Cookies and Shopping

Bo Bian[*], Michaela Pagel[†], Huan Tang[‡], and Emily Williams[§]

Abstract

This study analyzes the staggered adoption of cookie permission and compliance systems by US retailers in response to the European Union's data protection standards and the California Privacy Act. We combine information on when specific retailers implemented a cookie compliance system with individual-level bank and credit card transaction data from a US data aggregation and analytics provider. We first show that online spending decreases after retailers introduce cookie compliance systems. To address potential selection into who shops—and when—at retailers with enhanced privacy, we construct an enhanced-privacy exposure instrument using each consumer's pre-policy shopping basket shares and estimate reduced-form regressions that replace the endogenous post-policy indicator with this instrument. This reduced-form approach confirms the initial spending findings and further shows that individuals incur fewer overdraft, late, and cash-advance fees, roll over less credit card debt, and borrow less using other high-interest unsecured credit such as payday loans when treated with enhanced privacy. We discuss targeted advertising, third-degree price discrimination, and shopping convenience as the three main mechanisms driving these results. Our findings offer insights into the impact of data privacy regulations on consumers, informing policy considerations at both the state and federal levels.

Keywords: Data security, privacy regulation, web cookies, online shopping, financial well-being

JEL Codes: G5, G28, L86

[*]University of British Columbia, Sauder School of Business. E-mail: bo.bian@sauder.ubc.ca

[†]Washington University's Olin Business School, NBER, & CEPR. E-mail: mpagel@wustl.edu

[‡]University of Pennsylvania, The Wharton School & CEPR. E-mail: huan.ht.tang@gmail.com

[§]Harvard Business School & CEPR. E-mail: ewilliams@hbs.edu

We thank Hunt Allcott, Simona Abis, Tania Babina, Markus Mobius, David Rothschild, Steffen Meyer, Charline Uhr, and Matt Backus for valuable comments, as well as seminar participants at Barnard College, Aarhus University, UT Austin, OSU, UIUC, WashU Olin, Stockholm School of Economics, Federal Reserve Bank of Boston, HEC Paris, University of Heidelberg, DICE, Copenhagen Business School, University of Melbourne, Monash University, and the University of Missouri. We also thank Anton Tarasenko for providing outstanding research assistance, and Ansgar Walther, Antoine Uettwiller, and Tarun Ramadorai for sharing their data on cookies and web trackers.

1 Introduction

Over the past few decades, privacy has been significantly undermined by technological advancements and shifting societal norms. The ubiquity of digital devices and online services has led to unprecedented levels of data collection. At the same time, data storage capacity and analytical capabilities have advanced considerably. Improvements in individual identification and tracking have eroded anonymity, contributing to a complex interplay between convenience and privacy in the digital age.

Contributing to these developments was a significant acceleration in electronic commerce and retail sales, which was further amplified during the Covid-19 pandemic. Retail companies profit from consumers' private information by leveraging it for targeted advertising, personalized pricing, and operational optimization. These practices may benefit consumers by enhancing shopping convenience, but they may also harm consumers by enabling manipulation, exploiting behavioral vulnerabilities, and encouraging overspending. Additionally, third-degree price discrimination leads to a redistribution of surplus from consumers to producers and is widely perceived as unfair.

To protect user privacy and ensure transparent and secure online experiences, the most prominent data privacy regulations have specifically targeted web cookies—the primary tool used to track consumers' characteristics and behaviors online. The European Union passed the General Data Protection Regulation (GDPR) in May 2018, which required websites to obtain clear and informed consent from users before storing or accessing cookies on their devices. The California Consumer Privacy Act (CCPA) was enacted in January 2020 and established privacy standards for the collection of personal information, including data gathered through cookies.

In this paper, we analyze the staggered adoption of cookie compliance systems by US retailers that sought to comply with these two major privacy regulations. Our study holds direct policy relevance, particularly at the federal level in the United States, where regulatory bodies such as the Federal Trade Commission (FTC) and

the Consumer Financial Protection Bureau (CFPB) are currently considering new rules to safeguard consumer privacy and enhance data security.

Our objective is to evaluate and quantify the impact of strengthened data security standards on online shopping and individual financial well-being. To that end, we first examine how the introduction of cookie compliance systems—via cookie permission banners or pop-up windows—affects online shopping relative to offline shopping at the retailers implementing the privacy policy. We then assess whether the induced changes in spending translate into broader measures of individual financial well-being.

To obtain information on the adoption of cookie compliance systems, we collect data from BuiltWith, a web technology profiler. BuiltWith has gathered information about the technology stacks of websites through automated web crawling and code analysis since its launch in August 2007. To comply with cookie privacy regulations, retailers typically implement a cookie consent mechanism through third-party providers. BuiltWith detects the most commonly used of these providers: OneTrust, Cookiebot, and Osano. These companies' cookie compliance solutions are designed to align with the EU GDPR and CCPA. They automate cookie scanning, record user consent, and offer preference centers that allow users to manage their privacy settings. They also offer customizable banners, consent management platforms, and granular consent options.

The introduction of cookie compliance systems constitutes a policy change that enhances customer data privacy: some consumers decline cookies, which limits online retailers' ability to gather information and track them. According to survey data from Statista, approximately 50% of customers consistently reject cookie permission requests.[1] Additionally, the EU GDPR and CCPA explicitly mandate data minimization, which is implemented by the cookie compliance frameworks and enhances data security even for users who accept all cookies.[2] We provide direct evidence for this using actual cookie placement data obtained from Ramadorai et al. (2020).[3]

[1]See `https://www.statista.com/statistics/989896/users-consent-in-advertising-tracking/`.

[2]See `https://secureprivacy.ai/blog/understanding-cookie-compliance`.

[3]After matching their data with our data on firms' cookie compliance infrastructure dates by firms' URLs, we find significant negative effects for all cookie and tracking measures by retailers

We combine information on the adoption of cookie compliance systems with data on online and offline shopping from a transactions dataset provided by a prominent US data aggregation and analytics firm. This provider uses advanced data analytics to clean and categorize transaction data, which is then made available to researchers as well as institutional investors and investment managers in both aggregated and disaggregated forms. We use the de-identified transactions dataset, which includes all bank and credit card transactions—encompassing spending, income, and geographic information—for an unbalanced panel of approximately 50 million active consumers spanning January 2010 to August 2023.

We rely on merchant identifications provided by the data aggregator as well as their flag whether a transaction was physical (i.e., an in-person payment at a point-of-sale terminal) or non-physical (i.e., conducted remotely through digital or electronic means). With our large sample size and the substantial number of online retailers that implemented cookie compliance policies at different times, we have sufficient statistical power to detect even modest effects.

We first document that cookie permission banners lead to a reduction in online spending at the affected retailer, consistent with privacy policies diminishing shopping convenience and limiting retailers' ability to target advertisements or engage in price discrimination. Offline spending is reduced as well, though not as much, which is useful to quantify the targeted advertising channel relative to the online shopping convenience and price discrimination channels.

To address the concern that certain types of individuals select into shopping at retailers with enhanced data privacy at particular times—thereby confounding our results—we perform a reduced-form analysis that instruments for privacy exposure using each individual's pre-policy basket shares at each retailer. Specifically, we replace the endogenous post-policy indicator with a constructed enhanced-privacy exposure variable, yielding a reduced-form regression that identifies the effect of privacy exposure while avoiding selection bias. The validity of this approach rests on the exclusion restriction that pre-policy basket shares, interacted with the staggered timing of retailer adoption, affect spending only through their impact on individuals'

that implemented cookie compliance frameworks.

privacy environment. Several features of the design support this: the basket shares are fixed before any compliance adoption occurred, making them predetermined and immune to anticipation effects; adoption timing is driven by regulatory pressure from the GDPR and CCPA rather than retailer-specific demand shocks; and we show directly that compliance adoption dates are uncorrelated with the adoption of other website features that could independently shift the online-to-offline spending ratio. We confirm our initial spending results and also estimate a positive effect on financial well-being, as measured by overdraft fees, late fees, cash-advance fees, credit card interest payments, and utilization of other high-interest unsecured loans such as payday loans.

2 Conceptual Framework and Related Literature

2.1 Conceptual Framework

Our analysis centers on three channels through which cookie-based data collection affects consumer spending and financial well-being: targeted advertising, third-degree price discrimination, and shopping convenience. We discuss each in turn and derive the predictions we bring to the data.

Targeted advertising. Retailers use third-party cookies to build profiles of individual consumers and serve them personalized advertisements. A large advertising literature documents that targeted ads are substantially more effective than untargeted ones at inducing purchases (Johnson, 2013; Wernerfelt et al., 2022). By limiting cookie placement, cookie compliance frameworks reduce the precision with which retailers can target individual consumers, lowering the probability that any given ad induces a purchase. This channel predicts a reduction in online spending following the introduction of cookie compliance systems, primarily in product categories where persuasive advertising plays a larger role—typically discretionary goods rather than necessities. Crucially, targeted advertising is also the channel most likely to operate in the *offline* context: retailers use browsing data collected through cookies to run email campaigns, retargeting ads, and push notifications that drive traffic

to physical stores as well as websites. This means that the targeted advertising channel is the only one of the three with a cross-channel reach, and its contribution to the total spending effect should be visible in a decline in offline spending, not only online spending. The EU GDPR, which specifically targeted the use of third-party cookies for behavioral tracking and advertising, provides the most direct policy lever for this channel; US retailers who adopted compliance systems primarily in response to the GDPR are therefore particularly informative for identifying the advertising channel.

Third-degree price discrimination. Consumer data collected through cookies allows retailers to infer individual willingness to pay and, in turn, set personalized prices. Following Dubé and Misra (2023), who document large increases in firm profits and corresponding consumer surplus losses under machine-learning-based personalized pricing, we expect that reduced data availability lowers retailers' ability to price discriminate. The welfare implications are theoretically ambiguous: while some consumers who were previously charged high prices will benefit, the standard result under third-degree price discrimination is that total consumer surplus falls relative to uniform pricing (Bergemann et al., 2015). For inframarginal consumers—those who would purchase anyway—lower prices under a more competitive or uniform pricing environment could reduce average transaction values, potentially reducing measured spending even if quantities are unchanged. For marginal consumers who were induced to purchase only because of a tailored price offer, curtailing discrimination may reduce both the number of transactions and total spending. Empirically, we expect this channel to be more pronounced for high-value, discretionary purchases where willingness-to-pay heterogeneity is large, and for high-income consumers who were subject to the most aggressive personalized pricing. Unlike the advertising channel, price discrimination operates almost exclusively in the *online* channel, through the retailer's digital checkout and pricing engine, and should not generate a meaningful decline in offline spending. The CCPA's broader data minimization and data-sale restrictions, which constrain the data inputs to pricing algorithms, are particularly relevant for this channel.

Shopping convenience. Cookies also facilitate shopping convenience, for ex-

ample by enabling persistent shopping carts, personalized product recommendations, one-click checkout, and saved payment information. Reduced cookie use may therefore increase frictions in the online shopping experience, raising the effective cost of purchasing and lowering demand. Like price discrimination, this channel operates exclusively in the online context and should produce no decline in offline spending. Unlike price discrimination, the convenience channel predicts a broad reduction in online transaction volumes and values, with effects distributed relatively uniformly across transaction sizes and income groups, since the friction affects all online shoppers regardless of willingness to pay.

These three channels have a common prediction: the introduction of cookie compliance systems should reduce online spending relative to offline spending. Our difference-in-differences design, using offline spending at the same retailer as the control, is well suited to testing this joint prediction. To the extent that the channels differentially affect spending levels versus transaction counts, and discretionary versus non-discretionary categories, the category and transaction-size heterogeneity we explore in Section 4 can shed further light on their relative importance.

Finally, if reduced cookie use primarily curtails manipulation and overspending rather than genuinely welfare-improving transactions, we would expect improvements in downstream measures of financial health. Overspending induced by targeted advertising or supra-competitive prices should, on the margin, crowd out debt repayment and savings, increase the use of high-cost borrowing, and raise the likelihood of financial distress. Our financial well-being outcomes in Section 4 directly test this prediction.

2.2 Related Literature

Our paper contributes to several strands of the literature. We briefly survey each in turn; for comprehensive reviews of the economics of privacy, see Acquisti et al. (2016) and Goldfarb and Tucker (2019).

Theory of consumer data, targeted advertising, and price discrimination. Taylor (2004) develops a foundational model of consumer privacy in electronic

retailing, showing that the welfare effects of disclosing customer information depend critically on whether consumers anticipate data sales and on demand elasticity. Bergemann and Bonatti (2011) provide a model of targeting in advertising markets that characterizes the equilibrium implications of consumer data availability for the division of surplus between advertisers and publishers across online and offline media. Bergemann and Bonatti (2015) model the cookie market directly, characterizing when data intermediaries sell consumer information and how this affects advertiser profits and consumer welfare. Bergemann et al. (2015) derive welfare bounds on third-degree price discrimination, relevant to our price discrimination channel. Together, these theoretical papers motivate the three channels—targeted advertising, price discrimination, and shopping convenience—that we formalize in the conceptual framework above.

Value of consumer data to firms. Johnson (2013) shows that online retailers are willing to pay 52% more for ads with third-party cookies, consistent with a Google blog post (Bindra, 2019) reporting similar magnitudes. Wernerfelt et al. (2022) run a large-scale field experiment with hundreds of thousands of advertisers on Meta and estimate that losing access to offsite data raises the median cost per incremental customer by 37%, from \$43.88 to \$60.19. Todri (2022) finds that ad-blockers reduce total online consumer spending by approximately \$14.2 billion per year. Berman and Israeli (2021) exploit the staggered adoption of a retail analytics service across more than 1,000 e-commerce websites and find average revenue increases of 8–29% post-adoption. Dubé and Misra (2023) study personalized pricing at a large digital firm and find that it generates a 55% increase in profits relative to the firm's status quo, while reducing total consumer surplus by 23% relative to a uniform price. Their results quantify the surplus redistribution from consumers to firms that our reduced-form evidence is consistent with.

