Transparency is costly, but pays

Shumiao Ouyang — *University of Oxford*
Verina F. Que — *Nanyang Technological University*

August 2026

Abstract

Data consent is a central input into technological progress. To comply with privacy regulations, platforms increasingly offer privacy controls. But does being transparent about these controls help or hurt the platform? We study a platform's decision to make an existing privacy control transparent. In a randomized field experiment with 8,274 users of a large financial super-app, all users had access to a centralized list for reviewing and canceling third-party app authorizations. Treated users receive information highlighting centralized data-authorization tools; control users do not. Transparent information increases use of the highlighted privacy tools and raises downstream platform engagement and asset balances. In contrast, we detect no corresponding average change in the third-party app-authorization decisions and the acceptance rate. Heterogeneity patterns further suggest that tool use rises most among users with low baseline platform knowledge, whereas engagement responses are stronger among users who describe themselves as privacy-sensitive. Greater use of a privacy control can coexist with stronger platform activity, while the revenue-share condition determines when that combination benefits the firm.

Keywords: technology adoption; data consent; digital privacy; platform strategy; navigation friction; field experiment

JEL codes: D83, D91, L14, L86, O33

1 Introduction

Data consent is a central input into technological progress. Firms use data to improve prediction, personalization, fraud detection, matching, and the design of digital services. New artificial-intelligence systems make this input still more important by expanding the set of products that learn from users' actions and combine information across tasks. Yet firms do not own users' continuing willingness to supply data. Users must be willing to adopt data-intensive services, authorize information flows, and remain in the relationship as uses of their data evolve. Privacy and consent are therefore not compliance details appended to innovation. They are part of the organizational infrastructure through which firms deploy technology and sustain its use (Miller and Tucker, 2009; Acquisti et al., 2016; Jones and Tonetti, 2020; Bergemann et al., 2022; Goldfarb and Que, 2023).

To comply with privacy regulation, platforms increasingly provide tools through which users can review data access, withdraw prior authorizations, change personalization, or manage account relationships (European Union, 2016; Standing Committee of the National People's Congress of the People's Republic of China, 2021). Providing a control, however, is not the same as making it usable. A tool may exist while remaining buried behind menus, labeled in unfamiliar language, separated from the moment of need, or unknown to most users. The platform therefore makes two choices. It chooses whether to install a control, and it chooses how costly that control is to find and use. We call the second margin *visibility*. Visibility can change users' effective choice set without changing the formal choice set.

This implementation choice creates a fundamental tradeoff for the firm. Making an existing privacy control visible lowers users' search and navigation costs. That can increase review and cancellation of authorizations, threatening data flows, transactions, or commissions associated with the focal feature. At the same time, visible control can make adoption less risky from the user's perspective. A data decision that can be revisited is less irreversible; a user may consequently be more willing to remain active, try adjacent services, or deepen a financial relationship with the platform. Transparency can thus be costly on one margin and valuable on another. Whether it helps or hurts the platform is an empirical and organizational question, not a consequence of privacy preferences alone.

The tension is especially sharp in multi-product platforms. The immediate effect of a cancellation is local and easy to attribute to the feature or third-party app whose authorization is removed. The possible return to visible control is broader and may arrive later through continued participation, financial commitment, reputation, or adoption elsewhere in the ecosystem. A product team evaluated on short-run authorizations can prefer opacity even when the platform values a durable relationship. Conversely, a platform that internalizes activity across payments, financial services, and many third-party applications may accept a local loss in order to strengthen the ecosystem as a whole.

The visibility decision therefore depends on both the firm's revenue mix and the horizon over which it evaluates the user relationship.

The key business-model statistic is the percentage of platform revenue based on the focal feature. If a large revenue share depends on permissions that visibility places at risk, transparency is expensive. If that share is small and most revenue comes from continued activity elsewhere, the broader response can dominate. This logic explains why privacy design need not reduce to a conflict between users and firms. It also connects interface design to firm scope and organization: a diversified platform can internalize relationship value across services, whereas a specialized firm may depend much more heavily on the permission that a control allows users to cancel (Rochet and Tirole, 2003; Hagiu and Wright, 2015; Parker and Van Alstyne, 2018).

The incidence of transparency can also differ within the same ecosystem. A third-party app may bear the immediate cost when a user cancels a permission, while the host platform captures the return from greater activity or stronger financial engagement. In other settings, the host may share the loss because it earns commissions or data value from the app. Interface authority, exposure to the focal revenue stream, and ownership of the long-term relationship need not reside with the same actor. Who chooses visibility and who bears its cost are therefore questions about platform governance and firm boundaries, not only about interface design (Eisenmann et al., 2011; Hagiu and Wright, 2015).

Firm scale by itself does not resolve this tradeoff. A diversified platform can spread the fixed cost of designing and supporting a usable control across many users, and it can recover a local loss through activity elsewhere. A specialized firm may face a larger exposed revenue share and fewer margins on which the relationship pays. But large firms may also possess valuable data rents, more complex permission systems, or market power that weakens pressure to make controls usable. The relevant sufficient statistic is not size. It is the revenue exposure of the focal feature relative to the value of the broader user relationship.

This implementation margin is largely missing from standard accounts of technology adoption. Research on privacy and disclosure commonly studies the moment when a user first grants access: the default, framing, timing, or content of the consent prompt (Adjerid et al., 2019; Farronato et al., 2025). Research on attention, shrouding, and administrative burdens shows more generally that an option can be formally present yet behaviorally remote when finding or using it is costly (Gabaix and Laibson, 2006; Chetty et al., 2009; Finkelstein and Notowidigdo, 2019). We bring these ideas together by studying transparency about a control that already exists *after* initial consent. The firm does not change the underlying capability, the initial authorization prompt, or the default. It changes the cost of locating an installed tool for managing an ongoing digital relationship.

Ex-post control matters because permissions accumulate. Users authorize different applications at different times, may forget earlier grants, and may reassess the value of a service as circumstances

change. The initial prompt governs the flow of new data relationships; a centralized management tool governs the stock of existing authorizations. Lowering the cost of reviewing that stock can remove stale or unwanted access. It can also create an option value of future recourse that supports subsequent adoption. Measuring both the stock and the flow is therefore necessary to determine whether visible control simply contracts data use or instead changes the composition and durability of the platform relationship.

We study a platform's decision to make an existing privacy control transparent in a randomized field experiment with 8,274 users of a large Chinese financial super-app. The platform combines payments and financial products with an ecosystem of third-party mini-apps. All experimental users had access to a centralized Data Authorization List through which they could review and cancel third-party app authorizations. Treated users received information highlighting the centralized data-authorization tools; control users did not. The intervention did not create a new privacy capability or change the initial authorization interface. It made an installed control more visible.

The experiment holds formal access fixed and varies the information needed to use it. The treatment named the Data Authorization List, explained its centralized management function, and supplied a location cue. A user who did not know that the list existed, or expected it to be difficult to locate, thus faced lower expected navigation friction after receiving the information. We link 4,136 treatment and 4,138 control respondents to administrative records on visits to privacy tools, authorization cancellations, subsequent third-party authorization requests and decisions, requested data scopes, platform activity, and balances in platform-linked financial products. The setting lets us observe the control made visible, the data relationship that it can affect, and user activity elsewhere in the same firm.

The finance-and-mini-app architecture is particularly useful for the firm-side question. The highlighted tool manages authorizations held by third-party applications, whereas the host platform also provides payment and financial services. The focal feature and the wider user relationship are economically distinct within one ecosystem. Administrative records capture both routine engagement and asset balances, a costly measure of the financial relationship. The design can therefore follow a transparency intervention beyond the interface itself without assuming that every authorization has the same value to the host or to complementors.

Our empirical analysis follows the two sides of the platform's tradeoff. We first measure navigation to and use of the focal control: visits to the Data Authorization List, visits to its parent privacy setting, and recorded authorization cancellation. We next examine whether directing attention to privacy changes the flow of later third-party app authorizations, including request frequency, confirmation, the acceptance rate conditional on a request, and the number and sensitivity tier of requested data scopes. Finally, we measure the broader platform relationship through page views, active duration, and asset balances. The sequence separates the immediate privacy-control margin

from the wider participation margin on which visibility may produce value.

Transparent information increases use of the highlighted privacy tools. The treatment raises the probability of visiting the Data Authorization List by 2.2 percentage points, relative to a pre-treatment mean of 1.18 percent. It also raises the probability of recording at least one app-authorization cancellation by 1.8 percentage points, relative to a 3.69 percent baseline. These estimates correspond to 22 additional list visitors and 18 additional users with a recorded cancellation per 1,000 respondents. The 1.2 percentage point estimate for visits to the broader privacy setting is smaller and less precise. The largest response is thus attached to the tool highlighted by the treatment, and it extends from navigation to use.

Greater use of the privacy control coexists with stronger average activity on the platform. The treatment increases log page views by 0.024 and log active duration by 0.027. It also raises log asset balances by 0.064, although this estimate is less precise. Page views and duration capture ongoing platform engagement; balances capture a costly allocation within the financial relationship. These outcomes do not equal accounting revenue, and average effects on separate outcomes do not require that the same user both cancel an authorization and increase participation. They show that greater use of a privacy control can coexist with stronger platform activity and asset balances over the study horizon.

In contrast, we detect no corresponding average change in subsequent third-party authorization decisions or in the acceptance rate. The estimated change in observed request records is 0.0427 per user per two weeks, with a 95 percent confidence interval from -0.0127 to 0.0980. The acceptance rate among user-periods with a request changes by -0.29 percentage points, with an interval from -2.55 to 1.97 points. Requested-scope counts and maximum requested tiers likewise show no statistically detectable average change. The intervals do not establish exact equality, but they place the clearest response in management of the existing permission stock rather than a broad contraction in the observed flow of new authorizations.

Heterogeneity patterns further sharpen the economic interpretation. Use of the highlighted privacy tools rises most among users with low baseline platform knowledge, the users for whom an informational intervention should remove the most friction. Engagement responses are stronger among users who describe themselves as privacy-sensitive. These patterns need not reflect the same mechanism or the same users. Instead, they fit the distinction at the center of the paper: visibility can close a knowledge gap in the use of a particular tool while the existence of accessible control matters especially for continued activity among users who care more about privacy.

The experiment estimates the reduced-form effect of transparent information, not a single psychological mechanism. The message can teach users that the control exists, clarify its function, reduce expected path cost, attract attention to privacy, or communicate that the platform has invested in user control. The firm's decision can nevertheless be organized using observable behavioral margins.

Control visits and cancellation measure exposure on the focal feature; authorization decisions measure whether that exposure spills over to later data-sharing choices; page views, duration, and asset balances measure the broader relationship. This approach asks which economically relevant behaviors move before applying the platform's revenue weights.

The pattern of outcomes makes the firm-side tradeoff concrete. A response confined to list visits would show that information changes navigation without changing exercise. Cancellation without a participation response would expose the focal feature without evidence of an offset elsewhere. A decline in later requests or acceptance would show that the negative consequences of privacy attention spread from the stock of existing permissions to the flow of new ones. The observed combination—more use of the highlighted control, stronger platform activity and balances, and no detectable average change in subsequent authorization decisions—is precisely the case in which the firm's revenue mix determines whether transparency pays.

We formalize this logic with a parsimonious model whose elements map directly to the empirical setting. On the user side, a person uses an installed control when the expected value of reviewing prior authorizations exceeds the friction of finding and navigating to it. Treatment assignment measures visibility; pre-treatment platform knowledge and prior setting use proxy for baseline friction; privacy sensitivity captures heterogeneity in the value of control; and list visits and cancellations measure navigation and exercise. The model does not require additional latent channels to generate its main prediction: reducing navigation friction moves users near the action threshold toward the tool.

On the firm side, let s denote the percentage of platform revenue tied to the focal authorization-dependent feature. Let ℓ_H denote the proportional loss in focal-feature revenue and g_H the proportional gain in all other platform revenue over horizon H. The platform's proportional revenue change is $(1-s)g_H - s\ell_H$. When transparency both exposes the focal feature and strengthens activity elsewhere, it raises revenue if $s < g_H/(g_H + \ell_H)$. The cancellation outcomes diagnose the behavioral source of ℓ_H, while engagement and asset balances diagnose margins that can contribute to g_H. Internal revenue data are needed to monetize those responses and measure s. The condition separates short-run exposure from the longer-run relationship return without treating cancellations, page views, and balances as if they shared a unit.

