In Privacy We Trust: The Effect of Privacy Regulations on Data Sharing Behavior [*]

Ozge Demirci[†] Ayelet Israeli[‡] Eva Ascarza[§]

July 22, 2025

[*]We thank the anonymous company for providing the data and Noah Ahmadi for data preparation. We are grateful for valuable feedback from Catherine Tucker, and seminar participants at the International Industrial Organization Conference, the European Marketing Academy Conference, and Imperial College London Business School.

[†]Imperial College London Business School, odemirci@ic.ac.uk

[‡]Harvard Business School, aisraeli@hbs.edu

[§]Harvard Business School, eascarza@hbs.edu

In Privacy We Trust: The Effect of Privacy Regulations on Data Sharing Behavior

Please click here for the latest version.

Abstract

This paper studies the impact of privacy policies on consumer data-sharing behavior, focusing on policy changes in California and Virginia that took effect in 2023. Using data from a leading customer engagement app in the United States, where users upload shopping receipts in exchange for rewards, we find that privacy regulations led to a significant increase in both the volume and scope of receipt uploads, with the largest increases among users who were initially less inclined to share information. To investigate whether these policy changes had a broader impact beyond the platform, we analyze the nationally representative Consumer Expenditure Survey metadata, which details respondent interactions during the survey interviews. We find that respondents in treated states became more willing to share spending information after the policy. We further show that states where the new regulation was implemented experienced heightened privacy awareness, evidenced by an increase in privacy-related Google search activity and a decline in expressed privacy concerns during expenditure survey interviews. Together, these findings suggest that privacy regulations may encourage greater consumer participation by improving transparency and trust around data-sharing practices.

Keywords: Privacy, Privacy regulation, Data Sharing, Digital Platforms

1 Introduction

The rise of digital platforms has brought significant privacy concerns, as vast amounts of consumer data are continuously collected, processed, and utilized for various business objectives. This issue has become particularly urgent in recent years, with growing public awareness around data misuse, unauthorized access, and inadequate transparency in data-handling practices. Privacy policies have emerged as important regulatory responses to these concerns, but the question of how these policies influence consumer behavior itself, particularly their willingness to share data, still remains. Understanding these behavioral responses is essential not only for assessing the effectiveness of privacy protections but also for analyzing how such regulations reshape firm strategies, data markets, and the broader digital economy.

A priori, it is unclear whether privacy regulations would lead consumers to increase or decrease their willingness to share information. On the one hand, increasing transparency around data collection practices could heighten privacy concerns, leading individuals to become more cautious and reduce their willingness to share personal information (Goldfarb and Tucker 2011, Acquisti et al. 2016, Aridor et al. 2023). Greater awareness of potential risks may discourage engagement with platforms that collect user data. On the other hand, by enhancing consumer rights and providing stronger protections, privacy regulations could increase trust in how personal information is handled, thereby encouraging greater data sharing (Tucker 2014, Gupta et al. 2023). Furthermore, prior research highlights that consumer privacy preferences are not fixed and that the impact of privacy regulations on willingness to share information can be highly context-dependent and heterogeneous across individuals, influenced by how protections and trade-offs are perceived by individuals (John et al. 2011, Goldfarb and Tucker 2012, Miller and Tucker 2018). Whether the overall effect of privacy protections leads to reduced or increased information sharing is ultimately an empirical question, which is the focus of this paper.

This paper investigates how privacy regulations affect consumers' willingness to share their data. We focus on two recent policy changes in the United States, the California Privacy Rights Act and the Virginia Consumer Data Protection Act, which came into effect in January 2023. These regulations require platforms to provide greater transparency about data collection and to offer consumers the right to opt out from selling their data, thus attempting to strengthen individual control over personal information. We analyze the effects of these policies within the context of a leading customer engagement app where users upload their shopping receipts in exchange for monetary rewards. This unique setting enables us to observe user engagement with the platform over time, offering a unique opportunity to examine fine-grained changes in consumer data-sharing behavior following the implementation of the regulation.

Using a synthetic difference-in-differences (SDID) approach, we estimate the causal effects of these policies. Our findings indicate a significant 9.2 percent increase in receipt uploads to the platform in the states where the new privacy regulations entered into effect, California and Virginia (hereafter "treated states"). This increase corresponds to approximately 0.38 additional receipts per user per week based on the pre-policy mean. Beyond the overall increase in volume, we also find that users share a broader range of information after the policy changes, including receipts from more diverse types of retailers and product categories. These patterns suggest that increased transparency and consumer control over personal data can encourage greater openness and participation in data-sharing platforms. We also document meaningful heterogeneity in these effects; the largest behavioral response is observed among users who were least willing to share data before the policy.

To assess the implicit value users place on privacy protections, we compare the behavioral effects of the policy change to those of equivalent monetary incentives. Using variation in platform rewards, we estimate that a one-point increase in the average reward per offer (worth $0.001) leads to approximately 0.19 additional receipts uploaded per user per week. This implies that the platform would need to spend around $0.005 to generate one additional receipt. In comparison, the privacy policy increased sharing by about 0.38 additional receipts per user per week. This suggests that, in this setting, the privacy regulation had an effect on user behavior comparable to what could be achieved through a modest increase in financial incentives. These results provide a revealed-preference estimate of the value consumers place on privacy protections and demonstrate how structural interventions can shift behavior through mechanisms beyond economic rewards.

To assess whether the policy influenced users' behavior in settings not covered by the regulation, we leverage data from the nationally representative "Consumer Expenditure Survey." This survey tracks individuals over time and collects detailed information on spending behavior. Although the state-level privacy regulation changes do not apply to this national survey, we use the survey to investigate whether the changes in regulation changed consumer behavior in contexts beyond the direct scope of the regulation. The survey's metadata records respondent interactions, including the share of questions that were answered relative to those that were either refused or left unanswered, and whether privacy concerns were raised during the interview process. Consistent with our main results, respondents in treated states became more likely to report their spending details after the policy change. This increase is particularly pronounced among individuals who had lower baseline reporting rates prior to the policy, mirroring the heterogeneity we observe on the platform. These patterns suggest that the observed behavioral shifts extend beyond digital platforms and monetary incentives. We interpret this as evidence of a broader change in consumer privacy protection perceptions, reflecting broader changes in consumers' willingness to disclose information in response to strengthened privacy protections.

While the main results indicate a significant increase in data sharing following the introduction of privacy regulations, it is important to assess whether this behavioral shift reflects a direct response to the law itself or was influenced by contemporaneous changes in platform strategies. In particular, changes in reward structures, promotional incentives, or communications from the platform could have independently motivated users to share more data, complicating the interpretation of the observed effects. To investigate this possibility, we analyze platform-side incentive structures and find no evidence of adjustments in rewards or promotional campaigns around the time of the policy implementation. This strengthens the interpretation that the observed increase in data sharing reflects a behavioral response to strengthened privacy protections rather than to concurrent changes initiated by the platform.

Furthermore, to rule out the possibility that broader economic changes explain the observed patterns, we analyze underlying consumption behavior using the Consumer Expenditure Survey data. We find no significant differential changes in spending patterns between treated and control states, suggesting that the documented increase in data sharing is not driven by changes in consumers' purchasing activity, but rather by an increased willingness to share existing information.

To shed light on the mechanisms driving these behavioral changes, we analyze shifts in public attention to privacy and expressed privacy concerns. Using Google Search Volume Index data, we find that interest in privacy topics rose significantly in treated states around the time of policy implementation. While this heightened interest could, in principle, either increase caution or enhance perceived safety, further analysis of Consumer Expenditure Survey metadata clarifies the direction of the effect: respondents in treated states became significantly less likely to express privacy-related concerns during survey interviews after the policy. These results support the interpretation that strengthened privacy regulations improved consumers' perceptions of data safety, contributing to the observed increases in data-sharing behavior.

This study makes several contributions to the literature on the economics of privacy, expanding on work primarily focusing on the economic value and consequences of protecting and disclosing personal data. Previous research has largely focused on the economics of privacy (Acquisti et al. 2016, Goldfarb and Tucker 2012, Goldfarb and Que 2023) and how privacy regulations affect firms and market outcomes, including advertising effectiveness, innovation, market concentration, and profitability (Goldfarb and Tucker 2011, Johnson et al. 2020, Lefrere et al. 2020, Peukert et al. 2022, Johnson et al. 2023, Demirer et al. 2024, Goldberg et al. 2024). Recent studies also document the potential impacts of regulations on tracking and cookie restrictions for advertisers and platforms (Johnson et al. 2023, Aridor et al. 2024) and how privacy policies affect particular web traffic or clicks (Aridor et al. 2020, Schmitt et al. 2022, Goldberg et al. 2024). While this literature has provided valuable insights into the economic implications of privacy policies, most studies have focused on firm-level and market-level outcomes. We complement this work by examining how privacy regulations shape consumer behavior, with a particular focus on individuals' willingness to share their data. By shifting the lens to consumer-level responses, our study contributes to the economics of privacy literature by offering a clearer view of how individuals respond to regulatory efforts aimed at safeguarding personal information.

By using two complementary data sources to analyze consumer behavior after the regulations, our analysis examines changes in both instrumental and intrinsic privacy preferences. The platform data allows us to observe behavior driven by monetary incentives and trade-offs, while the Consumer Expenditure Survey captures privacy attitudes that are not connected to monetary consequences or directly to the regulation. By combining these two perspectives, we offer a comprehensive view of how privacy regulations affect consumers, both in terms of practical behavior and underlying attitudes. Building on existing work that measures revealed privacy preferences (Goldfarb and Tucker 2012, Acquisti et al. 2013, Athey et al. 2017, Kummer and Schulte 2019, Tang 2019, Lin 2022), we also estimate the economic value of privacy protections by translating behavioral responses into monetary terms. Consumers vary in their privacy preferences and in how they evaluate the trade-offs involved in sharing personal data (Goldfarb and Tucker 2012, Collis et al. 2021, Dubé et al. 2024, Lin and Strulov-Shlain 2025). We contribute to this literature by examining differences in baseline willingness to share information and showing that privacy regulations had the largest impact on those who were initially the most reluctant to share their data.

The structure of the paper is as follows: Section 2 introduces institutional details, including recent privacy regulations in the US and the customer engagement platform used in this study. Section 3 describes our main data source and sample construction. Section 4 examines the impact of privacy policies on consumer data-sharing behavior. Section 5 analyzes whether the platform adjusted its strategy in response to the policy. Section 6 explores potential mechanisms underlying the observed behavioral shifts. Section 7 discusses the welfare implications of privacy policies and provides concluding remarks.

2 Empirical Setting

2.1 Privacy Policy Regulations in the US

In recent years, privacy regulations worldwide have evolved significantly in response to rising concerns over consumer data protection and digital privacy. The European Union's General Data Protection Regulation, enacted in 2018, marked a milestone by establishing comprehensive data protection principles, such as the right to access, correct, and delete personal data, along with stricter consent requirements. These regulatory efforts are largely driven by the exponential growth of digital platforms and the increasing volume of personal data collected, raising concerns about consumer rights, data misuse, and transparency. In the United States, privacy discussions have primarily focused on creating state-led policies to address these challenges, aiming to build consumer trust in digital environments without comprehensive federal privacy legislation. The current privacy regulations vary at the state level.

The California Privacy Rights Act and the Virginia Consumer Data Protection Act, both effective as of January 1, 2023, represent two significant developments in data privacy regulation. The California law expands on earlier legislation by strengthening consumer rights and placing additional obligations on businesses that collect and process personal data. It enhances consumer control by granting individuals the right to correct inaccuracies in their data, limit the use of sensitive personal information, and opt out of data sales and targeted advertising based on activity across different websites or apps. The policy also requires businesses to disclose more clearly what categories of personal information they collect and how that data is used, a shift that may directly influence consumer perceptions of transparency and trust in data-sharing environments.[1]

While the California Consumer Privacy Act (CCPA), which became effective in 2020, marked an important early step in the U.S. privacy landscape, the California Privacy Rights Act (CPRA) introduced in 2023 significantly expanded and strengthened that framework. The CPRA added new consumer rights, including the ability to correct personal information and restrict the use of sensitive data, and provided a more robust and comprehensive data privacy framework. Importantly, our platform introduced new user-facing privacy controls around the time CPRA took effect, including options for requesting access to and deleting data, as well as opt-out requests. These changes indicate a significant institutional shift for the platform at the time of CPRA's implementation, helping us isolate the effects of the 2023 policy from earlier regulatory changes.[2]

Similarly, Virginia's law introduces key privacy protections by giving consumers the right to access, delete, and correct their data. It requires businesses to obtain consent before processing sensitive data, which includes precise geolocation, biometric data, and information related to children. It also allows users to opt out of targeted advertising, data sales, and profiling for automated decision-making. These legislative changes not only set new benchmarks for consumer rights but also challenge businesses to rethink their data management strategies. This regulatory environment encourages companies to adopt more transparent data practices, potentially leading to a more consistent approach to privacy across states as other legislatures consider similar laws. As more states explore comparable legislation,

[1] Although enforcement began on July 1, 2023, many businesses, including the platform analyzed here, began implementing the policy's requirements earlier to align with legal obligations and avoid future compliance risks.

[2] We also examined whether the introduction of the CCPA in 2020 had any effects on consumers' willingness to share expenditure data or express privacy concerns by using metadata from the Consumer Expenditure Survey. We found no significant effects attributable to the CCPA.

these laws establish a foundation that gradually shapes privacy practices across the country, influencing data protection approaches and shaping consumer expectations in the digital landscape.

