Experience and the Demand for Ad Personalization: Evidence from a Field Experiment*

Chiara Farronato
Harvard University, CEPR, NBER
cfarronato@hbs.edu

Andrey Fradkin
Boston University
fradkin@bu.edu

Tesary Lin
Boston University
tesary@bu.edu

Chuan Yu
University of North Carolina
chuanyu@unc.edu

November 3, 2026

Abstract

We study consumer preferences for ad personalization and whether these preferences are well informed. We conduct a field experiment using a browser extension that varies consumers' exposure to non-preferred ad settings and provides information on how to change the settings on Amazon, Meta, and Google. Most consumers initially prefer personalized advertising, but 27% of platform-level settings do not align with stated preferences, and half of consumers are misaligned on at least one platform. Experience affects preferences: consumers exposed to their non-preferred setting are 8.1 percentage points more likely to revise their preferences, and the valuation of their preferred setting increases by $4.34 on average. Information alone has limited effects: a substantial share of participants are unaware that their ad settings do not match their preferences, and few of them change their settings after receiving instructions, consistent with substantial switching costs. Finally, personalization improves perceived ad relevance and reduces distraction primarily among consumers who already prefer personalized ads. These findings suggest that consumers are imperfectly informed about their utility from personalized advertising, that observed settings reflect both imperfect information and switching frictions, and that experience can help consumers better align their settings with their preferences.

*This research was supported by the Internet Society Foundation and by Tesary Lin's Isabel Anderson Professorship at Boston University. Andrey Fradkin was employed by Amazon until July 1, 2026 and is no longer affiliated with the company. Amazon did not fund this project, did not provide data or other resources, nor do they have the right to review the paper. The authors have no other relevant financial relationships or conflicts of interest to disclose.

1 Introduction

Global advertising expenditures are projected to exceed US$1.2 trillion in 2026, with the vast majority spent on digital channels.[1] Digital advertising is a primary source of revenue for many firms, ranging from platforms such as Google, Meta, and Amazon to media and news publishers.[2] Its growth has been enabled in part by technologies that use consumer data to target and personalize ads.

Personalization can increase the value of advertising by allowing firms to tailor ads to individual consumers, but it also requires the collection and use of consumer data. While advertisers and hosting platforms benefit from personalization, its value to consumers is less clear. Personalized ads may increase relevance and reduce search costs, but they may also raise privacy concerns, increase temptation, or disrupt the browsing experience. Whether consumers benefit from personalization is therefore an empirical question.

We study consumers' demand for ad personalization and ask whether their ad settings reflect their preferences. Consumers' ad settings may differ from their preferences for three reasons. First, preferences may be poorly informed because consumers do not know how personalization affects their browsing experience, particularly if they have never experienced a platform under a different ad setting. Second, switching costs, such as the difficulty of locating and changing platform settings, may prevent consumers from implementing their preferences. Third, consumers may be unaware of their current ad settings and therefore may not know whether those settings match their preferences.

We design an experiment to separately study the roles of experience and switching costs in shaping consumers' choices. Using a browser extension, we elicit consumers' willingness to accept to switch to their less preferred ad personalization setting on three major digital platforms: Amazon, Google, and Meta, which together account for 60% of U.S. digital ad spending.[3] We then randomize participants into one of three groups: an information treatment that informs participants of their current ad settings and how to change them, a control group that experiences their preferred ad setting, and an experience treatment that exposes participants to their non-preferred setting. Comparing baseline and endline preferences and settings allows us to measure learning from experience and the frictions that prevent consumers from implementing their preferred choices.

We find that consumers' preferences for ad personalization are shaped by experience

[1] https://www.statista.com/outlook/amo/advertising/worldwide

[2] https://reutersinstitute.politics.ox.ac.uk/journalism-media-and-technology-trends-and-predictions-2026

[3] https://www.emarketer.com/content/triopoly-s-market-share-dip-exposes-new-digital-trends

rather than fully known ex ante. Although most consumers initially prefer personalized advertising, many have settings that do not reflect those preferences. Exposure to a non-preferred setting changes both the direction and intensity of preferences, whereas information about how to change settings has comparatively small effects. Together, these findings suggest that observed ad settings reflect both imperfectly informed preferences and switching costs. Finally, personalized advertising improves perceived ad relevance primarily among consumers who already prefer personalization.

Relevant Literature. Our paper contributes to three strands of literature. First, we contribute to the literature on consumer learning and experience goods. When consumers cannot fully evaluate alternatives before using them, observed choices may not reflect informed preferences. This problem has been studied in settings including antidepressants (Dickstein, 2021), school choice (Agte et al., 2024), and search engines (Allcott et al., ming). We extend this literature to personalized advertising, where consumers may have limited experience with alternative advertising environments and therefore imperfectly understand the utility they derive from each.

Second, we contribute to the literature on the welfare effects of personalized advertising on consumers. A large literature documents the effectiveness of personalized advertising (Johnson et al., 2017; Bleier and Eisenbeiss, 2015; Sahni and Zhang, 2024; Schnadower Mustri et al., 2023; Yu, 2025; Wernerfelt et al., 2025), which speaks to its relevance to consumers. However, the welfare effects of personalized advertising depend not only on relevance, but also on obtrusiveness (Goldfarb and Tucker, 2011) and consumers' intrinsic desire to limit data collection (Lin, 2022). Evaluating consumer welfare calls for a broader measure of preferences. Recent work directly measures consumers' willingness to pay (Beknazar-Yuzbashev et al., 2025; Brynjolfsson et al., 2025) or satisfaction (Deisenroth et al., 2026) with personalized advertising on Meta. We add to this literature by asking whether these preferences are well informed, studying the sources of the gap between preferences and settings, and extending the analysis to major digital advertising platforms beyond Meta.

Third, we contribute to the literature on information frictions, transparency, and privacy choices. Privacy policies are notoriously difficult for consumers to understand despite regulatory requirements (Gambato et al., ming). Collis et al. (2021) show that providing consumers with information about the economic value of their personal data leads many to revise their data valuations, while Aridor et al. (2026) study how information about websites' privacy practices affects consumers' online choices. However, it remains an open question which forms of information provision are most useful to consumers. Separately, Farronato et al. (2026) show that repeated consent prompts impose substantial costs on online brows-

ing, highlighting the frictions created by online privacy choice environments. We advance this literature by comparing two forms of information provision: direct experience with alternative ad settings and information about how to change those settings. This design allows us to distinguish preference learning from switching costs and to compare the effectiveness of different interventions in helping consumers implement their preferred choices.

The rest of the paper is organized as follows. The next section describes our study design. Section 3 presents a theoretical framework that maps the relevant economic forces to our experimental design and data. Section 4 presents our results, and Section 5 concludes.

2 Study Design

We conducted a three-week field experiment on desktop browsers and recruited participants through Prolific in 2026. To be eligible, respondents had to be living in the U.S. aged 18 or older, pass an attention check, use a personal desktop or laptop running Google Chrome, not routinely use an ad blocker or VPN, be willing to remain logged into their Google, Facebook/Instagram, and Amazon accounts throughout the study, and not use their browser solely for survey-related activities. Respondents who failed any of these screening criteria were excluded.[4]

After providing informed consent, participants installed a Chrome browser extension developed for the study (Farronato and Fradkin, 2026) and kept it installed throughout the experiment. Participants completed three surveys: an enrollment survey at the beginning of the study, a midline survey after one week, and an endline survey after three weeks. They received $5 upon successful enrollment and installation of the extension, $3 for completing the midline survey, and $5 for completing the endline survey, for a total guaranteed payment of $13. Participants could also receive incentive payments through the preference-elicitation mechanism described below.

The browser extension recorded participants' online advertising environments throughout the study. It tracked ad personalization settings on Amazon, Google/YouTube, and Facebook/Instagram;[5] the URL and timestamp of each webpage visited; and information about ads displayed on those pages. To identify ads, the extension captured screenshots of webpages once per minute during active browsing. These screenshots were processed by an automated image-classification script using GPT-5 and then permanently deleted.

[4]To ensure that participants understood the study procedures, we also asked comprehension questions covering the study duration, required tasks, compensation, and their ability to withdraw at any time. Participants could proceed only after answering all questions correctly.

[5]Google and YouTube shared the same ad personalization setting, as did Facebook and Instagram.

The extension also periodically prompted participants with brief questions while they were actively browsing. During the baseline and intervention periods, pop-up surveys appeared following the detection of ads and asked participants whether the ads on the current webpage were relevant or distracting, or whether no ads were present (see Appendix Figure A.1 for screenshots of the extension).[6]

The first week of the study served as the baseline period. During this week, the extension recorded participants' browsing behavior, advertising exposure, and platform ad settings while participants retained their existing ad personalization settings. At enrollment, we also collected participants' baseline ratings of the relevance and obtrusiveness of ads on each of five platforms—Amazon, Google, YouTube, Facebook, and Instagram—on a scale from 0 to 100. We elicited the same ratings at endline, allowing us to study how changes in ad personalization affect perceived ad relevance and distraction.[7]

The midline survey, administered after the first week, elicited participants' preferences for personalized advertising on Amazon, Google/YouTube, and Facebook/Instagram. Participants first indicated whether they preferred personalized or non-personalized advertising on each platform. To make these choices incentive compatible, we informed participants that the browser extension would implement their stated preferences. We also elicited their beliefs about their current ad personalization settings on each platform. Participants were then shown a summary of their actual settings, as recorded by the browser extension, and given an opportunity to revise their preferences before any experimental intervention.[8] This procedure allows us to distinguish misperceptions about current settings from preference changes induced by the intervention.