The paper makes four contributions. First, it places privacy governance inside the economics of technological adoption. Data-intensive products require not only technical capability and financing but also users' willingness to enter, maintain, and revise data relationships. Prior work shows that privacy protection can shape technology diffusion and firms' deployment of digital services (Miller and Tucker, 2009; Tucker, 2014). We show that making a reversal mechanism usable can accompany stronger platform participation even as users exercise a particular control more often. Privacy infrastructure is therefore an input into technology deployment, not simply a constraint imposed

after innovation.

Second, the experiment identifies an implementation margin distinct from both technological capability and legal availability. Firms design not only which controls exist but also the effective user cost of reaching them. Naming, explanation, placement, path length, and feedback are intangible product investments that can change behavior without changing the installed feature. This extends research on attention and shrouding from attributes observed at purchase to management of an ongoing digital relationship. It also shows why evaluating privacy controls only by their formal existence misses a margin on which firms continue to choose.

Third, we derive a firm-side revenue-share condition for the supply of transparency. Existing work emphasizes the productive value of data, the costs of privacy loss, and the allocation of surplus across platform sides (Acquisti et al., 2016; Goldfarb and Que, 2023; Jones and Tonetti, 2020). Our condition gives the visibility decision a measurable business-model foundation. The same user response can favor transparency at a diversified platform and opacity at a firm concentrated in the exposed feature. It also clarifies why a short-horizon feature team and a platform-wide governance group may rationally prefer different designs.

Fourth, we separate management of the stock of existing permissions from the flow of new authorization. Initial prompts determine whether new data relationships begin; centralized ex-post management determines whether an accumulated portfolio remains in place. Measuring both shows why greater use of control over the stock need not produce a comparable contraction in the later flow. The distinction applies beyond privacy to subscription management, account deletion, data portability, appeals, and permissions delegated to artificial-intelligence agents.

Together, these contributions speak directly to how firms organize, finance, and deploy technology. Visibility is a product-design choice, an investment in intangible implementation capital, and a governance decision about which unit internalizes the user relationship. Its return varies with revenue mix, organizational horizon, and platform scope. The title's claim is correspondingly precise: transparency is costly because it increases exercise of a focal privacy control; it pays when the broader relationship return exceeds that exposure, as determined by the revenue-share condition.

The remainder of the paper proceeds as follows. Section 2 develops the navigation-friction and revenue-share model. Section 3 positions the contribution in research on attention, privacy, technology adoption, and platform strategy. Section 4 describes the setting, treatment, data, and empirical approach. Section 5 reports the results and connects each outcome family to the model. Section 6 develops implications for firm organization, technology deployment, market structure, and policy. Section 7 concludes.

2 Navigation Friction and the Platform's Visibility Choice

This section presents a deliberately spare model. A user faces a cost of reaching an existing privacy control, and visibility lowers that cost. The platform compares any revenue lost on authorization-dependent activity with the revenue gained from the user's broader relationship with the platform. The model contains no latent trust, assurance, recourse, or welfare parameter. Each object is either an outcome in the experiment, a pre-treatment variable, or an accounting quantity that can be reported as a transparent calibration.

2.1 Users: visibility lowers navigation friction

Let $v \in \{0, 1\}$ indicate whether the Data Authorization List is made visible. In the experiment, v is represented by randomized assignment to the additional information item. The underlying list is available under both values of v; visibility changes the cost of finding and reaching it, not the set of actions that the list permits.

Let c_i be user i's pre-treatment navigation friction. Its empirical index uses only pre-treatment information: reported knowledge of where platform controls are located and pre-period visits to the Data Authorization List or its parent privacy setting. Higher knowledge and prior navigation imply lower c_i. Let b_i be the value of reviewing the permission portfolio, represented by the common pre-treatment measures of privacy concern, desired protections, and self-reported privacy sensitivity. Visibility reduces the effective navigation cost to

$$c_i(v) = c_i - \delta v, \qquad \delta \geq 0. \tag{1}$$

The randomized change in use of the highlighted tool represents δ; it is not interpreted as minutes, clicks, or a psychological mediator. User i uses the list when

$$M_i(v) = \mathbf{1}\{b_i \geq c_i(v)\}. \tag{2}$$

Here M_i is the observed post-assignment indicator for visiting the Data Authorization List. Equations (1)–(2) imply $M_i(1) \geq M_i(0)$: lowering navigation friction weakly increases use of an already installed control. The increase should be concentrated among users for whom b_i lies between $c_i - \delta$ and c_i. In this setting, that prediction is evaluated directly by interacting assignment with the pre-treatment platform-knowledge indicator. The broader privacy-setting visit is a second navigation outcome; recorded cancellation is an action taken after reaching a control, not a measure of navigation cost itself.

This prediction is intentionally narrow. The information item may also draw attention to data sharing. The model does not name that attention response as an unobserved mechanism. Instead, it

measures any economically adverse consequence in the authorization outcomes themselves: recorded cancellations, the number of new authorization confirmations, and the confirmation rate among requests. Thus greater use of the highlighted tool need not imply a decline in authorization-dependent activity.

2.2 The platform: a revenue-share condition

The platform chooses whether to make the control visible. Fix a common evaluation horizon H, normalize pre-policy revenue over that horizon to one, and let

$$s = \frac{\text{revenue tied to authorization-dependent mini-app activity}}{\text{total platform revenue}} \in [0, 1]. \tag{3}$$

Unlike an abstract marginal value of data, s is a percentage with a direct accounting interpretation. We state the platform condition for every $s \in [0, 1]$ and report its break-even value.

Visibility can move two observed activity margins. Let $g_H \geq 0$ denote the proportional gain in revenue from the rest of the platform relationship through H. In the experiment, g_H is represented by the treatment effects on page views, active duration, and platform-linked asset balances. Each outcome can be used separately, with an explicitly stated revenue conversion, rather than combined into an opaque index. Let $\ell_H \geq 0$ denote the proportional loss in focal- feature revenue through the same horizon if making data sharing salient reduces authorization-dependent activity:

$$\ell_H = \max\left\{0, -\frac{\mathbb{E}[A_i(1) - A_i(0)]}{\mathbb{E}[A_i(0)]}\right\}, \tag{4}$$

where A_i is one authorization-activity measure at a time, oriented so that larger values mean more focal-feature activity and normalized to the control mean. The empirical counterparts for A_i are the unconditional number of recorded confirmations, the conditional confirmation rate, and one minus the recorded- cancellation indicator. Platform accounts translate each activity into focal revenue. When an authorization outcome does not decline on average, its point calibration sets $\ell_H = 0$.

Let $R_H(v)$ denote total platform revenue through H. Once focal and nonfocal responses are expressed in revenue units, the revenue effect of visibility is

$$\frac{R_H(1) - R_H(0)}{R_H(0)} = (1 - s)g_H - s\ell_H. \tag{5}$$

The first term is the gain on the non-focal share of the platform; the second is the short-run cost on the share exposed to authorization salience. This is the platform's tradeoff. For $g_H + \ell_H > 0$, visibility pays weakly when

$$s \leq s^*_H \equiv \frac{g_H}{g_H + \ell_H}. \tag{6}$$

The result is immediate. A small focal-feature revenue share makes transparency more attractive because the platform risks little revenue on the highlighted margin while retaining the broader engagement gain. A large focal-feature share makes the same adverse authorization response more costly. If the data show $g_H > 0$ and no average decline in authorization activity, the point estimate places s_H^* at one; wider confidence intervals for ℓ_H yield a range of break-even shares. The experiment therefore supplies the reduced-form changes used to calibrate g_H and ℓ_H, while the platform's revenue share determines whether their combination pays. Any accounting application uses one common horizon H for the focal and nonfocal revenue responses.

3 Related Literature and Positioning

The implementation layer connects literatures that are often treated separately. Work on inattention and shrouding explains why formally available attributes or remedies may not enter choice. Privacy research studies how information, control, and context affect data-related decisions. Research on technology adoption shows that privacy can shape the diffusion and performance of data-intensive products. Platform economics clarifies why the host may internalize both permission-dependent value and ecosystem-wide relationship effects. Research on consent and choice architecture distinguishes initial authorization from ex-post management. We connect these domains by treating the visibility of an installed control as a firm investment that changes both user behavior and the return to a data-intensive technology.

3.1 Inattention, shrouding, and costly exercise

An option can exist without affecting behavior when consumers do not notice it or face costs in acting. Models of shrouded attributes show that firms may hide dimensions on which informed consumers would respond, particularly when attention is costly and competition does not fully unravel the strategy (Gabaix and Laibson, 2006). Salience can change the effective weight placed on a price or attribute without changing the underlying menu (Chetty et al., 2009; Bordalo et al., 2013). Empirical work on benefit take-up and administrative burdens shows that information, assistance, and procedural design can determine whether formally eligible households receive a benefit (Finkelstein and Notowidigdo, 2019). The same attention-cost logic applies when fees or product terms are difficult to monitor (Gabaix and Laibson, 2006).

Digital controls share this wedge but add three features. First, they are often separated from the transaction that creates the underlying relationship. Consent may occur inside an app, whereas ex-post management sits in a central setting. Second, the platform can change discoverability continuously through interface updates without changing formal product terms. Third, the platform may govern a portfolio of third-party relationships rather than a single bilateral contract. The

relevant burden is therefore not only reading an attribute at purchase. It is locating, understanding, and exercising a portfolio-management capability later.

Our model imports the attention-cost insight without presuming a pure shrouding equilibrium. Section 2 gives a simple threshold: a user reaches the installed control when its value exceeds the navigation cost of finding and using it. The experiment varies information about the control while holding the formal capability fixed. Its contribution is therefore an implementation contrast—visibility of an installed option—combined with measurement of the focal control, later authorization, and broader platform participation.

3.2 Privacy valuation, context, and control

The economics of privacy emphasizes both the productive value of information and the heterogeneous costs of data use (Acquisti et al., 2016; Goldfarb and Que, 2023). Individuals may dislike observation intrinsically, anticipate instrumental harms, value personalized services, or confront uncertainty about future use. Small changes in incentives, defaults, or context can alter disclosure, and stated concern need not map simply into behavior (Athey et al., 2017; Acquisti et al., 2012; Lin, 2022). Contextual integrity further emphasizes that the acceptability of a data flow depends on actors, information type, purpose, and transmission norms rather than a scalar taste for secrecy (Nissenbaum, 2010).

Control can enter these decisions through several channels. An available choice can have intrinsic value, even when the chosen allocation is unchanged (Bartling et al., 2014). Perceived control can change disclosure or response to personalization (Brandimarte et al., 2013; Tucker, 2014). Privacy communication can also prime concern, interrupt behavior, or alter interpretation of downstream choices (Burtch et al., 2015; Adjerid et al., 2019). These mechanisms have different welfare and policy implications, but the same behavioral pattern can be consistent with more than one of them.

The paper therefore treats visibility as a reduced-form implementation investment rather than equating it with trust. The Q6 item names the control, explains its function, and cues its location. These elements lower the expected cost of finding the tool while also making privacy salient. The experiment is designed to measure their combined economic effect on control use, authorization, and platform participation.

This distinction also changes how contextual-authorization outcomes are read. If future revocation has option value, visibility can make a new authorization less costly. If the intervention raises privacy salience, it can make the same authorization less attractive. New request exposure can change because platform activity and app mix change. A conditional confirmation rate then describes the selected request set, not a primitive privacy preference. The result is a theoretically ambiguous outcome family, not a placebo through which one mechanism can be selected.

3.3 Data as a platform input and strategic visibility

Digital platforms coordinate interactions across user and complementor sides (Rochet and Tirole, 2003; Hagiu and Wright, 2015; Parker and Van Alstyne, 2018). Data can improve matching, personalization, prediction, and product development, but market incentives for collection and protection need not coincide with social value (Jones and Tonetti, 2020; Bergemann et al., 2022; Acemoglu et al., 2022; Ichihashi, 2020). Privacy rules and platform architecture can consequently reallocate surplus across users, advertisers, developers, and the host (Goldfarb and Tucker, 2011; Goldfarb and Que, 2023). Platform data strategy and privacy regulation can also change competition and market structure (Bhargava et al., 2020; Campbell et al., 2015).