2.2 Customer Engagement Platform

The main dataset used in this study is provided by a leading customer engagement platform based in the U.S. The platform incentivizes users to upload receipts from everyday purchases, creating a rich dataset of shopping behaviors. The receipts include itemized transaction details including the store name, location, timestamp, individual products purchased, short product descriptions, and prices — both listed and discounted, when applicable. Users can upload their physical receipts by submitting photos of these receipts (most common), or they can upload digital receipts of their online purchases (by connecting their e-commerce or e-mail account to the platform's app). Users are rewarded with points for each receipt upload, with additional incentives available for purchasing specific products or reaching spending thresholds. Points accumulated through these activities can be redeemed for rewards, such as gift cards and discounts, creating a cycle of engagement that motivates consistent use and data sharing. The platform's business model is centered on partnerships with brands that pay to promote their products and gain insights into consumer behavior. This reward-based model encourages users to participate actively, providing a comprehensive view of their spending patterns across a range of retailers and product categories.

Beyond receipt uploads, the platform engages users through personalized promotions and in-app offers, often in partnership with brands. These promotions reward users for uploading receipts that include promoted products or meet specific purchase conditions. While these incentives serve to encourage ongoing participation, they also create natural variation in when and how users share information, offering a valuable perspective to observe patterns in data-sharing behavior. Importantly, the platform captures detailed user-level data on purchase behavior, including store types, product categories, and geographic locations. This level of granularity makes it well-suited for studying how external changes, such as new privacy regulations, affect individuals' willingness to share personal consumption information.

2.3 Implications of Privacy Policies on the Platform

The privacy regulations introduced in California and Virginia impose various requirements on businesses handling consumer data. However, not all provisions directly impact the platform studied here. For instance, restrictions on processing sensitive personal data, such as health or biometric information, are not relevant, as the platform primarily collects consumer shopping behavior data rather than highly regulated personal information. Similarly, provisions related to automated decision-making and profiling have limited applicability, given that the platform's data processing focuses on consumer purchase activity rather than predictive risk assessments.

The most significant regulatory changes affecting the platform are the requirement to provide transparency about the data collected and the requirement to offer consumers the ability to opt out of data sales and sharing. Under CPRA, "personal information" is defined broadly to include commercial data such as purchase histories, even if not directly tied to identifiable personal details. Since the platform monetizes data by sharing it with third-party advertisers and brand partners, it had to implement an opt-out mechanism, which is now available within its app, allowing users to restrict the sale of their personal information. Additionally, new transparency and disclosure requirements mandated clearer privacy policies and more explicit communication about data usage, leading the platform to update its terms of use to reflect consumer rights under the new laws. These adjustments ensure compliance while maintaining access to consumer data under the revised legal framework.

Although the interface update was shown uniformly across users in all states, the platform also introduced state-specific disclosures in line with the new privacy policy requirements. For example, California residents were directed to dedicated sections within the app and on the website that outlined their rights under CPRA and offered a streamlined opt-out process. These updates reflect a combination of universal design changes and jurisdiction-specific compliance measures in states that adopted new privacy laws.

3 Data

The main dataset used in this paper is from a customer engagement platform, which includes customer receipts data and partner offer data.

3.1 Data on Customer Behavior

We observe receipt-level data for 15,806 users, randomly drawn through stratified sampling, over 21 weeks from October 2022 to March 2023. The raw receipt uploads are aggregated at the user-week level, which serves as the primary unit of analysis. To ensure comparability across users, we restrict the sample to those who downloaded the app and uploaded their first receipt in the first week of October. This helps eliminate cohort-based differences and aligns all users on a common timeline.

The main variables in the analysis, summarized in Table 1, provide insights into user activity on the platform. Users uploaded an average of 3.69 receipts per week, and total weekly spending averaged $215.16, reflecting a range of engagement levels, with some users contributing higher spending volumes. Additionally, the average number of stores visited per week was 2.22, with users purchasing an average of 25.54 items weekly. These variables represent different features of user engagement, from purchase frequency to various retailers and items purchased, capturing a broad range of platform interactions.

We present summary statistics for additional engagement metrics — such as the number of unique stores, ZIP codes, and retailer categories visited or shared in a given week — in Web Appendix A (Table W1). These variables capture more nuanced aspects of activity, including the diversity of retail behavior and categories purchased.

Table 1: Summary Statistics for Main Variables

Variables	**Mean**	**SD**	**Median**
Number of receipts uploaded	3.69	6.25	0
Total spending	215.16	755.61	0
Number of stores visited	2.22	3.48	0
Number of items bought	25.54	59.80	0
Amount spent on awarded items	177.23	597.14	0
Points available / offers available	152.61	102.35	151.93
Recency (weeks since last upload)	2.83	3.82	1

Notes: This table presents mean, standard deviation, and median, respectively, for the variables used in the main analysis. The dataset is aggregated at the user-week level. Total spending and amount spent on awarded items are in US dollars.

3.2 Data on Platform Activity

For the platform-side analysis, we use a dataset covering all 2,529 offers available on the platform during the study period. For each offer, the data includes the number of users who received it in each U.S. state, the number of points assigned, and the offer's start and end dates. This structure allows us to observe both the timing and geographic reach of incentives for consumers to share their data, and examine whether there were any changes in these incentives following the policy change. The majority of the offers are distributed nationally, with over 80% of them available in nearly all states. As a result, variation in offer availability across states is limited. Instead, our analysis focuses on differences in user exposure, i.e., how many users received a given offer in each state, and the points each user was exposed to as proxies for the platform's targeting intensity and promotional strategy.

4 Impacts of Privacy Policies on Consumer Behavior

4.1 Identification challenges

Our goal is to measure the causal impact of privacy policy changes in California and Virginia on users' data-sharing behavior on the platform. This requires addressing two key identification challenges.

First, we must construct a valid counterfactual to assess what sharing behavior California and Virginia would have looked like in the absence of the policy. Second, we must have no other confounding events that coincided with the policy change and could have independently influenced user behavior.

Challenge 1: Constructing a valid counterfactual. To address the first challenge, we use users in other states as the control group in a quasi-experimental setting. We compare data-sharing behavior of users in treated and control states, adopting a difference-in-differences (DID) framework. A DID approach requires two assumptions: It assumes that trends in the outcome variable for treated and control units would have followed parallel paths in the absence of treatment (the "parallel trends" assumption). Second, it assumes that treatment in one group does not indirectly affect outcomes in the control group (i.e., no spillovers).

The second condition is likely satisfied in our context, as the policy applied only to users in the treated states. While national media coverage could have affected privacy awareness more broadly, any indirect effects in control states would likely reduce, not increase, data sharing, potentially biasing our estimates conservatively.[3]

The first condition — parallel trends — is more difficult to guarantee in this context. Pre-treatment receipt-sharing behavior varies across users due to geographical and behavioral

[3]Since users in control states were not protected by the policy, increased awareness could lead to greater caution and lower sharing. However, as discussed in Section 6.1, we find no evidence of heightened privacy-related search activity in control states around the policy date, which helps alleviate spillover concerns.

differences, which may violate the assumption. To address this, we implement the synthetic difference-in-differences (SDID) estimator developed by Arkhangelsky et al. (2021). SDID combines the strengths of traditional DID with synthetic control weighting to construct a control group that more closely matches the treated group's pre-treatment trends. This relaxes the standard parallel trends assumption to a conditional version, allowing for more credible identification when baseline differences exist (see Web Appendix B.1 and Web Appendix B.4 for diagnostic tests and event-study evidence).

Challenge 2: Ruling out confounding changes at the time of the policy change. The second challenge is to ensure that there were no other events that drove observed changes in behavior at the time of the policy change. At the state level, we are not aware of any such events. We therefore investigate if any platform-side events co-occurred with the policy change and could have impacted user behavior. For example, the platform may have adjusted its reward structure, launched new engagement campaigns, or changed how it communicated with users in response to the policy.

Although the policy required clearer disclosures and opt-out mechanisms, it did not prohibit the platform from collecting consumer data. Thus, the firm may have revised its strategies to preserve or boost engagement. If such changes occurred concurrently with the policy shift, they could confound our estimates.

To address this, we examine whether platform-side incentives changed following the policy change. Specifically, we track variation in user-facing offers — such as promotions, discounts, or bonus rewards — that differ across users and states. By comparing offer exposure in treated versus control states before and after the policy, we evaluate whether observed behavioral changes could be attributed to adjustments in platform incentives.

Finally, to further isolate the policy's effect from potential platform confounds, we analyze data from the Consumer Expenditure Surveys. This setting is free from platform-side incentives and enables us to assess whether individuals became more or less willing to share consumption information in a non-monetized, survey-based context. Moreover, metadata from the survey indicates whether respondents raised privacy concerns during interviews, allowing us to directly connect behavioral shifts to changes in privacy attitudes.

4.2 Empirical Strategy: Synthetic Difference-in-Differences

The SDID estimator is implemented by solving a weighted two-way fixed effects regression that minimizes the squared-error loss between observed outcomes and a model that includes additive unit and time effects, as well as a treatment indicator. Specifically, SDID solves:

$$(\hat{\tau}^{SDID}, \hat{\mu}, \hat{\alpha}, \hat{\beta}) =_{\tau^{SDID},\mu,\alpha,\beta} \sum_{i=1}^{N} \sum_{t=1}^{T} \left(y_{it} - \mu - \alpha_i - \beta_t - \tau^{SDID} W_{it}\right)^2 \hat{\omega}_i \hat{\lambda}_t \tag{1}$$

where y_{it} is the observed outcome (log number of receipts uploaded) for user i in week t, and W_{it} is an indicator equal to 1 if user i is in the treated group and the observation falls in the post-policy period. The model includes individual fixed effects α_i, and week fixed effects β_t to account for user-specific and temporal unobserved heterogeneity. The coefficient τ^{SDID} captures the average treatment effect. The key feature of the estimator lies in the weights $\hat{\omega}_i$ and $\hat{\lambda}_t$, which are estimated from the data to construct a comparison group of untreated units that closely matches the treated group's pre-treatment outcomes, while also aligning pre- and post-treatment periods. This doubly weighted design helps ensure better pre-treatment balance and focuses estimation on observations most comparable to the treated group and periods, improving robustness and precision of the estimated average treatment effect $\hat{\tau}^{SDID}$.[4]

One main advantage of SDID is its use of both unit and time weights, which enable it to align the pre-treatment outcome paths of treated units with a reweighted synthetic control group while also reweighting the timeline to give more emphasis to periods that better represent the counterfactual evolution. This dual weighting helps address potential imbalances that arise when treated users differ systematically from control users or when treatment

[4]We also adjust the SDID procedure by incorporating covariates through outcome residualization. For more details on this procedure, please refer to Web Appendix B.2.

effects are sensitive to time-varying factors. In contrast to standard DID estimators, which rely on an unweighted average of untreated units, SDID constructs a counterfactual that is more tailored to the treated sample, improving the credibility of the estimated effects.

At the same time, SDID incorporates individual and time fixed effects, much like a two-way fixed effects (TWFE) specification. This feature helps absorb persistent differences across users and general shocks across weeks, stabilizing estimates and facilitating interpretation. As shown in Arkhangelsky et al. (2021), the inclusion of fixed effects within the weighted least squares framework improves the efficiency of the estimator and allows for extrapolation beyond the pre-treatment period, which is particularly useful in longer post-treatment windows.

These properties make SDID especially suitable for our setting. Users' engagement with the platform, measured through receipt uploads, may vary systematically across individuals and over time, driven by differences in location, prior activity, shopping habits, or responsiveness to incentives. By combining fixed effects with synthetic control-style weighting, SDID allows us to estimate the causal effect of privacy policies while reducing bias from baseline differences and time trends.

Alternative Estimation Methods — We complement our main analysis with additional estimation strategies to assess the robustness of our findings. As a robustness check to our primary SDID estimates, we implement the Doubly Robust Difference-in-Differences (DRDID) estimator proposed by Sant'Anna and Zhao (2020). DRDID is particularly well-suited to settings like ours, with sharp treatment timing and a panel data structure. The estimator combines outcome regression with inverse probability weighting, improving the balance between treated and untreated units while preserving the main idea of the DID identification strategy. Whereas SDID focuses on matching units based on pre-treatment outcome trends, DRDID offers a complementary approach by balancing on observed covariates. This makes it especially useful in settings with rich baseline information, helping ensure that treated and control groups are comparable in terms of observed characteristics. Its implementation does not require staggered adoption or strong parametric assumptions, making it aligned with the structure of our dataset and policy change. By relying on a different weighting and adjustment approach than SDID, DRDID allows us to test whether our findings are robust to alternative estimation methods.

Alternative Outcome Specification — The primary outcome variable — receipt uploads — is a count data that can include zeros. In our main analysis, we apply a log transformation to reduce skewness and to facilitate interpretation of the results in percentage terms. However, recent work has highlighted interpretational challenges when log-transforming outcomes with zero values, particularly due to concerns about unit dependence (Chen and Roth 2023). To address this, we also estimate the SDID model using the raw (untransformed) outcome to verify that our results are not sensitive to functional form.[5] Further methodological details are provided in Web Appendix B.

4.3 Results

4.3.1 Effects on the Volume of Shared Data

In this section, we estimate the baseline regressions to assess the impact of privacy policies on the amount of information shared on the platform, specifically focusing on the number of receipts uploaded. Table 2 presents the results of this analysis, showing average treatment effects ($\hat{\tau}^{SDID}$ in Equation (1) and $\hat{\tau}^{DRDID}$ in Equation (B2)), across different estimation methods: SDID and DRDID. Across specifications, we control for time-varying controls including points per offer, recency of uploading, lagged spending, lagged award activity, lagged number of items bought, and lagged number of stores visited.