Following preference elicitation, we randomly assigned participants to one of three experimental groups: an information treatment (20%), an experience treatment in which participants were assigned to experience their non-preferred settings (40%), and a control group in which participants were assigned to experience their preferred settings (40%). Participants in the information treatment were informed that the extension would not modify their ad personalization settings. Instead, after we disclosed their actual ad settings to them, participants received detailed instructions on how to locate and modify those settings on each

[6] A pop-up was triggered after a screenshot was classified as containing ads, provided that more than 10 minutes had passed since the previous pop-up. We capped the number of pop-ups at 15 per day.

[7] These relevance and distraction ratings (hereafter, baseline and endline ratings) complement the pop-up ad surveys. The baseline and endline ratings are more granular and have high completion rates, but they capture participants' perceptions only of ads on those five platform. The pop-up surveys also capture advertising encountered elsewhere on the web.

[8] In practice, only 3.6% of preferences across Amazon, Google, and Facebook were revised after participants were shown their actual ad personalization settings. Throughout the paper, we use these updated preferences as participants' pre-intervention preferences.

platform themselves. By reducing the informational costs of changing settings while leaving implementation to participants, this treatment helps us assess the importance of switching frictions.

For the remaining 80% of participants, we elicited their monetary valuation of ad personalization using an incentive-compatible Becker–DeGroot–Marschak (BDM) mechanism. For participants who preferred either fully personalized or fully non-personalized advertising across platforms, we elicited their willingness to accept (WTA) to switch to the less-preferred configuration for the two-week intervention period. Participants with mixed preferences across platforms reported their WTA for switching to both the fully personalized and fully non-personalized configurations, after which one of these two configurations was randomly selected for the intervention. We also elicited participants' WTA for switching each individual platform to its non-preferred setting, although these platform-specific valuations were not designed to be incentive compatible.

Participants assigned to the experience treatment were offered $25 to adopt their non-preferred ad personalization configuration for the two-week intervention period. This offer amount was chosen based on pilot studies and exceeded the WTA of 96.3% of participants. Participants with WTA below $25 received the payment and experienced their non-preferred settings, whereas those with WTA above $25 experienced their preferred settings.

For participants assigned to the control group, we set their ad personalization settings to match their stated preferences. Importantly, some participants entered the experiment with settings that differed from their preferences. For these participants, assignment to the experience treatment left their settings *unchanged*, since their existing settings were already their non-preferred settings, while assignment to the control group switched them to their preferred settings and eliminated the initial mismatch. (For a share of mismatches between baseline preferences and settings, see Panel B of Table 1 below).

We define the experience and control groups relative to participants' preferred settings, rather than by whether their actual settings changed, as required by the incentive-compatible BDM mechanism. Random assignment generates exogenous variation in exposure to non-preferred advertising settings, allowing us to identify how such experience affects preferences for ad personalization. In our analysis, we consider two treatment measures: (1) assignment to the experience treatment and (2) actual exposure to an ad personalization setting that differs from the participant's pre-intervention setting. The former is exogenous by design, whereas the latter is endogenous. We instrument exposure to a different setting with assignment to the experience treatment.

At the end of the two-week intervention period, participants completed the endline

survey. We again elicited their preferred ad personalization settings and their ratings of ad relevance and obtrusiveness on each platform. Participants whose preferences differed from those reported at midline were asked to explain the reasons for their change. We then repeated the incentive-compatible BDM elicitation to measure changes in participants' WTA for switching to their less-preferred ad settings for a two-week period. To maintain incentive compatibility, the second BDM assigned the vast majority of participants to their preferred settings with zero compensation and a small share to their non-preferred settings with a $25 offer payment, and the extension implemented the resulting settings. As a result, the vast majority of participants left the experiment with ad settings that reflected their most recently expressed preferences.

2.1 Sample and Baseline Perceptions

We recruited a total of 1,805 eligible participants from Prolific for the study. Of these, 1,509 completed the midline survey and 1,182 completed the endline survey. Our main analysis sample consists of participants who completed all study milestones and had at least three active browsing days during the baseline week (first week).[9] These criteria leave a final analysis sample of 869 participants. Panel A of Table 1 summarizes participant retention, the sample selection criteria, and how they vary across treatment arms. Participants assigned to non-preferred settings have a higher retention rate, likely because the $25 offer gives them a stronger incentive to complete the study. Panel B summarizes the platform-level observations. In principle, each participant contributes three platform-level settings. In a small share of cases, the extension fails to record a setting, but the missing rate is highly similar across treatment arms.

Table 2 presents descriptive statistics for the main analysis sample. The average participant is 37 years old, 66% are female, 74% identify as White, 11% as Black, and 5% as Asian. Approximately 23% of participants are students, and 65% report being employed. During the baseline week, participants visited an average of 17.5 unique domains and spent roughly 80 minutes browsing per day. Nearly all participants report visiting Google at least weekly (98%), with substantial shares also visiting YouTube (80%), Facebook (71%), and Instagram (68%) at least weekly. Amazon is visited at least weekly by 40% of participants.

Figure 1 summarizes participants' baseline perceptions of advertising across the five platforms operated by Amazon, Google, and Meta. These ratings were elicited at enrollment, before the baseline observation period and any experimental intervention. Perceptions differ

[9] An active day is defined as a day with at least five unique domains visited and at least ten captured screenshots.

Table 1: Sample Flow

	Info	Preferred setting	Non-preferred setting	Total
Panel A: From intake to the analysis sample				
Completed intake, randomized	286	594	629	1,509
Completed endline survey	210	446	526	1,182
and active user (analysis sample)	157	326	386	869
Retention	54.9%	54.9%	61.4%	57.6%
Panel B: Platform-level observations				
Platform preference–setting pairs observed	447	921	1,093	2,461
of which mismatched	127	239	310	676
Average settings recorded per active user	2.85	2.83	2.83	2.83

Table 2: Descriptive Statistics

	N	Mean	Std. Dev.	10th Pct.	Median	90th Pct.
Demographics						
Age	869	37.30	12.62	21	36	55
Share Female	868	0.66	0.47	0	1	1
Share White	867	0.73	0.44	0	1	1
Share Black	867	0.11	0.31	0	0	1
Share Asian	867	0.05	0.22	0	0	0
Share Student	869	0.23	0.42	0	0	1
Share Employed	793	0.65	0.48	0	1	1
Browsing						
Domain Visited Daily	869	17.01	11.36	7	14	32
Daily Browsing Time (minutes)	869	72.54	63.00	19	53	156
Platform Usage						
Visit Google at least weekly	869	0.98	0.13	1	1	1
Visit YouTube at least weekly	869	0.80	0.40	0	1	1
Visit Facebook at least weekly	869	0.71	0.46	0	1	1
Visit Instagram at least weekly	869	0.68	0.47	0	1	1
Visit Amazon at least weekly	869	0.40	0.49	0	0	1

substantially across platforms. Amazon stands out as the platform where ads are perceived as most relevant and least distracting, while YouTube lies at the opposite extreme. On Google, ads are also perceived as somewhat more relevant than distracting. In contrast, on Instagram and Facebook, participants perceive ads as more distracting than relevant.

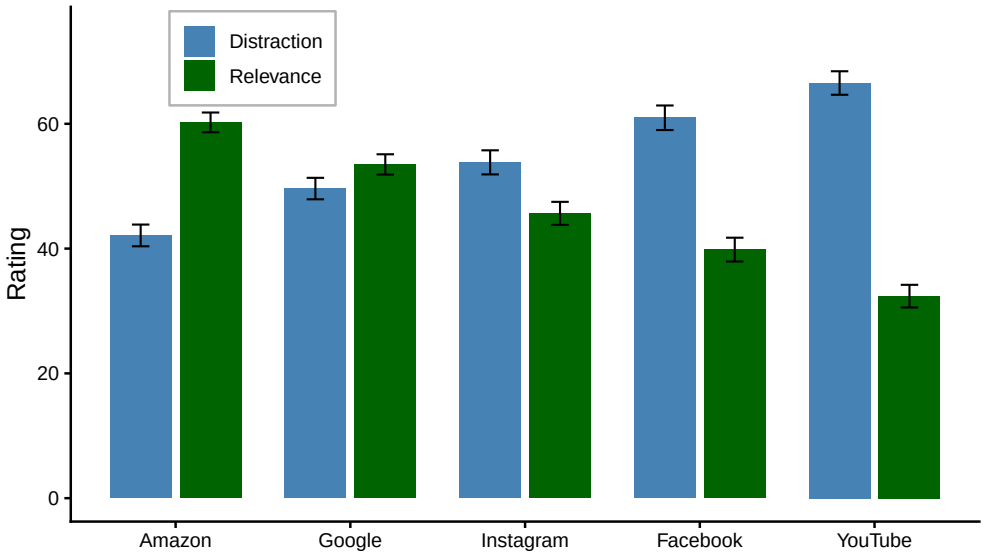


Figure 1: Ad Relevance and Distraction Across Platforms

Notes: The figure reports average ratings of ad distraction and relevance across platforms. After consenting to the study, participants were asked separately for each platform to indicate their agreement with the following statements: "ads impede my experience by cluttering my results with distracting content" and "ads improve my experience by showing me relevant products and services." Responses are scaled from 0 to 100, with higher values indicating stronger agreement. Error bars denote 95% confidence intervals.