The present setting adds an ex-post governance choice. A host platform can make third-party authorization management easier to find, potentially reducing some complementors' continuing access. At the same time, the host may benefit from users' willingness to remain active across the ecosystem. The host and a complementor can therefore face different incidence. Equation (5) represents this tension using the percentage of platform revenue tied to the focal feature, the proportional focal loss, and the proportional revenue gain elsewhere on the platform.

This condition gives the strategy contribution a measurable business-model foundation. The experiment supplies the behavioral margins underlying focal exposure and broad participation. Feature-level and platform-level accounts translate those responses into revenue and supply the focal share. A third-party permission-management control may expose a small percentage of a host's total payment and financial revenue even when it is important to the affected complementor. A diversified platform can consequently prefer visibility while a specialized firm facing the same user response prefers opacity. The prediction concerns revenue exposure rather than firm size by itself.

3.4 From ex-ante consent to ex-post governance

Most privacy choice occurs at the moment a service requests data. Research on disclosure and consent therefore studies defaults, framing, timing, and the architecture of an initial prompt (Adjerid et al., 2019; Tsai et al., 2011; Farronato et al., 2025). That margin is important, but digital relationships continue. Users accumulate permissions; app purposes and personal values may change; forgotten services may retain access; and a user may later want to review several grants together.

Centralized ex-post management differs in economic function from centralized ex-ante consent. The initial prompt governs a flow of new decisions. A management list governs a portfolio of existing authorizations. The latter can reduce the cost of review across multiple apps and create an option value of future recourse. Yet centralization alone does not guarantee behavioral access. A consolidated control that is unknown or hard to navigate remains formally available but behaviorally remote.

The experiment isolates neither centralization nor capability. Supplied materials state that both arms had the same Data Authorization List. The assignment adds information about that list after five common survey modules. The study therefore complements research that changes initial consent architecture rather than replicating it. It asks whether information about ex-post management changes the use of existing controls, the flow of later requests, and platform participation.

The legal motivation is adjacent but separate. For covered processing based on consent, Article 7(3) of the European Union's General Data Protection Regulation addresses ease of withdrawal, and Article 15 of China's Personal Information Protection Law calls for a convenient withdrawal method (European Union, 2016; Standing Committee of the National People's Congress of the People's Republic of China, 2021). The provisions differ in actor, basis, scope, force, and remedy. They motivate an economic question about usability; the experiment studies a product control and does not interpret its recorded cancellation event as a legal conclusion.

3.5 Disclosure, explanation, and the implementation layer

Information interventions can work through learning or salience. Labels and warnings can change demand when they reveal an attribute, alter attention, or provide a reference point (Bollinger et al., 2011; Araya et al., 2022; Adalja et al., 2023). In digital markets, privacy information can affect purchasing, disclosure, browsing, and response to personalization, but effects depend on baseline knowledge, placement, context, and the decision being measured (Tsai et al., 2011; Tucker, 2014; Burtch et al., 2015). A message that describes an available control is simultaneously informational and performative: it can teach a path, remind the user of risk, and communicate something about the platform.

This paper's term "implementation technology" should therefore not be read as "information only." Visibility can include naming, explanation, location, path shortening, or persistent access. The field experiment bundles a concise set of these elements in Q6. The treatment is a specific visibility investment, while the economic logic extends to other ways firms can reduce navigation friction.

The distinction is particularly useful for emerging digital controls. Account deletion, appeal systems, AI-agent permissions, model-memory controls, and data portability all require an implementation path after a formal capability is declared. These systems differ in agency, reversibility, provider boundaries, and propagation of action, but they share a firm choice over how costly the installed option is to exercise.

3.6 Contribution

The paper sits at the intersection of four questions. From consumer attention, it takes the idea that formal options may be inert when action is costly. From privacy economics, it takes the possibility that control, salience, service value, and heterogeneous concerns jointly determine behavior. From platform economics, it takes the host's need to internalize both permission-dependent value and cross-side relationship effects. From choice architecture, it adds a distinction between the flow of initial decisions and ex-post management of an existing portfolio.

The result is a paper about the economic implementation of digital control. The field design asks whether information makes an existing control behaviorally consequential. The platform model asks when a host benefits from that response. Combining the two reveals why transparency can activate a privacy control and strengthen platform activity at the same time, and why the firm's preference ultimately depends on the revenue share placed at risk.

4 Setting, Experimental Design, Data, and Empirical Strategy

The experiment changes the visibility of an existing privacy tool. It does not create a new privacy capability or alter the interface at which a user first decides whether to authorize a third-party app. This distinction is central to the design: treatment lowers the informational and navigational friction around an installed control while holding the underlying menu of controls fixed. We first describe the platform and its two authorization margins, then present the field experiment, administrative outcomes, and empirical specifications.

4.1 Institutional setting

4.1.1 A financial super-app and its third-party ecosystem

The Alipay platform is a large financial super-app that combines payments and financial products with an ecosystem of third-party mini-apps. Users can reach transportation, local commerce, public services, and many other functions without leaving the host platform. The host supplies the technical infrastructure and ecosystem rules, while individual mini-apps request access to particular user information when providing a service. A user can therefore hold many continuing app relationships and encounter additional authorization requests over time.

Data authorization is an important input into this production architecture. Permission to use information allows mini-apps to personalize service, complete transactions, and reduce verification costs. At the same time, persistent permissions can create privacy concerns that reduce users' willingness to engage with the ecosystem. Privacy controls are consequently part of the platform's technology: they shape both the terms on which data enter production and users' willingness to

maintain a long-lived platform relationship.

The platform's financial dimension makes downstream participation economically meaningful. Prior work shows how the platform's payment infrastructure produces information that can expand access to credit (Ouyang, 2026). Administrative records allow us to follow page views, active duration, and balances in platform-linked financial products. Page views and duration capture the breadth and intensity of platform activity. Asset balances capture a more consequential relationship margin because they require users to allocate funds to products inside the platform. None of these outcomes is, by itself, a measure of profit or welfare; together they describe whether transparency strengthens or weakens the user's broader relationship with the firm.

4.1.2 Two authorization margins

The setting contains two distinct authorization margins. At the *contextual-authorization margin*, a mini-app requests one or more data scopes when the user seeks a service. The user can confirm the request or leave no recorded confirmation. Service value, requested information, and app identity vary across encounters, so the number of requests reflects both user participation and the set of apps reached.

At the *portfolio-management margin*, the Alipay platform offers a Central Privacy Setting that contains a Data Authorization List. The list brings third-party app authorizations into one location and allows a user to review and cancel them. Both treatment and control users had access to this installed tool. The intervention changes whether the tool is brought to the user's attention, not whether it exists.

The two margins play different roles in the model. Visits to the list and recorded cancellations reveal whether visibility overcomes the cost of finding and using an existing control. Later authorization requests and confirmations measure a potential short-run downside: making privacy salient could discourage permission-dependent activity. Platform engagement and financial balances capture the broader relationship margin through which visibility can generate a longer-run return.

Throughout the paper, a "recorded authorization cancellation" refers to the administrative event observed by the platform. It captures the exercise of an installed control but does not imply deletion of data previously transferred to a third party. Similarly, "no recorded confirmation" is broader than an active rejection: it can include dismissal, abandonment, or timeout. These definitions keep the empirical claims aligned with the events recorded in the administrative data.

4.2 Experimental design

The experiment was embedded in a privacy survey delivered through the platform's front-page message box in May 2024. Figure 1 shows the survey period, assignment, linkage to administrative

records, and the event windows used in the analysis.

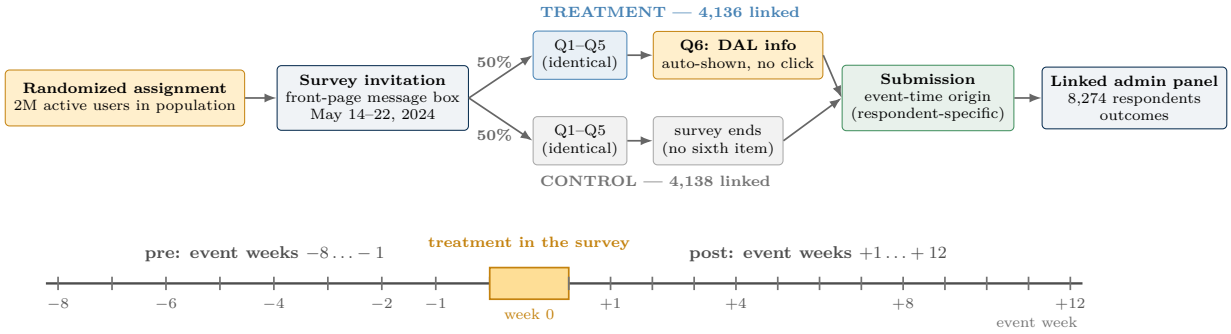


Figure 1: Experimental Timeline and Analysis Window

Notes: The survey was distributed from May 14 through May 22, 2024, through the Alipay platform's front-page message box. The broader rollout targeted two million accounts. The linked experimental sample contains 4,136 treated users and 4,138 control users. Event time is defined relative to each user's survey date. The engagement panels cover eight pre-treatment weeks and twelve post-treatment weeks; authorization panels use biweekly event periods. Period −1 is the omitted reference period in the event-study specifications.

4.2.1 Assignment and analysis sample

Accounts in the broader rollout were randomly assigned to one of two survey versions. Both versions contained the common Q1–Q5 modules and differed only in the additional Q6 information item. The linked respondent sample contains 8,274 users: 4,136 assigned to treatment and 4,138 assigned to control. Randomization produced nearly equal linked arm sizes and separated the visibility intervention from the underlying availability of the Data Authorization List.

The experiment was part of a much broader messaging campaign, while the paper's analysis population is the set of respondents linked to the administrative data. We therefore report treatment effects for these 8,274 users rather than extrapolating estimates to the full two-million-account rollout or to every user of the Alipay platform. The experimental sample is especially informative about people who engage with platform communications and can be followed across survey and administrative records.

Because outcomes are observed before and after the survey, the main estimates compare changes in the treatment group with changes in the control group. The pre-period absorbs persistent differences in users' activity and permission portfolios, while randomized assignment identifies the effect of receiving the additional information.

4.2.2 Information about an existing privacy tool

Both arms first completed five common privacy modules, Q1–Q5. The control survey ended after those modules. The treatment survey continued to Q6, which named the Data Authorization List,

explained that it could be used to centrally manage mini-app authorizations, and showed users where to find the tool. Figure 2 isolates this treatment–control difference.

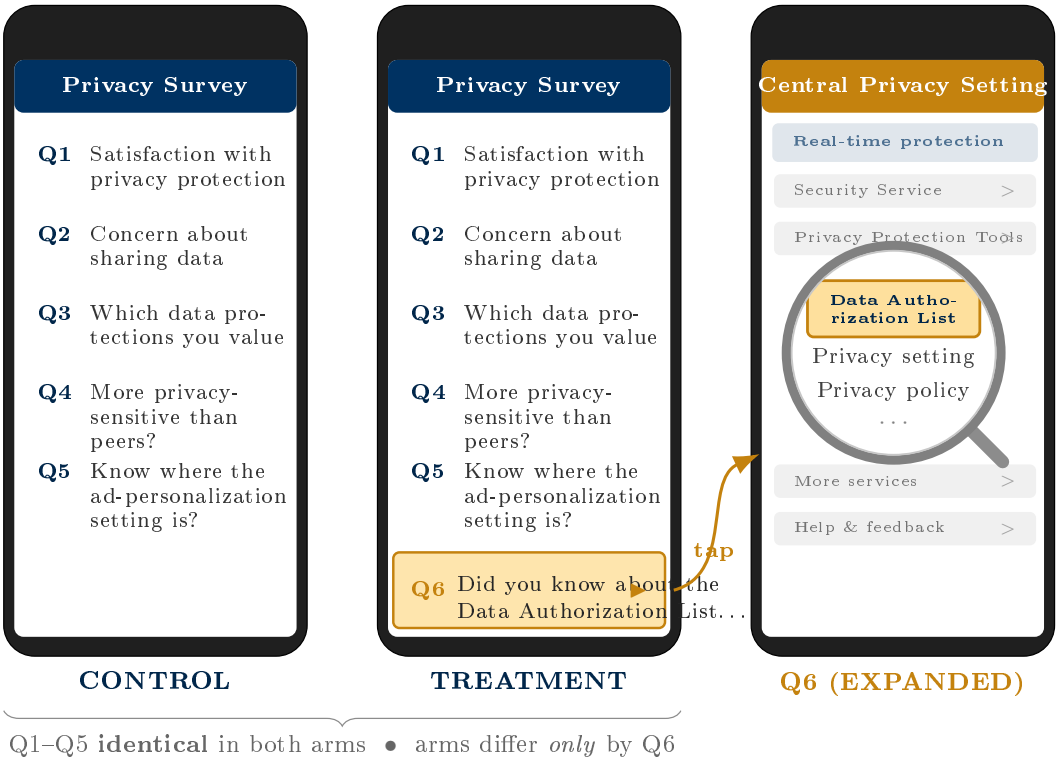


Figure 2: Information Treatment

Notes: Both groups answered the same five privacy questions. Treated users then received an additional item that named the Data Authorization List, described its centralized-management function, and provided a visual location cue. Control users did not receive this item. The intervention changed information about an existing control; it did not create the control or alter the initial mini-app authorization prompt. English text in the figure is a translation of the platform survey. Q6 was displayed automatically; the "tap" arrow in the schematic denotes expansion of the location illustration.