Column (1) reports results from the SDID estimator, using the logarithm of receipt uploads as the outcome. The coefficient of 0.088 implies an estimated increase of approximately

[5] While count models such as Poisson or negative binomial regressions are often used for this type of outcome, they are less appropriate in our setting due to the inclusion of high-dimensional fixed effects. These models can suffer from estimation difficulties and incidental parameter problems in such contexts. Using a linear model ensures computational tractability and allows for a consistent fixed effects structure across specifications.

Table 2: Changes in the number of receipts uploaded weekly

	(1) (SDID, logged)	(2) (DRDID)	(3) (SDID, raw)
Treated×Post	0.0877	0.0525	0.591
	(0.0024)	(0.0104)	(0.012)
	[p = 0.000]	[p = 0.000]	[p = 0.000]
Observations	316,806	316,806	316,806
Pre-treatment Mean	4.13	4.13	4.13
Week FE	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes

Notes: This table reports the estimated treatment effects of privacy policy implementation on the number of receipts uploaded per user per week by estimating Equation (1) in Columns (1) and (3) and Equation (B2) in Column (2). The dependent variable is the logarithm of receipt uploads in Columns (1) and (2) and the raw count of receipt uploads in Column (3). All models include user and week fixed effects and the same covariates. The 'Pre-treatment Mean' shows the average number of receipts uploaded per user-week in the control group before policy implementation. Columns (1) and (3) use the SDID estimator with standard errors based on bootstrap-based inference. Column (2) reports results from the DRDID with robust and asymptotic standard errors. Corresponding p-values (in square brackets) are reported below standard errors for each estimate.

9.2% in the number of receipts uploaded by treated users after the policy implementation, relative to the control group.[6] While this relative effect is modest, it corresponds to an increase of about 0.38 receipts per user per week, based on the pre-treatment mean of 4.13. This amounts to roughly 1.5 additional receipts per user per month after the policy implementation.

Column (2) presents results from the DRDID estimator. This method combines outcome regression with inverse probability weighting and estimates the average treatment effect on the treated within a two-period framework. The coefficient of 0.052 suggests a 5.4% increase in receipt uploads post-policy. This smaller magnitude, compared to SDID, may reflect differences in weighting strategies and sensitivity to treatment effect heterogeneity. Since DRDID effectively averages effects across covariate-adjusted comparisons, it may down-weight some of the stronger responses observed in SDID's reweighted panel framework.

[6] The percentage effect is calculated as $\exp(0.0877) - 1 \approx 0.0917$.

Column (3) shows results from an alternative SDID specification using the raw number of receipts uploaded as the outcome. The coefficient implies that treated users uploaded approximately 0.59 more receipts per week after the policy change. This estimate is larger than the implied 0.38 receipt increase in Column (1), which is consistent with known differences between log-transformed and raw count outcomes. Log transformation tends to compress variation and reduce the influence of large values and zeros, while linear models treat all changes equally in absolute terms. As such, the larger coefficient in Column (3) reflects that the linear specification captures absolute changes more directly, without re-scaling the outcome.

To examine the dynamics of the policy's effect over time, we estimate event-study specifications using the SDID framework. This approach allows us to assess how receipt uploads evolved in treated versus control states in the weeks before and after the policy implementation. Up until two weeks prior to the policy change, the plots show no evidence of pre-treatment differences in trends, supporting the validity of our identification strategy. The plots also confirm a sustained increase in receipt uploads after the policy takes effect. The modest rise in receipt uploading in treated states during the two weeks before the policy change is likely attributable to the anticipation and publicity of the upcoming policy changes. However, the more pronounced and persistent increase begins in January 2023, coinciding with the policy's implementation. The corresponding event-study plots, presented in Figure 1a and Figure 1b below, align with Columns (1) and (3), respectively.[7] The figures demonstrate that the SDID method generated a comparable set of control users before the policy change.

Overall, these results demonstrate a consistent and statistically significant positive effect of privacy policies on data-sharing behavior. Across all estimation methods, users in treated states increased the number of receipts uploaded after introducing privacy policies. The consistency of the effects across different estimation strategies reinforces the robustness and

[7] Because DRDID is designed for two-period settings, it does not allow for dynamic treatment effect estimations, so we do not construct an event-study plot for Columns (2).

Figure 1: Event Study Estimates of SDID for Receipt Upload Behavior

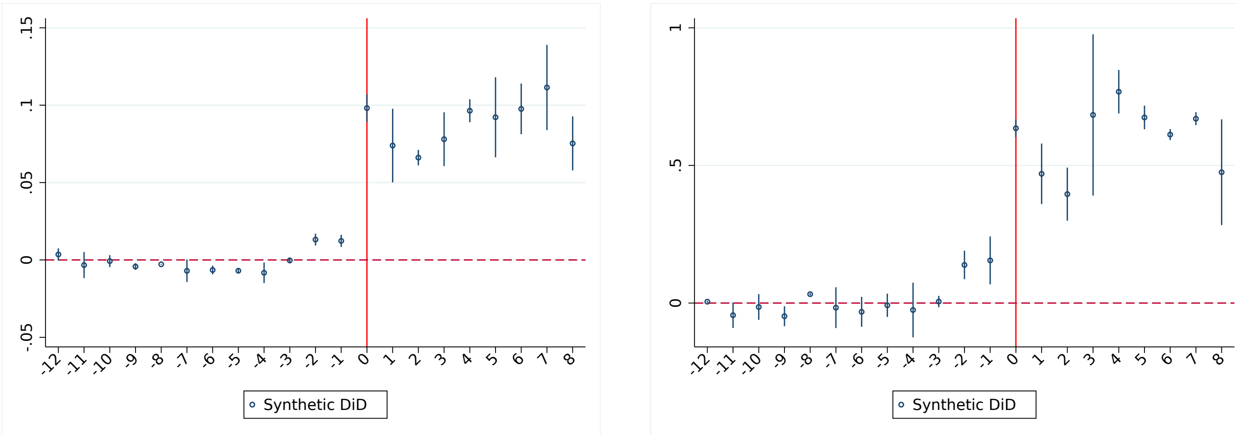

(a) Logged receipt count **(b)** Raw receipt count

Notes: Each panel plots time-varying treatment effects estimated using the SDID method, as described in Web Appendix B.4. Panel (a) uses log counts and Panel (b) uses raw counts of uploaded receipts as the outcome variable. Each point represents the estimated effect $\hat{\tau}^{\mathrm{SDID}}_{a,\ell}$ for week ℓ relative to the policy implementation date, marked by the red vertical line. Confidence intervals are constructed using bootstrap-based inference.

reliability of our findings.[8]

Subgroup analysis based on pre-treatment activity level — To explore heterogeneity in treatment effects, we estimate Equation (1) separately for user subgroups defined by their level of engagement pre-policy change. We classify users into below- and above-median activity groups based on the number of receipts uploaded before the policy change, using state-specific median thresholds. In Table 3, Column (1) reports the baseline estimates from Table 2, and Columns (2) and (3) report treatment effects for these subgroups, where treated users in each group are compared to control users from the same activity stratum. This approach allows us to examine whether the policy effect differs for users who were more or less engaged with the platform before the privacy update.

The results suggest meaningful differences in treatment effects based on users' pre-policy engagement levels. Column (2) shows that users below the median in pre-treatment receipt uploads experienced a statistically significant increase of 0.067 in log receipt uploads following

[8]As a robustness check, we replicate our main analysis using another random sample of 15,089 users. Results, reported in Web Appendix Table B1, are highly consistent with our main findings.

Table 3: Changes in the number of receipts by prior activity level

	(1) Full sample	(2) Below Median	(3) Above Median
Treated×Post	0.0877	0.0657	0.0059
	(0.0024)	(0.0187)	(0.0017)
	[p = 0.000]	[p = 0.000]	[p = 0.001]
Observations	316,806	156,366	157,080
Pre-treatment Mean	4.13	0.84	7.48
Week FE	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes

Notes: All columns report estimates from SDID regressions as specified in Equation (1). Subgroups in Columns (2) and (3) are based on users' pre-treatment upload activity relative to state-specific quartile thresholds and show below and above the median level active users, respectively. Control group units are matched with the same level of activity (i.e., below or above median) from control states in Columns (2) and (3). Standard errors computed via bootstrap replications. Corresponding p-values are listed in square brackets below each estimate.

the policy change. This corresponds to about 0.06 additional receipts per user per week, nearly a 7% increase relative to the subgroup's pre-policy mean of 0.84 receipts. Given the low baseline engagement, this is a relatively large effect in absolute terms.

In contrast, the estimated effect for users above the median in pre-treatment uploads is relatively small, with a coefficient of 0.006 as shown in Column (3). This suggests that the behavioral response to the privacy policy was more limited among users who were already active on the platform prior to the policy change. One possible explanation is that these users were already comfortable sharing data, leaving less room for additional behavioral adjustment. Overall, the pattern of results indicates that the policy may have had a more noticeable impact on users with previously lower levels of engagement.

This pattern is consistent with the idea that stronger privacy protections may have reduced perceived barriers or concerns for less engaged users, thereby encouraging greater data-sharing. For high-engagement users, who may already feel familiar with or trust the platform, the marginal impact of additional privacy disclosures appears limited. As each subgroup regression is estimated independently, the synthetic control weights are optimized separately for each group. This allows for meaningful within-subgroup comparisons but requires caution when comparing coefficients across subgroups directly. A detailed explanation of our subgrouping procedure and its implications is provided in Web Appendix B.6.

4.3.2 Effects on the Composition of Shared Data

Understanding not only how much data users share but also what types of data they are willing to disclose provides a richer perspective on the behavioral shifts affected by these privacy regulations. In this section, we examine whether the policies also influenced the type of information shared.

Variety of store visits — We extend our baseline SDID regression in Equation (1) by replacing the outcome variable with three alternative measures that capture distinct dimensions of informational variety: (1) the number of unique stores visited, reflecting store-level variety; (2) the number of unique ZIP codes of stores visited, capturing geographic variety; and (3) the number of distinct retailer categories, indicating categorical variety. Table 4 presents the results using each of these three outcomes.

Table 4: Changes in the type of data shared: Store visits

	(1)	(2)	(3)
DV:	Unique stores	Unique ZIPcodes	Retail Category
Treated×Post	0.0524	0.0347	0.0390
	(0.008)	(0.0199)	(0.0045)
	[p = 0.000]	[p = 0.081]	[p = 0.000]
Observations	316,806	316,806	316,806
Pre-treatment Mean	2.45	1.58	1.22
Week FEs	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes

Notes: Each column estimates logged numbers of unique store visits, unique ZIP codes visited, and number of retailer categories and reports results from a separate SDID regression with individual and week fixed effects. Standard errors computed via bootstrap replications. The coefficients represent average treatment effects on the treated. Pre-treatment refers to the average value of each outcome among control units in the pre-policy period. Corresponding p-values are reported in square brackets below each estimate.

Column (1) shows that treated users, on average, share information from a more diverse set of stores following the policy change. The estimated coefficient of 0.052 implies a 5.4% increase in the number of unique stores visited. Given the pre-treatment control mean of 2.45, this corresponds to an increase of approximately 0.13 stores per user per week. Column (2) reports a coefficient of 0.035, indicating a 3.5% increase in the number of unique ZIP codes shared, which translates to about 0.06 ZIP codes. Column (3) presents a coefficient of 0.04, suggesting a 4.0% increase in retailer category diversity, equivalent to roughly 0.05 additional categories disclosed per user per week.

To further investigate changes in the variety of store visits, Table 5 examines whether the policy change affected the types of retailers for which users share data. Columns represent food and beverage, department stores, health-related, and specialty retailers, respectively. Retailer classification was based on the 207 most common retailers we observe in the dataset (covering 75% of receipts), following classification methods from Dubé et al. (2018) and DellaVigna and Gentzkow (2019).

Table 5: Changes in the type of data shared: Retailer types

	(1)	(2)	(3)	(4)
DV:	Food & Beverage	Department Stores	Health	Specialty Stores
Treated×Post	0.0287	0.0633	0.0172	0.0043
	(0.0119)	(0.0096)	(0.0010)	(0.0052)
	[p = 0.016]	[p = 0.000]	[p = 0.000]	[p = 0.408]
Observations	316,806	316,806	316,806	316,806
Pre-treatment Mean	0.92	1.23	0.32	0.09
Week FE	Yes	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes	Yes

Notes: Outcomes are log counts of receipts from each retailer category. Specialty stores include apparel, home improvement, and entertainment retailers (e.g., clothing brands, hardware chains, cinemas). Standard errors calculated using bootstrap replications. Corresponding p-values are reported in square brackets below each estimate.

The table shows that treated users increased their sharing of food and beverage receipts by approximately 2.9% and department store receipts by about 6.5%, suggesting greater willingness to disclose core consumption activities. For health-related retailers, the estimated increase is smaller, around 1.7%, likely reflecting the more sensitive nature of medical and pharmaceutical purchases. In contrast, the effect on specialty store receipts is small (0.4%) and not statistically significant, suggesting that the policy had a more limited influence on the disclosure of discretionary or lifestyle-related purchases.

Digital receipts — On this platform, users can connect their e-commerce or email accounts to enable automatic uploading of digital receipts. Connecting these accounts to the platform requires logging in to those accounts through the platform's application by providing the usernames and passwords for those external accounts. Engaging in this type of activity can therefore be considered as a higher level of trust in the platform. This subsection examines changes in connecting additional digital accounts and uploading digital receipts following the policy change.