3 Theoretical Framework

In this section, we develop a theoretical framework for the value of ad personalization and the frictions that generate gaps between consumers' actual and optimal ad settings. We abstract from platform-specific preferences for simplicity, but introduce them when mapping the model to the data.

True and perceived value of personalization. Let $U(d)$ denote a consumer's fully informed utility under ad personalization setting $d \in \{0, 1\}$, where $d = 1$ denotes personalized ads. Define

$$v = U(1) - U(0)$$

as the consumer's fully informed value of personalized relative to non-personalized advertising. Before gaining experience with alternative settings, the consumer perceives this value to be $\widetilde{v}_0$. The resulting valuation error is $v - \widetilde{v}_0$.

The consumer's initial and fully informed binary preferences for ad personalization are

$$P_0 = \mathbf{1}\{\widetilde{v}_0 \geq 0\}\,, \qquad P^* = \mathbf{1}\{v \geq 0\}\,.$$

We are particularly interested in *sign errors*. Define

$$Q_0 = \mathbf{1}\{P_0 \neq P^*\},$$

so that $Q_0 = 1$ indicates that the consumer's initial preference differs from her fully informed preference. Such a consumer would reverse her preferred setting if experience led her to learn their fully informed value v.

Let $d_0 \in \{0, 1\}$ denote the consumer's initial setting, where $d_0 = 1$ denotes personalized ads. Initial settings may differ from both initial and fully informed preferences. Define the initial setting–preference mismatch rate m as

$$m = \Pr(M_0 = 1), \qquad M_0 = \mathbf{1}\{d_0 \neq P_0\}$$

Information about current settings and how to change them. A consumer's initial setting may differ from her initial preference for several reasons. Consumers may be unaware that their current setting differs from their preference; they may know that a mismatch exists but not know how to change the setting; or they may know how to change it but find doing so sufficiently costly.

Our information treatment addresses the first two frictions simultaneously. We reveal each consumer's current setting and provide a direct link to the page where it can be changed. A consumer with an initial mismatch ($M_0 = 1$) therefore learns both that a change is needed and how to make it. Under this interpretation, the only remaining obstacle is the residual cost of implementing the change, denoted by $C^I \geq 0$.

Let $Z^I \in \{0, 1\}$ denote randomized assignment to the information treatment, and let $S^I(z)$ indicate whether the consumer would switch under assignment z. For an initially mismatched consumer, the perceived gain from switching to the preferred setting is $w_0 = |\widetilde{v}_0|$. Under the information treatment, the consumer switches if and only if she is initially mismatched and this perceived gain is at least as large as the residual switching cost:

$$S^I(1) = M_0 \mathbf{1}\left\{|\widetilde{v}_0| \geq C^I\right\}.$$

We assume that consumers do not switch in the absence of treatment, $S^I(0) = 0$, and that the residual switching cost is the only reason an informed consumer does not switch. Because treatment assignment is randomized, the observed switching rate among treated consumers with an initial mismatch identifies

$$s = \Pr(S^I(1) = 1 \mid M_0 = 1) = \Pr(S^I = 1 \mid Z^I = 1, M_0 = 1).$$

The quantity s is the fraction of initial mismatches corrected by the information treatment. Thus, ms is the share of all consumers whose mismatch is corrected, while $m(1-s)$ is the share who remain mismatched after receiving the information treatment:

$$m = \underbrace{ms}_{\text{corrected by information treatment}} + \underbrace{m(1-s)}_{\text{remaining after information treatment}}.$$

This is an accounting decomposition of the initial mismatch rate. Under our maintained assumptions, the remaining mismatches are attributed to residual switching costs. By contrast, the corrected mismatches cannot be attributed to inattention alone, because the information treatment both reveals the consumer's current setting and reduces the cost of locating where to change it.

Learning through experience. Let $e \in \{0, 1\}$ denote randomized assignment to experience either the preferred or non-preferred ad setting, respectively. Following assignment to experience state e, the perceived value updates to:

$$\widetilde{v}_1(e) = \widetilde{v}_0 + \lambda_e\Big(v - \widetilde{v}_0\Big), \qquad 0 \le \lambda_0 \le \lambda_1 \le 1,$$

with endline preference

$$P_1(e) = \mathbf{1}\{\widetilde{v}_1(e) \ge 0\}.$$

The parameter λ_e captures the extent to which experience moves the consumer's perceived valuation toward her fully informed valuation. We assume that experience with the non-preferred setting ($e = 1$) generates at least as much learning as experience with the preferred setting ($e = 0$).

We observe both the baseline valuation $\widetilde{v}_0$ and the realized endline valuation $\widetilde{v}_1(e)$, allowing us to measure the valuation update directly:

$$\Delta\widetilde{v}(e) = \widetilde{v}_1(e) - \widetilde{v}_0 = \lambda_e\Big(v - \widetilde{v}_0\Big).$$

Experience may affect both the direction and intensity of preferences. A preference reversal occurs only when the updated valuation crosses zero, while $\Delta\widetilde{v}(e)$ also captures changes in preference intensity that leave the preferred setting unchanged. If consumers learn nothing from experiencing the same setting, then $\lambda_0 = 0$. If two weeks of exposure to a different setting are sufficient for full learning, then $\lambda_1 = 1$, and the endline valuation of the switchers reveals the fully informed utility v, up to measurement error.

Decomposition. For each experience state $e \in \{0, 1\}$, define the preference-revision indicator

$$R(e) = \mathbf{1}\{P_1(e) \neq P_0\},$$

and the revision rates among consumers who were initially aligned and mismatched:

$$r_A(e) = \Pr(R(e) = 1 \mid M_0 = 0), \qquad r_M(e) = \Pr(R(e) = 1 \mid M_0 = 1).$$

An initially mismatched consumer remains mismatched if she does not revise her preference, while an initially aligned consumer becomes mismatched if she does. Therefore, the following identity holds for either experience state:

$$\mathbf{1}\{d_0 \neq P_1(e)\} = M_0(1 - R(e)) + (1 - M_0)R(e).$$

Let $G(e) = \Pr(d_0 \neq P_1(e))$ denote the gap between the *initial* setting and the potential endline preference under experience state e. Using the same baseline population and the decomposition of m induced by the information treatment,

$$G(e) = ms + m(1 - s) + (1 - m)r_A(e) - mr_M(e).$$

The first two terms simply partition the initial mismatch rate m. Preference updating creates new gaps among initially aligned consumers, captured by $(1 - m)r_A(e)$, and resolves gaps among initially mismatched consumers, captured by $mr_M(e)$.

If there is no learning under same-setting experience, so that $\lambda_0 = 0$ and $\widetilde{v}_1(0) = \widetilde{v}_0$, then $P_1(0) = P_0$, $r_A(0) = r_M(0) = R(0) = 0$, and $G(0) = m$. If different-setting experience generates full learning, so that $\lambda_1 = 1$ and $\widetilde{v}_1(1) = v$, then $P_1(1) = P^*$, and $G(1)$ measures the gap between initial settings and fully informed preferences.

Causal decomposition using the randomized groups. The key empirical objects are the initial mismatch rate m, the information-treatment switch rate s, and the preference-revision rates in each group $R(e)$. Because we observe both initial alignment M_0 and initial preferences P_0, we can estimate these update rates separately within each baseline cell:

$$u_{ap}(e) = \Pr\Big(R(e) = 1 \mid M_0 = a, P_0 = p\Big), \qquad a, p \in \{0, 1\}.$$

The alignment-specific preference-revision rates are the corresponding weighted averages across baseline preference cells:

$$r_A(e) = \sum_{p\in\{0,1\}} u_{0p}(e)\Pr(P_0 = p \mid M_0 = 0),$$
$$r_M(e) = \sum_{p\in\{0,1\}} u_{1p}(e)\Pr(P_0 = p \mid M_0 = 1).$$

The endline gap is therefore

$$G(e) = m(1 - r_M(e)) + (1 - m)r_A(e)$$

The first term is the share of initial mismatches that remain, while the second is the share of initially aligned consumers who become mismatched.