The treatment is intentionally light-touch. It resembles the kinds of explanatory messages that product teams can add at low marginal cost without redesigning the underlying permission system. That feature makes the experiment useful for studying transparency as a deployment decision: the platform can change awareness and expected navigation cost while leaving the installed technology fixed.

The treatment should nevertheless be interpreted as an information bundle. Q6 names the tool, explains its function, supplies a location cue, and directs attention to privacy. The design therefore identifies the reduced-form effect of making the existing control salient and easier to find. It does not separately identify learning, navigation, reassurance, or privacy salience. The model uses these channels to organize the observable margins rather than treating any one of them as an independently estimated mediator.

4.2.3 Baseline survey measures

The five common modules record the user's baseline relationship with privacy and the platform:

Q1. satisfaction with the platform's privacy protection on a five-point scale;

Q2. concern about sharing data with the platform;

Q3. desired protection principles, including minimum collection, prevention of leakage, restrictions on third-party sharing, protection against price discrimination, and data retrieval;

Q4. relative self-reported privacy sensitivity and instrumental motives for sharing; and

Q5. knowledge of several platform privacy settings, including the central privacy setting, advertising-personalization controls, and app-authorization cancellation.

These measures serve two purposes. First, they describe the experimental sample and allow a check that treatment assignment is unrelated to observable baseline characteristics. Second, they capture theoretically relevant heterogeneity. Users with less knowledge of platform settings plausibly face higher initial search and navigation costs. Users who describe themselves as relatively privacy-sensitive may place greater value on the ability to monitor and revise permissions. The moderator analyses use these pre-treatment measures rather than outcomes elicited after the intervention.

4.3 Treatment contrast and estimand

Let Z_i equal one when user i is assigned to the treatment survey and zero otherwise. For outcome Y_i, define $Y_i(1)$ and $Y_i(0)$ as potential outcomes under the two survey versions. The primary estimand for the linked experimental sample $\mathcal{P}$ is the intention-to-treat effect

$$\tau_{\mathcal{P}} = \mathbb{E}\left[Y_i(1) - Y_i(0) \mid i \in \mathcal{P}\right]. \tag{7}$$

This estimand captures the effect of the platform's decision to show the Q6 information bundle. It does not condition on whether a user subsequently visits the list, cancels an authorization, or remembers the message. Those actions are outcomes of the intervention. The intention-to-treat contrast is therefore the policy-relevant effect of deploying the visibility message to an eligible respondent in this experimental sample.

The analysis assumes that one user's survey assignment does not change another sampled user's outcomes. Given the small share of the platform population in the linked sample and the individual nature of the message, large cross-user spillovers are unlikely over the study horizon. Any informal sharing of the tool with control users would tend to compress the measured treatment contrast.

4.4 Administrative data and outcome construction

The survey records are linked at the user level to platform activity, financial, and app-authorization logs. We group the outcomes into the three margins in the model. This organization is economic rather than chronological: it separates the immediate use of an existing control, the flow of new authorization decisions, and the user's broader relationship with the platform.

4.4.1 Use of the highlighted privacy tools

The proximal outcome family contains three binary measures over defined pre- and post-treatment windows: any visit to the Data Authorization List, any visit to the broader Central Privacy Setting, and any recorded app-authorization cancellation. The list visit is the closest behavioral counterpart to the information shown in Q6. The broader-setting visit captures navigation toward the parent privacy interface. Cancellation captures exercise of an authorization control after the user reaches a privacy-management decision.

Together these outcomes distinguish attention from consequential use. A list visit shows that the control entered the user's feasible choice set; a recorded cancellation shows that at least one existing authorization was changed. The data do not require every visit to culminate in cancellation: a user may inspect the portfolio and decide that existing permissions are worth keeping.

4.4.2 New authorization decisions

The second family measures the flow of authorization requests after treatment. We construct the number of observed requests, the number with a recorded confirmation, and the number without a recorded confirmation in each user-by-biweekly period. Periods with no request enter the count outcomes as zeros. We also calculate the recorded-confirmation rate—called the acceptance rate in the figures—among user-periods with at least one request.

The total number of requests is partly an activity outcome: a user who visits more mini-apps can encounter more prompts even if her response to a given prompt is unchanged. The conditional confirmation rate instead describes choices among observed request opportunities. Because treatment can affect which users and apps generate a request, we interpret the rate as the realized confirmation margin rather than a fixed-preference parameter.

For each observed user–app–day request, we also measure the number of distinct data scopes requested and the maximum tier on the study's four-level requested- scope classification. These outcomes test whether transparency changes the composition of permission demands encountered in the ecosystem. They describe requested fields, not the full stock of previously shared data.

4.4.3 Platform engagement and financial relationship

The third family contains log page views, log active duration, and log asset balances. Page views measure the breadth of platform activity; active duration measures time spent using the platform; and asset balances capture the strength of the platform-linked financial relationship. The engagement panel spans eight weeks before and twelve weeks after the survey and contains 165,480 user-week observations.

These outcomes map to the model's broader-benefit side. If usable privacy tools reduce the perceived cost of participating in a data-intensive ecosystem, users may engage more or hold more assets on the platform. The outcomes remain reduced-form measures: more page views need not always be better for users, and asset balances can change through deposits, withdrawals, returns, and portfolio reallocation. Their role is to reveal whether privacy-tool visibility spills over to economically meaningful platform activity.

4.5 Baseline characteristics

Figures 3 and 4 compare the treatment and control groups on survey responses and demographic characteristics measured before Q6. The distributions are similar across arms, as expected under random assignment.

The baseline figures also describe the population to which the experiment most directly speaks. Respondents vary substantially in privacy attitudes and knowledge of platform controls, leaving meaningful room to test the model's friction and privacy-sensitivity predictions. At the same time, all users had engaged enough with the platform to answer the survey and link to administrative records, a feature we return to when discussing external validity.

4.6 Empirical specification and inference

For outcomes observed before and after treatment, we estimate user fixed-effects difference-in-differences models of the form

$$Y_{it} = \alpha_i + \delta_t + \beta \left(\text{Post}_{it} \times Z_i\right) + \varepsilon_{it}, \tag{8}$$

where α_i absorbs time-invariant user characteristics, δ_t denotes period fixed effects, and β measures the treatment-induced change relative to control. Standard errors are clustered by user. For binary direct outcomes, the coefficient is reported in percentage-point units. For logged engagement outcomes, $100[\exp(\widehat{\beta}) - 1]$ gives the corresponding proportional change on the coefficient scale.

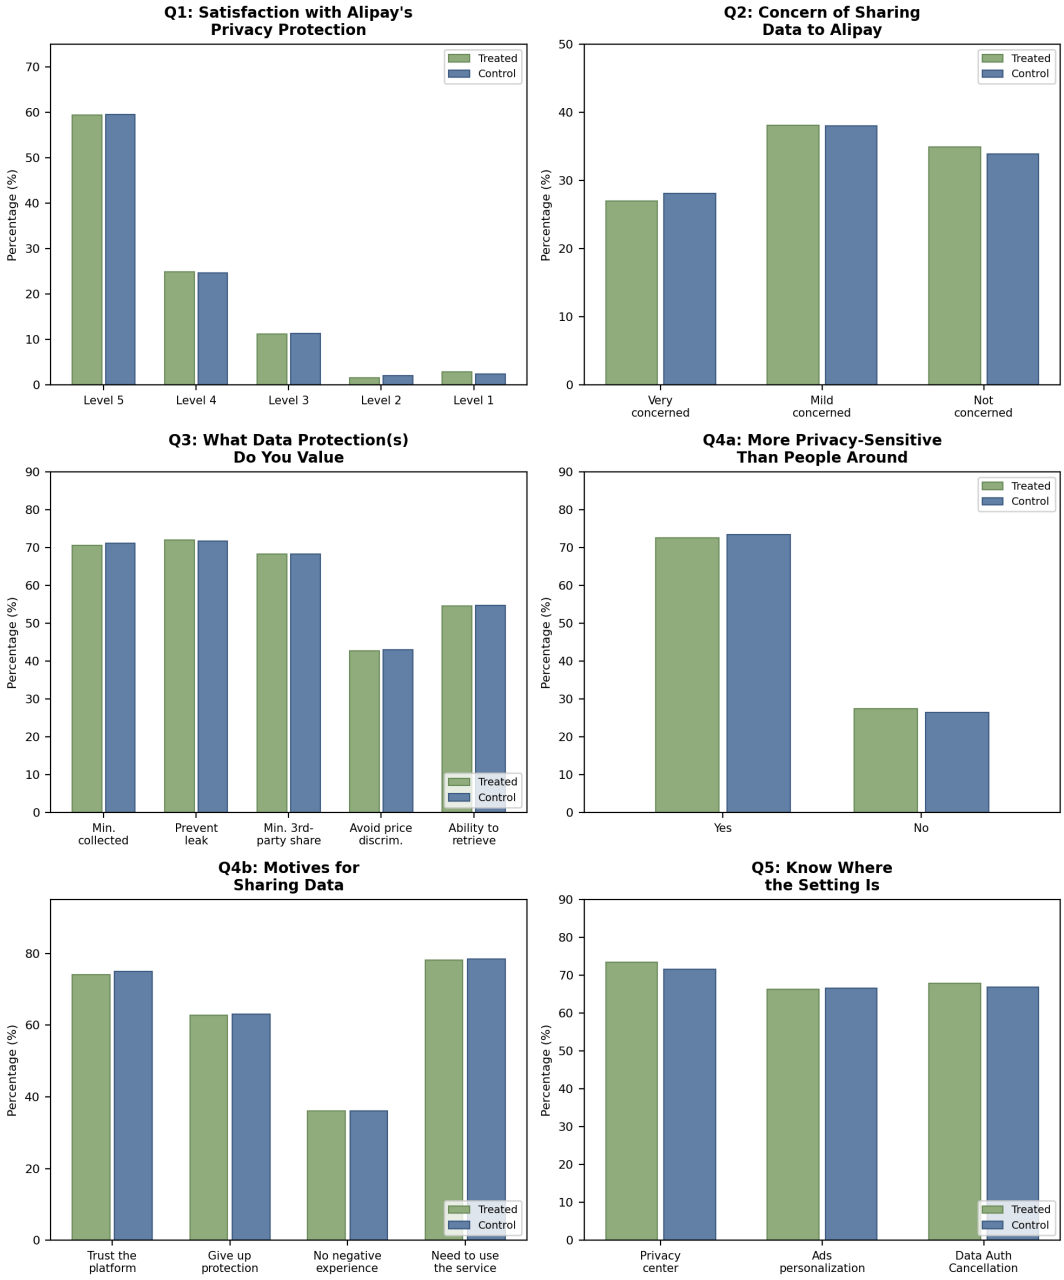


Figure 3: Baseline Survey Responses by Treatment Status

Notes: Bars report response shares for the five privacy modules that both groups completed before treated users saw Q6. The modules cover privacy satisfaction, concern about data sharing, desired protection principles, self-reported privacy sensitivity, motives for sharing, and knowledge of platform privacy settings. Treatment contains 4,136 users and control contains 4,138 users.