Effects of privacy regulation. Goldfarb and Tucker (2011) use responses from 3.3 million survey participants exposed to nearly 10,000 online banner advertising campaigns and find that EU privacy regulation substantially reduced advertising effectiveness. Goldberg et al. (2024) examine the GDPR's impact on website page views and revenue for over 1,000 online firms using Adobe Analytics data, finding

12% reductions in both metrics, with disproportionately larger effects on smaller sites. Aridor et al. (2021) show that the GDPR caused a 12.5% reduction in the number of consumers observable to online travel intermediaries, while raising the average value of trackable consumers to advertisers. Zhao et al. (2021) find evidence of increased search friction following the GDPR using consumer panels from four countries. Lukic et al. (2023) document that GDPR enforcement reduced the extent of online tracking. More broadly, the GDPR has been linked to changes in European web traffic (Goldberg et al., 2019), app market dynamics (Janssen et al., 2021), venture capital financing of technology firms (Jia et al., 2021), and firms' data monetization strategies (Aridor et al., 2020; Bessen et al., 2020; Peukert et al., 2021).

US privacy regulations and financial sector effects. Most closely related to our paper are two studies of US privacy regulations using consumer-level data. Bae et al. (2023) use a difference-in-differences design comparing Californians to non-Californians following the CCPA and find that Californians reduce purchases by 4.3%, increase returns by 3.0%, and reduce discretionary spending by $96. They also find evidence of increased search effort. Doerr et al. (2023) show that the CCPA increased loan applicants' willingness to share financial data with FinTech lenders and improved lender screening, enabling lower rates especially for underserved borrowers. Our paper complements these studies by focusing on retailer-level cookie compliance systems—a more granular treatment—and by directly linking privacy exposure to financial health outcomes. Babina et al. (2022) show that open banking regulations, which also govern consumer data sharing, spur FinTech investment. A related set of papers studies Apple's privacy initiatives: Bian et al. (2021) show that Apple's privacy labels reduced weekly app downloads by 14% and subscription revenue by 15% for iPhone apps; Kesler (2022) finds that the App Tracking Transparency (ATT) framework shifted app monetization toward paid models; and Bian et al. (2023) link consumer data privacy to the risk of financial fraud through data breaches.

3 Empirical Analysis

3.1 Privacy Policy Changes

Companies implement privacy policies and cookie permission systems to comply with evolving regulations and to address growing privacy concerns. To comply with the EU GDPR and CCPA in the domain of web cookies, companies partner with a cookie compliance provider—such as OneTrust—to help set up the infrastructure that controls the site's cookie policies. We obtained data on when exactly cookie compliance systems were implemented from BuiltWith, a web technology profiler that uses automated data analysis and web crawling to gather information about the technology features of websites. We collected the dates on which BuiltWith first detected the three most common cookie compliance providers: OneTrust, Osano, and CookieBot.

For example, McDonald's partnered with OneTrust, which was first detected by the BuiltWith crawler on August 9, 2019. Since then, the company maintains a privacy policy that outlines how it collects, uses, and protects consumer data. When a user visits the McDonald's website, they are asked to consent to the placement of cookies and similar tracking technologies. The specific cookie categories include:

- *Necessary Cookies:* Essential for basic website functionality; do not require explicit consent.
- *Performance Cookies:* Collect anonymized data on how users interact with the website, helping McDonald's improve site performance and user experience.
- *Functional Cookies:* Are cookies saved by a website to remember user preferences and settings, enhancing the browsing experience without tracking across sites.
- *Targeting Cookies:* Track users' browsing behavior to deliver personalized ads, promotions, and offers. Additionally, these cookies are the primary tools used by online retailers for consumer profiling by collecting data on users across websites to build interest profiles.

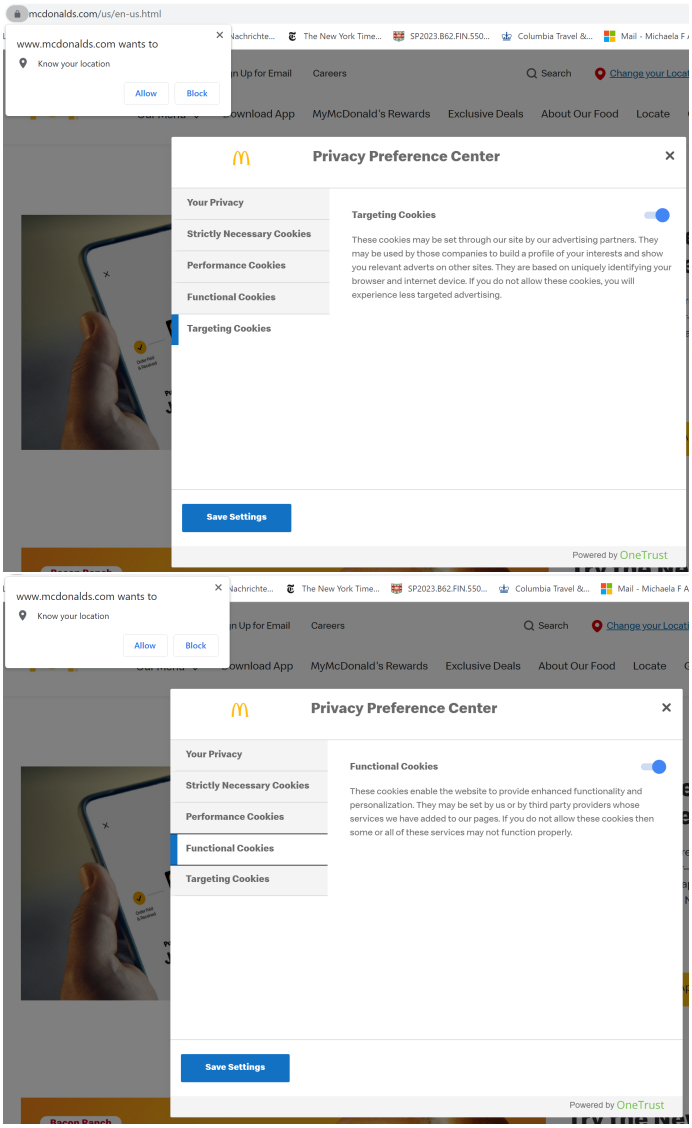
Users are presented with a cookie banner or pop-up on their first visit to the website and can manage their preferences by accepting or declining specific cookie categories. Figure 1 provides screenshots of McDonald's privacy center.

The premise of this paper is that when a retailer introduces a framework that makes its website compliant with privacy law, this constitutes a meaningful step toward enhancing customer privacy. Privacy enhancement can arise through two channels. First, some customers—anecdotally around 50% of users—reject cookies. Cookie compliance providers commit to withholding unnecessary cookies from these customers, thereby enhancing their privacy. Second, even among customers who accept cookies, the EU GDPR and CCPA explicitly mandate data minimization that the cookie compliance systems adhere to. Therefore, a retailer's investment in infrastructure compliant with privacy laws will result in less aggressive cookie placements and data collection overall.

We provide suggestive evidence for this using actual cookie placement data obtained from Ramadorai et al. (2020), who examine firms' data extraction practices and how they are shaped by privacy regulation. The authors collected firm-level data extraction measures using the "OpenWPM" crawler developed by Englehardt and Narayanan (2016). They shared with us data from two crawler runs: February 2019 and May 2023. We matched their data with our cookie compliance infrastructure dates by firms' URLs. We then regressed each cookie and tracking measure on firm fixed effects, crawling date fixed effects, and an indicator for whether a cookie compliance framework was in place at the time of the crawl (see Table 1). We find significant negative results for all cookie and tracking measures, reassuring us that the cookie compliance frameworks indeed enhanced customer privacy. Our measures include first-party (FP) cookies—primarily used for website functionality and showing less of a decline in response to the compliance framework—as well as third-party (TP) cookies and tracker requests, which are more substantially affected.

Figures 2 and 3 further illustrate the changes in all cookie and tracking measures before versus after firms introduced cookie compliance infrastructures. Figure 2 shows before-versus-after histograms for each tracking measure, and Figure 3 shows the distribution of within-firm differences for each cookie measure.

Figure 1: McDonald's Privacy Center



Notes: Screenshots of the cookie preference center displayed to users visiting the McDonald's website after the retailer adopted OneTrust as its cookie compliance provider (first detected by BuiltWith on August 9, 2019). The left panel shows the targeting cookies tab; the right panel shows the functional cookies tab. The interface illustrates how cookie compliance frameworks present users with granular consent options across cookie categories (Necessary, Performance, Functional, and Targeting).

3.2 Transaction-Level Bank Account Data

We combine our information on cookie compliance adoption dates at the retailer level with transaction-level data on individual bank and credit card accounts provided by a financial aggregation and analytics provider. For each individual, the dataset contains all income and spending transactions across a number of bank and credit card accounts, as well as information on overdraft fees, savings income, and other measures of financial well-being.

In this paper, we use bank account and credit card transaction data from a large U.S. account aggregator. Partnering with thousands of financial institutions and FinTech firms, they link multiple financial accounts to a single individual while stripping private individual information. While the overall sample includes 50 million individuals, we use a random subsample of about 5 thousand people for an unbalanced panel between 2010 and August 2023. This subsample size reflects data-access constraints; because the subsample is drawn uniformly at random from the full dataset, the demographic and geographic composition mirrors the broader population of account holders, and the large number of transactions per individual (Table 2 reports a mean of 123 transactions per user per month) provides ample within-individual variation to identify the effects of interest. Each observation consists of an individual identifier, a bank account identifier, a transaction amount, the day of the transaction, a merchant name, a description of the transaction, the zip code of the transaction location (if it is at a point-of-sale terminal), and a broad category name (i.e., "Groceries," "Rent," etc).

To look at financial well-being, we identify transactions that are overdrafts, non-sufficient fund fees, late fees, as well as cash-advance fees. We do so by manually searching within the description variable for keywords, such as "overdraft," "insufficient," "od fee," or "od item," to name a few. For overdraft, late, and non-sufficient fund fees, we restrict the relevant outgoing transactions to those below 100 USD. Given our interest in the spillover effects of bank policy changes on other indicators of financial well-being, we also do description searches for early wage access, payday loans, and other FinTech loans. Credit card interest charges are categorized by the

data provider as such.

3.3 Main Empirical Analyses

We first look at online and offline shopping after the dates of the cookie compliance policy implementation. Our regression specification is:

$$SpendPerTransaction_{i,m,t} = \alpha_i + \alpha_m + \alpha_t + \beta\, PostPolicy_{m,t} \times Online_{i,m,t} + \varepsilon_{i,m,t} \quad (1)$$

where i denotes individual, m denotes merchant, and t denotes day-of-sample. We use the raw transaction-level data, such that individual i may have zero, one, or multiple transactions at retailer m on day t. We control for individual, merchant, and day-of-sample fixed effects and cluster standard errors at the individual, merchant, and monthly levels. We also display results for a number of other fixed effects regimes, e.g., individual times online times day-of-sample fixed effects as well as individual times online and day-of-sample times merchant fixed effects.

Regression 1 can be run using just the raw data and results are qualitatively unchanged in that case. That said, in the interest of constraining the data to spending transactions of interest, we restrict to debit transactions in which the Retailer field is non-empty and the transaction categories are Electronics/General Merchandise, Personal/Family, Travel, Entertainment/Recreation, Restaurants, Groceries, Home Improvement, Utilities, Services/Supplies, Pets/Pet Care, Office Expenses, Gifts, Cable/Satellite/Telecom, and Automotive/Fuel. Additionally, to avoid long-term confounding variation biasing the estimate of β, the coefficient of interest, we restrict the sample to all observations that are transactions at retailers six months before to 13 months after they introduced their privacy policy. As the control group, we then use observations at all other retailers during those windows. This restricts the overall sample from about 20 to about 10 million observations.

Additionally, we winsorize the spending data at the 1% level and use both the logged values as well as the deviations from individual-level means as alternative outcome variables. Merchant fixed effects are included for merchants with at least

100 transactions in the cleaned spending dataset.

3.4 Selection and Privacy Exposure

The principal endogeneity concern in this analysis is that certain types of individuals may select into shopping at retailers with or without privacy policies at particular points in time. To address this selection problem, we construct a reduced-form privacy exposure variable—$EnhPrivacy_{i,t}$—using each individual's pre-policy basket shares at each retailer. This variable measures the share of a consumer's pre-period spending that is exposed to a cookie compliance framework at time t, and is orthogonal to the individual-level selection that confounds Specification 1. Crucially, Specification 2 is a *reduced-form regression*: it replaces the endogenous $PostPolicy_{m,t}$ indicator directly with the constructed exposure instrument, rather than instrumenting for it in a two-stage least squares framework. The coefficient β in Specification 2 therefore captures the intent-to-treat effect of privacy exposure on spending, and is directly interpretable as such. Its magnitude is not intended to be compared to the OLS coefficient in Specification 1, which estimates the average treatment effect conditional on the retailer having adopted a policy.

$$SpendPerTransaction_{i,m,t} = \alpha_i + \alpha_t + \alpha_m + \beta\, EnhPrivacy_{i,t} \times Online_{i,m,t} + \varepsilon_{i,m,t} \quad (2)$$

where i, m, and t denote individual, merchant, and day-of-sample as above. We again control for individual, merchant, and day-of-sample fixed effects and cluster standard errors at the individual, merchant, and monthly levels.

Validity of the instrument. The identifying assumption underlying Specification 2 is that the exposure instrument $EnhPrivacy_{i,t}$ affects spending only through its impact on individuals' privacy environment—the exclusion restriction. This requires that the cross-sectional variation in pre-policy basket shares, interacted with the time-series variation in which retailers have adopted cookie compliance, does not independently predict changes in online relative to offline spending for reasons unrelated to privacy. Several considerations support this assumption.

First, the basket shares used to construct $EnhPrivacy_{i,t}$ are fixed at the pre-policy period, before any cookie compliance adoption had occurred. They therefore reflect historical shopping patterns and cannot be endogenous to anticipated future compliance decisions by retailers. An individual who happened to shop more at retailers that later adopted early is not systematically different from one who shopped more at late adopters, because at the time of shopping neither individual knew which retailers would adopt first.

Second, the time-series variation in adoption dates is driven by regulatory pressure—specifically the GDPR and CCPA compliance deadlines—rather than by retailer-specific demand or supply shocks. The robustness check in Figures 13 and 14 confirms that compliance adoption dates are uncorrelated with the adoption of other website features that could independently affect online spending. This supports the view that the timing of adoption reflects regulatory compliance rather than concurrent operational changes.

Third, to the extent that the exclusion restriction could be violated, the most plausible threat is that retailers that adopted early also made other changes to their websites (e.g., redesigns, new payment options) that affected the online/offline spending ratio. The feature-date evidence directly addresses this concern. A less easily addressed threat is that early-adopting retailers were systematically different in their product offerings or pricing strategies in ways that coincided with changes in consumer demand. The individual fixed effects in Specification 2 absorb permanent individual-level differences in preferences for early-adopting retailers, and the day-of-sample fixed effects absorb common time trends. The residual identifying variation is therefore the interaction of an individual's pre-policy exposure weights with the staggered timing of adoption, which is unlikely to be correlated with unobservable demand shifts at specific retailers unless those shifts happened to occur precisely at the adoption date—which is tested and rejected by the flat pre-trends in the event-study figures.