To measure the causal contribution of experience, define the randomized treatment-control contrasts

$$\Delta r_A = r_A(1) - r_A(0), \qquad \Delta r_M = r_M(1) - r_M(0).$$

The gap in the absence of different-setting experience is

$$G(0) = ms + m(1 - s) + (1 - m)r_A(0) - mr_M(0).$$

Assuming that the information and experience mechanisms do not interact, the gap under the experience treatment admits the following signed causal decomposition:

$$\begin{aligned} G(1) = & \underbrace{ms}_{\text{corrected by information treatment}} + \underbrace{m(1-s)}_{\text{remaining after information treatment}} \\ & + \underbrace{(1-m)r_A(0) - mr_M(0)}_{\text{Control-group preference updating}} + \underbrace{(1-m)\Delta r_A - m\Delta r_M}_{\text{causal Experience effect}}. \end{aligned} \tag{1}$$

The first two terms provide an accounting decomposition of the initial mismatch rate. Under our maintained assumptions, the second term is attributed to residual switching costs. The third term captures preference updating that would also occur in the Control group. The final term captures the additional effect of exposure to a different setting: such experience creates mismatches among some initially aligned consumers while resolving mismatches among some initially misaligned consumers. The switching-cost component identifies the share of consumers deterred from switching by residual switching costs, rather than the monetary magnitude of those costs.

If different-setting experience fully reveals the value of personalization, $\lambda_1 = 1$, the causal Experience effect can be interpreted as the contribution of errors in consumers' initial valuations. Under incomplete learning, it instead measures the causal contribution of experience-induced preference updating.

Welfare decomposition. The gap $G(e)$ measures how often settings and full-informed preferences differ, but not the welfare consequences of those mismatches: a mismatch is more costly for a consumer with a large $|v|$ than for one who is nearly indifferent. The fully informed welfare loss associated with any setting d is

$$\ell(d) = U(P^*) - U(d) = |v|\mathbf{1}\{d \neq P^*\}.$$

Because the setting is binary, we have

$$\mathbf{1}\{d_0 \neq P^*\} = M_0(1 - Q_0) + (1 - M_0)Q_0.$$

Taking expectations gives

$$\mathcal{L}_0 = \underbrace{\mathbb{E}[|v|M_0(1 - Q_0)]}_{\text{implementation loss}} + \underbrace{\mathbb{E}[|v|(1 - M_0)Q_0]}_{\text{preference-learning loss}}.$$

The first term arises when the consumer's initial preference coincides with her fully informed preference, but her setting fails to implement it. The second arises when her setting implements her initial preference, but that preference differs from her fully informed preference. If $M_0 = Q_0 = 1$, these two errors offset: because the choice is binary, a setting that differs from an incorrectly signed preference coincides with the fully informed optimum and therefore generates no initial welfare loss.

We can further connect the implementation component to the information treatment. Since $S^I(1) \leq M_0$, we can partition initial welfare loss as:

$$\begin{aligned}\mathcal{L}_0 = & \underbrace{\mathbb{E}\Big[|v|S^I(1)(1 - Q_0)\Big]}_{\text{loss removable by information treatment}} \\ & + \underbrace{\mathbb{E}\Big[|v|\{M_0 - S^I(1)\}(1 - Q_0)\Big]}_{\text{residual implementation loss}} \\ & + \underbrace{\mathbb{E}[|v|(1 - M_0)Q_0]}_{\text{preference-learning loss}}. \end{aligned} \tag{2}$$

This is the welfare analogue of the mismatch-rate decomposition in equation (1), with each

term weighted by its welfare stakes $|v|$. It also makes clear why switching is not necessarily welfare improving. Let $d^I = S^I(1)P_0 + \{1 - S^I(1)\}d_0$ denote the post-information setting. The gross welfare effect of the information treatment is

$$\begin{aligned}\Delta\mathcal{W}_I &= \mathbb{E}\Big[U(d^I) - U(d_0)\Big] \\ &= \underbrace{\mathbb{E}\Big[|v|S^I(1)(1-Q_0)\Big]}_{\text{correct switches}} - \underbrace{\mathbb{E}\Big[|v|S^I(1)Q_0\Big]}_{\text{switches away from the optimum}} .\end{aligned}$$

Information increases welfare when it helps consumers implement an initial preference that coincides with their fully informed preference, but it reduces welfare when it facilitates implementation of an incorrectly signed preference. If C^I represents a real resource or utility cost rather than a purely behavioral wedge, the expected realized switching cost, $\mathbb{E}[S^I(1)C^I]$, must also be subtracted from this gross welfare effect.

The next section presents our empirical findings and the key treatment effects that we use to quantify the value of ad personalization and the roles of preference learning and switching frictions. A full mapping between the model and the data, together with an evaluation of the welfare effects, is work in progress.

4 Results

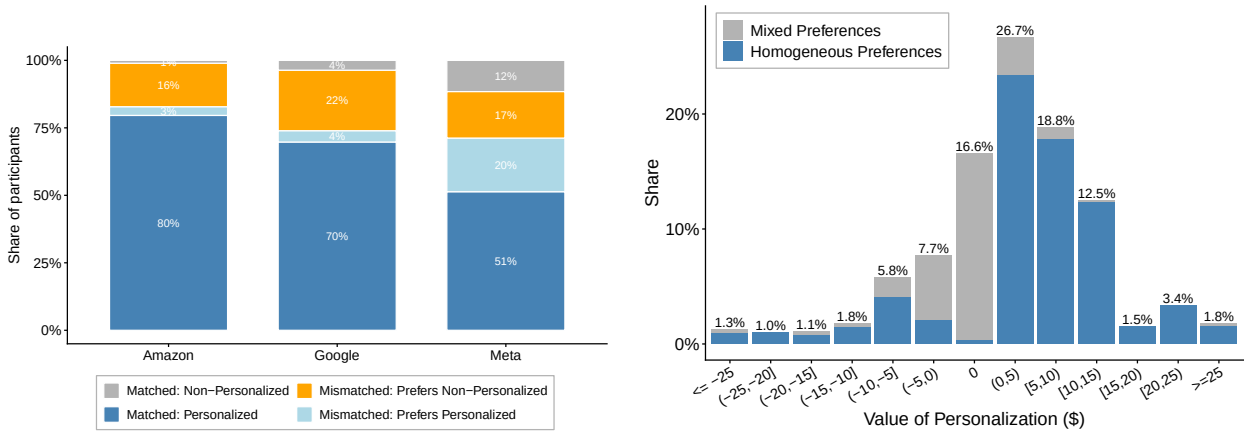
This section presents our main empirical results. We first document participants' baseline preferences and settings, and then examine the effects of experience and information separately.

4.1 Baseline Preferences and Settings

We begin by documenting participants' baseline preferences and ad settings, as well as the mismatches between them. Figure 2a compares participants' preferences for ad personalization with their actual settings on Amazon, Google, and Facebook before the experimental intervention.[10] Overall, participants are more likely to have personalized ad settings than to prefer personalized advertising, suggesting that some users who would prefer non-personalized ads do not have their preferred configuration. On Amazon, 83% of participants state a preference for personalized advertising, while 96% have personalized settings. A similar pattern emerges on Google, where 74% prefer personalized advertising but 92% have personalized settings. Facebook is an exception: 71% of participants prefer personalized

[10]For an additional breakdown of preferences by platform, see Appendix Figure A.2.

advertising, while 68% have personalized settings.



(a) Alignment Between Preferences and Ad Personalization Settings at Baseline

(b) Distribution of Valuations for Personalization

Figure 2: Preferences and Valuations for Personalized Advertising

Notes: Panel (a) compares participants' stated preferences for personalized advertising with their observed ad personalization settings on Amazon, Google, and Facebook at the beginning of the study. Each bar reports the share of participants whose preferences and settings are aligned or misaligned. We were unable to observe baseline ad personalization settings for 1.8% of participants on Amazon, 5.3% on Google, and 9.7% on Facebook. Panel (b) reports the distribution of participants' willingness to accept (WTA) compensation to switch advertising settings. Positive values indicate a preference for personalized advertising, while negative values indicate a preference for non-personalized advertising. The bars distinguish participants with homogeneous and mixed preferences across the three platforms. WTA was elicited only from participants assigned to the control or experience treatments; Panel (b) therefore includes only the 712 participants in those treatment groups. The corresponding post-intervention distribution is reported in Appendix Figure A.3.

Figure 2a also shows the extent to which participants' preferences align with their observed settings. Alignment is highest on Amazon, where 81% of participants have settings consistent with their preferences, largely because both preferences and settings tend to favor personalization. On Google, 74% of participants have settings that match their preferences. Meta exhibits the greatest disagreement, since 37% of participants have settings inconsistent with their initial preferences. Unlike on Amazon and Google, these mismatches are approximately symmetric on Meta, with similar shares of participants preferring more personalization than they receive and less personalization than they receive.

Participants are often uncertain or mistaken about their current ad-personalization settings. Across participant-platform pairs, 19.8% report being unsure of their current settings. Among those who provide a definite response, 74.2% correctly identify their setting. More importantly, among participant-platform pairs for which the actual setting does not match the participant's stated preference and the participant provides a definite response, 45% mistakenly believe that their preferred setting is already in place. These consumers therefore

have no perceived reason to switch. The remaining 55% correctly recognize the mismatch but have not corrected it, a pattern consistent with a role for implementation frictions.[11]

Figure 2b reports the distribution of participants' valuations for ad personalization before any experimental intervention. Positive values indicate an overall preference for personalized ads, while negative values indicate an overall preference for non-personalized ads across Amazon, Google, and Meta. The blue portions correspond to participants whose preferences are consistent across all three platforms: 59.8% prefer personalized advertising on all three platforms, while 11.3% prefer non-personalized advertising on all three. The gray portions correspond to the 28.9% of participants with mixed preferences across platforms. For these participants, the plotted valuation is the difference between their WTA to switch to a fully non-personalized configuration and their WTA to switch to a fully personalized configuration.[12]

The figure shows that most participants value personalized advertising, consistent with the preference patterns in Figure 2a. At the same time, the intensity of these preferences is generally modest. More than 96% of participants are willing to switch to their less-preferred advertising configuration for less than $25, the predetermined payment used in our experiment. As a result, the experience treatment successfully induced most participants to adopt their non-preferred ad setting. Finally, 16.7% of participants are indifferent between the fully personalized and fully non-personalized configurations. Virtually all of these participants report the same WTA for switching to the two configurations, rather than a WTA of zero for both.