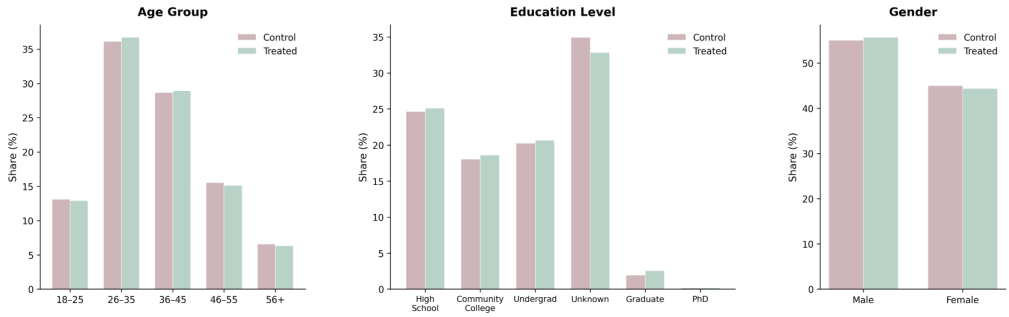


Figure 4: Baseline Demographic Characteristics by Treatment Status

Notes: Bars report the distributions of gender, age, and education in the linked experimental sample, separately by randomized treatment status. Treatment contains 4,136 users and control contains 4,138 users.

To examine timing, we estimate

$$Y_{it} = \alpha_i + \delta_t + \sum_{k \neq -1} \beta_k \left[Z_i \times \mathbf{1}\{t = k\} \right] + \varepsilon_{it}, \tag{9}$$

with event period -1 omitted. The event studies allow us to see whether group differences emerge after the visibility intervention and how they evolve over the observed horizon. The figures report point estimates and 95% confidence intervals based on user-clustered standard errors.

Moderator specifications interact treatment and post-treatment indicators with pre-treatment platform knowledge or relative privacy sensitivity. We report the effect for the lower-knowledge reference group and the incremental difference for higher-knowledge users. These analyses test whether treatment responses are concentrated where the simplified model predicts: among users with greater initial navigation friction or stronger demand for privacy control. We treat the patterns as effect heterogeneity rather than as proof of a unique psychological channel.

5 Results

The results follow the economic sequence in Section 2. The visibility treatment first changes the effective cost of finding and using an existing privacy control. That change can then affect two margins relevant to the platform: a short-run permission margin, on which greater privacy salience may reduce authorization-dependent activity, and a broader relationship margin, on which usable control may encourage engagement. We examine these margins in turn and then use baseline knowledge and privacy sensitivity to study which users respond most.

The empirical mapping is direct. Assignment to Q6 is the visibility change v. A visit to the Data Authorization List is the observed management action M_i, while baseline platform knowledge proxies for initial navigation friction c_i. New requests, confirmations, and cancellations determine

the focal activity-loss term ℓ_H. Page views, active duration, and asset balances represent the broader relationship-gain term g_H. This organization lets the data speak to both sides of the firm's tradeoff without treating any single downstream outcome as a complete measure of profit or welfare.

5.1 Use of the highlighted privacy tools

We begin with the behaviors closest to treatment. Naming the Data Authorization List, describing its function, and showing where to find it should increase use when users otherwise face awareness or navigation friction. The Data Authorization List visit is the most proximal outcome. Visits to the broader Central Privacy Setting reveal whether attention spills over to the parent interface, and recorded cancellation shows whether some users go beyond inspection and change an existing app authorization.

Table 1 reports the estimates. Transparent information increases Data Authorization List visits by 2.2 percentage points (standard error 0.4 percentage points) and recorded app-authorization cancellations by 1.8 percentage points (standard error 0.6 percentage points). The estimate for visits to the broader central privacy setting is 1.2 percentage points and is not statistically distinguishable from zero at conventional levels. The named-list estimate is larger and more precisely estimated than the parent-setting estimate. Because the table does not report their covariance, we do not interpret the difference in individual significance as a formal test that the two effects differ.

Table 1: Direct Effects of the Information Treatment

	Visited Data Authorization List	Visited central privacy setting	Canceled an app authorization
	(1)	(2)	(3)
Post	0.005**	0.047***	0.027***
	(0.002)	(0.006)	(0.004)
Post × Treat	0.022***	0.012	0.018***
	(0.004)	(0.009)	(0.006)
Observations	16,548	16,548	16,548
Individual FE	Yes	Yes	Yes
Clustered error	user	user	user
Pre-treatment avg.	1.18%	10.56%	3.69%

Notes: User-level pre/post difference-in-differences regressions with user fixed effects. Standard errors are clustered by user. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

The magnitudes are economically large relative to low baseline use. The 2.2-point increase in list visits equals about 186 percent of the 1.18 percent pre-treatment average. The 1.8-point increase in cancellation equals about 49 percent of the 3.69 percent pre-treatment average. In frequency terms, the estimates correspond to 22 additional list visitors and 18 additional users with a recorded cancellation per 1,000 treated respondents. A small information change thus shifts an installed

control from rarely used toward economically meaningful use.

The combination of visits and cancellation helps distinguish visibility from a purely attitudinal response. Treated users do not merely report greater concern; they navigate to the named interface and are more likely to take an action recorded in the platform's authorization system. At the same time, the two outcomes need not move one-for-one. Many users may inspect their portfolio and decide that current permissions remain valuable, while others may learn enough from the message to cancel without repeatedly visiting the parent setting.

These findings line up with the model's friction margin. An installed tool has little economic effect when the user does not know it exists or expects the path to be costly. Visibility lowers the effective cost $c_i(v)$, bringing some users across the use threshold in Equation (2). The treatment does not reveal the exact weight placed on search, comprehension, and navigation, but it shows that the implementation layer matters even when the underlying capability is unchanged.

A recorded cancellation should be interpreted at the same level as the data. It shows that a user exercised a platform control over a third-party authorization. It does not reveal the value of the cancelled service, the amount of data previously shared, or the user's welfare. Those quantities are not needed for the first result: transparency increases use of an existing privacy tool.

5.2 New authorization decisions

Greater use of an ex-post control could be accompanied by fewer new authorizations. The message may make privacy risks salient and discourage users from confirming later requests. Conversely, a visible cancellation option can make a new decision feel more reversible and thereby support experimentation with useful services. The average effect on new authorization is consequently ambiguous in the model and is the key observable margin on which transparency could impose an immediate cost on permission-dependent activity.

We detect no corresponding average change in the frequency of third-party app authorization decisions. The estimated effect on total request records is 0.0427 per user per two-week period, with a 95% confidence interval from -0.0127 to 0.0980. Relative to the control pre-treatment mean of 0.8919, the point estimate is 4.8 percent, but the interval includes zero. Counts with a recorded confirmation increase by 0.0128 and counts without a recorded confirmation increase by 0.0299; neither estimate is statistically distinguishable from zero.

The conditional confirmation rate yields the same qualitative conclusion. Among user-biweekly periods with at least one observed request, the estimate is -0.29 percentage points, with a 95% confidence interval from -2.55 to 1.97 percentage points. The data therefore provide no evidence of an average fall in acceptance after users learn about the centralized privacy tool. Because this rate is defined only when a request occurs, it summarizes realized choices among observed opportunities

rather than an unconditional preference for sharing.

Requested-scope composition also remains similar across groups. The estimated change in the number of distinct scopes requested is −0.0162, with a 95% confidence interval from −0.0463 to 0.0139. The estimated change in the maximum requested tier is 0.0107 on the four-level scale, with an interval from −0.0381 to 0.0594. These results do not suggest that the apps reached by treated users systematically ask for broader or more intensive data access.

The absence of a detectable average authorization response is substantively important because it locates the firm's short-run downside. Transparency raises the use of controls over existing authorizations, but it does not produce a corresponding average contraction in the flow of new requests or confirmations over the observed horizon. The result is not a claim that every individual or app is unaffected, nor is it an exact-equivalence statement. It says that the mean permission-flow margin is small relative to the clear proximal increase in privacy-tool use.

In the platform condition, any loss ℓ_H on this margin is weighted by the share s of the firm's revenue tied to the focal third-party authorization feature. When that share is small, even a modest authorization response creates a limited direct cost to the host. When the share is large, the same behavioral response can matter considerably. The experiment measures the behavioral input to Equation (5); the revenue share determines how strongly the firm values it.

5.3 Platform engagement and asset balances

We next examine the benefit side of the platform's tradeoff. If a visible privacy control reduces users' perceived cost of participating in a data-intensive ecosystem, it can increase activity beyond the privacy interface itself. Table 2 reports effects on log page views, log active duration, and log asset balances.

Table 2: Effects on Platform Engagement and Asset Balances

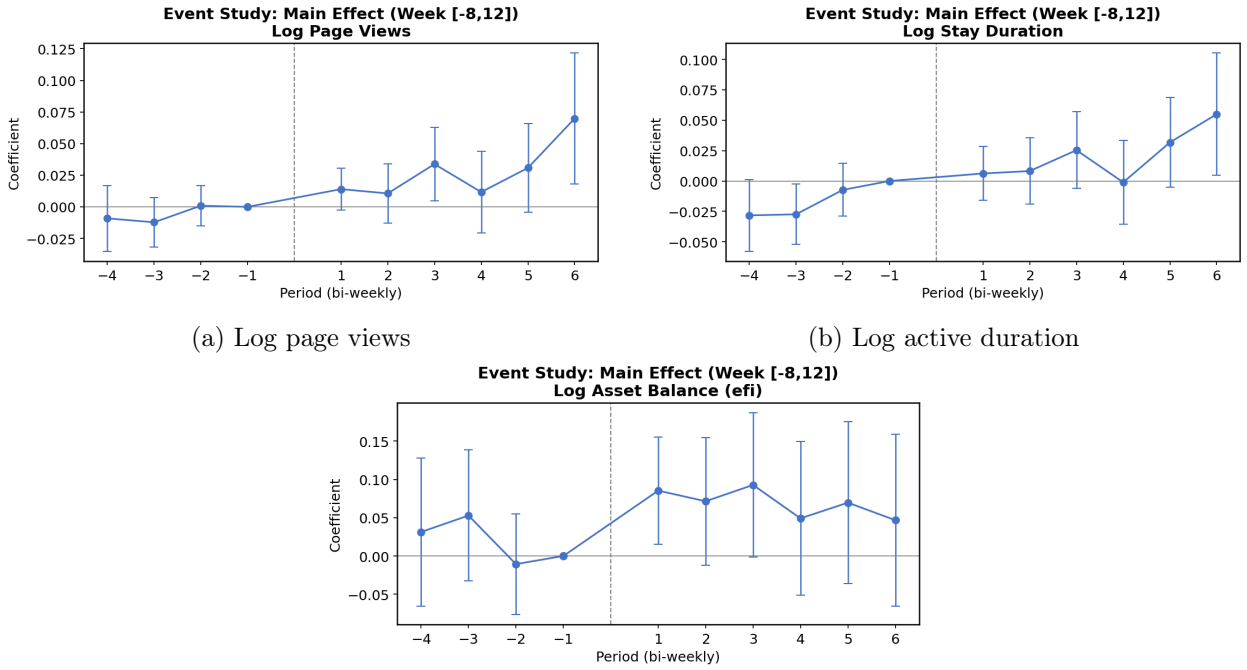
	Log page views	Log active duration	Log asset balance
	(1)	(2)	(3)
Post × Treat	0.024**	0.027**	0.064*
	(0.012)	(0.012)	(0.038)
Observations	165,480	165,480	165,480
Individual FE	Yes	Yes	Yes
Week FE	Yes	Yes	Yes
Clustered error	user	user	user

Notes: User-week difference-in-differences regressions with user and week fixed effects. Standard errors are clustered by user. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Transparent information increases all three outcomes. The coefficient on log page views is 0.024 (standard error 0.012), and the coefficient on log active duration is 0.027 (standard error 0.012).

On the coefficient scale, these values correspond to increases of approximately 2.4 percent and 2.7 percent. The coefficient on log asset balances is 0.064 (standard error 0.038), corresponding to approximately 6.6 percent. The balance estimate is less precise, but its magnitude suggests that the response extends beyond low-cost clicks or time spent reading a survey.

Figure 5 plots the dynamic estimates. Page views and active duration become positive after treatment and tend to grow over the later post-treatment periods. Asset-balance estimates are also positive throughout the displayed post-treatment window, although their confidence intervals are wider. The gradual pattern is consistent with a relationship response that develops as users return to the platform, rather than an effect confined to the moment at which they receive Q6.