As described above, we restrict spending outcomes to outgoing transactions with assigned retailers in assigned categories and to fixed windows six months before until 13 months after each retailer's privacy policy adoption.

We also aggregate the data to construct a balanced panel in which zeros can be assigned on days when individuals do not spend. Specifically, we aggregate to the individual × day × channel level, where channel indicates whether a transaction was online or offline. For each individual i on date t, we sum all spending and transaction counts separately for the online and offline channels, and fill in zeros for channel-days on which the individual had no transactions within their observed date range. The fraction of spending occurring at retailers that have already adopted a cookie compliance framework, $PostCookie_{i,t,c}$, is computed as a spending-weighted mean across transactions and is set to zero on zero-spending days. The privacy exposure instrument $EnhPrivacy_{i,t}$ is constructed from pre-policy basket shares and retailer adoption dates, and is carried forward on zero-spending days. We then estimate:

$$Spend_{i,t,c} = \alpha_i + \alpha_t + \beta\, PostCookie_{i,t,c} \times Online_{i,t,c} + \varepsilon_{i,t,c} \tag{3}$$

where i, t, and c denote individual, day-of-sample, and channel (online/offline). We control for individual and day-of-sample fixed effects and cluster standard errors at the individual and monthly levels.

Instead of using spending as an outcome in Regression 3, we can use measures of financial well-being, such as overdraft fees, cash-advance fees, and interest paid on credit card debt. For these outcomes, we aggregate to the individual × day level—rather than the individual × day × channel level used for spending—because financial health outcomes such as overdraft fees, credit card interest, and payday loan usage reflect the individual's overall financial position on a given day and are not naturally attributable to online versus offline spending. We fill in zeros for individual-days on which no such fee transaction occurred, within each individual's observed date range. Snapshots of checking account balances are available for part of the sample, and we use those to construct a daily balance by taking the snapshot as the starting point and summing up spending and income transactions before and after.

3.5 Summary Statistics

We conduct all analyses on a random subsample of 5,000 individuals drawn from the full dataset. Table 2 provides summary statistics of the transaction-level data.

Figure 4 shows the share of transactions occurring at retailers that already have privacy policies and cookie compliance systems in place, plotted over the sample period.

Figure 5 shows the average value of our enhanced privacy instrument over the sample period.

4 Results

4.1 Main Regression Results

Our analysis reveals that the introduction of cookie compliance systems leads to a significant decrease in consumer spending and, specifically, online spending. This is consistent with our proposed channels: restrictions on targeted advertising, consumer profiling, price discrimination, and shopping convenience should each exert downward pressure on spending.

As discussed, Table 3 uses the raw transactions sample and estimates Specifications 1 and 2. Specification 1 is the baseline OLS difference-in-differences; Specification 2 is the reduced-form regression that replaces the post-policy indicator with the pre-policy-basket-share exposure instrument. The key advantage of Specification 2 over Specification 1 is that it addresses selection bias: $PostPolicy_{m,t}$ in Specification 1 is an endogenous indicator that switches on for the specific retailers where an individual chooses to shop post-adoption. Individuals who are already inclined to spend less online may disproportionately patronize retailers with stronger privacy practices, biasing the OLS estimate toward zero (or away from zero if privacy-conscious individuals are heavy spenders). The exposure instrument $EnhPrivacy_{i,t}$ is constructed from each individual's pre-period basket shares and is therefore predetermined and orthogonal to post-adoption changes in shopping behavior. It varies across individuals because people shopped at different retailers before any policy was

introduced, and it varies over time because retailers adopted cookie compliance at different dates. Critically, the instrument can be thought of as a Bartik-style shifter: the cross-sectional weights (pre-policy basket shares) are fixed before the policy and the time-series variation (which retailers have adopted by date t) is driven by regulatory pressure that is plausibly exogenous to any individual's spending decisions. We find spending decreases in both specifications, with the reduced-form estimates in Specification 2 serving as our preferred identification strategy. Table 4 contains the same analysis but adds more stringent fixed effects regimes.

Tables 5 and 6 break down the main regression results by merchant category. The negative effects are concentrated in retail spending on electronics, general merchandise, and travel. By contrast, categories such as groceries show little effect. This pattern is consistent with price discrimination and targeted advertising being more prevalent in discretionary rather than essential spending.

In order to assess whether the spending decreases are associated with improvements in financial health, we examine additional outcomes using the enhanced privacy instrument in Table 7. We find reductions in the incidence of overdraft fees, credit card payments, rolled-over credit card debt (as measured by interest charges), and the prevalence of other high-interest unsecured borrowing.

These results are consistent with two distinct welfare mechanisms that our design cannot fully disentangle. The first is a *behavioral overspending* channel: targeted advertising exploits cognitive biases and present-bias to induce purchases consumers later regret, which crowd out debt repayment and savings. Under this interpretation, the financial health improvements reflect the removal of a manipulation that was systematically harmful to consumer welfare, and the welfare gains from reduced privacy intrusion are unambiguously positive. The second is a *price discrimination* channel: consumers who were previously charged supra-competitive personalized prices faced higher effective prices, reducing their purchasing power and increasing recourse to costly credit. Under this interpretation, the financial health improvements reflect the pass-through of lower prices to debt-constrained consumers. While the welfare implications of price discrimination are theoretically ambiguous—some consumers benefit from tailored discounts while others are overcharged—the financial

health evidence suggests that the net effect of data-enabled price discrimination on financially vulnerable consumers was harmful. Both mechanisms predict the same direction of effect on financial health outcomes, and the evidence is consistent with both operating simultaneously. What the results collectively establish is that the financial harm from lax data privacy standards was real and measurable, regardless of the precise channel through which it operated. Note that, in these regressions, we use the raw transaction data, where one row in the data is a transaction. To construct the outcome variable, e.g., overdraft fees, we define a new variable that is either positive, if the transaction is an overdraft fee payment, or zero, if not.

In turn, we aggregate the spending data to the individual $\times$ day $\times$ channel level, filling in zeros for channel-days on which an individual did not transact, as described in Section 3. Tables 8 and 9 present results analogous to the previous regressions on spending in this aggregated dataset. For financial well-being outcomes, we aggregate instead to the individual $\times$ day level, since fee incurrences reflect overall daily financial position rather than a specific spending channel; Table 10 presents these results.

4.2 Mechanism Evidence

The heterogeneity results presented in this subsection speak directly to which of the three channels—targeted advertising, price discrimination, and shopping convenience—drive the main spending effects. We exploit variation in income, spending intensity relative to income, and the regulatory era of adoption to generate differential predictions across channels.

4.2.1 Income Heterogeneity and the Price Discrimination Channel

Table 13 splits the raw transaction sample by tercile of regular income. The negative effects on online spending grow monotonically with income. This gradient is both economically large and statistically robust, and provides the most direct cross-sectional evidence in the paper for the price discrimination channel.

The logic is as follows. Retailers infer willingness to pay from consumer data—browsing history, past purchase behavior, geographic and demographic signals—and use this information to set personalized prices. Because willingness to pay is positively correlated with income, high-income consumers are precisely those who would be charged the highest personalized prices under a regime of data-enabled price discrimination. When cookie compliance frameworks limit the data available to retailers, their ability to identify and target high-willingness-to-pay consumers is curtailed, and the largest reductions in spending should therefore occur among those who were being most aggressively price-discriminated against. This is exactly what we find.

The income gradient is harder to reconcile with the other two channels. The shopping convenience channel has no natural prediction for an income gradient: the friction imposed by a cookie banner or the loss of convenience features does not vary systematically with income. The targeted advertising channel, if anything, might predict effects in the opposite direction—lower-income and more financially constrained consumers may be more susceptible to persuasive advertising and impulse purchases induced by personalized promotions, which would imply larger effects at the bottom of the income distribution rather than the top. The fact that we find the opposite—with effects growing sharply with income—points specifically to price discrimination as the channel driving the income heterogeneity, and suggests that the welfare costs of data-enabled pricing were disproportionately borne by higher-income consumers who were systematically overcharged relative to what they would have paid under uniform pricing. Table 12 confirms this income gradient on the aggregated panel with zeros, which additionally captures the extensive margin of spending.

4.2.2 Spending Intensity and Mechanism Evidence

Table 16 splits the raw transaction sample by tercile of transaction size—a proxy for spending intensity at the transaction level. The effects are absent in the smallest transaction-size tercile and concentrated in medium- and large-value transactions. Table 11 confirms this pattern on the aggregated panel, splitting instead by each individual's mean spending relative to mean income—a cleaner individual-level mea-

sure of spending intensity—and similarly finds effects growing monotonically with spending intensity, with the largest effects concentrated among the highest-spending-intensity tercile and near-zero effects in the lowest tercile.

This pattern is informative about which channels are operative. The shopping convenience channel predicts effects that are roughly uniform across consumer types: the friction imposed by a cookie compliance banner, or the loss of saved preferences such as a persistent shopping cart or one-click checkout, raises the cost of transacting online irrespective of how much an individual typically spends relative to their income. The near-zero effect in the lowest spending-intensity tercile is therefore difficult to reconcile with shopping convenience as the primary operative channel.

The price discrimination and targeted advertising channels, by contrast, both predict effects concentrated among high-spending-intensity consumers. For price discrimination, individuals who spend more relative to their income are likely higher-value customers for retailers and therefore face more aggressive personalized pricing. Curtailing the data that enables this pricing should produce the largest spending reductions precisely among these consumers. For targeted advertising, high-spending-intensity individuals may be more susceptible to advertising-induced purchases that push spending above what their income would otherwise support—meaning that removing targeted advertising has a disproportionately large effect on their spending. The concentration of effects in the highest spending-intensity tercile is therefore consistent with both the price discrimination and targeted advertising channels, and provides evidence against shopping convenience as the primary mechanism. This interpretation is reinforced by the category results in Tables 5 and 6, where effects are concentrated in discretionary categories such as electronics, general merchandise, and travel.

4.2.3 GDPR-Era versus CCPA-Era Adoptions and the Targeted Advertising Channel

Table 14 splits the sample by tercile of the sample period. The effects are statistically insignificant in the earliest time tercile, when very few retailers had yet adopted

compliance systems, and grow substantially in the second and third terciles, consistent with the staggered nature of the treatment and the concentration of compliance adoptions after 2019. Table 15 restricts attention to the post-2019 subsample and finds similarly large effects.

The timing of this growth is informative about the relative importance of the targeted advertising and price discrimination channels. The EU GDPR, enforced from May 2018, was directed almost exclusively at cookie consent for behavioral tracking and targeted advertising—it is the primary regulatory instrument behind the advertising channel. US retailers who adopted compliance systems in the earliest period were therefore doing so predominantly in response to GDPR pressure, meaning that early-period variation in our treatment primarily activates the advertising channel.

In contrast, the CCPA, enacted in January 2020, has a broader mandate: in addition to cookie consent, it restricts the sale and sharing of personal data more generally, which constrains the data inputs that feed personalized pricing algorithms and data broker markets. Retailers adopting in the later periods were therefore increasingly responding to both sets of requirements, layering the price discrimination channel on top of the advertising channel.

Under this interpretation, the growing effect size across time terciles reflects in part the accumulation of the price discrimination channel as CCPA-era compliance adds data restrictions that curtail personalized pricing beyond the advertising restrictions already in place from GDPR-era compliance. This interpretation is consistent with the income gradient result in Table 13: if the price discrimination channel grows in importance over time, one would expect effects among high-income consumers—who face the most aggressive personalized pricing—to be disproportionately concentrated in the later period. We acknowledge that growing effect sizes over time could also partly reflect the expanding reach of the instrument as more retailers adopt compliance systems, which increases the variation available for identification.

4.3 Dynamics

We examine the dynamics of online and offline spending around the retailer-specific cookie compliance introduction dates. We first use the raw transaction data with winsorized, logged, or individual-mean-demeaned spending as the outcome, including individual, day-of-sample, and their interaction as fixed effects. We also present results on the aggregated individual × day × channel panel with zeros. The results are presented in Figures 6–12.

Each event-study figure plots two series. The black line (labeled "Baseline") traces the evolution of offline spending in event time, with the month of policy introduction normalized to zero. The blue line (labeled "Interaction with online") plots the estimated interaction coefficient between the event-time indicator and the online dummy: it therefore captures the *differential* change in online spending relative to offline spending in each event-time bin. A negative interaction coefficient does not mean that online spending is negative in absolute terms; it means that online spending declined by *more* than offline spending following the introduction of the cookie compliance framework.

The offline baseline falls—and what this implies for the channels. Importantly, the baseline (offline) series itself trends downward after the policy introduction. This finding warrants careful interpretation and provides direct evidence on the relative importance of the three channels discussed in Section 2. Of the three channels, only the targeted advertising channel has a natural mechanism for reducing offline spending: retailer cookies are used not only to personalize the online shopping experience, but also to run cross-channel email campaigns, retargeting ads, and push notifications that drive consumers to physical stores. When cookie compliance frameworks limit data collection, these advertising tools are curtailed across all channels, reducing store traffic as well as website visits. The price discrimination and shopping convenience channels, by contrast, are essentially online-only in their operation: personalized pricing operates through the online checkout process and retailers' dynamic pricing algorithms, while convenience features (saved carts, one-click checkout, personalized product feeds) are features of the digital interface.

Neither has a direct mechanism for reducing offline spending.

This asymmetry is empirically useful. Since targeted advertising is the primary channel that operates in the offline context, the decline in the baseline offline series provides a rough lower bound on the aggregate contribution of the targeted advertising channel. The interaction coefficient—the differential decline in online relative to offline—then captures the *additional* effects that are concentrated in the online channel: specifically, the price discrimination and shopping convenience channels. Our identification strategy uses offline spending as the control to absorb common time trends, which is valid under the parallel trends assumption in the pre-period. However, the post-policy decline in offline spending means that the interaction coefficient *understates* the total effect of cookie compliance on online spending: the full impact on online spending equals the interaction coefficient (the additional online-specific decline) plus the baseline offline decline (the advertising component that affects both channels equally). The estimated interaction effects should therefore be interpreted as lower bounds on the total impact of cookie compliance on online spending. This lower-bound interpretation requires, however, that the observed offline decline is attributable entirely to the advertising channel rather than to substitution from online to offline shopping within the same retailer; to the extent that any offline decline reflects such substitution, the lower-bound logic would be partially offset.