Table 6: Changes in the type of data shared: Digital receipts

	(1)	(2)
DV:	Digital Account Connection	Number of Digital Receipts
Treated×Post	0.0035	0.0538
	(0.0004)	(0.0005)
	[p = 0.000]	[p = 0.000]
Observations	316,806	316,806
Pre-treatment Mean	0.028	0.67
Week FE	Yes	Yes
Individual FE	Yes	Yes

Notes: Column (1) estimates changes in an indicator for whether the user has linked at least one digital source (e.g., e-commerce or email account). Column (2) estimates changes in the logged number of digital receipts shared. Standard errors are computed using bootstrapping. Corresponding p-values are reported in square brackets below each estimate.

In Table 6, Column (1) focuses on changes in whether users chose to link their digital accounts to the platform. The outcome is a binary variable equal to one if a user had connected at least one digital source, and zero otherwise. The estimated coefficient of 0.0035 suggests a 0.35 percentage point increase in the likelihood of digital account connection. Given a pre-policy connection rate of only 3%, this represents a relatively small absolute effect, but a noticeable increase in relative terms. Column (2) estimates the effect of privacy policies on the logged number of digital receipts uploaded per user per week. The coefficient of 0.054 corresponds to an approximate 5.4% increase relative to the pre-treatment mean, indicating a modest but statistically significant rise in digital data sharing following the policy change. Together, these results suggest that the policy changes may have encouraged slightly greater engagement with digital data-sharing features on the platform. Although the changes are modest in magnitude, they are consistent with the broader pattern observed in physical receipt uploads.

Overall, the results in Section 4.3.2 show that privacy policies not only boost the volume of shared data but also expand its scope across variety of stores, consumption categories, and channels (digital vs. paper receipts). This suggests that users are more comfortable sharing information about their regular shopping habits. Diversifying shared data may enhance the platform's ability to infer consumer behavior and to personalize offerings, while reflecting users' growing trust in data handling under strengthened privacy regulations.

4.3.3 Supporting Evidence from the Consumer Expenditure Survey

To assess whether the behavioral shifts we observe on the platform extend beyond that setting, we draw on complementary evidence from the Consumer Expenditure (CE) Survey between 2019 and 2024. The CE Survey is a nationally representative dataset providing quarterly household-level data on expenditures, as well as detailed metadata on the interview process itself, including whether respondents expressed concerns about privacy and the extent to which they answered questions when asked about their spending. This analysis helps us examine whether changes in consumer privacy perceptions triggered by new legislation can affect behaviors even in settings that are not legally governed by the CPRA, such as government-administered surveys. Using these data, we assess whether the implementation of privacy policies influenced consumers' willingness to share expenditure details in a broader setting beyond the platform analyzed in this paper and its subscribers.[9]

[9]For more details about the CE Survey used in this study, please refer to Web Appendix Table W2.

We leverage the CE Survey data in two ways. First, we use it to rule out an alternative explanation for our main findings— namely, that the observed increase in receipt uploads reflects changes in underlying consumption in treated states rather than greater willingness to share data. Second, we use the survey to examine whether the same pattern of increased data-sharing among individuals in treated states emerges in a broader, non-platform setting, which is not governed by the changes in regulation.

Ruling out changes in consumption patterns in treated state — Using the CE Survey, we examine whether overall consumption patterns in treated states changed contemporaneously with the policy implementation. A potential concern is that the increase in receipt uploads observed in the platform data may reflect changes in actual spending behavior rather than a shift in data-sharing behavior. For instance, January is a period when consumer behavior may be influenced by seasonal factors such as New Year's resolutions or broader economic trends. Additionally, if consumers in treated states systematically differ from those in control states in terms of spending habits, it could raise concerns about whether the observed increase in data sharing is driven by underlying changes in consumption rather than privacy policy effects.

To assess this, we use the CE data, focusing on total expenditures as well as spending in food, beverages, and apparel, which are the categories that saw the largest increase in receipt uploads in the platform data. (Results are presented in Web Appendix Table B2). Across all specifications, we find no significant changes in total consumption or spending on these key categories following the policy implementation in treated states relative to control states. These findings suggest that the observed increase in receipt uploads is not driven by a shift in underlying consumption behavior but rather by changes in consumer engagement with the platform, reinforcing the interpretation that privacy regulations influenced willingness to share data rather than overall spending patterns.

Impacts of the policy on sharing in the CE survey — We construct an indicator of sharing expenditure data in the CE Survey, which is the proportion of questions answered relative to the total number of questions asked, including those that were left unanswered due to respondents refusing to answer or stating that they did not recall the information. A lower value for this measure suggests a greater reluctance to disclose expense details, whereas a higher value indicates a greater degree of engagement in sharing expenditure information. To estimate the impact of privacy policies on expenditure-sharing behaviors in the CE Survey, we use a DID framework, in which users from California and Virginia ("treated") are compared to those from other states ("control").[10] We estimate the following specification:

$$y_{it} = \gamma \text{Post}_t \times \text{Treat}_i + \beta_c \text{X}_{it} + \phi_i + \theta_t + \epsilon_{it} \tag{2}$$

where y_{it} represents the share of expenditure-related questions answered by respondent i at time t. The variable $\text{Post}_t \times \text{Treat}_i$ is an interaction term that takes the value of one if respondent i is from a treated state and the interview occurs after January 1, 2023, and zero otherwise. Household fixed effects (ϕ_i) control for time-invariant respondent characteristics, and quarter fixed effects (θ_t) account for seasonality in survey responses. Additionally, we control for income (X_{it}) when estimating the ratio of shared expenditures, as income levels may influence individuals' likelihood of having more expenditure details.

Before presenting the main results, we validate the parallel trends assumption using an event-study analysis, detailed further in Web Appendix B.7. Specifically, we estimate the event-study specification provided in Equation (B6), which allows us to examine the evolution of respondents' willingness to share expenditure information in treated states relative to control states over time. Results, presented in Web Appendix Figure B2, indicate that before the policy implementation, differences between treated and control states are small,

[10] We do not apply the SDID estimator to the CE Survey analysis because the survey has a repeated cross-sectional structure. Each household is observed every quarter in a year, but only within a single calendar year, and the sample changes across years. Since SDID requires a balanced panel structure to construct unit and time weights based on pre-treatment trends, it is not well-suited to this data setting.

statistically insignificant, and stable, supporting the validity of our parallel trends assumption. After the introduction of privacy policies, however, there is a clear upward shift in treated respondents' engagement in expenditure reporting, indicating a positive impact of the policies on respondents' willingness to disclose information.

Table 7 reports the results for the ratio of shared expenditures among all questions. Respondents in treated states increased the proportion of expenditure questions they answered relative to the total number of questions asked. The median split is calculated at the state level using respondents' pre-policy average share of expenditures reported. Column (1) indicates that the ratio of shared expenditures increased by 1.45 percentage points post-policy, representing a meaningful increase relative to the high baseline (around 96%). This increase is particularly pronounced among respondents who initially shared fewer details (below-median sharers), for whom the ratio increased by 3.3 percentage points, while highly-engaged respondents exhibited no statistically significant changes.[11]

Table 7: Changes in the ratio of shared expenses overall responses

DV: Ratio of shared expenses	Full Sample	Below Median	Above Median
Treated×Post	0.0145	0.033	0.002
	(0.003)	(0.008)	(0.002)
	[p = 0.000]	[p = 0.000]	[p = 0.317]
Observations	69,646	25,686	31,449
R-squared	0.664	0.669	0.599
Household FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes

Notes: The table presents estimation results for Equation (2) by using the share of reported expenditures as an outcome variable. Each regression controls for household income, consumer unit, and quarter time fixed effects. Subgroups in Columns (2) and (3) are based on users' pre-treatment median of share of expenditures in each state. Heteroskedasticity robust standard errors are reported in parentheses. Corresponding p-values are listed in square brackets below each estimate.

[11] Two states in our control group, Colorado and Connecticut, enacted new privacy laws that came into effect on July 1, 2023, which overlap with the timeline of our CE survey data. To ensure the robustness of our results, we re-estimated the analysis excluding these two states from the control group. The findings remain consistent with those presented in the table and are robust to this alternative sample.

Overall, the findings from the CE Survey reinforce the patterns observed in the platform data, suggesting that privacy regulations positively influence individuals' willingness to share information. The increase in expenditure-sharing behavior is particularly pronounced among respondents who were initially less willing to disclose information, mirroring the stronger behavioral response observed among less active users on the platform. Together, these results indicate that the effects of privacy policies extend beyond digital interactions to broader patterns of information disclosure.

Although the CE Survey is not directly subject to state-level privacy laws, the behavioral shifts we document are likely shaped by broader changes in consumer sentiment and perceived data protection norms. Regulatory interventions can increase the salience of privacy and signal a heightened institutional emphasis on data protection. These effects may spill over into contexts like government surveys, where legal frameworks remain unchanged but psychological perceptions of safety and control improve. As a result, individuals who were previously more cautious may become more willing to share personal information even outside commercial platforms, reflecting a broader shift in norms around data sharing. These results support the interpretation that privacy legislation can affect behavior through both direct legal mechanisms and broader shifts in perceived norms. We further explore this explanation in Section 6.

5 Impacts of Privacy Policies on Platform Strategy

After establishing that the change in privacy policies led to more consumer data sharing, we next examine whether the platform made any changes to its data-sharing incentives that could explain the observed increase in user engagement. By doing so, we aim to assess whether the rise in data sharing can be attributed to platform-level strategic responses rather than consumer-driven reactions to privacy protections.

Specifically, the new privacy regulations introduced stricter requirements for data transparency and consumer control, but they did not directly restrict the platform's data collection, as the platform does not process sensitive personal information. The most immediate impact came from the requirement to provide an opt-out mechanism and clearer disclosures on collected data, prompting the platform to update its terms of use and adjust its compliance practices. These changes can influence how the platform structures data-sharing incentives, as well as how it engages with its business partners.

From the platform's perspective, stricter data regulations may change the marginal value of consumer data for the platform. By increasing compliance obligations and making data collection more transparent, the policy can lead the platform to reconsider its approach to maintaining user engagement. If regulatory requirements raise the cost of managing consumer data, the platform may adjust its incentive structure to ensure continued participation. Offering greater rewards for data-sharing activities can help sustain engagement levels and offset any potential decline in data availability due to opt-outs. Alternatively, since the policy can increase transparency and potentially build greater consumer trust, it can lead to more data being shared on the platform (as shown in the previous section). Even if the platform does not anticipate a significant change in data collection, it can still modify incentives to reinforce engagement and ensure a steady flow of high-quality data, particularly as users become more aware of their privacy rights and how their information is being used.

5.1 Empirical Strategy and Results

To assess whether the platform responded to the privacy policy by adjusting its data-sharing incentives, we use the platform offer data that shows how many users received each offer in each U.S. state. Most offers in the dataset were distributed nationally, with over 80% of the offers available in every state. Because of this limited variation in where offers were shown, we focus on how many users receive each offer in each state. This lets us examine whether the platform changed the intensity of offer exposure in treated states after the policy.

In particular, we estimate the following DID specification:

$$y_{it} = \beta(\text{Post}_t \times \text{Treat}_i) + \alpha_i + \gamma_t + \epsilon_{it} \tag{3}$$

where y_{it} is the outcome for state i in week t, Treat_i is an indicator for whether the state is California or Virginia, and Post_t is a post-policy indicator that equals 1 after January 1, 2023. All models include state fixed effects (α_i) and week fixed effects (γ_t). In addition to DID, we also provide SDID estimates for additional robustness.

We examine three outcomes. The first is the number of users who received each offer, capturing total exposure at the state-offer level. The second is the average number of users per offer in each state-week, reflecting the intensity of offer targeting. The third is the average number of points assigned per offer, serving as a proxy for the reward value attached to the platform's promotions.

We validate the parallel trends assumption using an event-study analysis, detailed in Web Appendix C, using both the DID and SDID methods, plotting the dynamics of the treatment effect over time. The plots show no systematic differences in the variables of interest between treated and control states before the policy implementation. We find no evidence of differential pre-treatment trends, supporting the identification strategy used in our DID and SDID estimations. Table 8 reports the results of our DID and SDID analyses.

We find no statistically significant changes in any of the three outcomes examined. The number of users receiving offers, the average number of users per offer, and the average number of points assigned per offer all remain statistically unchanged in treated states relative to the control states following the policy. Given the large pre-treatment means, these small estimates suggest that the platform did not systematically adjust its targeting or reward structure in response to the policy. Overall, the results provide no evidence that the platform expanded its data-sharing incentives in reaction to the regulatory change.

These results help rule out the possibility that the observed increase in data sharing was driven by changes in the platform's promotional strategy. We find no evidence that the

Table 8: Changes in data-sharing incentives

	(1) #Users with offer (DID)	(2) Users per offer (DID)	(3) Points per offer (DID)	(4) #Users with offer (SDID)	(5) Users per offer (SDID)	(6) Points per offer (SDID)
Post × Treat	19,900	14,350	15.92	64.48	-11,188	-24.66
	(13,689)	(9,843)	(9.81)	(4421.6)	(7596.7)	(18.76)
	[p = 0.146]	[p = 0.145]	[p = 0.105]	[p = 0.988]	[p = 0.141]	[p = 0.189]
Observations	1,113	1,113	1,113	1,113	1,113	1,113
Pre-treatment Mean	64,385	68,856	2,167	64,385	68,856	2,167
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each column reports estimates from a separate regression using a state-week level panel dataset. Columns (1)–(3) use DID with state and week fixed effects; Columns (4)–(6) use the SDID estimator with bootstrap-based standard errors. The outcomes include: number of users receiving offers (Columns 1 and 4), users per offer (Columns 2 and 5), and points per offer (Columns 3 and 6). Users per offer is defined as the number of users divided by the number of offers in each state-week; points per offer is the total number of points assigned divided by the number of offers. In Columns (1) and (4), we additionally control for the total number of offers available in each state-week. Standard errors clustered at the state level for DID, and are computed using bootstrapping for SDID; they are shown in parentheses. Corresponding p-values (in square brackets) are reported below standard errors for each estimate.

platform expanded its targeting intensity or reward values in response to the policy. Instead, the absence of significant changes in exposure to offers or incentive magnitude supports the interpretation that the behavioral shifts documented earlier reflect users' direct responses to the privacy regulations rather than strategic adjustments by the platform during the same period.