(a) Log page views

(b) Log active duration

(c) Log asset balance

Figure 5: Dynamic Effects on Platform Engagement and Asset Balances

Notes: Panels report treatment-by-event-period coefficients from user-level event-study regressions. Period -1 is omitted. Event periods are biweekly, with four pre-treatment and six post-treatment periods. Whiskers are 95% confidence intervals based on standard errors clustered by user. The embedded "Week $[-8, 12]$" headings describe the source week range, which is plotted in biweekly bins on the horizontal axis.

Each outcome captures a distinct dimension of the platform relationship. Page views can rise because users access a wider set of services. Active duration can rise because they spend more time in those services. Asset balances require a financial allocation decision and therefore provide a higher-stakes complement to engagement. The consistency of signs across the three outcomes makes a narrow explanation based solely on visiting the privacy page less compelling.

The estimates map to the model's broad relationship return g_H. A user who knows that

permissions can be reviewed and cancelled may be more willing to remain active in the ecosystem even if she never exercises the tool again. Greater perceived reversibility can make experimentation less costly, and greater control can reduce the privacy friction associated with holding assets or using additional services. The experiment does not separately identify these channels, but it shows that making the control visible has consequences beyond the control itself.

The engagement results do not directly measure revenue. Rather, they identify the observable base over which the platform earns its broader relationship return. Whether the increase is sufficient to make transparency profitable depends on monetization of engagement and balances, the revenue share attached to the focal third-party permission channel, and the cost of deploying the visibility feature. That condition links the field estimates to the firm's technology-deployment decision.

5.4 Who responds? Baseline knowledge and privacy sensitivity

The simplified model predicts that visibility should matter most when initial friction is high or control is especially valuable. We use two pre-treatment survey measures to examine these margins. Knowledge of platform privacy settings proxies for the user's initial ability to navigate the platform. Relative privacy sensitivity identifies users who are likely to place greater value on control over data access.

Table 3 reports direct effects by baseline platform knowledge. The reference group consists of users who did not report knowing the location of the platform's advertising-personalization controls. For this lower-knowledge group, treatment increases Data Authorization List visits by 3.3 percentage points, central-setting visits by 2.3 points, and recorded cancellation by 3.3 points. The treatment-by-knowledge interactions are negative in all three columns, indicating smaller responses among users who already know more about platform settings.

Adding the interaction to the reference-group effect gives implied effects for the higher-knowledge group of 1.7 percentage points for list visits, 0.8 points for parent-setting visits, and 1.0 point for cancellation. The estimates remain positive in magnitude but are much smaller than those for users with low baseline knowledge. Visibility therefore narrows a behavioral knowledge gap: it changes use most for users who initially face the greatest difficulty finding privacy controls.

The knowledge measure is deliberately broader than prior awareness of the Data Authorization List itself. That breadth is useful for the model because it captures familiarity with navigating privacy features across the platform. The result is not simply that users repeat information they were shown. It is that an information intervention produces larger administrative behavior changes among users with lower initial platform knowledge.

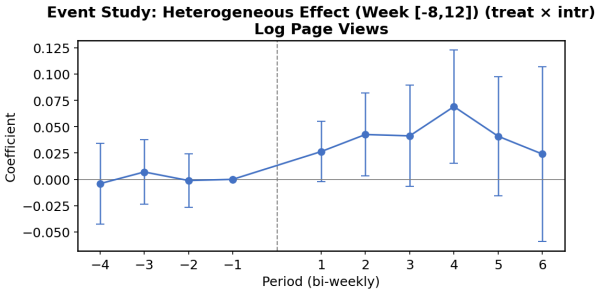
Figure 6 turns to relative privacy sensitivity. The panels plot the incremental treatment response for users who describe themselves as more privacy-sensitive than people around them. The post-

Table 3: Direct Effects by Baseline Platform Knowledge

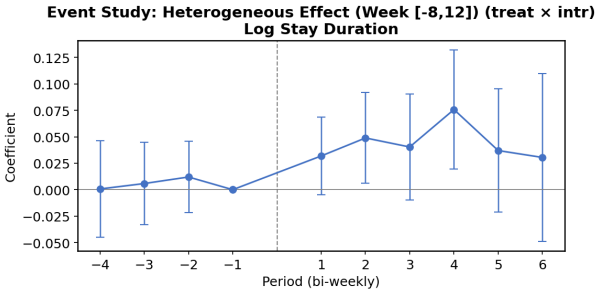
	Visited Data Authorization List	Visited central privacy setting	Canceled an app authorization
	(1)	(2)	(3)
Post	0.005	0.041***	0.024***
	(0.003)	(0.006)	(0.005)
Post × Treat	0.033***	0.023**	0.033***
	(0.008)	(0.011)	(0.010)
Post × Knowledge	−0.004	−0.008	−0.010
	(0.004)	(0.007)	(0.007)
Post × Treat × Knowledge	−0.016*	−0.015	−0.023*
	(0.009)	(0.013)	(0.012)
Observations	16,548	16,548	16,548
Individual FE	Yes	Yes	Yes
Clustered error	user	user	user

Notes: User-level pre/post difference-in-differences regressions. "Knowledge" equals one for users who report knowing the location of advertising-personalization controls. The lower-knowledge group is the omitted category. Standard errors are clustered by user. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

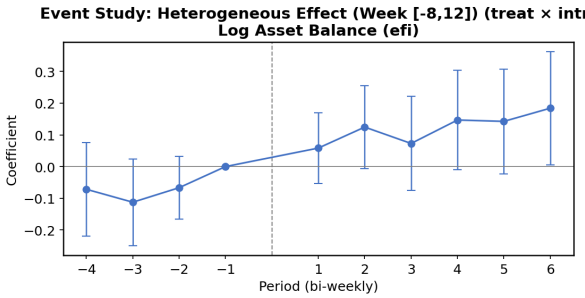
treatment differentials are generally positive for page views, active duration, and asset balances, with the largest and most persistent differences in asset balances. Thus the broader engagement response is stronger among users for whom control over data use is likely to be especially valuable.



(a) Log page views



(b) Log active duration



(c) Log asset balance

Figure 6: Dynamic Effects by Baseline Privacy Sensitivity

Notes: Panels report the incremental treatment-by-event-period effect for users who describe themselves as more privacy-sensitive than people around them. Period −1 is omitted. Event periods are biweekly. Whiskers are 95% confidence intervals based on standard errors clustered by user. Total effects for privacy-sensitive users combine the common treatment effect and the incremental coefficient shown here. The embedded week range is plotted in biweekly bins on the horizontal axis; "intr" denotes the baseline privacy- sensitivity indicator.

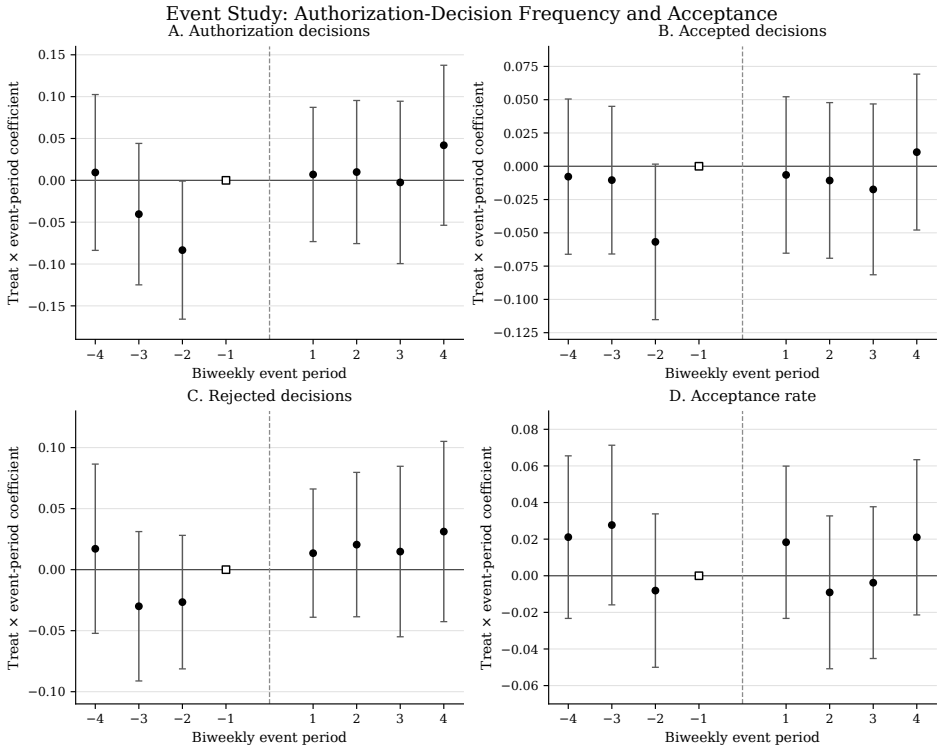
The two moderator patterns illuminate different sides of user friction. Low platform knowledge predicts a stronger response at the privacy tool itself, where finding and navigation costs are most immediate. Privacy sensitivity predicts a stronger response in downstream engagement and balances, where the value of having control can shape the user's willingness to deepen the platform relationship. The patterns are consistent with the conceptual framework even though they do not isolate a single psychological mechanism.

These results also help rule out a simple generic-attention account. If Q6 only added another survey item, one would not naturally expect tool use to be concentrated among low-knowledge users or downstream activity to be concentrated among privacy-sensitive users. The alignment between baseline characteristics and the outcomes they most directly predict supports an interpretation centered on navigation friction and the value of control.

5.5 Dynamics of authorization requests and requested data scope

Figures 7 and 8 show the dynamic authorization results. These plots complement the average estimates by asking whether any short-lived decline in permission-dependent activity is hidden by

aggregation over the post-treatment period.



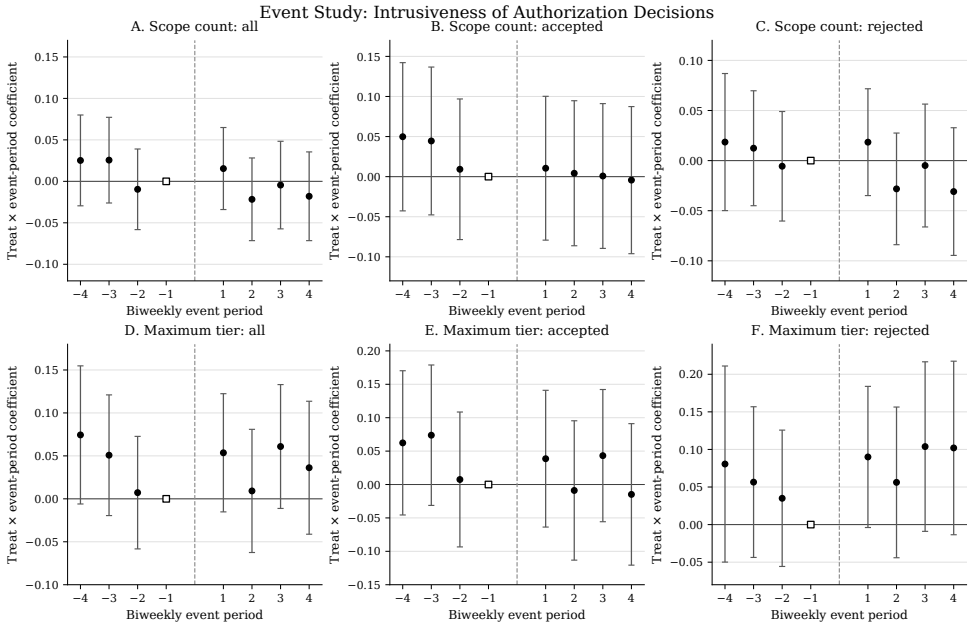
Notes: Bars are 95% confidence intervals from user-clustered standard errors. Period -1 is the omitted reference; the dashed line marks the treatment boundary. There is no period 0 in the code.

Figure 7: Dynamic Effects on Authorization Decisions and Confirmation

Notes: Panels report treatment-by-event-period coefficients from user-level event-study regressions. In the embedded labels, "accepted" denotes a request with recorded confirmation; "rejected" denotes a request without a recorded confirmation. The acceptance rate is defined among user-biweekly periods with at least one request. Period -1 is omitted. Whiskers are 95% confidence intervals based on standard errors clustered by user.