The pre-policy coefficients (months -6 through -1) are small in magnitude and show no systematic trend for either the baseline or the interaction series, consistent with parallel pre-trends. The online and offline spending paths evolve similarly in the months leading up to the policy introduction, lending credibility to the identifying assumption that, absent the cookie compliance intervention, online and offline spending would have continued to evolve in parallel.

4.4 Robustness

To address the concern that companies introduced cookie compliance systems concurrently with other website features that independently affect online shopping, we show that the adoption dates of other common website features are uncorrelated

with the adoption of cookie compliance systems. Figure 13 shows this for features that would likely affect online shopping, and Figure 14 for features less likely to do so. All subplots share the same axis ranges; if the adoption dates were correlated, we would expect mass along the 45-degree line, which is not observed.

The sample splits by income tercile, sample period, and spending intensity relative to income — which serve both as robustness checks and as mechanism evidence — are presented and discussed in Section 4.2. The spending-intensity and income patterns hold up on both the raw transaction sample (Tables 13 and 16) and the aggregated panel with zeros (Tables 12 and 11).

4.5 Placebo

As a placebo exercise, we randomly reshuffle the cookie compliance dates among retailers that implemented cookie compliance frameworks, and rerun our main dynamic regression analyses using the aggregated specification with winsorized, as well as logged spending and individual-mean-demeaned spending as outcome variables. In this placebo exercise, shown in Figures 15 to 17, we do not find a decreasing pattern after the policy implementation.

5 Discussion and Conclusion

Our findings raise important considerations regarding the trade-off between personalized marketing and consumer privacy. While targeted advertising and price discrimination can increase online retailers' profits, our results suggest they also impose substantial economic costs on consumers through overspending and financial distress.

Our analysis shows that the introduction of cookie permission banners reduces online spending at the affected retailer. To address potential selection bias, we construct a privacy exposure instrument based on each individual's pre-policy shopping basket shares and estimate reduced-form regressions that confirm the direction and significance of the spending effects while avoiding the endogeneity of shopping at

privacy-compliant retailers. These specifications further reveal a positive impact on financial well-being, reflected in lower overdraft fees, higher checking account balances, reduced credit card interest payments, and decreased utilization of high-interest unsecured loans.

Several strands of the heterogeneity evidence shed light on which channels drive these results. The decline in offline spending in the event-study figures points to targeted advertising as an operative channel: it is the only mechanism with a plausible cross-channel reach, and its contribution to the total effect is reflected in the falling baseline. The concentration of effects in medium- and large-value transactions in the raw data (Table 16), confirmed on the aggregated panel (Table 11), are more consistent with price discrimination and targeted advertising than with shopping convenience, which would predict roughly uniform effects across consumer types. The monotonic increase in effect size across income terciles provides the most direct evidence for the price discrimination channel specifically: high-income consumers, who face the highest personalized prices, experience the largest reductions in spending when those prices are curtailed; this gradient is established on the raw transaction sample (Table 13) and confirmed on the aggregated panel (Table 12). The comparison of GDPR-era and CCPA-era adopters further suggests that the advertising channel is most directly activated by the GDPR's cookie-consent requirements, while the broader data-minimization provisions of the CCPA may additionally restrict personalized pricing. Together, these patterns support the interpretation that all three channels are operative, with price discrimination and targeted advertising contributing more than shopping convenience.

Our research highlights the importance of data privacy regulation in safeguarding consumer interests. Existing regulations such as the GDPR and state-level privacy acts play a vital role in protecting users. Policymakers should also consider the potential benefits of federal-level regulations that enhance data security, with a view toward protecting consumers from overspending and the associated adverse financial well-being consequences. A promising direction for future work is a formal decomposition of effects by the regulatory trigger of each retailer's adoption—distinguishing GDPR-driven from CCPA-driven compliance—which would provide sharper evidence

on the relative contributions of the advertising and price discrimination channels and may better inform which regulatory instruments are most effective.

References

Acquisti, A., C. R. Taylor, and L. Wagman (2016). The economics of privacy. *Journal of Economic Literature 54*(2), 442–492.

Aridor, G., Y.-K. Che, and T. Salz (2020). The economic consequences of data privacy regulation: Empirical evidence from GDPR. NBER Working Paper 26900, Columbia University, Massachusetts Institute of Technology.

Aridor, G., Y.-K. Che, and T. Salz (2021). The effect of privacy regulation on the data industry: Empirical evidence from gdpr. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pp. 93–94.

Babina, T., G. Buchak, and W. Gornall (2022). Customer data access and fintech entry: Early evidence from open banking. *Available at SSRN*.

Bae, D., R. Mayya, and T. Nian (2023). Privacy regulation and its unintended consequence on consumption behaviors: Evidence from ccpa.

Bergemann, D. and A. Bonatti (2011). Targeting in advertising markets: Implications for offline versus online media. *RAND Journal of Economics 42*(3), 417–443.

Bergemann, D. and A. Bonatti (2015). Selling cookies. *American Economic Journal: Microeconomics 7*(3), 259–294.

Bergemann, D., B. Brooks, and S. Morris (2015). The limits of price discrimination. *American Economic Review 105*(3), 921–957.

Berman, R. and A. Israeli (2021). The added value of data-analytics: Evidence from online retailers. Technical report, working paper.

Bessen, J. E., S. M. Impink, L. Reichensperger, and R. Seamans (2020). GDPR and the importance of data to AI startups. Working paper, New York University, Boston University.

Bian, B., X. Ma, and H. Tang (2021). The supply and demand for data privacy: Evidence from mobile apps. *Available at SSRN*.

Bian, B., M. Pagel, and H. Tang (2023). Consumer surveillance and financial fraud. Technical report, mimeo.

Bindra, C. (2019). Next steps to ensure transparency, choice and control in digital advertising.

Doerr, S., L. Gambacorta, L. Guiso, and M. Sanchez del Villar (2023). Privacy regulation and fintech lending. *Available at SSRN 4353798*.

Dubé, J.-P. and S. Misra (2023). Personalized pricing and consumer welfare. *Journal of Political Economy 131*(1), 131–189.

Englehardt, S. and A. Narayanan (2016). Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 1388–1401.

Goldberg, S., G. Johnson, and S. Shriver (2019). Regulating privacy online: The early impact of the gdpr on european web traffic & e-commerce outcomes. *Available at SSRN 3421731*.

Goldberg, S. G., G. A. Johnson, and S. K. Shriver (2024). Regulating privacy online: An economic evaluation of the GDPR. *American Economic Journal: Economic Policy 16*(1), 325–358.

Goldfarb, A. and C. Tucker (2019). Digital economics. *Journal of Economic Literature 57*(1), 3–43.

Goldfarb, A. and C. E. Tucker (2011). Privacy regulation and online advertising. *Management Science 57*(1), 57–71.

Janssen, R., R. Kesler, M. Kummer, and J. Waldfogel (2021). GDPR and the lost generation of innovative apps. NBER Working Paper 146409, University of Zurich,

University of Minnesota, University of East Anglia, Georgia Institute of Technology.

Jia, J., G. Z. Jin, and L. Wagman (2021). The short-run effects of the general data protection regulation on technology venture investment. *Marketing Science*, forthcoming.

Johnson, G. (2013). The impact of privacy policy on the auction market for online display advertising.

Kesler, R. (2022). The impact of Apple's app tracking transparency on app monetization. *Available at SSRN 4090786*.

Lukic, K., K. M. Miller, and B. Skiera (2023). The impact of the general data protection regulation (gdpr) on online tracking. *Available at SSRN*.

Peukert, C., S. Bechtold, M. Batikas, and K. Tobias (2021). Regulatory spillovers and data governance: Evidence from the GDPR. *Marketing Science*, forthcoming.

Ramadorai, T., A. Uettwiller, and A. Walther (2020). Privacy policies and consumer data extraction: Evidence from u.s. firms. *Available at SSRN 3352175*.

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics 35*(4), 631–650.

Todri, V. (2022). Frontiers: The impact of ad-blockers on online consumer behavior. *Marketing Science 41*(1), 7–18.

Wernerfelt, N., A. Tuchman, B. Shapiro, and R. Moakler (2022). Estimating the value of offsite data to advertisers on meta. *University of Chicago, Becker Friedman Institute for Economics Working Paper* (114).

Zhao, Y., P. Yildirim, and P. K. Chintagunta (2021). Privacy regulations and online search friction: Evidence from gdpr. *Available at SSRN 3903599*.

Tables

Table 1: Cookie Compliance Frameworks Reduce Cookie and Tracker Placement

	(1) FP cookies	(2) TP tracking cookies	(3) Tracker requests	(4) TP requests	(5) Cookies	(6) Total cookies
Post cookie policy	-1.283	-4.232***	-8.574***	-24.169***	-12.924***	-5.515***
	(0.954)	(0.717)	(1.741)	(6.230)	(2.559)	(1.496)
Mean of dependent variable	7.76	2.94	9.48	36.69	15.83	10.70
URL/firm FEs	✓	✓	✓	✓	✓	✓
Crawling date FEs	✓	✓	✓	✓	✓	✓
Observations	7,734	7,734	7,734	7,734	7,734	7,734
Adjusted R^2	0.63	0.54	0.56	0.51	0.52	0.59

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is a matched panel of retailer URLs observed in two web-crawling runs (February 2019 and May 2023) from Ramadorai et al. (2020), merged with the BuiltWith cookie compliance adoption dates from this paper. Each column reports an OLS regression of the indicated cookie or tracker measure on URL/firm fixed effects, crawling-date fixed effects, and a binary indicator equal to one if a cookie compliance framework (OneTrust, Osano, or CookieBot) was active at the time of the crawl. Standard errors are in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2: Summary Statistics

	Mean	Standard deviation	Median	5th percentile	25th percentile	75th percentile	95th percentile
Retail spending	52.02	330.7	21.30	2.490	9.590	49	172.5
Online spending	62.89	405.2	22.44	2	9.990	59.48	208.4
Restaurant spending	29.78	147.9	15.93	3.160	8.460	31.13	80.38
Grocery spending	39.12	141.8	19.11	1.650	7.400	42.13	130.1
Regular income	1588.8	2184.6	1182.5	83.12	500	2043	4064.3
Other income	1077.4	4182.0	166.9	3	35.25	748.3	3470
Overdraft fees	0.0537	1.574	0	0	0	0	0
Late fees	28.39	7.791	27	15	25	35	39
NSF fees	67.19	277.3	35	5	14.06	50.69	200
Cash advance fees	25.70	132.4	7.340	0.0900	1.510	22	100
Interest charges	54.51	68.96	29.44	0.630	8.610	80.22	178.3
Payday/EWA/FinTech loans	1140.2	3152.5	347.8	11.28	75	1153.1	4387.7
Transactions per user per month	123.4	103.4	107	28	69	156	261
Number of users per month	1584.5	919.1	1581	170	782	2396	3022
Observations (total)	9,395,105						

Notes: The sample consists of 5,000 randomly selected individuals from the full de-identified transactions dataset, covering January 2014 to August 2023. The unit of observation is a transaction. Spending and income variables are measured in dollars per transaction. Fee variables are measured in dollars per transaction occurrence; because the summary statistics are computed over the raw transaction-level data—where each row corresponds to an actual transaction—all fee observations are positive by construction (there are no zero-fee rows). The reported means therefore reflect the typical dollar amount of a fee conditional on it being charged. Payday/EWA/FinTech loans are measured in dollars. We restrict spending transactions to debit transactions in which the Retailer field is non-empty and the transaction categories are Electronics/General Merchandise, Personal/Family, Travel, Entertainment/Recreation, Restaurants, Groceries, Home Improvement, Utilities, Services/Supplies, Pets/Pet Care, Office Expenses, Gifts, Cable/Satellite/Telecom, and Automotive/Fuel. We restrict the sample to all observations that are transactions at retailers six months before to 13 months after they introduced their privacy policy. As the control group, we then use observations at all other retailers during those windows. Transaction counts and user counts are measured at the monthly level.

Table 3: Effect of Cookie Compliance on Spending: OLS and Reduced-Form Estimates

	(1) Spending winsorized	(2) Spending log	(3) Spending deviation from mean	(4) Spending winsorized	(5) Spending log	(6) Spending deviation from mean
Post cookie policy =1	2.729***	0.038***	0.076***			
	(0.442)	(0.009)	(0.014)			
Online =1	15.592***	0.243***	0.407***	15.713***	0.249***	0.408***
	(0.932)	(0.015)	(0.027)	(0.945)	(0.015)	(0.027)
Post cookie policy =1 × Online=1	-6.641***	-0.044**	-0.162***			
	(1.011)	(0.021)	(0.031)			
Enhanced privacy				0.395**	0.011***	0.009**
				(0.152)	(0.003)	(0.004)
Online =1 × Enhanced privacy				-1.550***	-0.024***	-0.029***
				(0.350)	(0.007)	(0.010)
Individual FEs	✓	✓	✓	✓	✓	✓
Day-of-sample FEs	✓	✓	✓	✓	✓	✓
Merchant FEs	✓	✓	✓	✓	✓	✓
Observations	4,288,970	4,288,970	4,288,970	4,288,970	4,288,970	4,288,970
Adjusted R^2	0.33	0.41	0.11	0.33	0.41	0.11

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample consists of 5,000 randomly selected individuals from the full de-identified transactions dataset. We then restrict to spending transactions in select retailer categories as described in the caption of Table 2. The sample period is restricted to fixed windows 6 months before until 13 months after each retailer's privacy policies, with control observations kept if they occured within those windows. Columns (1)–(3) estimate Specification 1: $SpendPerTransaction_{i,m,t} = \alpha_i + \alpha_m + \alpha_t + \beta\, PostPolicy_{m,t} \times Online_{i,m,t} + \varepsilon_{i,m,t}$. Columns (4)–(6) estimate Specification 2, replacing the endogenous post-policy indicator with the reduced-form privacy exposure instrument $EnhPrivacy_{i,t}$. The dependent variable is winsorized spending per transaction in Columns (1) and (4), log spending in Columns (2) and (5), and spending deviation from the individual mean in Columns (3) and (6). All specifications include individual, day-of-sample, and merchant fixed effects. Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 4: Effect of Cookie Compliance on Spending: Additional Fixed Effects Specifications