6 Mechanism Analysis

With the platform-side explanations ruled out, we next investigate the underlying mechanism on the consumer side. Because we observe behavioral changes both on the platform and in the CE Survey, we believe that the underlying mechanism extends beyond platform disclosures and platform-specific design changes. In particular, we investigate whether the changes in policies shaped public perceptions of privacy and data safety, which in turn would also explain the changes in CE Survey responses. Section 6.1 examines changes in public attention to privacy issues using Google search behavior, while Section 6.2 assesses whether privacy concerns declined in treated states using survey metadata from the CE Survey.

6.1 Changes in Interest in Privacy: Evidence from Google Search

We begin by examining whether the observed increase in data-sharing behavior was accompanied by heightened public attention to privacy-related issues. One potential mechanism is that the policy drew greater public interest in privacy protections, which in turn may have influenced individuals' perceptions of data safety. On the one hand, increased attention could have amplified concerns about data misuse, reducing individuals' willingness to share. On the other hand, the publicity surrounding strengthened regulations may have reassured consumers, fostering greater trust and openness. To explore this, we analyze data from the Google Search Volume Index (SVI), which tracks interest in specific topics across U.S. states. We focus on search frequency for the term "privacy" from October 2022 to March 2023, aligning with the period covered by our platform data. However, since the SVI is reported monthly, this timeframe provides only a few pre-policy observations. To create a more reliable comparison baseline, we also collect data for the same months one year earlier — October 2021 to March 2022. This extended, seasonal matched comparison period allows us to better assess shifts in privacy-related search interest leading up to and following the policy change.

Google SVI reflects the relative search interest in privacy-related terms within each state, based on their proportion of overall search volume. The search volumes for each state are then defined relative to each other, allowing for meaningful comparisons within the selected set of states. Google Trends allows search volume comparisons across a maximum of five states at a time. Merging results from different comparisons is not feasible in this case, as Google dynamically normalizes search data based on varying reference points across queries. To ensure consistency, we focus our analysis on a set of large U.S. states: California, Virginia, Texas, Florida, and New York. These states not only capture diverse regional and demographic characteristics but also offer a robust comparison between treated and control groups, minimizing the influence of small-sample fluctuations and ensuring meaningful cross-state analysis.

We use the Google SVI data to answer three key questions: (i) whether treated states (California and Virginia) exhibited greater interest in privacy-related topics around policy implementation relative to control states, (ii) whether privacy-related searches increased over time within treated states, particularly as the policy implementation approached, and (iii) whether there was a differential increase in search interest for privacy between treated and control states, accounting for both over-time and cross-state patterns.

It is important to note that interpreting these results requires caution, as this setting differs from previous analyses. First, privacy policies received national media attention, meaning that individuals in control states may have also been exposed to discussions surrounding these policies, potentially influencing their search behavior. As a result, assuming that control states remained entirely unaffected may not be realistic. Second, defining a clear 'pre' and 'post' period is more complex. Given that discussions and media coverage of the policy likely intensified in the weeks leading up to implementation, we set the 'post-treatment' period to begin one month before the implementation date to capture shifts in search behavior leading up to the policy change.

To visually assess search trends, Web Appendix Figure D1 presents the unconditional average Google SVI values for the pre- and post-policy periods, separately for treated and control states. We observe that, while treated states already exhibited a greater interest in privacy-related topics before the policy, this interest further increased around the time of policy implementation. In contrast, the control group did not experience a similar rise; if anything, privacy-related search activity slightly declined. The higher search volume in treated states suggests public engagement with privacy issues, reinforcing the idea that individuals in these states were more attuned to privacy matters at the time of policy introduction.

To formally test for differential trends in privacy-related searches, we estimate the following regression model:

$$\mathrm{SVI}_{st} = \beta(\mathrm{Post}_t \times \mathrm{Treat}_s) + \gamma_s + \delta_t + \alpha_m + \varepsilon_{st} \quad (4)$$

where SVI_{st} represents the search interest for privacy topics in state s at month t. Post_t is a binary variable equal to one for the post-policy period (one month before implementation onward), and Treat_s is an indicator for whether a state was subject to privacy regulation (California or Virginia). The model includes state fixed effects (γ_s) to control for time-invariant state characteristics, month fixed effects (δ_t) and year fixed effects (α_m) to account for time trends in privacy-related searches.

Regression results, reported in Table 9, show that treated states exhibited significantly greater search activity for privacy-related topics compared to control states around the time of policy implementation. Column (1) estimates the effect of treated status on search interest by focusing on only one month prior to the policy onward. It finds that, on average, treated states had 6.55 indexed units higher Google SVI than control states during this period. Column (2) examines changes within treated states over time. While Web Appendix Figure D1 shows an increase in search activity unconditionally on controls, the regression estimate turns negative when state and month fixed effects are included. This suggests that the observed increase may reflect broader time trends that are accounted for in the regression. Finally, Column (3) confirms that the increase in privacy searches was 1.37 indexed units higher in treated states compared to control states around policy implementation. This reinforces the idea that privacy regulations heighten public engagement with privacy-related issues.[12]

[12]We believe that most of the observed activity is due to individuals' interest in privacy. While it is possible that some Google search activity may reflect business-side interest, we believe this is unlikely to be a major driver of the observed "privacy" search trends. First, the privacy regulations examined in this study apply only to relatively large firms. Specifically, it applies to businesses with over $25M in annual revenue or those handling data on 100,000 or more consumers and deriving at least 50% of revenue from data sales. This makes the number of firms affected in each state relatively small compared to the overall population of consumers. Second, we examined comparative Google Trends indices for search terms more likely to reflect firm compliance activity, such as "privacy compliance", "data protection" and "data security". These terms consistently show very low search volume (indices close to 0–15) relative to "privacy" (indices ranging from 40 to 100) over the same period. This suggests that the "privacy" search term predominantly reflects consumer-driven interest rather than firm-side compliance research.

Table 9: Changes in Google SVI for privacy

	(1)	(2)	(3)
Treated	6.554		
	(0.243)		
	[p = 0.000]		
Post		-2.088	
		(0.238)	
		[p = 0.000]	
Treated×Post			1.377
			(0.203)
			[p = 0.000]
Observations	595	728	1,820
Pre-treatment Mean	51.26	53.39	51.38
R-squared	0.833	0.893	0.914
Month FE	Yes	Yes	Yes
Year FE	Yes	No	Yes
State FE	No	Yes	Yes

Notes: Column (1) restricts to the post-period and compares treated vs control states with month and year fixed effects. Column (2) compares pre- vs post-period search activity within treated states month and using state fixed effects. Column (3) shows the estimation results for Equation (4) and includes both month, state and year fixed effects. The data used in this analysis is aggregated at the monthly level. Heteroskedasticity robust standard errors are reported in parentheses. Corresponding p-values (in square brackets) are reported below standard errors for each estimate.

These findings suggest that the states that adopted privacy regulations experienced a meaningful rise in public interest in privacy-related issues, particularly around the policy implementation date. This heightened attention could have had two opposing effects: it could have led to increased concern and reduced sharing if consumers became more wary, or it could have reduced concern and fostered greater sharing if consumers perceived the regulatory environment as safer. Our subsequent analysis of CE Survey data on privacy concerns shows that the latter option appears to dominate.

6.2 Changes in Privacy Concerns: Evidence from CE Survey

To explore the observed behavioral responses to privacy regulations, we analyze metadata from the CE Survey that tracks respondents' expressed concerns about privacy during the survey process. Specifically, the survey records instances where respondents explicitly cited privacy as a reason for hesitation, expressed general concerns about data collection, or refused to disclose information due to privacy reasons. Analyzing these data allows us to test whether improved perceptions of data safety due to strengthened privacy protections may have contributed to users' increased willingness to share personal information.

To estimate the impact of privacy policies on expressed privacy concerns, we use a DID framework as outlined in Equation 2. In this specification, the outcome variable is a binary indicator equal to one if the respondent expressed any privacy-related concern during the CE Survey interview. The coefficient on the interaction term captures the differential change in privacy concerns for respondents in treated states (California and Virginia) relative to those in control states, before and after the policy implementation. The model includes consumer unit fixed effects to control for time-invariant respondent characteristics and quarter fixed effects to account for common time shocks.

We validate the parallel trends assumption by conducting an event-study analysis, as detailed in Web Appendix D.2. The event study plot in Web Appendix Figure D2 demonstrates that there are no significant differences in privacy concerns between treated and control states before the policy implementation. Following the introduction of privacy regulations, however, respondents in treated states exhibit a sharp and sustained decline in expressed privacy concerns, suggesting that strengthened protections improved perceptions of data safety.

Table 10 presents the results examining changes in privacy concerns. Column (1) shows that, following the policy implementation, respondents in treated states were 5 percentage points less likely to express privacy-related concerns during the survey compared to respon-

Table 10: Changes in privacy concerns

DV: Privacy concern	Full Sample	Below Median	Above Median
Treated×Post	-0.0503	-0.0805	-0.0294
	(0.012)	(0.024)	(0.015)
	[0.000]	[0.001]	[0.050]
Observations	69,646	25,686	31,449
R-squared	0.594	0.626	0.533
Household FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes

Notes: The table presents estimation results for Equation (2) by using privacy concern indicator variable as an outcome. Each regression includes consumer unit and quarter time fixed effects. Subgroups in Columns (2) and (3) are based on users' pre-treatment median of share of expenditures in each state. Heteroskedasticity robust standard errors are reported in parentheses. Corresponding p-values (in square brackets) are reported below standard errors for each estimate.

dents in control states. This is a substantial decline given that the baseline prevalence of such concerns was around 14%. Columns (2) and (3) show that the reduction is concentrated among respondents who were initially less forthcoming with their expenditure information: privacy concerns declined by 8 percentage points among those respondents who previously shared below the median group, compared to a smaller 3 percentage point drop among those who previously shared above the above median. This pattern is consistent with our earlier finding that below-median users on the platform and in the CE Survey increased their data-sharing behavior after the policy. Together, the results suggest that the observed rise in information disclosure is likely driven by reduced privacy apprehension, especially among users who were initially more cautious.

Overall, combining the results from the Google SVI data and the CE survey, we find that in states that adopted the privacy policy, there is an increased interest in privacy-related information, and that people are less concerned about their privacy. This suggests a mechanism by which we observe an increase in data sharing in these states — users are less concerned about privacy due to increased awareness and protection. Together with the findings in Section 4, our analyses suggest that these changes in privacy perceptions impact user behavior both in contexts directly affected by the policy changes (the platform) and more broadly in contexts beyond the scope of the policy changes (the CE survey).

7 Concluding Remarks

This study provides new evidence on how privacy regulations shape consumer data-sharing behavior. Focusing on the implementation of the California Privacy Rights Act and the Virginia Consumer Data Protection Act in 2023, we find that these regulations led to a significant increase in the number of receipts shared on a customer engagement platform. In addition to the overall increase in volume, consumers also began sharing a broader variety of information across store types, geographic locations, product categories, and channels, suggesting that the regulations encouraged a more open and diversified engagement with the platform.

These behavioral shifts were not uniform across consumers. The largest increases in data sharing came from consumers who were least inclined to share before the policy change, highlighting that privacy protections may be especially effective in reducing reluctance among more hesitant individuals. Complementary analysis of the Consumer Expenditure Survey reveals similar patterns. Individuals in treated states were less likely to express privacy concerns and more likely to report their expenditures in expenditure surveys after the policies were implemented. These shifts also occurred primarily among respondents who had previously shared less, suggesting that privacy policies help close the participation gaps between different types of users. Importantly, these behavioral shifts appear to be driven by changes in perceived data protection, indicating that changes in consumer sentiment may extend the influence of privacy regulation beyond its formal scope.

Welfare Implications To interpret the economic relevance of our findings, we benchmark the effect of privacy regulation against users' responsiveness to monetary incentives on the same platform. In supplementary analysis (Web Appendix E), we show that a one-point increase in the average reward per offer, equivalent to roughly \$0.001, raises receipt uploads by 0.047 log points per user per week. This implies that the platform would need to increase payments by approximately \$0.005 per offer to generate one additional receipt upload. By comparison, the privacy policy led to an average increase of 0.38 additional receipts per user per week. This is equivalent to what the platform could achieve by raising per-offer rewards by around two points, a sizable behavioral shift in response to a structural policy change.

These back-of-the-envelope calculations suggest that privacy protections can be as influential as financial incentives in encouraging user engagement. Unlike monetary rewards, however, privacy policies may have lasting effects by shaping the institutional environment in which data-sharing decisions occur, potentially building trust and reducing perceived risks over time. Our analysis provides a revealed-preference estimate of how users value such protections, highlighting the potential of regulatory tools to promote participation in digital ecosystems without relying solely on financial incentives.