The authorization series show no discrete post-treatment decline. Point estimates for total requests, requests with confirmation, requests without confirmation, and the conditional confirmation rate move around zero, and all displayed post-treatment confidence intervals include zero. The fourth post-treatment period is positive for both total requests and the acceptance rate, but it is imprecisely estimated. The dynamic evidence therefore reinforces the average result: greater use of existing privacy controls is not accompanied by a detectable contraction in new authorization activity.

The figure also makes clear why request frequency and acceptance should be read together. Increased platform participation can create more authorization opportunities, while privacy salience can affect the probability of confirming a given request. A stable confirmation rate alongside stable or weakly higher request counts is consistent with transparency encouraging platform use without materially changing the average realized authorization decision.



Notes: Bars are 95% confidence intervals from user-clustered standard errors. Each observation is a user-app-day authorization decision. Period -1 is the omitted reference; the dashed line marks the treatment boundary.

Figure 8: Dynamic Effects on Requested Data Scope

Notes: Each observation is an observed user–app–day request. Scope count is the number of distinct data scopes requested. Maximum tier is measured on the study's four-level classification. In the embedded labels, "accepted" denotes a request with recorded confirmation and "rejected" denotes a request without recorded confirmation. Period −1 is omitted. Whiskers are 95% confidence intervals based on standard errors clustered by user.

There is likewise no systematic change in the composition of requested data. Across all requests, requests with confirmation, and requests without recorded confirmation, post-treatment coefficients for scope count and maximum tier are small relative to their confidence intervals. Every displayed post-treatment interval includes zero. The visibility intervention changes users' access to an ex-post management tool without inducing a detectable average change in the intensity of new data requests they encounter.

These composition results sharpen the platform interpretation. A message that raises privacy-tool use could in principle prompt users to migrate toward apps that request fewer fields or induce a different mix of authorization outcomes. We do not observe such a reallocation on average. The aggregate request and confirmation estimates are not statistically distinguishable from zero, and a downstream cost is likewise not apparent in requested-scope intensity.

5.6 Control Use and Platform Activity

Four findings form a coherent pattern. First, information about an existing privacy control substantially increases visits to the highlighted tool and recorded cancellation. Second, tool use rises most among users with low baseline knowledge of platform settings, consistent with a reduction in

search and navigation friction. Third, we detect no corresponding average change in new authorization decisions, the confirmation rate, or requested data scope. Fourth, platform engagement and asset balances rise, with stronger downstream responses among users who describe themselves as privacy-sensitive.

This combination is the central empirical result. Greater use of a privacy control can coexist with stronger activity on the same platform. Transparency does not mechanically push users away from a data-intensive technology. In this setting, making the control easier to find expands effective user choice while also strengthening the broader platform relationship.

The findings map directly to the firm's short-run and long-run incentives. The immediate cost of transparency comes from the possibility that salience and easier cancellation reduce the permission-dependent value of the focal feature. The benefit comes from greater engagement, deeper financial relationships, and the durability of the user–platform relationship. Our estimates show the behavioral margins on both sides: more control use, no detectable average contraction in new authorization, and more downstream activity.

Whether this combination pays for the firm depends on the focal feature's share s of platform revenue. When little revenue is tied to third-party app authorizations, the cost created by additional control use is limited and the broader engagement return can dominate. When a large share of revenue depends on retaining those authorizations, the same visibility policy can be privately costly even if users become more active elsewhere. The break-even condition in Equation (6), rather than a universal preference for transparency or opacity, determines the platform's incentive.

The threshold gives a transparent percentage test. If the monetized proportional focal loss equals the monetized broad gain, transparency pays when the focal feature generates less than 50 percent of platform revenue. If the focal loss is twice the broad gain, the threshold is one-third; if it is half the broad gain, the threshold is two-thirds. The authorization-flow point estimates imply $\ell_H = 0$ for confirmations and therefore a 100 percent break-even share on that margin. The positive cancellation response implies a strictly lower threshold for the existing-permission margin, with its exact level determined by the revenue attached to a cancelled authorization. The experimental estimates identify the behaviors to monetize; applying one common horizon H to focal and nonfocal accounts supplies g_H, ℓ_H, and the firm-specific percentage s^*_H.

The experiment does not itself convert engagement or balances into dollars, so it does not estimate a realized profit effect. Its contribution is to identify the user-side responses that enter the platform decision and to show that the two sides need not move in opposition. This is precisely why privacy control is part of technology strategy: data consent affects the inputs available to firms, while the way firms deploy consent-management technology affects adoption, participation, and the value of the wider ecosystem.

The result also explains the title's tension. Transparency is costly because it draws attention

to privacy and makes revocation easier. It can nevertheless pay because usable control reduces an important friction in the adoption and use of digital services. The balance depends on what the focal permission contributes to current revenue relative to what transparency contributes to the continuing relationship. The field experiment shows that both forces are empirically relevant and can be studied within the same platform.

6 Discussion

The experiment reveals a firm decision that is easy to miss when privacy is treated only as a legal constraint. A platform can install a control and still choose how costly it is for users to locate and exercise. That implementation decision affects the effective terms on which users adopt data-intensive services. It can expose a focal permission-dependent activity to cancellation while strengthening participation elsewhere in the firm. The empirical pattern—more use of the highlighted control, higher platform engagement and asset balances, and no statistically detectable average change in subsequent authorization or acceptance—shows why the two sides of this decision must be measured together.

6.1 Privacy controls as technology infrastructure

Data consent is an input into technological production. Digital services use data for matching, prediction, fraud detection, personalization, and product development; artificial intelligence expands both the range of productive uses and the difficulty of anticipating those uses at the moment of initial consent (Jones and Tonetti, 2020; Bergemann et al., 2022; Goldfarb and Que, 2023). Firms therefore need more than the technical capacity to collect and process information. They also need users who are willing to begin, continue, and revise data relationships as services evolve.

Usable ex-post control is part of that infrastructure. A user may accept an authorization today while remaining uncertain about the app's future value or the consequences of continued access. A centralized list lowers the cost of reviewing the accumulated stock of authorizations. Making that list visible can reduce the practical irreversibility of an earlier choice, even though the experiment does not change the authorization prompt or the underlying control. In this sense, the information treatment changes the implementation of an installed technology, not the technology's formal menu.

The distinction between the stock and flow of permissions is important for firms. Product teams often focus on the flow because initial authorization is close to conversion. Users, however, accumulate a stock of permissions across applications and over time. Some become stale, some outlive the service that motivated them, and some are reassessed as users learn. Centralized management can prune that stock while leaving the later flow of useful relationships largely intact. The experiment's outcome pattern illustrates this separation: recorded cancellation rises, while the estimated changes

in subsequent request frequency, conditional confirmation, and requested-scope composition are not statistically distinguishable from zero on average.

Privacy governance is therefore neither simply a brake on innovation nor a costless assurance device. It is an organizational input that determines how reversible data-intensive adoption is in practice. Product naming, placement, path length, explanation, and feedback are forms of implementation capital. They can change the effective user cost of exercising a control without altering formal access, and their return depends on how the firm values the focal feature and the wider relationship.

6.2 When transparency pays

The platform's problem is not "privacy versus growth" in the abstract. It is a revenue-incidence problem. Greater visibility can increase cancellation or otherwise reduce activity tied to the focal authorization-dependent feature. At the same time, it can increase activity in the rest of the platform. The relevant business-model statistic is therefore the percentage of baseline revenue tied to the focal feature.

Equation (5) makes this comparison explicit. The focal revenue share s weights the proportional focal loss ℓ_H, and the remaining share $1 - s$ weights the proportional broad gain g_H over a common horizon H. When visibility reduces focal revenue but raises revenue elsewhere, Equation (6) gives the break-even share: transparency pays when the focal feature's revenue share lies below $g_H/(g_H + \ell_H)$. The condition is deliberately expressed in percentages so that every term can be constructed from firm data.

The experiment measures the behavioral inputs to this calculation. Visits and recorded cancellations show that the focal control becomes economically active. Authorization requests, confirmations, and requested scopes show whether the response spreads into the subsequent flow of permission-dependent activity. Page views, active duration, and asset balances capture broader relationship margins. Platform accounts can then map those behaviors into focal and nonfocal revenue and supply s. This separation is preferable to treating a page view, a cancellation, and a dollar of balance as if they shared a unit.

The estimated pattern is favorable to transparency for a broad set of revenue shares, but it is not itself a profit estimate. The treatment increases list visits by 2.2 percentage points and recorded cancellation by 1.8 points, which reveals potential exposure on the focal margin. It also increases log page views by 0.024, log active duration by 0.027, and log asset balances by 0.064, although the balance estimate is less precise. Because the later-authorization estimates show no statistically detectable average change, the most visible focal response concerns management of existing permissions rather than a broad contraction in new app relationships. Whether the measured

broad response is large enough in revenue units depends on the platform's monetization and focal share.

This accounting perspective sharpens the paper's title. "Costly" describes the exposure created when visibility prompts more users to exercise the focal control. "Pays" describes the positive response in wider platform activity and the revenue-share condition under which that response exceeds the focal loss. Transparency is not free, and its return is not universal. Its value is an empirical property of the firm's business model.

6.3 Short- and long-run platform incentives

The two sides of the tradeoff need not arrive at the same time. A cancellation is immediate, salient in an operating dashboard, and readily assigned to the feature or third-party app whose authorization was removed. The broader return can accrue through repeated use, deeper financial participation, cross-service adoption, and retention. A short evaluation window can therefore make transparency appear less attractive than a longer one, even with the same users and technology.

This timing creates an organizational wedge. A mini-app team evaluated on retained permissions or current conversion may emphasize the focal loss. A platform-wide product, privacy, or governance team may internalize activity across payments, finance, and other services. The teams can rationally prefer different interface designs because they face different accounting boundaries and horizons. Making the control visible is then not merely a communications choice; it is a decision about which unit owns the user relationship and which outcomes enter its objective.

The same logic helps interpret the field setting. The highlighted list governs third-party mini-app authorizations, while the host operates a much broader financial super-app. The host can capture relationship value outside the permission exposed by the control, whereas an affected complementor may bear a more concentrated cost. This architecture makes the platform a natural place to observe a local control response alongside broader activity. It also illustrates why host and complementor incentives can diverge even when both are firms in the same ecosystem.

Evaluation horizon should consequently be treated as part of the investment decision. Firms can calculate s, g_H, and ℓ_H at multiple horizons and ask whether the break-even condition changes. If broader revenue accumulates while the focal loss stabilizes, transparency becomes more attractive over time. If cancellation produces continuing losses or the engagement response is transitory, the opposite can occur. A common horizon for behavior and revenue keeps the decision internally consistent.

6.4 Firm boundaries, market structure, and incidence

The revenue-share condition gives a sharper comparative prediction than firm size alone. A diversified platform can have a small focal share and many channels through which a durable relationship creates value. It can also spread the fixed costs of control design, interface engineering, support, and monitoring across a large user base. These features can make usable controls a productive investment rather than a pure compliance expense.

Specialized firms face a different calculation. If most revenue depends on the permission placed at risk, s is high and the immediate cost receives a large weight. The firm also has fewer nonfocal activities on which a stronger relationship can pay. A control with the same behavioral effect can therefore be voluntarily surfaced by a diversified platform and obscured by a firm whose production is concentrated in the exposed feature.

Scale by itself does not determine the sign. Large platforms may derive substantial rents from data, operate more complex permission graphs, or face weak competitive discipline. A small firm may differentiate on credible privacy governance. The model instead points to observable dimensions: revenue concentration, contribution margins, implementation cost, evaluation horizon, and the allocation of interface authority. Cross-firm comparisons that measure these objects can explain variation in transparency more directly than a simple large-firm indicator.

Interface authority adds a further incidence issue. The host can decide which permissions are easy to review even when a complementor owns the revenue stream. It may surface third-party authorizations while leaving controls over its own core data less prominent, or it may centralize controls to improve the quality of the whole ecosystem. Comparing visibility across permission categories within the same platform can reveal how governance and business model interact. In this way, privacy implementation can influence production concentration through a channel beyond data scale: diversified firms may be better able to internalize broad relationship benefits and finance the fixed cost of usable control.

6.5 Policy and external validity

Policy often asks whether a control exists. The experiment shows why the cost of reaching it is a separate economic margin. Both treatment and control users had the same centralized authorization list, yet a short message naming and locating it changed navigation and recorded exercise. Formal availability did not make the interface behaviorally irrelevant. Product audits can therefore measure discovery, path length, successful navigation, task completion, and exercise in addition to checking whether the capability is installed.