			(1) Spending winsorized	(2) Spending winsorized	(3) Spending winsorized	(4) Spending winsorized	(5) Spending winsorized
Post cookie policy		=1	0.126	0.479	-0.987*	-0.161	54.745
			(0.630)	(0.626)	(0.563)	(0.498)	(67.791)
Online		=1	0.000	0.000	0.000	0.000	0.000
			(0.000)	(0.000)	(.)	(.)	(0.000)
Post cookie policy	=1 × Online	=1	-13.932***	-14.412***	-10.772***	-13.902***	-40.666***
			(2.178)	(2.155)	(2.146)	(2.562)	(2.929)
Individual FEs			✓				
Day-of-sample FEs			✓		✓		
Merchant FEs			✓	✓	✓	✓	
Online FEs			✓	✓			
Individual times day-of-sample FEs				✓			
Individual times online FEs					✓		
Individual times online times day-of-sample FEs						✓	
Individual times online and day-of-sample times merchant FEs							✓
Observations			11,841,261	11,834,673	11,841,250	9,563,423	10,017,615
Adjusted R^2			0.41	0.43	0.42	0.49	0.43

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is as described in Table 3. This table replicates the main spending regressions under more stringent fixed effects regimes, including individual × online × day-of-sample fixed effects and individual × online combined with day-of-sample × merchant fixed effects, to absorb additional sources of unobserved heterogeneity. The dependent variable is winsorized spending per transaction in all Columns (1)–(5). Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Effect of Cookie Compliance on Spending by Merchant Category: OLS Estimates

		(1) Spending winsorized	(2) Spending winsorized	(3) Spending winsorized	(4) Spending winsorized	(5) Spending winsorized	(6) Spending winsorized
		Electronics general merchandise	Personal and family	Home improvement	Travel	Restaurants	Groceries
Post cookie policy	=1	3.784**	-3.119**	8.304	7.473***	-1.466***	4.455**
		(1.725)	(1.378)	(11.185)	(1.434)	(0.332)	(1.719)
Online	=1	12.033***	6.510***	65.677***	14.943***	13.601***	25.359***
		(2.043)	(1.792)	(7.219)	(4.351)	(2.068)	(7.428)
Post cookie policy	=1 × Online=1	-29.961***	-2.427	-3.742	-5.930***	-4.357**	6.371
		(4.885)	(3.388)	(11.587)	(1.533)	(2.096)	(5.751)
Individual FEs		✓	✓	✓	✓	✓	✓
Day-of-sample FEs		✓	✓	✓	✓	✓	✓
Merchant FEs		✓	✓	✓	✓	✓	✓
Observations		1,967,384	556,298	311,424	766,906	2,588,458	1,953,711
Adjusted R^2		0.13	0.34	0.30	0.25	0.29	0.31

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is as described in Table 3. Each column estimates Specification 1 on the subsample of transactions in the indicated merchant category. The dependent variable is winsorized spending per transaction in all Columns (1)–(6), with each column corresponding to a distinct merchant category: Column (1) Electronics/General Merchandise, Column (2) Personal and Family, Column (3) Home Improvement, Column (4) Travel, Column (5) Restaurants, and Column (6) Groceries. All specifications include individual, day-of-sample, and merchant fixed effects. Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6: Effect of Cookie Compliance on Spending by Merchant Category: Reduced-Form Estimates

	(1) Spending winsorized	(2) Spending winsorized	(3) Spending winsorized	(4) Spending winsorized	(5) Spending winsorized	(6) Spending winsorized
	Electronics general merchandise	Personal and family	Home improvement	Travel	Restaurants	Groceries
Enhanced privacy	3.020***	0.002	-1.475	0.755	-0.429**	0.065
	(0.569)	(0.833)	(2.543)	(0.694)	(0.203)	(0.299)
Online =1	11.666***	7.247***	65.301***	14.580***	13.930***	25.565***
	(1.919)	(1.814)	(7.280)	(4.342)	(2.177)	(7.819)
Online=1 × Enhanced privacy	-8.303***	-5.180***	1.525	-1.346	-2.333**	-0.770
	(1.031)	(1.103)	(2.748)	(0.854)	(1.125)	(2.453)
Individual FEs	✓	✓	✓	✓	✓	✓
Day-of-sample FEs	✓	✓	✓	✓	✓	✓
Merchant FEs	✓	✓	✓	✓	✓	✓
Observations	1,967,384	556,298	311,424	766,906	2,588,458	1,953,711
Adjusted R^2	0.13	0.34	0.30	0.25	0.29	0.31

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is as described in Table 3. Each column estimates Specification 2 on the subsample of transactions in the indicated merchant category, replacing the endogenous post-policy indicator with the privacy exposure instrument. The dependent variable is winsorized spending per transaction in all Columns (1)–(6), with each column corresponding to a distinct merchant category: Column (1) Electronics/General Merchandise, Column (2) Personal and Family, Column (3) Home Improvement, Column (4) Travel, Column (5) Restaurants, and Column (6) Groceries. All specifications include individual, day-of-sample, and merchant fixed effects. Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7: Effect of Cookie Compliance on Financial Well-Being: Reduced-Form Estimates

	(1) Overdraft fees	(2) Late fees	(3) NSF fees	(4) Cash advance fees	(5) Interest charges	(6) Payday/EWA /FinTech loans	(7) Checking account balance log
Enhanced privacy	-0.026***	-0.003***	-0.006	-0.004*	-0.054***	-0.277***	0.027**
	(0.004)	(0.001)	(0.007)	(0.002)	(0.007)	(0.097)	(0.012)
Mean of dependent variable	0.05	0.01	0.05	0.02	0.20	0.88	8.51
Individual FEs	✓	✓	✓	✓	✓	✓	✓
Day-of-sample FEs	✓	✓	✓	✓	✓	✓	✓
Observations	9,394,374	9,394,374	9,394,374	9,394,374	9,394,374	9,394,374	1,967,011
Adjusted R^2	0.02	0.00	0.00	0.00	0.02	0.01	0.83

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is as described in Table 3. Each column estimates Specification 2 with a financial well-being outcome as the dependent variable, replacing the endogenous post-policy indicator with the privacy exposure instrument. Financial well-being outcomes are measured as the dollar amounts of each incurrence and set to zeros for all non-occurrence transactions. The dependent variables are: Column (1) overdraft fees, Column (2) late fees, Column (3) NSF fees, Column (4) cash-advance fees, Column (5) credit card interest charges, Column (6) Payday/EWA/FinTech loan amounts, and Column (7) log checking account balance. Snapshots of checking account balances and, therefore, the constructed balances, are only available for part of the sample. All specifications include individual and day-of-sample fixed effects as indicated. Standard errors, in parentheses, are clustered at the individual and month levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 8: Effect of Cookie Compliance on Spending: Aggregated Panel with Zeros, OLS Estimates

	(1) Spending winsorized	(2) Number of transactions	(3) Spending deviation from mean	(4) Spending winsorized	(5) Number of transactions	(6) Spending deviation from mean
Post cookie policy	58.473***	2.050***	4.298***	58.440***	2.039***	4.243***
	(1.679)	(0.045)	(0.184)	(1.661)	(0.046)	(0.182)
Online =1	-14.093***	-0.472***	-0.842***	-14.076***	-0.472***	-0.840***
	(0.474)	(0.016)	(0.034)	(0.474)	(0.016)	(0.034)
Online=1 × Post cookie policy	-37.676***	-1.092***	-2.000***	-35.813***	-1.025***	-1.862***
	(2.448)	(0.082)	(0.335)	(2.467)	(0.083)	(0.337)
Individual FEs	✓	✓	✓			
Day-of-sample FEs	✓	✓	✓			
Individual times month-of-sample FEs				✓	✓	✓
Day-of-week FEs				✓	✓	✓
Week-of-month FEs				✓	✓	✓
Observations	7,490,641	7,490,641	7,488,349	7,489,073	7,489,073	7,486,950
Adjusted R^2	0.17	0.25	0.02	0.17	0.25	0.02

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is aggregated to the individual × day × channel (online/offline) level, filling in zeros for channel-days on which an individual did not transact within their observed date range. This allows the extensive margin of spending to be captured. The fraction of spending at post-cookie retailers ($PostCookie_{i,t,c}$) is set to zero on zero-spending days. Columns (1)–(3) include individual and day-of-sample fixed effects; Columns (4)–(6) replace day-of-sample fixed effects with individual × month-of-sample, day-of-week, and week-of-month fixed effects. The dependent variables are winsorized spending, number of transactions, and spending deviation from the individual mean. Standard errors, in parentheses, are clustered at the individual and month levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 9: Effect of Cookie Compliance on Spending: Aggregated Panel with Zeros, Reduced-Form Estimates

	(1) Spending winsorized	(2) Number of transactions	(3) Spending deviation from mean	(4) Spending winsorized	(5) Number of transactions	(6) Spending deviation from mean
Enhanced privacy	20.188***	0.688***	1.149***	20.726***	0.707***	1.191***
	(0.639)	(0.028)	(0.046)	(0.630)	(0.029)	(0.043)
Online =1	-12.787***	-0.421***	-0.778***	-12.775***	-0.421***	-0.777***
	(0.557)	(0.019)	(0.040)	(0.556)	(0.019)	(0.040)
Online=1 × Enhanced privacy	-15.297***	-0.564***	-0.820***	-15.114***	-0.560***	-0.804***
	(0.818)	(0.031)	(0.057)	(0.823)	(0.031)	(0.057)
Mean of dependent variable	22.97	0.57	1.50	22.97	0.57	1.50
Individual FEs	✓	✓	✓			
Day-of-sample FEs	✓	✓	✓			
Individual times month-of-sample FEs				✓	✓	✓
Day-of-week FEs				✓	✓	✓
Week-of-month FEs				✓	✓	✓
Observations	7,490,452	7,490,452	7,488,160	7,488,781	7,488,781	7,486,658
Adjusted R^2	0.17	0.26	0.02	0.17	0.26	0.02

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is aggregated as described in Table 8. Each set of columns estimates Specification 3 replacing the *PostCookie* indicator with the reduced-form privacy exposure instrument $EnhPrivacy_{i,t}$. Columns (1)–(3) include individual and day-of-sample fixed effects; Columns (4)–(6) replace day-of-sample fixed effects with individual × month-of-sample, day-of-week, and week-of-month fixed effects. The dependent variables are winsorized spending, number of transactions, and spending deviation from the individual mean. Standard errors, in parentheses, are clustered at the individual and month levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Effect of Cookie Compliance on Financial Well-Being: Aggregated Panel with Zeros

	(1) Overdraft fees	(2) Late fees	(3) NSF fees	(4) Cash advance fees	(5) Interest charges	(6) Payday/EWA/ FinTech loans	(7) Checking account balance log
Enhanced privacy	-0.174***	-0.009***	0.048	-0.007	-0.149***	-0.297	0.016
	(0.066)	(0.003)	(0.043)	(0.014)	(0.037)	(0.459)	(0.016)
Mean of dependent variable	0.70	0.04	0.24	0.11	0.80	2.93	8.65
Individual FEs	✓	✓	✓	✓	✓	✓	✓
Day-of-sample FEs	✓	✓	✓	✓	✓	✓	✓
Observations	4,612,695	4,612,695	4,612,695	4,612,695	4,612,695	4,612,695	4,476,017
Adjusted R^2	0.01	0.01	0.04	0.00	0.04	0.01	0.74

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Financial well-being outcomes are aggregated to the individual × day level (not the individual × day × channel level used for spending outcomes), because fee incurrences and credit usage reflect overall daily financial position rather than a specific spending channel. For each individual-day, all fee transactions of the indicated type are summed and a zero is filled in on days with no such transaction within the individual's observed date range. The dependent variables are: Column (1) overdraft fees, Column (2) late fees, Column (3) NSF fees, Column (4) cash-advance fees, Column (5) credit card interest charges, Column (6) Payday/EWA/FinTech loan amounts, and Column (7) log checking account balance. The checking account balance is aggregated by construction from account snapshots and running transaction totals; snapshots are available only for part of the sample. All specifications include individual and day-of-sample fixed effects. Standard errors, in parentheses, are clustered at the individual and month levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 11: Effect of Cookie Compliance on Spending by Income Tercile

	(1) Spending winsorized	(2) Spending deviation from mean	(3) Spending winsorized	(4) Spending deviation from mean	(5) Spending winsorized	(6) Spending deviation from mean
	1st income tercile		2nd income tercile		3rd income tercile	
Post cookie policy =1	-1.312*	-0.000	0.963	0.007	-0.620	-0.018***
	(0.743)	(0.006)	(0.945)	(0.006)	(1.379)	(0.006)
Online =1	35.863***	0.223***	48.784***	0.260***	59.844***	0.216***
	(2.276)	(0.014)	(3.378)	(0.015)	(3.783)	(0.013)
Post cookie policy=1 × Online=1	-8.203***	-0.051**	-14.115***	-0.058**	-22.199***	-0.043
	(2.639)	(0.025)	(2.834)	(0.023)	(4.676)	(0.038)
Median of dependent variable	49.73		57.01		76.20	
Individual FEs	✓	✓	✓	✓	✓	✓
Day-of-sample FEs	✓	✓	✓	✓	✓	✓
Merchant FEs	✓	✓	✓	✓	✓	✓
Observations	3,692,359	3,692,359	3,822,546	3,822,546	3,701,822	3,701,822
Adjusted R^2	0.40	0.21	0.40	0.24	0.47	0.12

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is as described in Table 3, split by tercile of average regular income. Each pair of columns estimates Specification 1 on the indicated income tercile subsample. Columns (1)–(2) report results for the lowest income tercile, Columns (3)–(4) for the middle tercile, and Columns (5)–(6) for the highest tercile. Within each pair, the odd column reports winsorized spending per transaction and the even column reports spending deviation from the individual mean. All specifications include individual, day-of-sample, and merchant fixed effects. Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 12: Effect of Cookie Compliance on Spending by Income Tercile: Aggregated Panel with Zeros

	(1) Spending winsorized	(2) Number of transactions	(3) Spending winsorized	(4) Number of transactions	(5) Spending winsorized	(6) Number of transactions
	1st income tercile		2nd income tercile		3rd income tercile	
Post cookie policy	49.067***	2.078***	58.464***	2.054***	66.596***	1.988***
	(2.330)	(0.072)	(2.129)	(0.066)	(3.137)	(0.073)
Online =1	-12.921***	-0.488***	-15.819***	-0.543***	-14.462***	-0.439***
	(0.617)	(0.022)	(0.731)	(0.025)	(0.839)	(0.025)
Online=1 × Post cookie policy	-28.324***	-1.123***	-37.938***	-1.031***	-38.826***	-0.900***
	(3.263)	(0.109)	(3.072)	(0.105)	(4.341)	(0.115)
Mean of dependent variable	18.69	0.53	25.27	0.64	27.12	0.61
Individual times month-of-sample FEs	✓	✓	✓	✓	✓	✓
Day-of-week FEs	✓	✓	✓	✓	✓	✓
Week-of-month FEs	✓	✓	✓	✓	✓	✓
Observations	2,188,451	2,188,451	2,460,015	2,460,015	2,259,319	2,259,319
Adjusted R^2	0.17	0.25	0.17	0.26	0.17	0.26