4.2 Learning through Experience

To assess whether experience with settings that differ from a participant's stated preferences causes her to update those preferences, we restrict the sample to participants assigned to the experience and control groups. The theoretical framework motivates three dimensions of preference updating: whether a participant changes her stated preference for personalization; whether experience shifts her valuation, on average, toward or away from personalization; and whether experience changes the intensity of her preferences. We therefore construct three outcome variables.

[11] Although beliefs about settings were elicited before participants were shown their actual settings, for consistency we define baseline stated preferences throughout the analysis using preferences reported after this revelation. Only 3.6% of participants revised their stated preferences, suggesting that the revelation primarily corrected beliefs about existing settings rather than changing preferences, and that this timing distinction is quantitatively unimportant.

[12] For participants with mixed preferences, one of the two bundles—fully personalized or fully non-personalized—was selected at random for the experimental intervention.

First, we define Switch_i as an indicator equal to one if participant i changes her stated preference on at least one of the three platforms:

$$\text{Switch}_i \equiv \mathbf{1}\left\{P_{ik1} \neq P_{ik0} \text{ for at least one } k \in \mathcal{K}\right\}, \tag{3}$$

where $\mathcal{K} = \{\text{Amazon}, \text{Google}, \text{Facebook}\}$.

Second, we measure the signed change in the participant's valuation of personalization across the three platforms:

$$\Delta v_i \equiv \widetilde{v}_{i1} - \widetilde{v}_{i0}. \tag{4}$$

A positive value indicates updating toward greater perceived value from personalization, while a negative value indicates updating toward greater value from non-personalization.

Third, we measure the change in the intensity of the participant's preference:

$$\Delta|v_i| \equiv |\widetilde{v}_{i1}| - |\widetilde{v}_{i0}|. \tag{5}$$

A positive value indicates that the participant's preference becomes stronger, regardless of whether she prefers personalized or non-personalized advertising. We cap both valuation-based outcomes at the 1st and 99th percentiles of their respective distributions to limit the influence of extreme changes.[13]

In our first set of results, we define treatment in terms of exposure to the participant's non-preferred ad setting. As discussed in Section 2, experiencing the non-preferred setting does not necessarily imply experiencing a setting different from the participant's pre-experiment setting. Some participants entered the experiment with settings that differed from their stated preferences, so assignment to the experience treatment would leave their settings unchanged. Later, we consider an alternative treatment measure that captures whether the participant actually experienced a change in her advertising setting.

For each outcome Y_i, we first estimate the reduced-form specification

$$Y_i = \alpha + \beta Z_i + \varepsilon_i, \tag{6}$$

where Z_i indicates assignment to the experience treatment. The coefficient β therefore captures the intent-to-treat effect of assignment to the experience treatment, which is designed to expose participants to their non-preferred setting. We then estimate the effect of actually experiencing the non-preferred setting, using treatment assignment Z_i as an instrument for realized exposure. Under the standard IV assumptions, this specification identifies

[13] Capping the outcomes at -25 and 25, or excluding outliers altogether, yields similar results.

the effect of experiencing the non-preferred setting for participants whose exposure is induced by treatment assignment. All regressions are estimated at the participant level with heteroskedasticity-robust standard errors.

The first three columns of Table 3 report intent-to-treat effects, while the final three columns report the corresponding IV estimates. Assignment to the experience treatment substantially increases preference updating. It raises the probability that a participant revises her preference on at least one platform by 7.8 percentage points, relative to a control-group mean of 36.5 percent.[14] The IV estimate indicates that actually experiencing the non-preferred setting increases the probability of revising at least one platform preference by 8.1 percentage points. Both estimates are statistically significant at the 5% level.

Table 3: Effect of Non-Preferred Setting Experience on Preference Updating

Estimator		OLS			IV	
Outcome	Switch	Value update	Intensity update	Switch	Value update	Intensity update
	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Treatment	0.078**	2.39***	4.16***			
	(0.037)	(0.875)	(0.719)			
Experience				0.081**	2.49***	4.34***
				(0.038)	(0.910)	(0.741)
Constant	0.365***	-0.284	-0.160	0.365***	-0.284	-0.160
	(0.027)	(0.554)	(0.461)	(0.027)	(0.554)	(0.461)
Observations	712	712	712	712	712	712
R^2	0.006	0.010	0.043	0.006	0.010	0.043

Notes: The estimation sample includes participants assigned to the experience and control groups. The unit of observation is the participant. The switch outcome is an indicator equal to one if the participant changes her stated preference for ad personalization on at least one of Amazon, Google, or Facebook between the pre-intervention and endline preference elicitation. The value-update outcome is $\Delta v_i = \widetilde{v}_{i1} - \widetilde{v}_{i0}$, where positive values indicate updating toward greater value from personalization. The intensity-update outcome is $\Delta|v_i| = |\widetilde{v}_{i1}| - |\widetilde{v}_{i0}|$, where positive values indicate a strengthening of preferences. Both valuation outcomes are capped at the 1st and 99th percentiles of their respective distributions. Columns (1)–(3) report reduced-form estimates of the effect of assignment to the experience treatment. Columns (4)–(6) report IV estimates in which experience-treatment assignment instruments for realized exposure to the non-preferred setting. Heteroskedasticity-robust standard errors are reported in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Experience also changes participants' valuations of personalization. Assignment to the experience treatment increases the signed valuation of personalization by \$2.39 relative to

[14] Some participants in the control group also experienced changes in their settings because their existing settings did not match their stated preferences. The control-group switching rate may therefore partly reflect updating induced by these setting changes. We discuss the effect of experiencing an ad setting switch in Table 5.

the control group, indicating that participants update, on average, toward greater value from personalization. The corresponding IV estimate is similar, at \$2.49. Assignment also increases preference intensity by \$4.16, with a corresponding IV estimate of \$4.34. Thus, experience does more than induce participants to reconsider their platform-level preferences: it shifts valuations toward personalization on average and leads them to hold more intense preferences. The similarity between the reduced-form and IV estimates is consistent with the high rate at which treatment assignment induced actual exposure to the non-preferred setting.

Participant–Platform Outcomes. We next disaggregate the analysis to the participant–platform level. For each participant i and platform p, we construct analogues of the three participant-level outcomes. The first, Switch_{ip}, is an indicator equal to one if participant i changes her stated preference on platform p. The second is the signed valuation update, $\Delta v_{ip} = \widetilde{v}_{ip1} - \widetilde{v}_{ip0}$. As before, $\widetilde{v}_{ipt}$ is defined so that positive values indicate a preference for personalization: it equals the WTA to switch away from personalization when the participant prefers personalized advertising, and the negative of the WTA to switch away from non-personalization when she prefers non-personalized advertising. The third outcome is the change in platform-level preference intensity, $\Delta|v_{ip}| = |\widetilde{v}_{ip1}| - |\widetilde{v}_{ip0}|$. As above, the two valuation outcomes are capped at the 1st and 99th percentiles of their respective distributions.

We repeat the reduced-form and IV specifications described above at the participant–platform level. The pooled specifications include platform fixed effects, and standard errors are clustered at the participant level to account for correlation across the three platform observations contributed by each participant.

Columns (1)–(4) of Table 4 show that experience with the non-preferred setting increases preference updating at the platform level. Assignment to the experience treatment increases the probability of revising a platform-level preference by approximately 8.7 percentage points. The corresponding IV estimates are similar in magnitude. The estimates are statistically significant at the 1% level and are stable across specifications with and without platform fixed effects.

Columns (5)–(8) show that experience also shifts platform-level valuations toward personalization. Assignment to the experience treatment increases the signed valuation of personalization by \$1.88, indicating that participants update, on average, toward greater value from personalization. The corresponding IV estimate implies that actually experiencing the non-preferred setting increases the signed valuation of personalization by \$1.97. Both estimates are statistically significant at the 5% level.