The results also caution against assuming that stronger control necessarily reduces adoption of data-intensive services. In this setting, more exercise of an ex-post privacy control coexists with

greater average platform activity and asset balances, while subsequent authorization and acceptance show no statistically detectable average change. This pattern does not imply that every transparency intervention expands adoption. It establishes that user control and firm activity need not be opposites, particularly when the control affects a limited share of a multi-product platform's revenue.

The estimates pertain to respondents on one finance-and-mini-app platform and to a survey-delivered information treatment. Persistent dashboards, search redesigns, push notifications, or regulatory notices can produce different attention and recipient composition. A control tied directly to core advertising, payment, or model-training data can also expose a larger revenue share and produce a different firm choice. These differences are naturally organized by the same model: the intervention changes friction, while the business model determines how the behavioral response enters revenue.

The conceptual structure extends beyond privacy. Subscription cancellation, account deletion, data portability, appeals, and permissions delegated to AI agents all separate the formal presence of an option from the user's cost of exercising it. In each case, the firm trades immediate feature retention against the credibility and adoption of a broader technology. The relevant friction, focal share, and evaluation horizon differ, but the firm problem is the same.

6.6 Applying the revenue test

The revenue condition translates the field experiment into an accounting protocol. For the focal side, the platform combines permission-level cancellation and authorization histories with mini-app transactions, commissions, advertising, or other authorization-dependent revenue. This measures focal revenue under each visibility policy and therefore ℓ_H. For the broad side, the platform connects page views, time, product use, and asset balances to revenue outside the affected mini-app channel, which supplies g_H. Baseline accounts then supply the focal share s.

This measurement is informative at several horizons and organizational levels. A feature-level calculation reveals the local incentive; a platform-wide calculation incorporates cross-service gains; and a combined host-complementor calculation shows incidence across firms. Repeating the exercise for permission categories with different revenue exposure shows which controls have the strongest private case for visibility and which expose enough revenue to create internal resistance.

The same accounting logic disciplines the choice among visibility components. Naming a control, explaining its function, showing its location, shortening the path, and making the destination persistent lower different portions of navigation friction and carry different implementation costs. Product teams can compare each component's broad return per unit of focal exposure over short and long horizons. This makes visibility comparable to other investments in product design: the firm measures its cost, its revenue incidence, and the durability of the response.

The central lesson is already visible in the present experiment. A firm can make a consequential privacy control easier to use without inducing a broad average retreat from the measured flow of new authorizations, and the same intervention can strengthen activity elsewhere on the platform. The firm-side decision is then governed by revenue exposure and horizon, not by a generic presumption that transparency and technological adoption conflict.

7 Conclusion

Data consent is a central input into technological progress because firms need both information and a durable relationship with the people who generate it. Privacy controls shape that relationship after an initial authorization has been granted. Yet installing a control is not the end of the firm's decision. The firm also chooses whether users can find and exercise it at reasonable cost. This paper studies that implementation margin.

In a randomized field experiment with 8,274 users of a large financial super-app, all users had access to a centralized list for reviewing and canceling third-party app authorizations. The treatment supplied information that named the list, explained its centralized-management function, and made its location salient. The intervention therefore held the formal capability fixed while reducing the friction associated with finding it.

The treatment increased visits to the Data Authorization List by 2.2 percentage points and raised recorded app-authorization cancellation by 1.8 points. These effects are economically large relative to pre-treatment rates of 1.18 percent and 3.69 percent. Visibility changed both navigation to the control and exercise of a consequential privacy action. The heterogeneity estimates reinforce the navigation-friction interpretation: tool use rose most among users with lower baseline knowledge of platform settings.

Greater use of the control coexisted with stronger activity elsewhere on the platform. The treatment increased log page views by 0.024 and log active duration by 0.027. The coefficient on log asset balances was 0.064 and positive, although estimated less precisely. At the same time, the experiment detected no statistically significant average change in subsequent authorization requests, confirmed requests, the acceptance rate conditional on a request, or the composition of requested data scopes. The engagement response was stronger among users who described themselves as privacy-sensitive. Together, the results show that making an existing privacy control visible can activate management of past permissions without producing a broad average retreat from observed new authorization or from the host platform.

The model translates this coexistence into a firm decision. Users act when the value of reaching the control exceeds navigation friction; visibility lowers that friction. For the platform, the behavioral response has two sides. More exercise can reduce revenue tied to the focal authorization-dependent

feature, while broader participation can raise revenue elsewhere. If s is the percentage of baseline revenue tied to the focal feature, ℓ_H is the proportional focal loss, and g_H is the proportional gain in other revenue, transparency raises revenue when

$$s \leq \frac{g_H}{g_H + \ell_H}.$$

The condition makes the title precise. Transparency is costly because it exposes a focal permission-dependent activity to greater exercise of control; it pays when the wider relationship response outweighs that exposure. The experiment measures the behavioral margins, while firm accounts provide the revenue share and monetization needed to apply the condition.

The revenue-share condition also explains why platform incentives vary across firms and over time. A diversified platform can have a small exposed share, spread implementation costs across many users, and recover a local loss through payments, finance, or other services. A specialized firm may depend heavily on the permission placed at risk. Likewise, a feature team evaluated on immediate cancellation or conversion can prefer less visibility even when the platform as a whole benefits through a longer user relationship. Transparency is thus a product-design investment, an internal-accounting problem, and a question of firm boundaries.

The broader implication is that privacy governance belongs inside the economics of technology adoption. Data-intensive services, including artificial intelligence, create value from information whose future uses can be difficult to anticipate. A credible ability to revisit earlier choices can make that relationship more sustainable. Firms that measure only retained permissions will miss the wider return to usable control; policymakers that measure only formal availability will miss the friction that prevents exercise.

The distinction extends beyond data authorization. Account deletion, subscription cancellation, portability, appeals, and AI-agent permissions can all exist formally while remaining behaviorally remote. In each setting, a firm chooses an implementation layer between an installed option and effective use. The evidence here shows why that layer matters: transparency can lead users to exercise control and, at the same time, deepen their activity on the technology that provides it. Whether that combination benefits the firm is not a matter of intuition alone. It is a measurable consequence of the focal revenue share, the size of the broader response, and the horizon over which the firm values the relationship.

References

Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar. Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4):218–256, 2022. doi: 10.1257/mic.20200200.

Alessandro Acquisti, Leslie K. John, and George Loewenstein. The impact of relative standards on the propensity to disclose. *Journal of Marketing Research*, 49(2):160–174, 2012. doi: 10.1509/jmr.09.0215.

Alessandro Acquisti, Curtis Taylor, and Liad Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–492, 2016. doi: 10.1257/jel.54.2.442.

Aaron Adalja, Jura Liaukonyté, Emily Wang, and Xinrong Zhu. GMO and non-GMO labeling effects: Evidence from a quasi-natural experiment. *Marketing Science*, 42(2):233–250, 2023.

Idris Adjerid, Alessandro Acquisti, and George Loewenstein. Choice architecture, framing, and cascaded privacy choices. *Management Science*, 65(5):2267–2290, 2019. doi: 10.1287/mnsc.2018.3028.

Sebastián Araya, Andrés Elberg, Carlos Noton, and Daniel Schwartz. Identifying food labeling effects on consumer behavior. *Marketing Science*, 41(5):982–1003, 2022.

Susan Athey, Christian Catalini, and Catherine E. Tucker. The digital privacy paradox: Small money, small costs, small talk. NBER Working Paper 23488, National Bureau of Economic Research, 2017.

Björn Bartling, Ernst Fehr, and Holger Herz. The intrinsic value of decision rights. *Econometrica*, 82(6):2005–2039, 2014. doi: 10.3982/ECTA11573.

Dirk Bergemann, Alessandro Bonatti, and Tan Gan. The economics of social data. *RAND Journal of Economics*, 53(2):263–296, 2022. doi: 10.1111/1756-2171.12407.

Hemant K. Bhargava, Olivier Rubel, Elizabeth J. Altman, Ramnik Arora, Jörn Boehnke, Kaitlin Daniels, Timothy Derdenger, Bryan Kirschner, Darin LaFramboise, Pantelis Loupos, Geoffrey Parker, and Adithya Pattabhiramaiah. Platform data strategy. *Marketing Letters*, 31(4):323–334, 2020. doi: 10.1007/s11002-020-09539-3.

Bryan Bollinger, Phillip Leslie, and Alan Sorensen. Calorie posting in chain restaurants. *American Economic Journal: Economic Policy*, 3(1):91–128, 2011.

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer. Salience and consumer choice. *Journal of Political Economy*, 121(5):803–843, 2013.

Laura Brandimarte, Alessandro Acquisti, and George Loewenstein. Misplaced confidences: Privacy and the control paradox. *Social Psychological and Personality Science*, 4(3):340–347, 2013. doi: 10.1177/1948550612455931.

Gordon Burtch, Anindya Ghose, and Sunil Wattal. The hidden cost of accommodating crowdfunder privacy preferences: A randomized field experiment. *Management Science*, 61(5):949–962, 2015. doi: 10.1287/mnsc.2014.2069.

James Campbell, Avi Goldfarb, and Catherine Tucker. Privacy regulation and market structure. *Journal of Economics & Management Strategy*, 24(1):47–73, 2015. doi: 10.1111/jems.12079.

Raj Chetty, Adam Looney, and Emmanuel Saez. Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177, 2009.

Thomas Eisenmann, Geoffrey Parker, and Marshall Van Alstyne. Platform envelopment. *Strategic Management Journal*, 32(12):1270–1285, 2011.

European Union. Regulation (EU) 2016/679: General data protection regulation. Official Journal of the European Union, 2016. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679.

Chiara Farronato, Andrey Fradkin, and Tesary Lin. Designing consent: Choice architecture and consumer welfare in data sharing. Working Paper 34025, National Bureau of Economic Research, 2025.

Amy Finkelstein and Matthew J. Notowidigdo. Take-up and targeting: Experimental evidence from SNAP. *Quarterly Journal of Economics*, 134(3):1505–1556, 2019. doi: 10.1093/qje/qjz013.

Xavier Gabaix and David Laibson. Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics*, 121(2):505–540, 2006. doi: 10.1162/qjec.2006.121.2.505.

Avi Goldfarb and Verina F. Que. The economics of digital privacy. *Annual Review of Economics*, 15: 267–286, 2023. doi: 10.1146/annurev-economics-082322-014346.

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

Andrei Hagiu and Julian Wright. Multi-sided platforms. *International Journal of Industrial Organization*, 43:162–174, 2015. doi: 10.1016/j.ijindorg.2015.03.003.

Shota Ichihashi. Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–595, 2020. doi: 10.1257/aer.20181052.

Charles I. Jones and Christopher Tonetti. Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–2858, 2020. doi: 10.1257/aer.20191330.

Tesary Lin. Valuing intrinsic and instrumental preferences for privacy. *Marketing Science*, 41(4): 663–681, 2022. doi: 10.1287/mksc.2022.1368.

Amalia R. Miller and Catherine E. Tucker. Privacy protection and technology diffusion: The case of electronic medical records. *Management Science*, 55(7):1077–1093, 2009.

Helen Nissenbaum. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press, 2010.

Shumiao Ouyang. Cashless payment and financial inclusion. *Journal of Financial Economics*, 180: 104277, 2026. doi: 10.1016/j.jfineco.2026.104277.

Geoffrey G. Parker and Marshall W. Van Alstyne. Innovation, openness, and platform control. *Management Science*, 64(7):3015–3032, 2018. doi: 10.1287/mnsc.2017.2757.

Jean-Charles Rochet and Jean Tirole. Platform competition in two-sided markets. *Journal of the European Economic Association*, 1(4):990–1029, 2003. doi: 10.1162/154247603322493212.

Standing Committee of the National People's Congress of the People's Republic of China. Personal information protection law of the people's republic of china. Official Chinese text, National People's Congress, 2021. URL https://www.npc.gov.cn/npc/c2/c30834/202108/t20210820_313088.html.

Janice Y. Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. The effect of online privacy information on purchasing behavior: An experimental study. *Information Systems Research*, 22 (2):254–268, 2011. doi: 10.1287/isre.1090.0260.

Catherine E. Tucker. Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research*, 51(5):546–562, 2014. doi: 10.1509/jmr.10.0355.