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is aggregated as described in Table 8 and split by tercile of average regular income. Each set of columns estimates Specification 3 on the indicated income tercile subsample. As in Table 13 above, the effects grow monotonically with income, consistent with the price discrimination channel. Columns (1)–(2) report results for the lowest income tercile, Columns (3)–(4) for the middle tercile, and Columns (5)–(6) for the highest tercile. Within each pair, the odd column reports winsorized spending and the even column reports the number of transactions. All specifications include individual and day-of-sample fixed effects. Standard errors, in parentheses, are clustered at the individual and month levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 13: Effect of Cookie Compliance on Spending by Transaction Size Tercile

		(1) Spending winsorized	(2) Spending log	(3) Spending winsorized	(4) Spending log	(5) Spending winsorized	(6) Spending log
		1st spent tercile		2nd spent tercile		3rd spent tercile	
Post cookie policy	=1	0.237***	0.026***	0.636***	0.027***	-1.291	0.002
		(0.042)	(0.007)	(0.076)	(0.003)	(2.254)	(0.008)
Online	=1	-0.381***	-0.077***	1.279***	0.050***	56.010***	0.180***
		(0.048)	(0.008)	(0.072)	(0.003)	(2.875)	(0.007)
Post cookie policy	=1 × Online=1	-0.002	0.013	-0.416***	-0.011*	-15.893**	-0.061***
		(0.068)	(0.012)	(0.149)	(0.006)	(6.287)	(0.018)
Median of dependent variable		6.94		26.80		189.16	
Individual FEs		✓	✓	✓	✓	✓	✓
Day-of-sample FEs		✓	✓	✓	✓	✓	✓
Merchant FEs		✓	✓	✓	✓	✓	✓
Observations		4,393,235	4,393,235	4,388,846	4,388,846	3,057,756	3,057,756
Adjusted R^2		0.23	0.28	0.14	0.15	0.41	0.40

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is as described in Table 3, split by tercile of transaction size (dollar amount). Each pair of columns estimates Specification 1 on the indicated transaction-size tercile subsample. The absence of an effect in the smallest tercile and its concentration in medium and large transactions is consistent with price discrimination and targeted advertising—rather than shopping convenience—as the primary operative channels. Columns (1)–(2) report results for the smallest transaction-size tercile, Columns (3)–(4) for the middle tercile, and Columns (5)–(6) for the largest tercile. Within each pair, the odd column reports winsorized spending per transaction and the even column reports log spending. All specifications include individual, day-of-sample, and merchant fixed effects. Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 14: Effect of Cookie Compliance on Spending by Spending Intensity: Aggregated Panel with Zeros

	(1) Spending winsorized	(2) Number of transactions	(3) Spending winsorized	(4) Number of transactions	(5) Spending winsorized	(6) Number of transactions
	1st spent tercile		2nd spent tercile		3rd spent tercile	
Post cookie policy	57.108***	1.683***	50.929***	1.790***	53.205***	2.116***
	(2.865)	(0.049)	(1.904)	(0.044)	(2.385)	(0.068)
Online =1	-3.655***	-0.158***	-15.058***	-0.510***	-23.426***	-0.766***
	(0.387)	(0.010)	(0.539)	(0.017)	(0.745)	(0.026)
Online=1 × Post cookie policy	-32.991***	-0.862***	-27.880***	-0.776***	-38.307***	-1.172***
	(3.535)	(0.072)	(2.915)	(0.071)	(3.030)	(0.092)
Median of dependent variable	10.01	0.23	24.25	0.60	35.25	0.90
Individual FEs	✓	✓	✓	✓	✓	✓
Day-of-sample FEs	✓	✓	✓	✓	✓	✓
Observations	2,082,783	2,082,783	2,422,975	2,422,975	2,403,369	2,403,369
Adjusted R^2	0.08	0.16	0.15	0.25	0.21	0.29

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is aggregated as described in Table 8 and split by tercile of each individual's mean retail spending divided by mean regular income—a measure of spending intensity capturing the share of resources devoted to retail purchases. As in Table 16 above, effects grow monotonically with spending intensity, consistent with higher-intensity consumers facing more aggressive price discrimination and being more susceptible to targeted advertising, and inconsistent with the shopping convenience channel, which predicts uniform effects across consumer types. Columns (1)–(2) report results for the lowest spending-intensity tercile, Columns (3)–(4) for the middle tercile, and Columns (5)–(6) for the highest tercile. Within each pair, the odd column reports winsorized spending and the even column reports the number of transactions. All specifications include individual and day-of-sample fixed effects. Standard errors, in parentheses, are clustered at the individual and month levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 15: Effect of Cookie Compliance on Spending by Sample Period Tercile

		(1) Spending winsorized	(2) Spending deviation from mean	(3) Spending winsorized	(4) Spending deviation from mean	(5) Spending winsorized	(6) Spending deviation from mean
		1st time tercile		2nd time tercile		3rd time tercile	
Post cookie policy	=0	0.000	0.000				
		(.)	(.)				
Online	=1	58.187***	0.275***	56.784***	0.281***	37.815***	0.189***
		(2.534)	(0.012)	(2.707)	(0.013)	(2.403)	(0.011)
Post cookie policy	=0 × Online=1	0.000	0.000				
		(.)	(.)				
Post cookie policy	=1			3.052*	0.021	0.693	0.003
				(1.664)	(0.014)	(0.712)	(0.004)
Post cookie policy	=1 × Online=1			-20.729	-0.323	-9.996***	-0.034**
				(14.705)	(0.198)	(2.027)	(0.014)
Median of dependent variable		57.10		59.28		67.36	
Individual FEs		✓	✓	✓	✓	✓	✓
Day-of-sample FEs		✓	✓	✓	✓	✓	✓
Merchant FEs		✓	✓	✓	✓	✓	✓
Observations		3,758,563	3,758,563	4,022,221	4,022,221	4,059,587	4,059,587
Adjusted R^2		0.40	0.23	0.44	0.22	0.43	0.14

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is as described in Table 3, split into three equal-length subperiods of the full sample. Each pair of columns estimates Specification 1 on the indicated time tercile subsample. Columns (1)–(2) report results for the earliest time tercile, Columns (3)–(4) for the middle tercile, and Columns (5)–(6) for the latest tercile. Within each pair, the odd column reports winsorized spending per transaction and the even column reports spending deviation from the individual mean. All specifications include individual, day-of-sample, and merchant fixed effects. Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 16: Effect of Cookie Compliance on Spending by Sample Period Tercile: Post-2019 Subsample

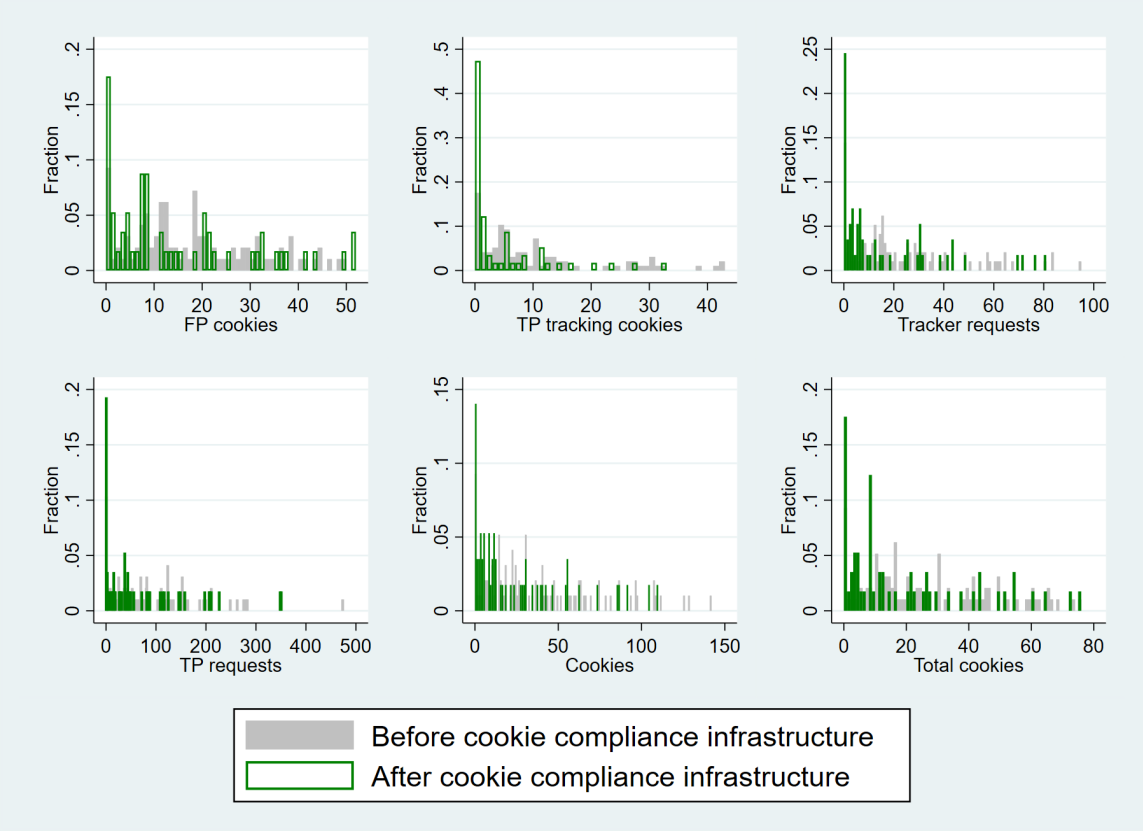
	(1) Spending winsorized	(2) Spending deviation from mean	(3) Spending winsorized	(4) Spending deviation from mean	(5) Spending winsorized	(6) Spending deviation from mean
	1st time tercile		2nd time tercile		3rd time tercile	
Post cookie policy =1	2.404	0.020	2.236**	0.018***	8.408***	0.037**
	(1.370)	(0.011)	(0.777)	(0.005)	(2.044)	(0.014)
Online =1	35.168***	0.163***	36.371***	0.193***	41.009***	0.220***
	(3.456)	(0.014)	(2.874)	(0.015)	(3.504)	(0.022)
Post cookie policy=1 × Online=1	-6.855**	-0.033	-22.456***	-0.096***	-26.761***	-0.125***
	(2.946)	(0.021)	(3.988)	(0.024)	(4.012)	(0.022)
Median of dependent variable	66.05		67.17		73.62	
Individual FEs	✓	✓	✓	✓	✓	✓
Day-of-sample FEs	✓	✓	✓	✓	✓	✓
Merchant FEs	✓	✓	✓	✓	✓	✓
Observations	1,095,864	1,095,864	1,096,747	1,096,747	1,056,948	1,056,948
Adjusted R^2	0.45	0.23	0.43	0.09	0.45	0.17

Standard errors in parentheses
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The sample is restricted to transactions occurring after January 2019, capturing the period of heaviest cookie compliance adoption and overlapping both GDPR-era and CCPA-era compliance rollouts. The restricted sample is then split into three equal-length subperiods. Each pair of columns estimates Specification 1 on the indicated time tercile subsample. Columns (1)–(2) report results for the earliest post-2019 tercile, Columns (3)–(4) for the middle tercile, and Columns (5)–(6) for the latest tercile. Within each pair, the odd column reports winsorized spending per transaction and the even column reports spending deviation from the individual mean. All specifications include individual, day-of-sample, and merchant fixed effects. Standard errors, in parentheses, are clustered at the individual, month, and merchant levels. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

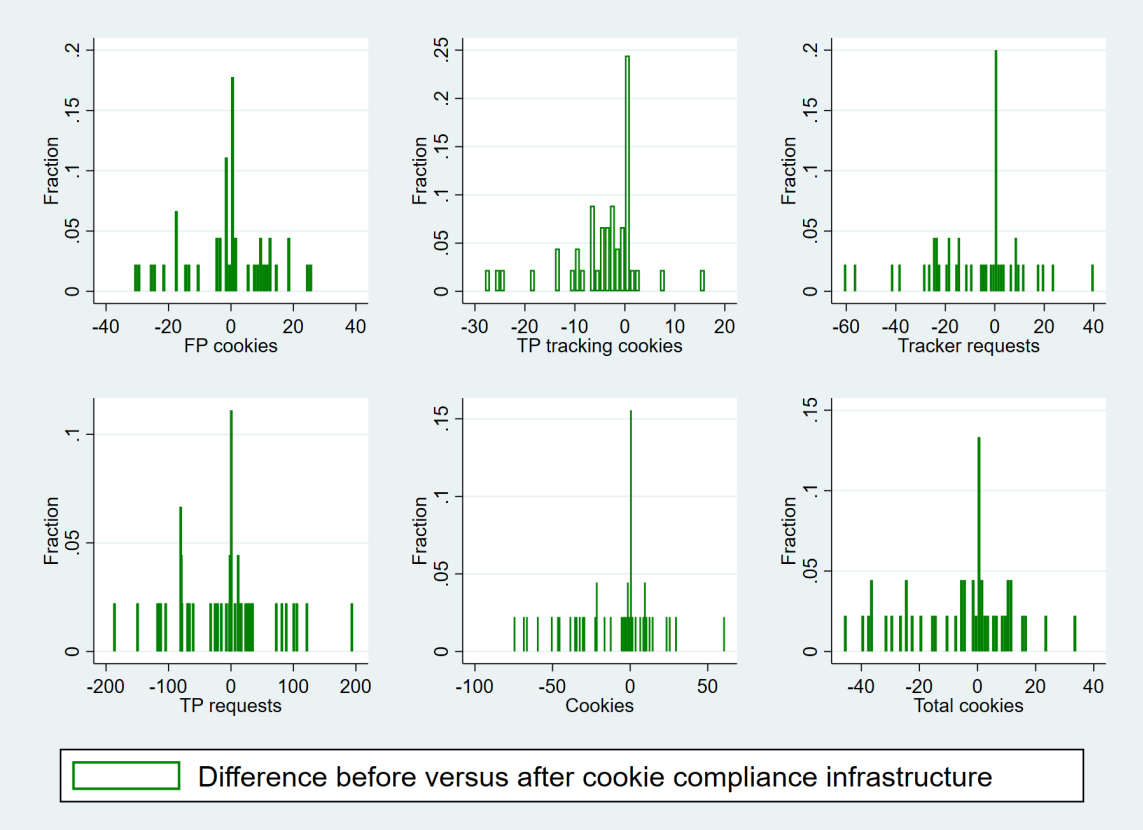
Figures

Figure 2: Cookie Compliance Frameworks Reduce Cookie and Tracker Placement: Before vs. After Histograms