Broader Implications Our findings highlight the potential of privacy regulation not just to restrict firms' use of data but to increase consumers' willingness to share data when they feel less concerned about their data privacy. Institutional mechanisms, such as transparency requirements, opt-out controls, and well-communicated privacy measures, can foster more durable trust than short-term financial incentives. Prior research suggests that when firms adopt fair and transparent data practices, consumers are more likely to disclose personal information and accept its use in service personalization, reinforcing the idea that perceived fairness can be a powerful driver of data-sharing behavior (Culnan and Armstrong 1999). In this sense, privacy regulation may serve as a complement or even a substitute for monetary tools to promote participation in data ecosystems.

This pattern of results also highlights important variation in how individuals respond to privacy protections. The largest behavioral changes came from users who were initially least inclined to share, suggesting that regulatory safeguards can be especially effective in lowering participation barriers for more hesitant users. In doing so, privacy policies may help increase participation among a broader set of users, contributing to more balanced patterns of data sharing across the user base.

More broadly, this study advances our understanding of the behavioral effects of privacy regulations by focusing on individual responses rather than firm outcomes. In doing so, it highlights that privacy policies can shape both the volume and diversity of data shared, offering platforms and policymakers a clearer picture of how trust-enhancing regulations can support more inclusive and engaged participation in the data economy.

Limitations and Future Directions While our study provides robust evidence on the impact of privacy regulations, several limitations warrant consideration. First, our analysis centers on a specific form of data-sharing behavior — voluntary receipt uploads in exchange for rewards — which may not generalize to contexts involving more sensitive information, such as health, location, or financial data. Understanding whether and how privacy regulations influence sharing in these domains remains an important area for future research. Second, our study is limited to several months after the policy change, which allows us to investigate only the short-term impact of the changes in privacy regulation. Third, although we employ a synthetic difference-in-differences approach to strengthen causal identification, we cannot fully rule out the influence of concurrent events or unobserved factors that may affect data-sharing behavior. Fourth, while the Consumer Expenditure Survey offers valuable external evidence, its indirect measures of privacy concern may not capture the full spectrum of attitudinal shifts. Further studies using direct measures of trust and concerns could yield deeper insights into the underlying mechanisms.

Building on these limitations, several broader questions emerge. As data-sharing increases in the long term, it will be important to understand how platforms will respond, whether through changes in design, targeting strategies, or monetization models, and how such responses shape the broader market equilibrium. In addition, as privacy policies become more widespread across U.S. states, future research can explore how policy saturation influences both user expectations and platform adaptation.

References

Acquisti A, John LK, Loewenstein G (2013) What is privacy worth? *The Journal of Legal Studies* 42(2):249–274.

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

Aridor G, Che Y, Salz T (2023) The effect of privacy regulation on the data industry: empirical evidence from gdpr. *RAND Journal of Economics* 54(4):695–730, ISSN 0741-6261, URL http://dx.doi.org/10.1111/1756-2171.12455, publisher Copyright: © 2023 The Authors. The RAND Journal of Economics published by Wiley Periodicals LLC on behalf of The RAND Corporation.

Aridor G, Che YK, Hollenbeck B, McCarthy D, Kaiser M (2024) Evaluating the impact of privacy regulation on e-commerce firms: Evidence from apple's app tracking transparency .

Aridor G, Che YK, Salz T (2020) *The economic consequences of data privacy regulation: Empirical evidence from GDPR* (National Bureau of Economic Research Cambridge, MA, USA).

Arkhangelsky D, Athey S, Hirshberg DA, Imbens GW, Wager S (2021) Synthetic difference-in-differences. *American Economic Review* 111(12):4088–4118.

Athey S, Catalini C, Tucker C (2017) The digital privacy paradox: Small money, small costs, small talk. Technical report, National Bureau of Economic Research.

Chen J, Roth J (2023) Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics* .

Clarke D, Pailañir D, Athey S, Imbens GW (2023) Synthetic difference-in-differences estimation .

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

Culnan MJ, Armstrong PK (1999) Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization science* 10(1):104–115.

DellaVigna S, Gentzkow M (2019) Uniform pricing in us retail chains. *The Quarterly Journal of Economics* 134(4):2011–2084.

Demirer M, Hernández DJJ, Li D, Peng S (2024) Data, privacy laws and firm production: Evidence from the gdpr. Technical report, National Bureau of Economic Research.

Dubé JP, Bergemann D, Demirer M, Goldfarb A, Johnson G, Lambrecht A, Lin T, Tuchman A, Tucker CE, Lynch JG (2024) The intended and unintended consequences of privacy regulation for consumer marketing: A marketing science institute report. *Available at SSRN 4847653* .

Dubé JP, Hitsch GJ, Rossi PE (2018) Income and wealth effects on private-label demand: Evidence from the great recession. *Marketing Science* 37(1):22–53.

Goldberg SG, Johnson GA, Shriver SK (2024) Regulating privacy online: An economic evaluation of the gdpr. *American Economic Journal: Economic Policy* 16(1):325–358.

Goldfarb A, Que VF (2023) The economics of digital privacy. *Annual Review of Economics* 15(1):267–286.

Goldfarb A, Tucker C (2012) Shifts in privacy concerns. *American Economic Review* 102(3):349–353.

Goldfarb A, Tucker CE (2011) Privacy regulation and online advertising. *Management science* 57(1):57–71.

Gupta R, Iyengar R, Sharma M, Cannuscio CC, Merchant RM, Asch DA, Mitra N, Grande D (2023) Consumer views on privacy protections and sharing of personal digital health information. *JAMA Network Open* 6(3):e231305–e231305.

John LK, Acquisti A, Loewenstein G (2011) Strangers on a plane: Context-dependent willingness to divulge sensitive information. *Journal of consumer research* 37(5):858–873.

Johnson GA, Shriver SK, Du S (2020) Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science* 39(1):33–51.

Johnson GA, Shriver SK, Goldberg SG (2023) Privacy and market concentration: intended and unintended consequences of the gdpr. *Management Science* 69(10):5695–5721.

Kummer M, Schulte P (2019) When private information settles the bill: Money and privacy in google's market for smartphone applications. *Management Science* 65(8):3470–3494.

Lefrere V, Warberg L, Cheyre C, Marotta V, Acquisti A, et al. (2020) The impact of the gdpr on content providers. *The 2020 workshop on the economics of information security.*

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

Lin T, Strulov-Shlain A (2025) Choice architecture, privacy valuations, and selection bias in consumer data. *Marketing Science* .

Miller AR, Tucker C (2018) Privacy protection, personalized medicine, and genetic testing. *Management Science* 64(10):4648–4668.

Peukert C, Bechtold S, Batikas M, Kretschmer T (2022) Regulatory spillovers and data governance: Evidence from the gdpr. *Marketing Science* 41(4):746–768.

Sant'Anna PH, Zhao J (2020) Doubly robust difference-in-differences estimators. *Journal of econometrics* 219(1):101–122.

Schmitt J, Miller KM, Skiera B (2022) The impact of privacy laws on online user behavior. *HEC Paris Research Paper* .

Tang H (2019) The value of privacy: Evidence from online borrowers. *Available at SSRN* .

Tucker CE (2014) Social networks, personalized advertising, and privacy controls. *Journal of marketing research* 51(5):546–562.

Web Appendix

Table of Contents

These materials have been supplied by the authors to aid in the understanding of their paper. The AMA is sharing these materials at the request of the authors.

Web Appendix A Additional summary statistics

Table W1: Summary Statistics — Platform Data

Variable	**Mean**	**SD**	**Median**
Number of unique stores visited	2.22	3.49	0
Number of unique ZIP codes visited	1.43	2.02	0
Number of unique retailer categories visited	1.09	1.49	0
Number of digital receipts uploaded	0.62	2.59	0
Digital accounts connection status	0.25	0.43	0
Number of food retailers visited	0.85	2.01	0
Number of department stores visited	1.08	2.49	0
Number of health-related stores visited	0.08	0.49	0
Number of entertainment related stores visited	0.04	0.28	0

Notes: This table uses platform data and presents summary statistics for the additional variables used in the analysis. Digital accounts connection status represents a dummy variable taking the value one if a user connected her digital shopping accounts to the platform, zero otherwise.

Table W2: Summary Statistics — Consumer Expenditure Survey

Variable	**Mean**	**SD**	**Median**
Any privacy concern	0.123	0.329	0
Privacy concern	0.091	0.287	0
Data management concern	0.022	0.149	0
Personal question concern	0.017	0.130	0
Number of expenditures	34.35	13.01	33
Number of questions refused	0.103	0.839	0
Share of questions answered	0.969	0.078	1

Notes: This table presents summary statistics for the Consumer Expenditure Survey between 2019 and 2024.

Web Appendix B Details on empirical strategy and results

Web Appendix B.1 Event Study Plot for DID Estimator

To assess the plausibility of the parallel trends assumption required by a difference-in-differences (DID) framework, we implement an event-study analysis using log number of receipt uploads per week as the outcome. This approach allows us to track differences in sharing behavior between treated and control states over time, both before and after the policy implementation in January 2023. Specifically, we estimate the following model:

$$y_{it} = \sum_{w=-12}^{-1} \beta_w \text{Pre}_w \times \text{Treat}_i + \sum_{w=0}^{8} \beta_w \text{Post}_w \times \text{Treat}_i + \gamma X_{it} + \phi_i + \theta_t + \epsilon_{it} \tag{B1}$$

where y_{it} denotes the number of receipts uploaded by user i in week t. The variables Pre_w and Post_w are week-specific event-time indicators for w weeks before or after the policy, with the reference period omitted ($w = -1$). Treat_i is an indicator for whether user i is in a treated state. X_{it} is the control variables including points available per offer and recency (weeks since the last upload) and lagged variables from the previous week, such as weekly spending, the amount spent on awarded items, the number of items bought, and the number of stores visited. User fixed effects (ϕ_i) control for time-invariant differences in user behavior, while week fixed effects (θ_t) capture time trends and shocks.

Figure B1 plots the estimated coefficients β_w over time. We observe a deviation in the first three pre-treatment weeks, where treated users display significantly higher receipt uploads than control users. Given this, we adopt the SDID estimator in our main analysis to mitigate concerns of imbalance and pre-trend violations. By assigning weights to control users and pre-treatment periods, SDID improves covariate balance and better aligns the treated and synthetic control groups prior to treatment.

Web Appendix B.2 Covariate Adjustment in SDID

Synthetic Difference-in-Differences (SDID) matches treated and control units by reweighting observations to minimize differences in pre-treatment outcomes, using estimated unit and time weights. As described by Arkhangelsky et al. (2021), the baseline SDID estimator does not incorporate covariates directly into the weighting optimization. Instead, covariates can be incorporated by adjusting the outcome variable before estimation: regressing outcomes on covariates and fixed effects and then applying SDID to the residualized outcomes. Under this

Figure B1: Event study plot for difference-in-differences estimator

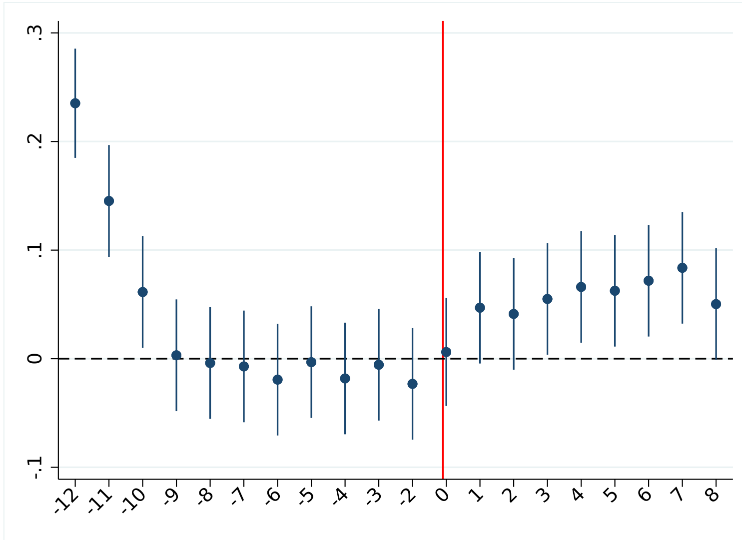

Notes: The figure plots β_w coefficients from an event-study regression using a traditional difference-in-differences framework. The red vertical line marks the policy implementation week ($w = 0$). Confidence intervals are based on robust standard errors clustered at the user level.

approach, covariates improve the precision of the estimated treatment effects without compromising consistency, provided that pre-treatment trends are conditionally parallel given covariates.

Following this methodology and its applications proposed by Clarke et al. (2023), we account for observed heterogeneity by including covariates directly in the residualization step. Specifically, we control for points available per offer and recency (weeks since the last upload) and lagged variables from the previous week, such as weekly spending, the amount spent on awarded items, the number of items bought, and the number of stores visited. These covariates capture various time-variant dimensions of user engagement and purchasing behavior that could influence receipt-uploading activity independently of the privacy policy intervention.

Web Appendix B.3 Doubly Robust DID Estimator

We include the doubly robust difference-in-differences (DRDID) estimator as a robustness check to our main specification. The method is designed for two-period panel data and combines outcome regression with inverse probability weighting (IPW) to estimate the average treatment effect on the treated (ATT). It is called "doubly robust" because it yields consistent estimates if either the outcome model or the treatment model is correctly specified.

Let Y_{i1} and Y_{i0} denote post- and pre-treatment outcomes for unit i, and define the difference $\Delta Y_i = Y_{i1} - Y_{i0}$. The DRDID estimator takes the form:

$$\hat{\tau}^{\mathrm{DRDID}} = \frac{1}{n}\sum_{i=1}^{n}\left[w_i \cdot (\Delta Y_i - \hat{\mu}_{0,1}(X_i) + \hat{\mu}_{0,0}(X_i))\right] \tag{B2}$$

Here, w_i are observation-specific weights that depend on treatment assignment and estimated propensity scores, and $\hat{\mu}_{0,t}(X_i)$ denotes the predicted outcome for control units at time t based on covariates. These components are combined to adjust for differences in covariates and pre-treatment trends between treated and control groups.