Table 4: Platform-Level Effects of Non-Preferred Setting Experience on Preference Updating

Outcome	Preference revision				Signed value update				Preference-intensity update			
Estimator	OLS		IV		OLS		IV		OLS		IV	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Variables												
Treatment	0.087***	0.087***			1.88**	1.88**			5.47***	5.47***		
	(0.022)	(0.022)			(0.893)	(0.893)			(0.778)	(0.778)		
Experience			0.090***	0.090***			1.97**	1.97**			5.71***	5.71***
			(0.023)	(0.023)			(0.929)	(0.930)			(0.800)	(0.801)
Constant	0.171***		0.171***		0.054		0.054		0.096		0.096	
	(0.014)		(0.014)		(0.569)		(0.569)		(0.519)		(0.519)	
Platform FE		Yes		Yes		Yes		Yes		Yes		Yes
Observations	2,136	2,136	2,136	2,136	2,133	2,133	2,133	2,133	2,133	2,133	2,133	2,133
R^2	0.011	0.013	0.011	0.013	0.004	0.004	0.004	0.004	0.060	0.060	0.060	0.060

Notes: The estimation sample includes participants assigned to the experience and control groups. The unit of observation is a participant–platform pair. Columns (1)–(4) examine preference revision, defined as an indicator equal to one if the participant changes her stated preference on the platform between the pre-intervention and endline preference elicitations. Columns (5)–(8) examine the signed change in the participant's platform-level valuation of personalization, $\Delta v_{ip} = \widetilde{v}_{ip1} - \widetilde{v}_{ip0}$. Columns (9)–(12) examine the change in platform-level preference intensity, $\Delta|v_{ip}| = |\widetilde{v}_{ip1}| - |\widetilde{v}_{ip0}|$. The valuation outcomes are capped at the 1st and 99th percentiles of their respective distributions. Within each set of four columns, the first two report reduced-form estimates of the effect of assignment to the experience treatment, and the final two report IV estimates in which treatment assignment instruments for actual exposure to the non-preferred setting. The second and fourth columns within each outcome group include platform fixed effects. Standard errors are clustered at the participant level. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

Finally, columns (9)–(12) show a larger effect on preference intensity. Assignment to the experience treatment increases the absolute strength of platform-level valuations by \$5.47, while the corresponding IV estimate is \$5.71. Both estimates are statistically significant at the 1% level, and the estimates are essentially unchanged by the inclusion of platform fixed effects. Taken together, these platform-level results reinforce the participant-level findings: experience with the non-preferred setting makes participants more likely to revise their preferences, shifts valuations toward personalization on average, and leads them to hold more strongly differentiated preferences.

Appendix Table A.1 disaggregates these effects by platform. The effects on preference revision and preference intensity are broadly consistent across platforms: treatment assignment increases the probability of preference revision by 6.9 to 10.1 percentage points and increases preference intensity by approximately \$5.5 on each platform. The signed valuation effects in Panel B, however, show suggestive heterogeneity across platforms. Assignment to the experience treatment increases the value of personalization by \$2.52 on Amazon and \$2.10 on Google, compared with a much smaller and statistically insignificant estimate of \$1.03 on Meta. The corresponding IV estimates display the same pattern, at \$2.62, \$2.19, and \$1.08, respectively. Thus, the average shift toward personalization appears to be driven primarily

by updating on Amazon and Google. Appendix Figures A.4–A.6 show that these estimates remain virtually unchanged after adding demographic, browsing, and baseline setting-belief controls.

Effects of Platform-Level Setting Changes. We next use the participant–platform panel to estimate the effect of changing a participant's ad setting, rather than simply exposing the participant to her non-preferred setting. Let

$$S_{ip} \equiv \mathbf{1}\left\{d_{ip}^{E} \neq d_{ip0}\right\} \tag{7}$$

indicate that the setting implemented by the browser extension on platform p differs from participant i's baseline setting. For each of the three platform-level outcomes—preference revision, signed valuation updating, and change in preference intensity—we estimate

$$Y_{ip} = \beta \widehat{S}_{ip} + \gamma_p + \varepsilon_{ip}, \tag{8}$$

where S_{ip} is instrumented by assignment to the experience treatment, Z_i. The coefficient β therefore captures the effect of experiencing a change in the ad setting for participant–platform pairs whose setting change is induced by treatment assignment. All specifications include platform fixed effects, and standard errors are clustered at the participant level.

Column (4) of Table 5 shows a strong first stage: assignment to the experience treatment increases the probability of a platform-level setting change by 43.8 percentage points. The coefficient is below one because some participants entered the experiment with settings that differed from their stated preferences. For these participants, the experience treatment leaves settings unchanged, while the control group induces a change. Nevertheless, the large positive first-stage coefficient indicates that such cases are less common: overall, assignment to the experience treatment is substantially more likely to induce a change in settings than assignment to the control group.

The IV estimates indicate that a treatment-induced setting change increases the probability of revising the stated preference for that platform by 19.7 percentage points. It also increases the signed valuation of personalization by \$4.86, indicating that participants update, on average, toward greater value from personalization, and increases preference intensity by \$12.7. The effects on preference revision and intensity are statistically significant at the 1% level, while the signed valuation effect is significant at the 5% level. The magnitudes are all larger than those in Table A.1, as they capture the effect of actually experiencing a change in settings rather than simply experiencing the non-preferred setting. For participants whose

Table 5: Effect of Platform-Level Setting Changes on Preference Updating

Stage	Second stage			First stage
Outcome	Preference revision	Signed value update	Intensity update	Setting change
	(1)	(2)	(3)	(4)
Variables				
Setting change	0.197***	4.86**	12.7***	
	(0.053)	(2.04)	(1.88)	
Treatment assignment				0.438***
				(0.025)
Platform FE	Yes	Yes	Yes	Yes
Observations	2,014	2,011	2,011	2,014
R^2	0.013	0.006	0.062	0.191

Notes: The estimation sample includes participants assigned to the experience and control groups, and the unit of observation is a participant–platform pair. A setting change is an indicator equal to one if the setting implemented during the intervention differs from the participant's baseline setting on that platform. Columns (1)–(3) report second-stage IV estimates for, respectively, platform-level preference revision, the signed change in the valuation of personalization, and the change in preference intensity, using assignment to the experience treatment as an instrument for the setting change. The valuation outcomes are capped at the 1st and 99th percentiles of their respective distributions. Column (4) reports the corresponding first-stage effect of experience-treatment assignment on the probability of a platform-level setting change. All specifications include platform fixed effects. Standard errors are clustered at the participant level. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.

baseline settings already differ from their preferences, the latter may involve no change in their advertising environment and therefore little new experience.

Taken together, these results indicate that treatment-induced setting changes lead participants both to reconsider which setting they prefer and to hold more strongly differentiated preferences afterward.

4.3 Information Treatment

To assess whether participants in the information treatment changed their ad-personalization settings, we compare each participant's pre-intervention setting with all subsequent settings recorded by the browser extension. We classify a participant–platform pair as having switched if, at any point after the midline survey, the extension records a setting that differs from its pre-intervention setting. This definition is intentionally generous: a single observed deviation is sufficient to classify the pair as having switched. Even under this liberal definition, switching is rare. Among the 157 participants assigned to the information treatment, only 4% of participant–platform pairs ever exhibit a switch, corresponding to 8.9% of participants changing the setting on at least one platform.

Switching is concentrated among participant–platform pairs whose baseline settings are

misaligned with stated preferences. Among pairs with an initial mismatch, 9.4% switch settings during the study, compared with only 2.2% of pairs whose baseline settings are aligned with participants' preferences.[15]

Switching also varies substantially across platforms. Conditional on an initial mismatch between preferences and settings, participants are more likely to switch on Amazon (18.2%) than on Google (6.8%) or Facebook (6%). This pattern may reflect cross-platform differences in both the value participants place on personalized advertising and the frictions they face in changing their settings.

5 Conclusion

Observed ad personalization settings are an imperfect measure of what consumers want. We study demand for ad personalization in a three-week field experiment spanning Amazon, Google, and Meta. Combining observed platform settings with incentive-compatible preference elicitation and randomized exposure to information and non-preferred ad settings allows us to distinguish learning through experience from frictions that prevent consumers from implementing their preferred choices.

The results reveal substantial gaps between settings and fully informed preferences. Although most participants initially prefer personalized advertising, mismatch rates range from 19% to 37% across platforms. Exposure to non-preferred ad settings increases the probability of revising preferences by 8.1 percentage points and increases the valuation of the preferred setting by $4.34. On average, valuations shift toward personalization. By contrast, information alone produces limited switching: even after participants are shown their actual settings and given direct instructions for changing them, only 9.4% of initial participant–platform mismatches are corrected.

Taken together, these findings caution against interpreting a consumer's current advertising setting as either a stable preference or a fully informed welfare choice. Consumers face frictions in implementing the preferences they already hold, while experience with alternative advertising environments can change those preferences themselves. Effective choice environments may therefore need to facilitate both learning and implementation—for example, by making alternative settings easy to try, compare, and reverse, rather than relying solely on one-time disclosures or defaults.

Several remaining analyses will provide a more complete assessment of the effects of

[15]The vast majority of recorded switches (63.2%) occur on the same day that participants complete the survey containing the switching instructions. This timing provides reassurance that the observed switches were induced by the information intervention.

ad personalization. First, we will use the browser-extension data to examine how different personalization settings affect objective outcomes, including browsing behavior, platform engagement, advertising exposure, and downstream behavior. Second, we will map the empirical moments, including baseline setting–preference mismatches, switching in the information treatment, and experience-induced changes in preferences and valuations, to the theoretical framework. This exercise will allow us to quantify the relative importance of imperfectly informed preferences and implementation frictions. Finally, we will combine these estimates with participants' incentive-compatible valuations to decompose welfare losses into components attributable to preference learning and implementation frictions, and to evaluate the welfare effects of alternative ad settings and choice environments.