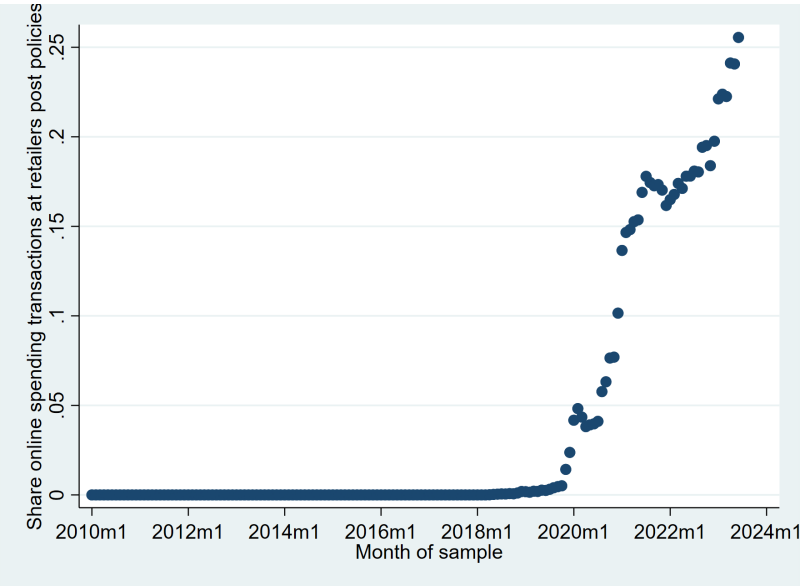
Notes: The sample is as described in Table 1. Each panel shows the distribution of the indicated cookie or tracking measure across all retailer-crawl observations, separately for observations before (gray) and after (green) the retailer adopted a cookie compliance framework.

Figure 3: Cookie Compliance Frameworks Reduce Cookie and Tracker Placement: Within-Firm Differences



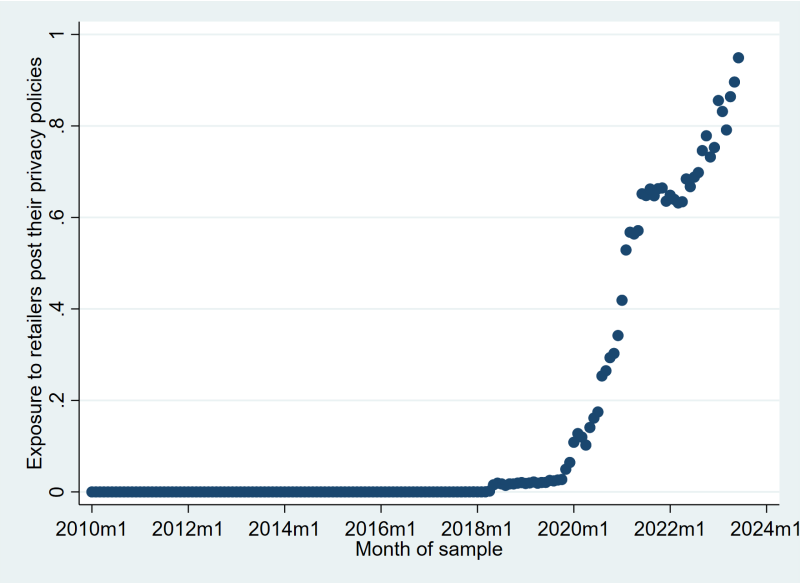
Notes: The sample is as described in Table 1. Each panel shows the distribution of the within-firm (before minus after) difference in the indicated cookie or tracking measure across all retailers that were observed in both pre- and post-adoption crawls.

Figure 4: Rollout of Cookie Compliance Frameworks over Time



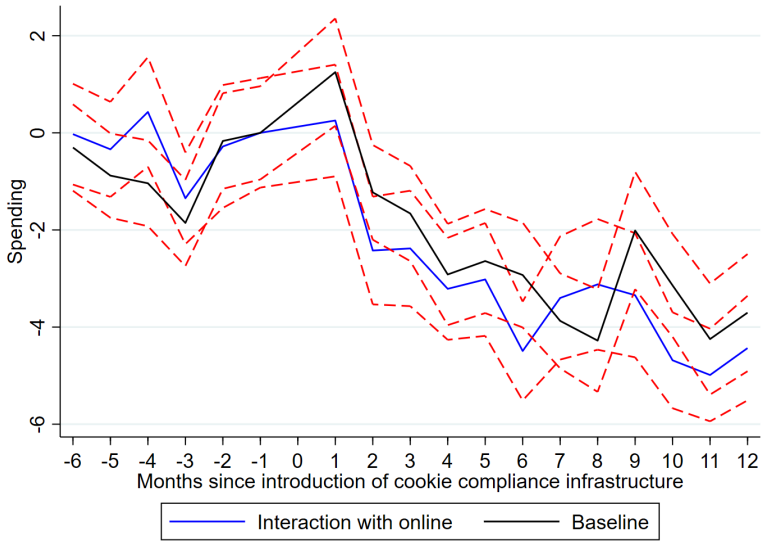
Notes: The figure plots the monthly share of online retail transactions in the sample that occur at retailers which have already implemented a cookie compliance framework (OneTrust, Osano, or CookieBot), as detected by BuiltWith. The sample covers January 2010 to August 2023.

Figure 5: Privacy Exposure Instrument over Time



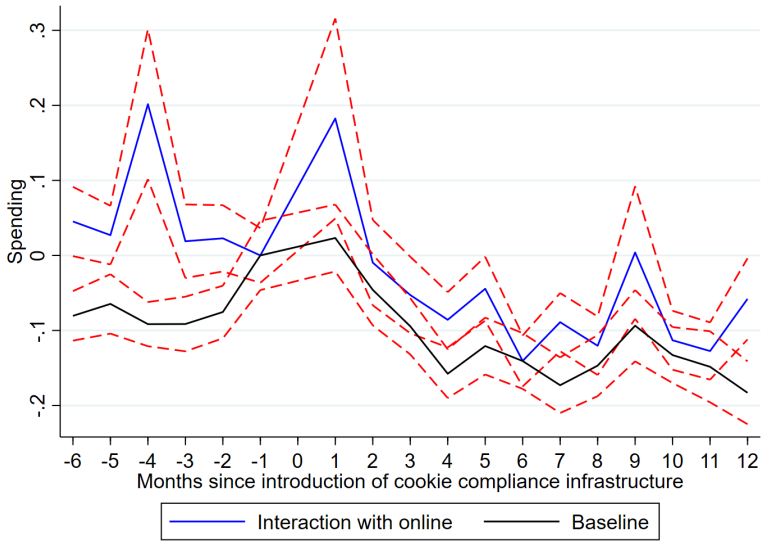
Notes: The figure plots the monthly cross-sectional average of the privacy exposure instrument, $EnhPrivacy_{i,t}$, constructed as the sum over retailers of an individual's pre-policy spending basket share at retailer m times an indicator for whether retailer m has adopted a cookie compliance framework by month t. The sample covers January 2010 to August 2023.

Figure 6: Dynamic Effects of Cookie Compliance on Online vs. Offline Spending



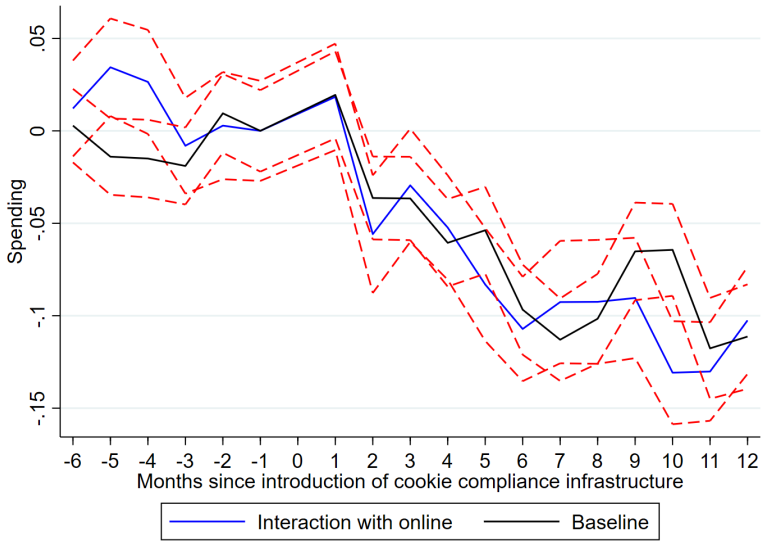
Notes: The figure plots event-study coefficients from Specification 3 estimated on the aggregated panel. The x-axis measures months relative to the retailer's cookie compliance adoption date (month 0). The black line ("Baseline") shows the coefficient on each event-time indicator, capturing the evolution of offline spending. The blue line ("Interaction with online") shows the coefficient on the interaction of each event-time indicator with the online dummy, capturing the *differential* decline in online relative to offline spending. A negative interaction coefficient indicates that online spending fell by more than offline spending. The dashed red lines show 95% confidence intervals around the interaction series. The dependent variable is winsorized spending per day. Individual and day-of-sample fixed effects are included. Standard errors are clustered at the individual and month levels.

Figure 7: Dynamic Effects of Cookie Compliance on Online vs. Offline Spending: Deviations from Individual Means



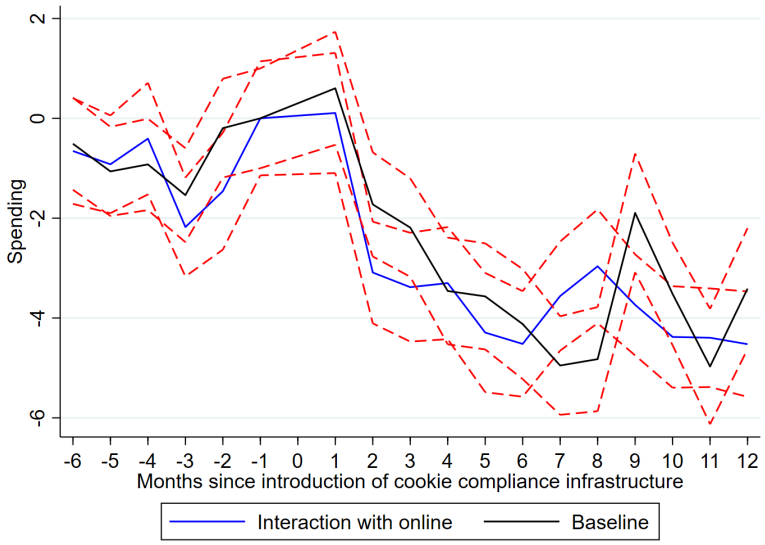
Notes: As Figure 6, but with spending measured as deviation from the individual's time-series mean rather than in levels. Individual and day-of-sample fixed effects are included. Standard errors are clustered at the individual and month levels.

Figure 8: Dynamic Effects of Cookie Compliance on Online vs. Offline Spending: Log Specification



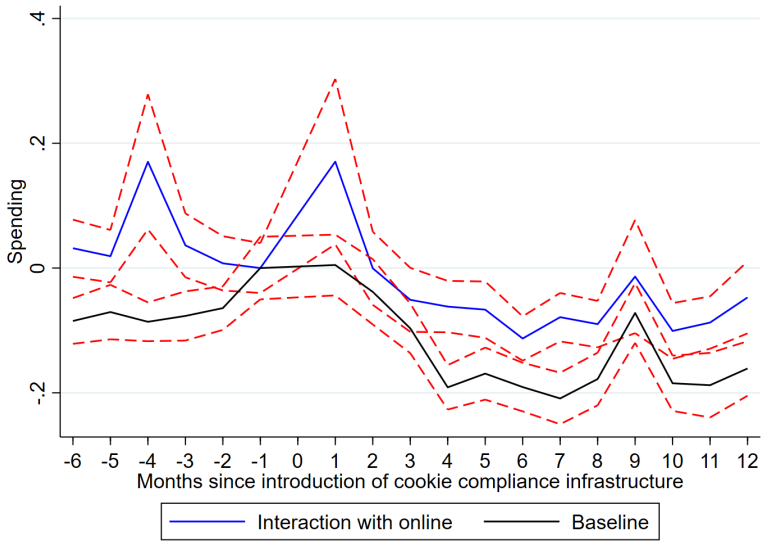
Notes: As Figure 6, but with log spending as the dependent variable. Individual and day-of-sample fixed effects are included. Standard errors are clustered at the individual and month levels.

Figure 9: Dynamic Effects of Cookie Compliance on Online vs. Offline Spending: High-Dimensional Fixed Effects



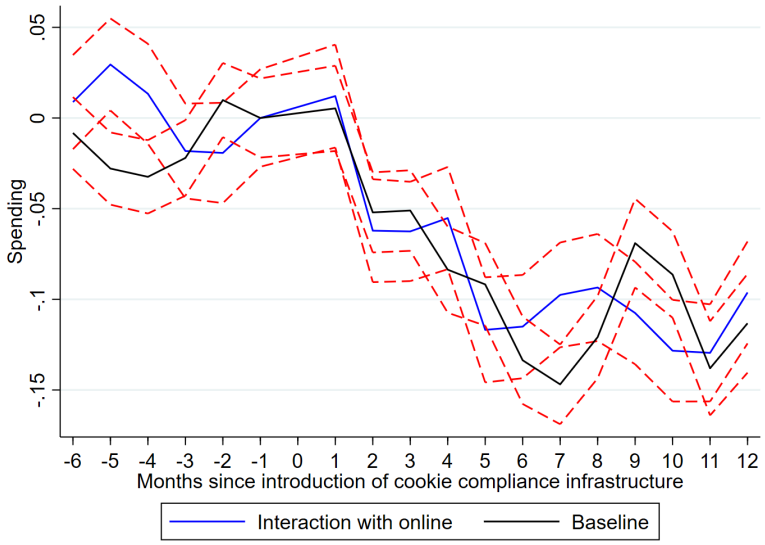
Notes: As Figure 6, but replacing individual and day-of-sample fixed effects with individual $\times$ month-of-sample, day-of-week, and week-in-month fixed effects. The dependent variable is winsorized spending per day. Standard errors are clustered at the individual and month levels.