Because DRDID is designed for a two-period setting, it does not allow for dynamic treatment effect estimation. Therefore, we do not construct an event-study plot based on this method.

Web Appendix B.4 Event Study Plot for SDID

To explore the dynamics of the policy effect over time, we disaggregate the average treatment effect from the SDID estimator into a sequence of weekly event-time effects. This approach follows the methodology outlined in Clarke et al. (2023), adapted to our setting with a single treatment cohort and a single treatment time a.

The SDID estimator for the average treatment effect can be expressed as:

$$\hat{\tau}_a^{\mathrm{SDID}} = \underbrace{\frac{1}{T}\sum_{t=a}^{T}\left(\frac{1}{N_{\mathrm{treat}}}\sum_{i\in\mathcal{T}} Y_{i,t} - \sum_{i=1}^{N_{\mathrm{co}}}\hat{\omega}_i Y_{i,t}\right)}_{\text{Post-treatment average gap}} - \underbrace{\sum_{t=1}^{a-1}\hat{\lambda}_t\left(\frac{1}{N_{\mathrm{treat}}}\sum_{i\in\mathcal{T}} Y_{i,t} - \sum_{i=1}^{N_{\mathrm{co}}}\hat{\omega}_i Y_{i,t}\right)}_{\text{Weighted pre-treatment gap}} \tag{B3}$$

where $\mathcal{T}$ and N_{treat} refer to the set and number of treated units, while N_{co} is the number of control units. The weights $\hat{\omega}_i$ are assigned to control units to match the pre-treatment outcome trajectory of the treated group, and the weights $\hat{\lambda}_t$ assign relative importance to each pre-treatment time period when constructing the synthetic control.

To estimate the effect ℓ weeks after treatment, we define the event-time SDID estimate as:

$$\hat{\tau}_{a,\ell}^{\mathrm{SDID}} = \left(\frac{1}{N_{\mathrm{treat}}}\sum_{i\in\mathcal{T}} Y_{i,a-1+\ell} - \sum_{i=1}^{N_{\mathrm{co}}}\hat{\omega}_i Y_{i,a-1+\ell}\right) - \sum_{t=1}^{a-1}\hat{\lambda}_t\left(\frac{1}{N_{\mathrm{treat}}}\sum_{i\in\mathcal{T}} Y_{i,t} - \sum_{i=1}^{N_{\mathrm{co}}}\hat{\omega}_i Y_{i,t}\right) \tag{B4}$$

Each $\hat{\tau}_{a,\ell}^{\text{SDID}}$ represents the treatment effect ℓ periods after the policy change, adjusted by the weighted average of pre-treatment differences. These dynamic estimates form the basis of the event-study plot in the main text.

By construction, the average treatment effect from SDID equals the average of the post-treatment event-time effects:

$$\hat{\tau}_a^{\text{SDID}} = \frac{1}{T-a+1} \sum_{\ell=1}^{T-a+1} \hat{\tau}_{a,\ell}^{\text{SDID}} \tag{B5}$$

Standard errors and confidence intervals are obtained via a bootstrap procedure, in which units are resampled with replacement and the SDID estimator is re-estimated across multiple replications. This approach accounts for both sampling variability and the structure of the estimator.

Figure 1a and Figure 1b present the results of the event study analysis, where the outcomes are the logarithm of the number of receipts uploaded and the raw count of receipts, respectively. Each point in the figures represents the estimated weekly treatment effect $\hat{\tau}_{a,\ell}^{\text{SDID}}$ relative to the policy implementation date, with confidence intervals based on bootstrap-based inference. While the SDID estimator does not require a strict assumption of parallel trends, the flat pattern of pre-treatment estimates supports the credibility of the identification strategy. The upward shift in post-treatment estimates is consistent with a behavioral response to the privacy policy.

Web Appendix B.5 Robustness of Main Results with an Alternative Sample

Table B1: Changes in the number of receipts uploaded weekly

	(1) (SDID, logged)	(2) (DRDID)	(3) (SDID, raw)
Treated×Post	0.0875	0.0608	0.518
	(0.0008)	(0.0107)	(0.0049)
	[p = 0.000]	[p = 0.000]	[p = 0.000]
Observations	316,806	316,806	316,806
Pre-treatment Mean	4.19	4.19	4.19
Week FE	Yes	Yes	Yes
Individual FE	Yes	Yes	Yes

Notes: This table reports the estimated treatment effects of privacy policy implementation on the number of receipts uploaded per user per week by estimating Equation 1 in Columns (1) and (3) and Equation B2 in Column (2). The dependent variable is the logarithm of receipt uploads in Columns (1) and (2) and the raw count of receipt uploads in Column (3). All models include user and week fixed effects and the same covariates. The 'Pre-treatment Mean' shows the average number of receipts uploaded per user-week in the control group before policy implementation. Columns (1) and (3) use the SDID estimator with standard errors based on bootstrap-based inference. Column (2) reports results from the DRDID using the doubly robust approach, with robust and asymptotic standard errors estimated. Corresponding p-values (in square brackets) are reported below standard errors for each estimate.

To assess the robustness of our findings to random sample variation, we replicate the main analysis using a separately drawn random sample of 15,089 users. Table B1 presents the corresponding results. Across all specifications (Synthetic DID with the logged outcome, DRDID with the logged outcome, and Synthetic DID with the raw count outcome) the estimated treatment effects remain positive, statistically significant, and similar in magnitude to those obtained with the original sample. These findings reinforce the robustness of our results and suggest that they are not driven by sample-specific idiosyncrasies.

Web Appendix B.6 Subgroup Analysis Using SDID

To analyze variation in policy effects by baseline user engagement, we implement a subgroup analysis using SDID. Specifically, we estimate separate SDID regressions for users with below- and above-median pre-treatment activity levels, where activity is defined as the number of receipts uploaded prior to the policy change. These subgroup classifications are constructed within each state to account for state-level variation in user behavior. Users are assigned to the below- or above-median subgroup based on whether their average pre-policy upload rate falls below or above the median within their own state.

Unlike earlier versions of the analysis, where we used all untreated users as the control group for each subgroup, we now construct subgroup-specific control groups. That is, treated users in the below-median group are compared only to control users with similarly low pre-treatment activity, and the same applies to the above-median subgroup. This adjustment ensures that the control group provides a more appropriate counterfactual by aligning baseline engagement levels across treated and control users within each subgroup.

Each subgroup model is estimated separately using the SDID procedure described in Equation 1, which re-optimizes unit and time weights to match the pre-treatment outcome trajectories of the subgroup of interest. While this allows for a flexible and credible estimation of treatment effects within each activity group, it also implies that the synthetic control differs across subgroups. As a result, treatment effect magnitudes across columns are not directly comparable in a structural sense and should be interpreted as within-subgroup effects.

This matched subgroup approach improves interpretability by holding constant baseline behavior across treatment and control users within each comparison, and helps reduce bias that could arise from compositional differences. It also preserves the identification strength of SDID by allowing the estimator to adapt separately to the time dynamics and composition of each subgroup.

Web Appendix B.7 Event Study Plot: Consumer Expenditure Survey for Sharing Expenditures

To validate the parallel trends assumption underlying our DID analysis of the Consumer Expenditure (CE) Survey data, we implement an event-study approach. This allows us to examine the evolution of survey response behavior among respondents in treated states (California and Virginia) relative to control states, both before and after the privacy policy implementation on January 1, 2023. Specifically, we estimate the following event-study regression:

$$y_{it} = \sum_{q=-15}^{-1} \beta_q \text{Pre}_q \times T_i + \sum_{q=0}^{5} \beta_q \text{Post}_q \times T_i + \phi_i + \theta_t + \epsilon_{it} \tag{B6}$$

where y_{it} represents the survey response outcome (ratio of shared expenditures) for consumer unit i at quarter t. The variables Pre_q and Post_q are indicators equal to 1 if observation

t occurs q quarters before or after the policy implementation, respectively, with the reference period being the quarter immediately before implementation ($q = -1$, corresponding to 2022Q4). The indicator T_i equals 1 for treated respondents (California and Virginia) and 0 otherwise. Consumer unit fixed effects (ϕ_i) control for time-invariant household characteristics, and quarter fixed effects (θ_t) account for common seasonal and temporal shocks.

Figure B2: Changes in Response Behavior in CE Survey

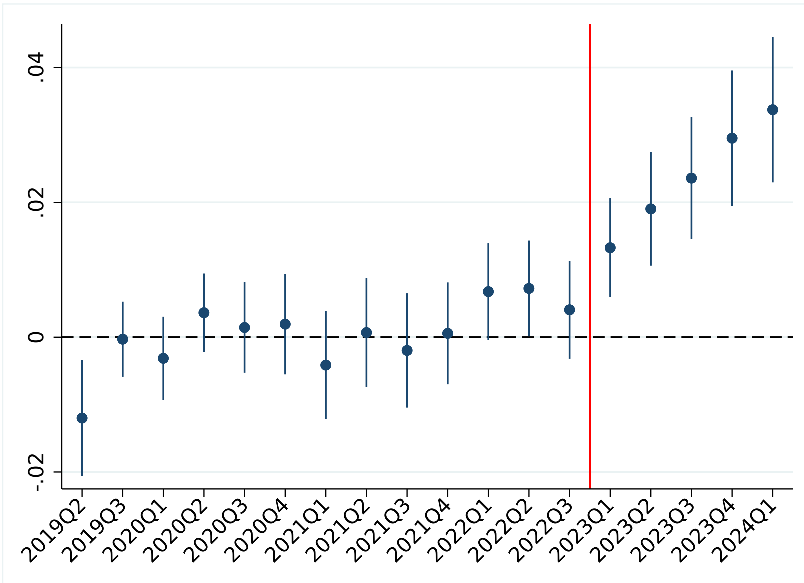

Notes: The figures plot β_k estimated from Equation B6. The red vertical line marks January 1st, 2023, when the privacy policy became effective. The outcome variable, ratio of shared expenditures, is calculated as the ratio of the number of expenditures stated in the survey to all questions asked, including refused or left unanswered for various reasons. Heteroskedasticity-robust standard errors are used for the event study regression.

Figure B2 presents the event-study estimates, plotting coefficients β_q along with their 95% confidence intervals. The vertical red line marks January 1, 2023, when privacy policies took effect. Before policy implementation, the estimated coefficients hover around zero and show no clear trend or significant differences between treated and control states. This supports the validity of the parallel trends assumption required for causal inference. Following policy implementation, however, we observe a clear and significant upward shift in the ratio of expenditures shared by respondents in treated states. This divergence indicates that the privacy policy had a measurable and positive impact on respondents' willingness to disclose expenditure information in the CE Survey context.

Web Appendix B.8 Analysis on Changes in Consumption Patterns

Table B2: Robustness Analysis for Consumption Patterns

	Total	Food	Beverage	Apparel
Treated×Post	-1.63	-1.386	-1.227	0.977
	(45.38)	(38.27)	(5.64)	(13.81)
	[p = 0.971]	[p = 0.971]	[p = 0.828]	[p = 0.944]
Observations	77,645	77,645	77,645	77,645
R-squared	0.844	0.831	0.737	0.595
Quarter FE	Yes	Yes	Yes	Yes
Household FE	Yes	Yes	Yes	Yes

Notes: This study uses data from Consumer Expenditure Surveys between 2019Q1 and 2024Q1, focusing on total expenditures and expenditures on food, beverage, and apparel. Each column reports the diff-in-diff coefficient for the regression of different types of expenditures (total, food, beverages, and apparel) on the interaction term of treated and post, household income. household and quarter fixed effects. Heteroskedasticity robust standard errors are reported in parentheses. Exact p-values (in square brackets) are reported below standard errors for each estimate.

Web Appendix C Additional results: Impacts on platform strategy

To test the validity of the parallel trends assumption and further examine the dynamics of platform behavior around the policy implementation, we first estimate an event study specification of the form:

$$y_{st} = \sum_{w=-12}^{-1} \beta_w \, \mathrm{Pre}_w \times T_s + \sum_{w=0}^{8} \beta_w \, \mathrm{Post}_w \times T_s + \phi_s + \theta_t + \epsilon_{st} \tag{C1}$$

where y_{st} represents the outcome of interest (e.g., user exposure or points per offer) for state s in week t. The variables Pre_w and Post_w are indicators for event time (weeks relative to the policy implementation), with $w = -1$ (the week just before policy) omitted as the reference period. The variable T_s equals 1 for treated states (California and Virginia) and 0 otherwise. State fixed effects (ϕ_s) absorb time-invariant differences across states, and week fixed effects (θ_t) control for common weekly shocks affecting all states. The coefficients β_w trace out the dynamics of treatment effects over time relative to the baseline week.

We examine three outcomes related to platform strategy. The first is the number of users who received each offer, capturing total exposure at the state-offer level. The second is the average number of users per offer in each state-week, measuring the intensity of offer targeting. The third is the average number of points assigned per offer, which is a proxy for the reward value attached to the offers.

To validate and complement the results from this DID-based event study, we also estimate an event study using the synthetic difference-in-differences (SDID) approach, described in Web Appendix B.4. SDID constructs a synthetic control by reweighting control states and pre-policy periods to better match the treated states' pre-policy behavior. This allows for a more robust comparison of post-policy changes, particularly in settings with possible baseline differences or modest violations of parallel trends.