References

Agte, P., Allende, C., Kapor, A., Neilson, C., and Ochoa, F. (2024). Search and biased beliefs in education markets. Technical report, National Bureau of Economic Research.

Allcott, H., Castillo, J. C., Gentzkow, M. A., Musolff, L., and Salz, T. (Forthcoming). Sources of market power in web search: Evidence from a field experiment. *American Economic Review*.

Aridor, G., Goldberg, S., and Li, J. (2026). What is online privacy worth? evidence from an information provision experiment. Working paper.

Beknazar-Yuzbashev, G., Jiménez-Durán, R., Simonov, A., and Stalinski, M. (2025). Social media advertising loads as prices. *Columbia Business School Research Paper*, (5952337).

Bleier, A. and Eisenbeiss, M. (2015). Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science*, 34(5):669–688.

Brynjolfsson, E., Collis, A., Deisenroth, D., Garro, H., Kutzman, D., Liaqat, A., and Wernerfelt, N. (2025). The consumer welfare effects of online ads: Evidence from a nine-year experiment. *American Economic Review: Insights*, 7(4):447–462.

Collis, A., Moehring, A., Sen, A., and Acquisti, A. (2021). Information frictions and heterogeneity in valuations of social media data. *Available at SSRN 3974826*.

Deisenroth, D., Manjeer, U., and Wernerfelt, N. (2026). The consumer value of ad personalization: Evidence from meta. *Available at SSRN 6924098*.

Dickstein, M. J. (2021). Efficient provision of experience goods: Evidence from antidepressant choice.

Farronato, C. and Fradkin, A. (2026). Webmunk 2: A tool for studying online behavior and digital platforms. White paper.

Farronato, C., Fradkin, A., and Lin, T. (2026). Designing consent: Choice architecture and consumer welfare in data sharing. *Management Science*.

Gambato, J., Ganglmair, B., and Krämer, J. (forthcoming). Effective regulation and firm compliance: The case of german privacy policies. In Gong, R., Hotz, V. J., and Schmutte, I. M., editors, *Data Privacy Protection and the Conduct of Applied Research: Methods, Approaches and New Findings*. University of Chicago Press.

Goldfarb, A. and Tucker, C. (2011). Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404.

Johnson, G. A., Lewis, R. A., and Nubbemeyer, E. I. (2017). Ghost ads: Improving the economics of measuring online ad effectiveness. *Journal of Marketing Research*, 54(6):867–884.

Lin, T. (2022). Valuing intrinsic and instrumental preferences for privacy. *Marketing Science*, 41(4):663–681.

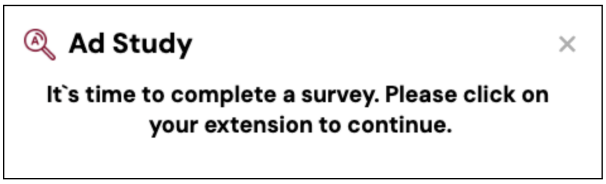
Sahni, N. S. and Zhang, C. (2024). Are consumers averse to sponsored messages? the role of search advertising in information discovery. *Quantitative Marketing and Economics*, 22(1):63–114.

Schnadower Mustri, E., Adjerid, I., and Acquisti, A. (2023). Behavioral advertising and consumer welfare. *Available at SSRN 4398428*.

Wernerfelt, N., Tuchman, A., Shapiro, B. T., and Moakler, R. (2025). Estimating the value of offsite tracking data to advertisers: Evidence from meta. *Marketing Science*, 44(2):268–286.

Yu, C. (2025). The welfare effects of sponsored product advertising. *Available at SSRN 4817542*.

A Additional Figures and Tables

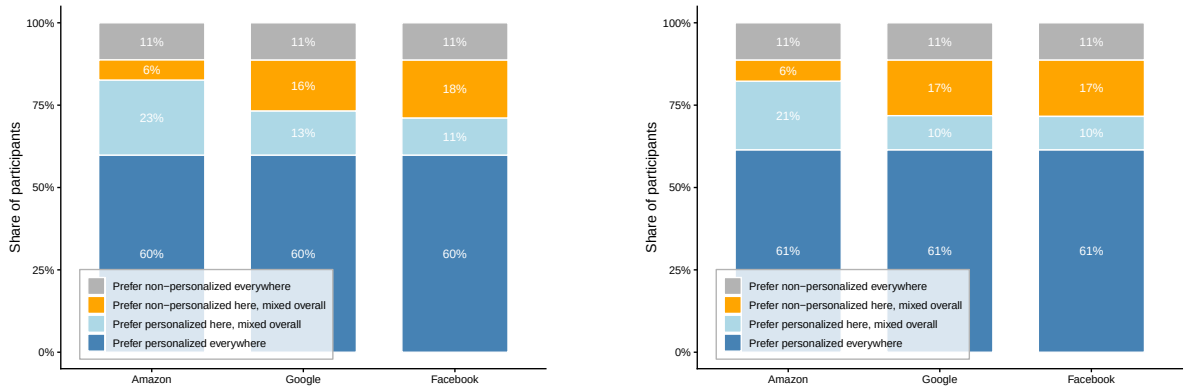


(a) Browser extension reminder prompting participants to complete the scheduled survey.



(b) Pop-up asking participants to evaluate the ads on the current webpage.

Figure A.1: Screenshots from the study browser extension. The extension periodically reminded participants to complete scheduled surveys and displayed brief pop-up questions during browsing to collect real-time measures of participants' advertising experience.

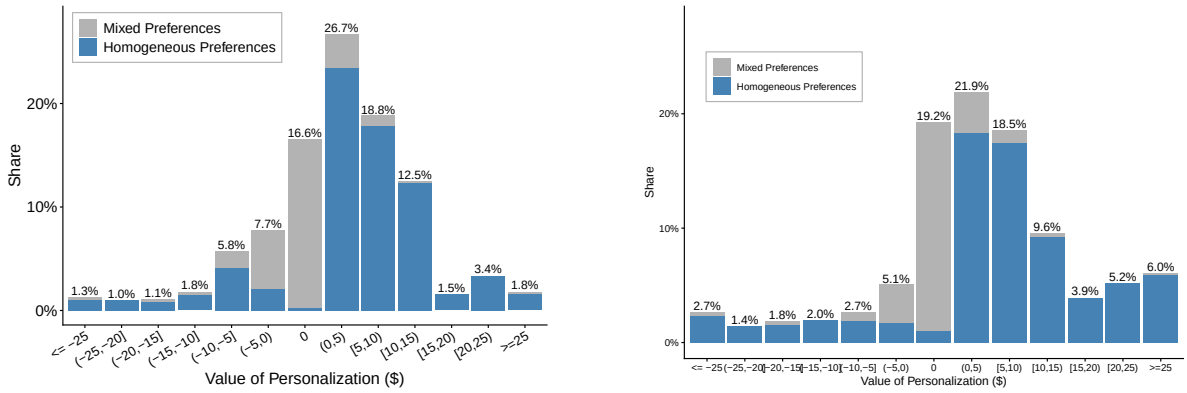


(a) Pre-intervention preferences

(b) Post-intervention preferences

Figure A.2: Preferences for Personalized Advertising Across Platforms

Notes: Each bar reports the share of participants who prefer personalized advertising on the indicated platform. The bars are decomposed according to participants' overall preference profiles across Amazon, Google, and Facebook. The dark blue segment corresponds to participants who prefer personalized advertising on all three platforms, while the gray segment corresponds to participants who prefer non-personalized advertising on all three platforms. The light blue and orange segments correspond to participants with mixed preferences, distinguishing those who prefer personalized advertising on the focal platform from those who prefer non-personalized advertising on the focal platform. Panel (a) reports preferences elicited before the experimental intervention, while panel (b) reports preferences elicited at the end of the study.