Figure 10: Dynamic Effects of Cookie Compliance on Online vs. Offline Spending: Deviations from Individual Means, High-Dimensional Fixed Effects



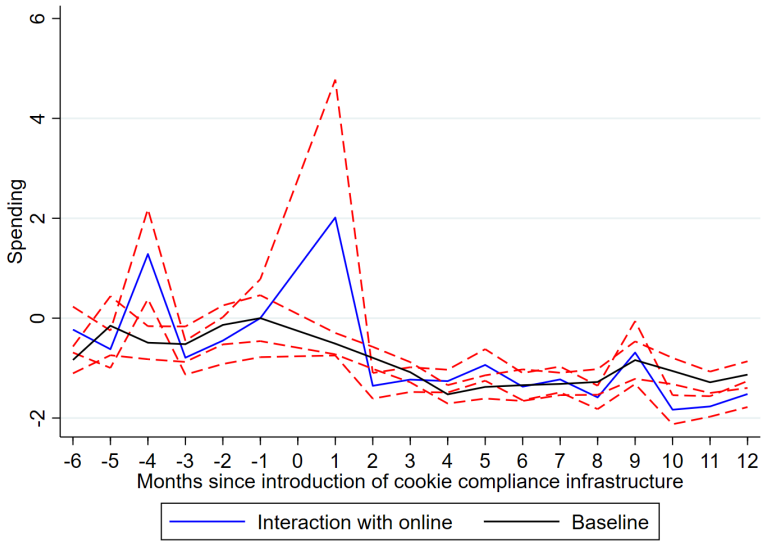
Notes: As Figure 7, but replacing individual and day-of-sample fixed effects with individual × month-of-sample, day-of-week, and week-in-month fixed effects. Standard errors are clustered at the individual and month levels.

Figure 11: Dynamic Effects of Cookie Compliance on Online vs. Offline Spending: Log Specification, High-Dimensional Fixed Effects



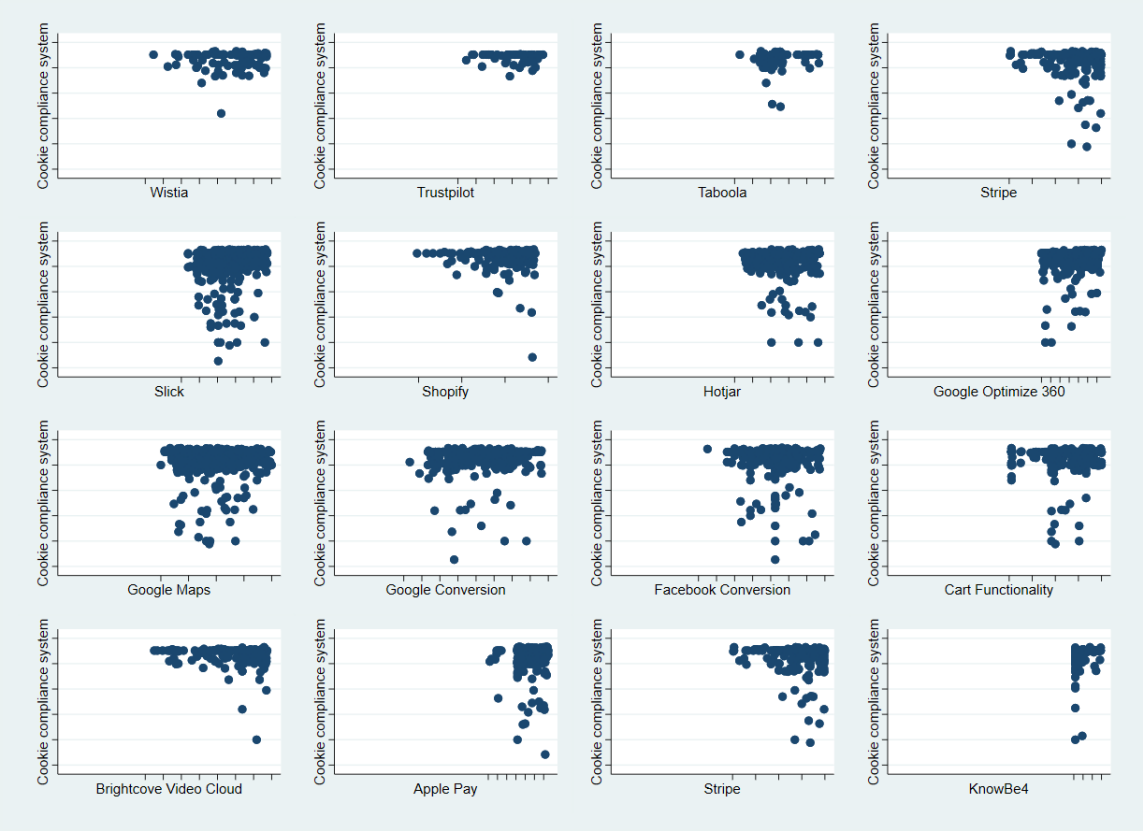
Notes: As Figure 8, but replacing individual and day-of-sample fixed effects with individual × month-of-sample, day-of-week, and week-in-month fixed effects. Standard errors are clustered at the individual and month levels.

Figure 12: Dynamic Effects of Cookie Compliance on Online vs. Offline Spending: Including Zero-Spending Days



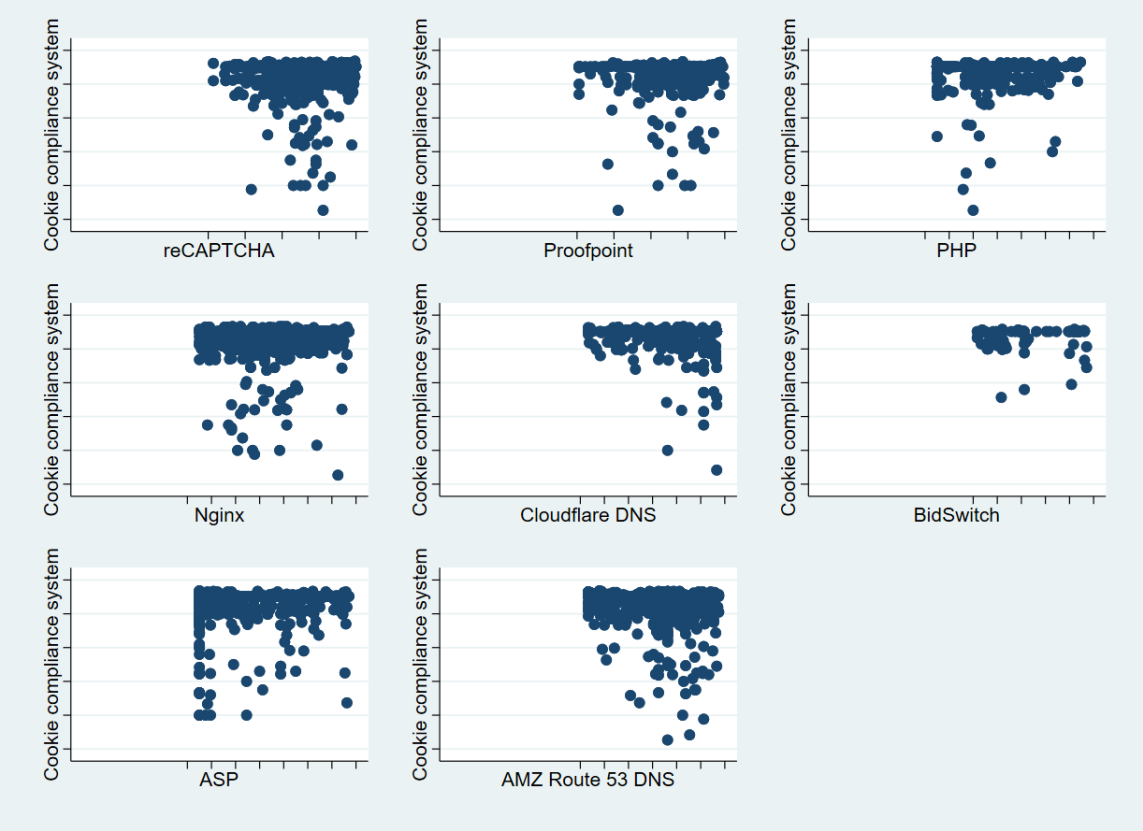
Notes: As Figure 10, but estimated on the aggregated panel that fills in zero spending on days when an individual did not transact at any retailer in the respective channel and policy status. This specification captures the extensive margin of spending. Standard errors are clustered at the individual and month levels.

Figure 13: Cookie Compliance Adoption Dates Are Uncorrelated with Other Shopping-Related Feature Adoption Dates



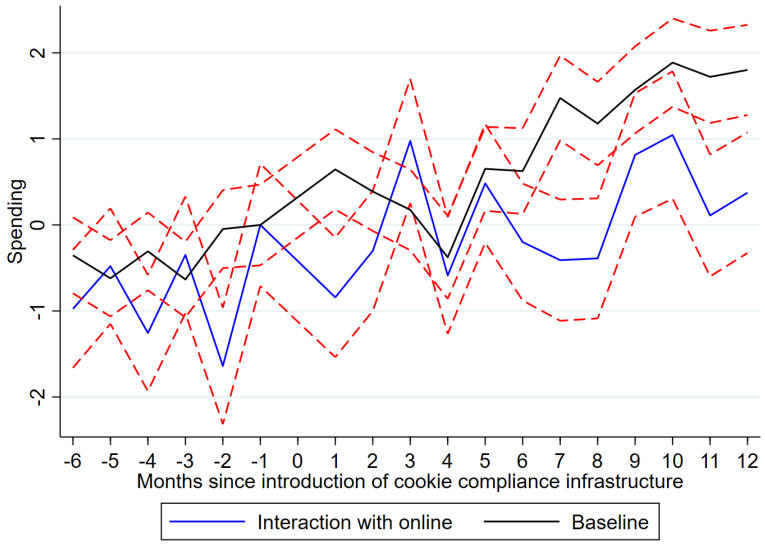
Notes: Each panel plots the date on which a retailer adopted a cookie compliance framework (y-axis) against the date on which the same retailer adopted an alternative website feature that could plausibly affect online shopping behavior (x-axis). Both axes span May 2001 to June 2023. If adoption dates were correlated, observations would cluster along the 45-degree line; the absence of such clustering supports the interpretation that cookie compliance adoption was not driven by, or contemporaneous with, broader website upgrades.

Figure 14: Cookie Compliance Adoption Dates Are Uncorrelated with Adoption Dates of Features Unlikely to Affect Shopping



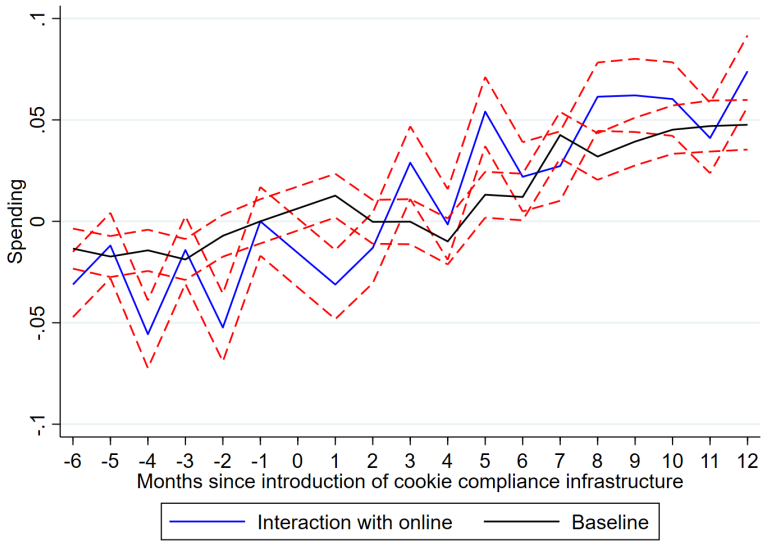
Notes: As Figure 13, but for website features that are unlikely to affect online shopping behavior (e.g., security and infrastructure tools). The absence of clustering along the 45-degree line confirms that cookie compliance adoption was not part of a coordinated technology overhaul.

Figure 15: Placebo Test: Reshuffled Cookie Compliance Dates, Winsorized Spending



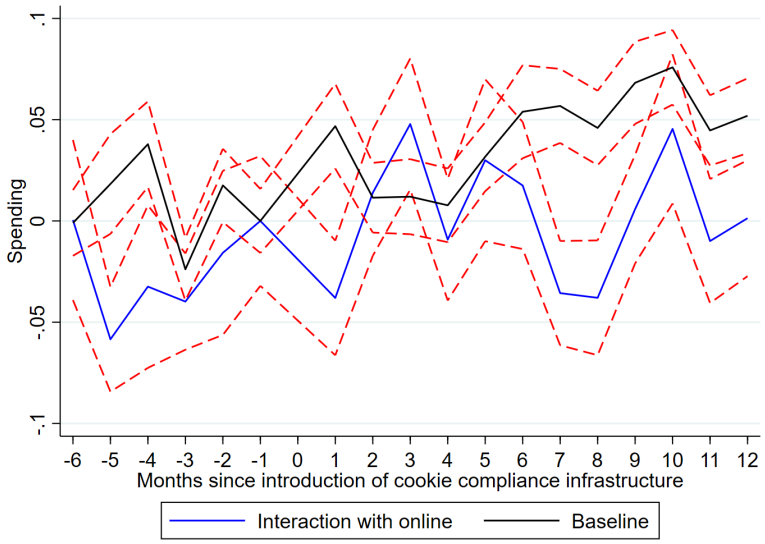
Notes: Placebo exercise in which cookie compliance dates are randomly reshuffled across retailers. The figure then replicates the event-study specification of Figure 6 on the placebo-assigned dates. The absence of a systematic post-event decline confirms that the findings in Figure 6 are not driven by spurious time trends.

Figure 16: Placebo Test: Reshuffled Cookie Compliance Dates, Log Spending



Notes: As Figure 15, but with log spending as the dependent variable. Corresponds to the placebo version of Figure 8.

Figure 17: Placebo Test: Reshuffled Cookie Compliance Dates, Deviations from Individual Means



Notes: As Figure 15, but with spending deviations from the individual mean as the dependent variable. Corresponds to the placebo version of Figure 7.