Figure C1 display the event-study results for each outcome from both methods. In each panel, we observe no systematic pre-policy divergence between treated and control states, and no evidence of a meaningful change in any of the three outcomes following the policy. These results reinforce the main findings in Table 8, suggesting that the increase in data sharing was not driven by contemporaneous changes in platform-side incentives or promotional intensity.

Figure C1: Event study: Platform-side outcomes over time

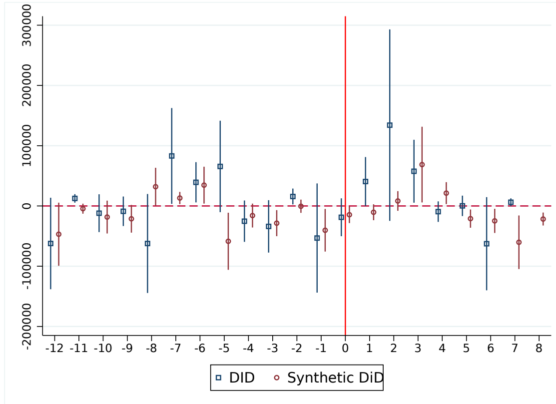

(a) Number of users with offers

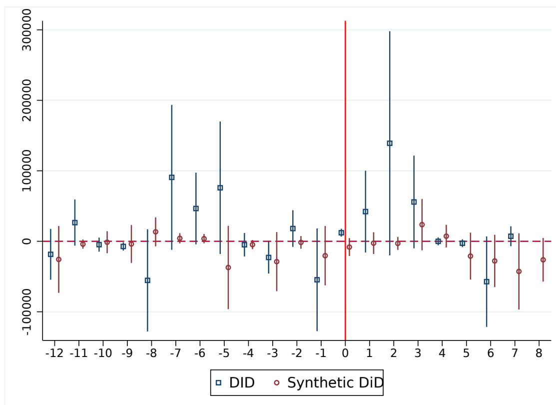

(b) Users per offer

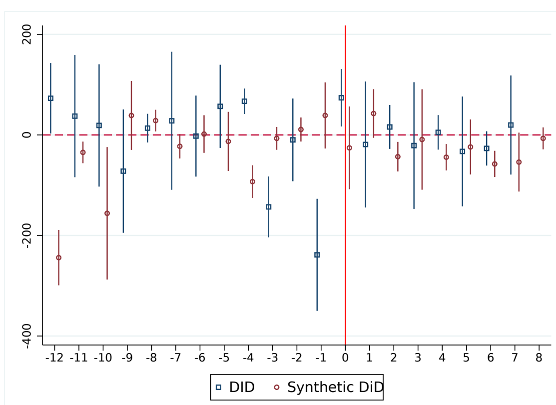

(c) Points per offer

Notes: Each panel shows coefficients from event-study regressions estimated using both DID and Synthetic DID methods. The y-axis represents the difference in the outcome variable between treated and control states relative to the week before policy implementation. Vertical red lines mark the policy implementation week. Confidence intervals for DID estimates are based on robust standard errors clustered at the state level. SDID confidence bands are constructed using bootstrap-based inference.

Web Appendix D Mechanism analysis

Web Appendix D.1 Analysis using Google SVI

Figure D1: Google Trends SVI for control and treated states

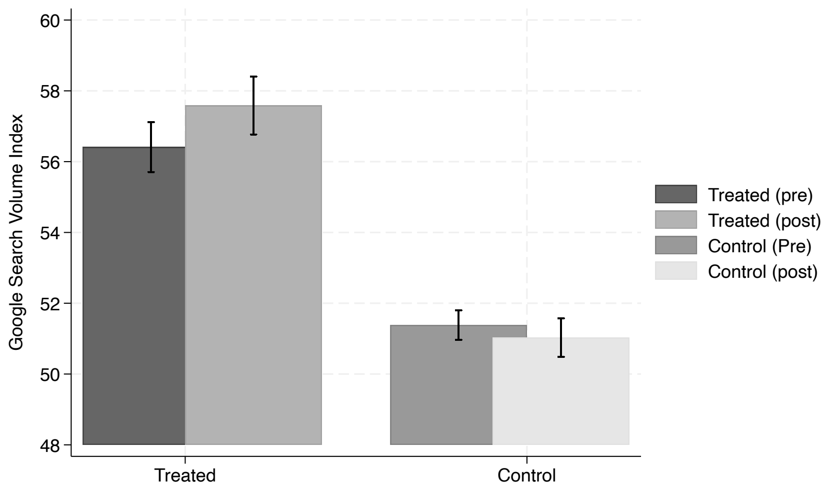

Web Appendix D.2 Event Study Plot: Consumer Expenditure Survey for Privacy Concerns

To analyze the mechanisms driving the increase in data-sharing, we also estimate the event-study specification in Equation B6 using a different outcome variable: an indicator for whether the respondent expressed any privacy concerns during the interview process. This variable captures instances in which privacy was explicitly cited as a reason for not answering expenditure questions or as a broader concern about data collection.

Figure D2 presents the results of this analysis. The estimates show no significant differences in privacy concerns between treated and control states in the pre-policy period, supporting the validity of the parallel trends assumption. Starting in the first quarter of 2023, however, we observe a sharp and sustained decline in the likelihood of respondents in treated states expressing privacy-related concerns. This post-policy divergence suggests that privacy regulations may have reduced apprehension about data sharing, particularly among those previously sensitive to privacy issues. These findings provide support for the interpretation that increased data sharing following the policy change is at least partially driven by improved perceptions of data safety.

Figure D2: Changes in Expressed Privacy Concerns during CE Survey

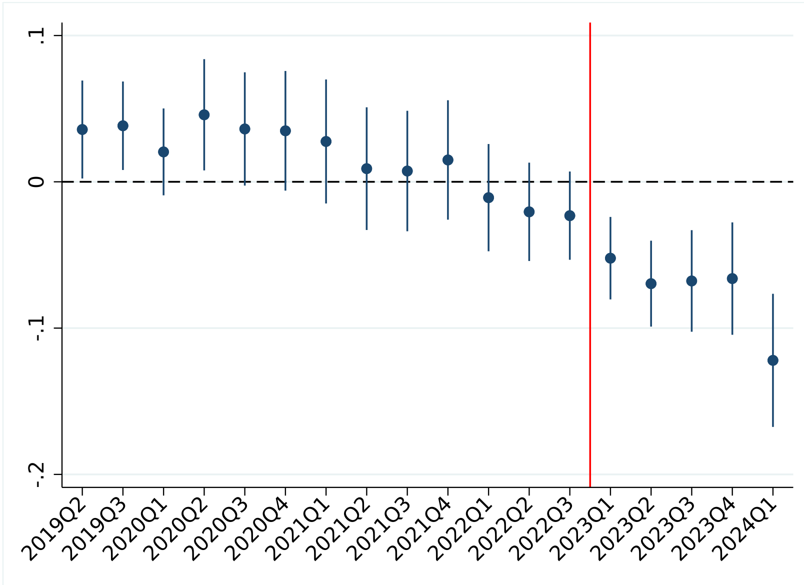

Notes: The figures plot β_k estimated from Equation B6. The red vertical line marks January 1st, 2023, when the privacy policy became effective. The outcome variable is privacy concerns which represents a binary indicator for any expressed privacy concern during the survey process during the interview process. Heteroskedasticity-robust standard errors are used for the event study regression.

Web Appendix E What predicts sharing more information?

In this section, we explore how monetary incentives influence data-sharing behavior on the platform. Specifically, we focus on predicting users' weekly receipt uploads using the average number of points per offer as the key independent variable. Since more active users may be rewarded with higher-value offers, the relationship between incentives and data sharing may suffer from endogeneity. To address this concern, we employ an instrumental variable (IV) strategy using 'treatment exposure' as an instrument for points per offer. Treatment exposure is defined as the ratio of the count of offers available to users in the treated group relative to the total offers available. It can be a plausible instrument because it captures variations in exposure to offers tied to the treatment but is unlikely to be directly related to users' endogenous choices regarding data-sharing behavior. Our IV estimation allows us to isolate the causal effect of marginal increases in points per offer on the volume of data users share.

The two-stage instrumental variable estimation is specified as follows:

$$\text{Stage 1:} \quad (\text{points/offers})_{it} = \beta_0(\text{treatment exposure})_{it} + \beta_1 X_{i(t-1)} + +\beta_2 X_{it} + \alpha_i + \gamma_t + \epsilon_{it} \tag{E1}$$

$$\text{Stage 2:} \quad \log(\text{receipts}_{it}) = \beta_0(\widehat{\text{points/offers}})_{it} + \beta_1 X_{i(t-1)} + \beta_2 X_{it)} + \alpha_i + \gamma_t + \epsilon_{it} \tag{E2}$$

In the first stage represented by Equation E1, we regress the points available per offer on treatment exposure and the control variables $X_{i(t-1)}$ and X_{it}, which include spending, the amount awarded, items, stores, and recency. User fixed effects (α) and week fixed effects (γ) are also included to control for unobserved heterogeneity. In Stage 2 in Equation E2, the log-transformed number of receipts uploaded is regressed on the fitted values of points available per offer from Stage 1, along with the same controls and fixed effects.

Table E1: Results: Instrumenting with treatment exposure

	(1) points/offers	(2) log(receipt)
Spending at t-1	-2.127	0.0241
	(0.361)	(0.0180)
	[p = 0.000]	[p = 0.087]
Amount spent on awarded at t-1	2.741	-0.0975
	(0.322)	(0.0163)
	[p = 0.000]	[p = 0.000]
Number of items bought at t-1	-0.657	0.0420
	(0.340)	(0.0165)
	[p = 0.056]	[p = 0.012]
Number of stores visited at t-1	0.526	0.108
	(0.527)	(0.0256)
	[p = 0.321]	[p = 0.000]
Recency	-0.519	-0.0752
	(0.0706)	(0.0036)
	[p = 0.000]	[p = 0.000]
Points / Offers		0.0475
		(0.00158)
		[p = 0.000]
Treatment exposure	27.76	
	(0.906)	
	[p = 0.000]	
Observations	316,806	316,806

Notes: This table uses platform data to estimate Equation E1 in Column (1) and Equation E2 in Column (2). All control variables except recent and points per offer are lagged and in logarithmic versions. The recency variable is censored at 10 weeks and more. Points/offers represent points available divided by the number of offers available per user per week. Corresponding p-values (in square brackets) are reported below standard errors for each estimate.

Table E1 presents the results from this instrumental variable estimation. Column (1) reports the first-stage results, showing the strong predictive power of treatment exposure on points per offer. The coefficient of treatment exposure is 27.76, indicating a significant relationship and confirming the relevance of the instrument. In Column (2), the second-stage results reveal a positive and significant association between points per offer and receipt uploads, with a coefficient of 0.0475.

This result suggests that, controlling for endogeneity, an increase in points per offer continues to have a positive effect on the likelihood of receipt uploading. Specifically, a one-unit increase in points per offer is associated with approximately a 4.8% increase in weekly receipt uploads. This instrumental variable approach helps address potential endogeneity concerns, yielding estimates that are less biased by reverse causality or omitted variable issues. By using treatment exposure as an instrument, we obtain a clearer understanding of the causal relationship between platform incentives and user data-sharing behavior, confirming that monetary incentives play a meaningful role in encouraging consumer engagement.

Web Appendix E.1 How much do consumers value privacy protections?

To interpret the economic significance of the privacy policy effect, we compare it to the behavioral response to monetary incentives. The IV analysis estimates how users respond to changes in the average number of points offered per data-sharing opportunity. We proceed in three steps:

Step 1. Calculating the marginal value of points per offer

The second-stage IV estimate in Table E1 shows that a one-unit increase in points per offer leads to a 0.0475 increase in the log of weekly receipts uploaded. This can be represented as:

$$\frac{d\log(\text{receipts})}{d(\text{points per offer})} = 0.0475$$

To convert this into a marginal change in the number of receipts (in levels), we multiply by the average number of receipts uploaded per user per week prior to the policy, which is 4.13:

$$\frac{d(\text{receipts})}{d(\text{points per offer})} = 0.0475 \times 4.13 \approx 0.196$$

This implies that each 1-point increase in reward per offer results in approximately 0.196 additional receipts shared per user per week.

Step 2. Translating the marginal effect into a monetary cost

One point is worth approximately \$0.001 in platform currency. Therefore, increasing the reward by 1 point per offer costs the platform \$0.001 and yields 0.196 additional receipts. The implied cost per additional receipt is:

$$\frac{\$0.001}{0.196} \approx \$0.0051$$

This provides a lower-bound estimate of how much users value their data-sharing: they require, on average, approximately 0.51 cents in reward incentives to upload one additional receipt.

Step 3. Comparing this to the effect of privacy protections

According to our main results, the introduction of privacy policies increased receipt uploads by approximately 9.2% (based on a coefficient of 0.088). Given the pre-policy average of 4.13 receipts per week, this corresponds to:

$$0.092 \times 4.13 \approx 0.38 \text{ additional receipts per user per week}$$

To achieve a similar increase in data sharing through monetary incentives alone, the platform would need to generate 0.38 additional receipts through higher rewards. Based on the marginal cost per additional receipt, this would require:

$$0.38 \times \$0.0051 \approx \$0.00194 \text{ per user per week}$$

Since one point costs $0.001, this corresponds to increasing the average reward per offer by approximately:

$$\frac{0.38}{0.196} \approx 1.94 \text{ points per offer}$$

Thus, the privacy policy had an effect equivalent to increasing the average reward per offer by about 2 points. Although modest in absolute monetary terms, this comparison illustrates that structural regulatory interventions can meaningfully influence consumer behavior in ways comparable to direct financial incentives, providing an alternative and potentially cost-effective lever for encouraging user data sharing.