surveys, since counting them would flag participants for answering our own questionnaires. Under the majority rule used here, seven participants spend more than seventy-five percent of their sessions on survey work, and twenty more than half. For the participants above half, of 14,503 merchant pageviews, 4.1 percent are immediately preceded within the session by a survey website and 0.7 percent by Swagbucks specifically, against 0.5 percent for a random sample of 200 other participants. Most of their merchant pageviews are preceded by a page of the same merchant (9,943 on Amazon, 1,493 on Walmart and 327 on Target), and the cash-back shopping portals operated by these platforms, through which survey work could lead to shopping, account for 4 of the 14,503 pageviews. In other words, these participants do paid tasks and shop in separate sessions.

Second, we re-estimate the main results without the flagged participants at two thresholds (Table A10). The coefficient on time per product changes from $+0.195$ to $+0.204$ and $+0.200$, and the other outcomes are unchanged. The last column applies a broader rule, in which a session counts as survey work if it contains a single survey pageview: this rule removes 102 participants and 5,740 intents, most of them participants with one such pageview in a long shopping session, and the coefficient falls to $+0.162$ ($p = 0.015$). We report this column as a limit case, since the rule flags a session of two hundred shopping pages because of one visit to a survey website. The share of flagged participants is similar

across arms at every threshold, so the exclusion is not correlated with treatment.

TABLE A8. Worked example of the intent assignment: one participant's product events over three days and the intent assigned to each. The language model reads the participant's chronological event stream (product names, timestamps, and the page each product appeared on), groups events that pursue the same shopping goal, and writes a label for each episode; episodes remain open across days and can interleave. The participant's activity on other intents in the same days is omitted. Assigned intents: **A** = "CR1025 Battery"; **B** = "Impact Driver Purchase".

Time	Merchant	Event	Intent
Day 1, 03:39	amazon	results pages, 35 products listed	A
Day 1, 03:41	amazon	opens product page	A
Day 1, 03:44	amazon	results pages, 28 products listed	B
Day 1, 03:52	amazon	opens product page: 20V MAX Impact Driver, Cordless, 3-Speed,...	B
Day 1, 03:52	amazon	results pages, 34 products listed	A
Day 1, 17:15	temu	opens product page	B
Day 1, 21:54	amazon	results pages, 4 products listed	A
Day 1, 21:55	amazon	**purchase**: PKCELL 5-Pack CR1025 Lithium Battery 3V...	A

TABLE A9. Within-subject analysis: search time on merchants where advertising could not be blocked. OLS on log time with participant and goal-category fixed effects (columns 3 and 4 add merchant fixed effects), Opt-In arm as reference; participants with intents on both types of merchant (343 participants, 29,532 intents). "Unblocked merchant" is an intent whose main merchant is eBay, Etsy, Temu, Target, AliExpress or Home Depot. The first row is the difference between the two types of merchant for Opt-In participants, who see sponsored listings on both; it is absorbed by the merchant fixed effects in columns 3 and 4. The interaction with Ad-Blocking is the additional time of an Ad-Blocking participant on the merchants where they see sponsored listings. SE clustered by participant in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Participant + category FE		+ merchant FE	
	Time per product	Time per intent	Time per product	Time per intent
Unblocked merchant (Opt-In)	-0.366^{***}	-0.188^{**}		
	(0.089)	(0.094)		
Unblocked × Ad-Blocking	$+0.204^{*}$	$+0.213^{*}$	$+0.193^{**}$	$+0.214^{*}$
	(0.108)	(0.118)	(0.087)	(0.110)
Unblocked × Anti-tracking	$+0.117$	$+0.092$	$+0.102$	$+0.081$
	(0.103)	(0.112)	(0.077)	(0.097)

TABLE A10. Headline results with professional survey takers excluded. Each column is a separate sample; the specification is the one in Table 4. A session counts as survey work when at least half its pageviews are on a survey host, except in the last column, where one such pageview suffices. Coefficients are Opt-In relative to Ad-Blocking, SE clustered by participant. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Full sample	Drop $\geq 75\%$	Drop $\geq 50\%$	Loose rule $\geq 50\%$
Search time per product	$+0.195^{***}$	$+0.204^{***}$	$+0.200^{***}$	$+0.162^{**}$
	(0.061)	(0.062)	(0.062)	(0.067)
Search time per intent	$+0.226^{***}$	$+0.234^{***}$	$+0.230^{***}$	$+0.189^{***}$
	(0.062)	(0.063)	(0.063)	(0.071)
Products evaluated	$+0.023$	$+0.023$	$+0.023$	$+0.017$
Merchants visited	$+0.002$	$+0.003$	$+0.002$	$+0.004$
P(purchase)	$+0.023$	$+0.026$	$+0.028$	$+0.018$
Participants	699	692	679	597
Intents	31,049	30,820	30,621	25,309