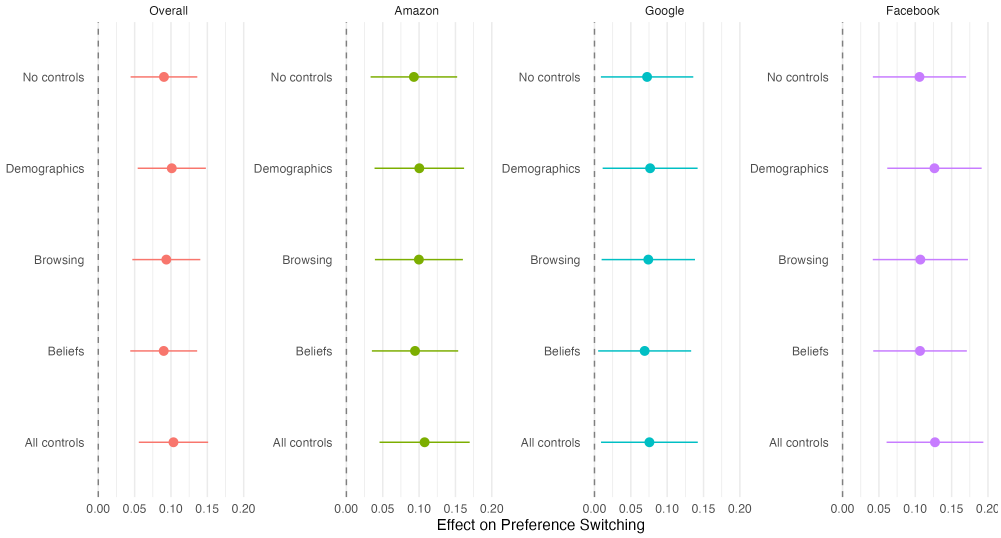
(a) Pre-intervention WTA

(b) Post-intervention WTA

Figure A.3: Distribution of Valuations for Personalized Advertising

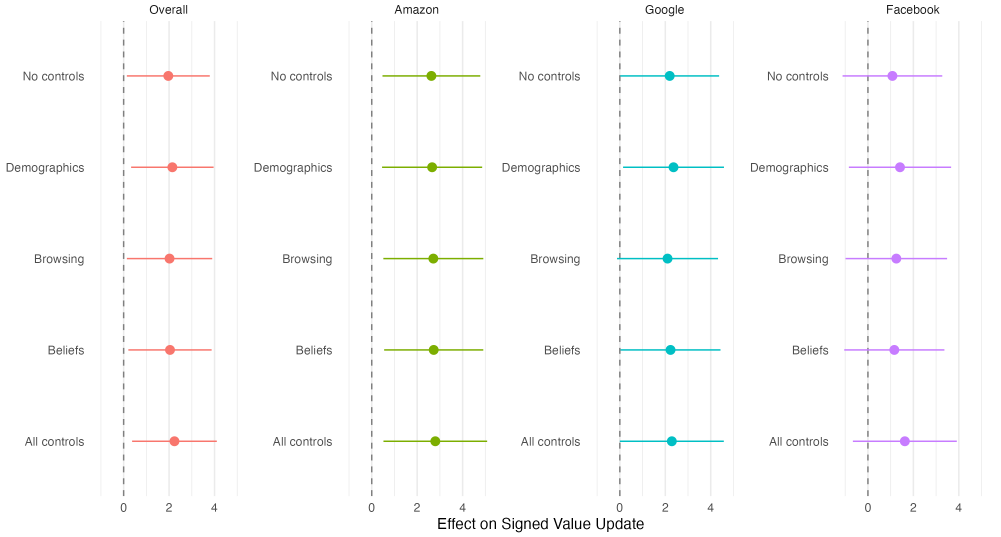
Notes: Each bar reports the share of participants in a given willingness-to-accept bin. Positive values indicate that participants prefer personalized advertising and would need to be compensated that amount to switch to non-personalized advertising (and vice versa for negative values). The bars are decomposed according to participants' overall preference profiles across Amazon, Google, and Facebook. The gray segment corresponds to participants with mixed preferences, and the blue segment corresponds to participants with homogeneous preferences (i.e., they prefer personalization or non-personalization across all three platforms). Panel (a) reports valuations elicited before the experimental intervention, while panel (b) reports valuations elicited at the end of the study.

Figure A.4: Robustness of the Effect of Experience on Platform-Level Preference Revision



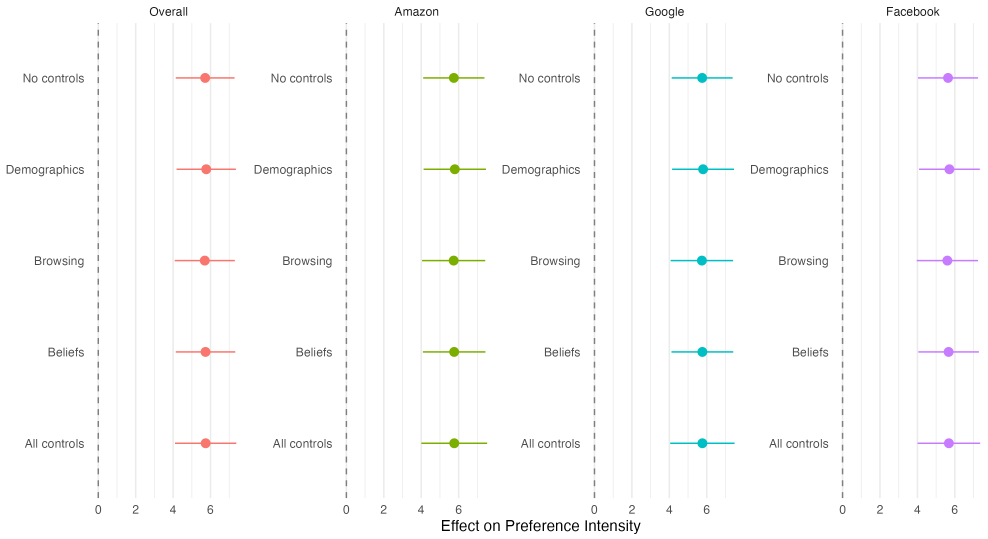
Notes: The figure reports IV estimates of the effect of experiencing the alternative setting on the probability of revising a platform-level preference. Experience-treatment assignment instruments for actual experience. The "Overall" estimates pool participant–platform observations and include platform fixed effects; the remaining estimates are obtained separately for Amazon, Google, and Facebook. The specifications sequentially include no controls, demographic controls for age group, sex, ethnicity, and employment status; browsing controls for the number of domains visited and total browsing time; baseline setting-belief controls indicating whether the participant correctly identified her setting on Amazon, Google, and Facebook; and all three sets of controls jointly. Error bars indicate 95 percent confidence intervals based on standard errors clustered at the participant level.

Figure A.5: Robustness of the Effect of Experience on Platform-Level Signed Value Updating



Notes: The figure reports IV estimates of the effect of experiencing the alternative setting on the signed change in participants' platform-level valuations of personalization. Otherwise, the figure is identical to Figure A.4.

Figure A.6: Robustness of the Effect of Experience on Platform-Level Preference Intensity



Notes: The figure reports IV estimates of the effect of experiencing the alternative setting on the change in participants' platform-level preference intensity. Otherwise, the figure is identical to Figure A.4.

Table A.1: Effect of Non-Preferred Setting Experience on Preference Updating By Platform

Panel A: Preference revision

Estimator		OLS			IV	
Platform	Amazon	Google	Meta	Amazon	Google	Meta
	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Treatment	0.089***	0.069**	0.101***			
	(0.029)	(0.031)	(0.031)			
Experience				0.093***	0.072**	0.106***
				(0.030)	(0.032)	(0.033)
Constant	0.144***	0.187***	0.181***	0.144***	0.187***	0.181***
	(0.019)	(0.022)	(0.021)	(0.019)	(0.022)	(0.021)
Observations	712	712	712	712	712	712
R^2	0.01265	0.00685	0.01416	0.01265	0.00685	0.01416

Panel B: Signed value update

Estimator		OLS			IV	
Platform	Amazon	Google	Meta	Amazon	Google	Meta
	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Treatment	2.52**	2.10**	1.03			
	(1.05)	(1.06)	(1.07)			
Experience				2.62**	2.19**	1.08
				(1.10)	(1.11)	(1.12)
Constant	-0.251	-0.122	0.536	-0.251	-0.122	0.536
	(0.702)	(0.696)	(0.684)	(0.702)	(0.696)	(0.684)
Observations	711	711	711	711	711	711
R^2	0.00776	0.00529	0.00125	0.00776	0.00529	0.00125

Panel C: Preference-intensity update

Estimator		OLS			IV	
Platform	Amazon	Google	Meta	Amazon	Google	Meta
	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Treatment	5.50***	5.52***	5.40***			
	(0.810)	(0.806)	(0.792)			
Experience				5.74***	5.75***	5.64***
				(0.834)	(0.829)	(0.817)
Constant	0.082	0.169	0.038	0.082	0.169	0.038
	(0.555)	(0.542)	(0.522)	(0.555)	(0.542)	(0.522)
Observations	711	711	711	711	711	711
R^2	0.05985	0.06028	0.05957	0.05985	0.06028	0.05957

Notes: The estimation sample includes participants assigned to the experience and control groups. Each column reports a separate regression for the indicated platform. In Panel A, the dependent variable is an indicator equal to one if the participant changes her stated preference on the platform between the pre-intervention and endline preference elicitations. In Panel B, the dependent variable is the signed change in the participant's platform-level valuation of personalization, $\Delta v_{ip} = \widetilde{v}_{ip1} - \widetilde{v}_{ip0}$. In Panel C, the dependent variable is the change in platform-level preference intensity, $\Delta|v_{ip}| = |\widetilde{v}_{ip1}| - |\widetilde{v}_{ip0}|$. The valuation outcomes are capped at the 1st and 99th percentiles of their respective distributions. Columns (1)–(3) report reduced-form estimates of the effect of assignment to the experience treatment for Amazon, Google, and Facebook, respectively. Columns (4)–(6) report the corresponding IV estimates, in which treatment assignment instruments for actual exposure to the non-preferred setting. Standard errors are clustered at the participant level. $^{***}p < 0.01$, $^{**}p < 0.05$, and $^{*}p < 0.10$.