How Do Past Privacy Choices Shape the Future?*

Verina F. Que†

This draft: Dec 1, 2024

Abstract

Consumers today often face frequent and consecutive digital privacy choices, but they do not necessarily make independent decisions each time. This paper demonstrates that past privacy choices affect consumers' current privacy choices. Such state-dependent choices suggest that privacy choices can have externalities within a platform in which one app's data requests can affect the ability of other apps to collect data. Specifically, I use an individual-level consumer panel to investigate data consent decisions by consumers on Alipay, a major digital platform that connects users and third-party apps. Leveraging a natural experiment that encourages users to accept data requests, I find that the probability of rejecting the next request declines 15%. This effect decays over time, and is larger when the next data consent decision is relatively harder on deliberating value-cost tradeoff (i.e. weak preference context), categorized by large language models (LLMs). The effect does not differ by whether the specific data requested in consecutive data consent decisions is the same nor by the underlying privacy preference of a consumer. Overall, I interpret these results to suggest that the externalities arising from state-dependent data consent choices are temporary. Nevertheless, they can positively or negatively impact the ability of apps to collect data, and therefore platforms have incentives to encourage apps to provide consumer-friendly data request designs.

Keywords: digital privacy, data sharing, state-dependence, platforms

*Acknowledgements: I am grateful to Avi Goldfarb for his guidance and support throughout this project. My appreciation extends to my dissertation committee members: Heski Bar-Isaac, Matthew Osborne, and Dilip Soman, for their insightful feedback. Additionally, I am thankful for helpful comments from Nitin Mehta and Liying Yang; and participants at the 9th Canadian E&T Symposium, the Mapping and Governing Online World conference, the Sheth AMA Doctoral Consortium, the Rotman brown-bag seminars, and the econ study-group meetings at Rotman. I extend my gratitude to Luohan Academy and AntGroup for the dissertation fellowship and access to the data.

†Rotman School of Management, University of Toronto, Toronto, ON M5S 3E6. Email: f.que@rotman.utoronto.ca.

1 Introduction

There is often a tradeoff between privacy protection and the sharing of personal data (Acquisti et al. (2016)). Firms trade off the competing goals of data acquisition, which generates value for both firms and consumers (Farboodi and Veldkamp (2022); Jones and Tonetti (2020)), and privacy protection. There is increasing attention to consumer data consent decisions in the literature (Tucker (2023); Lin (2022); Kummer and Schulte (2019)). There is also increasing attention to the role of the externalities that these consumer decisions create (Bonatti (2023)). Within a platform environment, consumer privacy decisions with one service provider might affect their decisions with other service providers if privacy decisions are state dependent. In other words, there is a connection between data consent decisions, state dependence, and externalities in a platform environment when consumers make repeated privacy decisions.

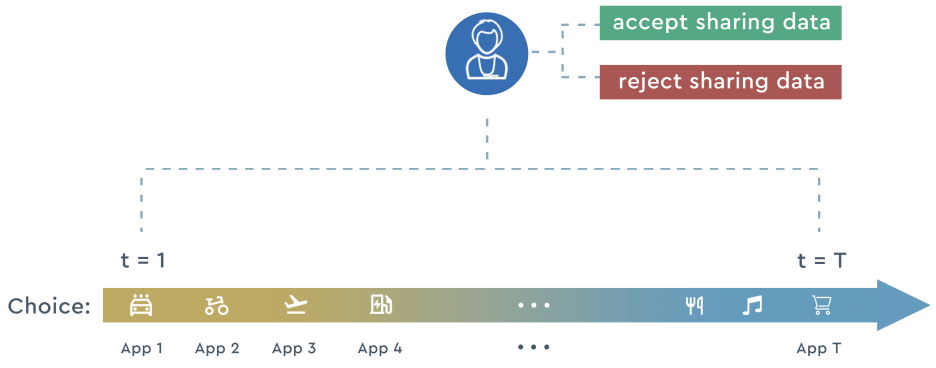
My research examines whether consumer privacy choices affect their subsequent privacy choices by using data from Alipay, a major digital platform that connects users and third-party apps. That is, it examines whether structural state dependence exists in consumer data-consent decisions in this online sequential choice setting. In the marketing literature, structural state dependence is the phenomenon where there is a causal effect of past behaviors on current behaviors (Heckman (1981); Dubé et al. (2010)). Some prior work uses stockout-style variation to identify structural state dependence (Goldfarb (2006a); Sudhir and Yang (2014); Levine and Seiler (2022)). I identify structural state dependence by leveraging a quasi-experiment in which consumers are highly likely to accept data requests. Specifically, on the Alipay platform, there exists a particular category of apps designed for consumers who need to charge their devices away from home when they run out of battery. These charging app brands deploy their charging stations at shopping malls, bars, restaurants, hotels, tourist attractions, and other public locations. Notably, the average rejection rate for data requests from these portable power bank apps is low, below 1 percent. In contrast, the rejection rates for data requests from other types of applications on the platform are much higher, ranging from 5 percent to 50 percent.

Leveraging this quasi-experiment, I document that structural state dependence exists. The state dependence between choices on the platform indicates an interaction between apps. Specifi-

cally, if consumers are more willing to accept data requests from one app, this would make them more likely to allow other apps to collect data as well. This state dependence suggests that a consumer's data consent decision in one context can have a positive or negative spillover effect on the broader ecosystem.

I examine privacy choices using proprietary data from Alipay, a major digital platform that serves over 1.3 billion users[1], that covers approximately 358,000 randomly sampled users from March 2020 to September 2022. On this platform, consumers repeatedly make decisions regarding data consent for functionality and advertising purposes. Figure 1 illustrates this sequence of choices where consumers make data-consent decisions on different apps with different use cases in different contexts over time. The unit of observation is the user-app-decision. Only by accepting the whole bundle of data types can the consumers use the services. Consumers' data consent decisions are binary: "accept" or "reject"; that is, in this setting consumers do not have customization options.

Figure 1: Consumers make sequential data-consent decisions



Outside of the particular Alipay context, such choices are not unusual. For example, a similar binary accept-or-reject decision of both data terms and use of a service is a "cookie wall". A cookie wall is a cookie banner that requires users to 'agree' or 'accept' all cookies and trackers in order

[1]From 2020 statistics, which is the starting year of my panel. Source: https://en.wikipedia.org/wiki/Alipay.

to access a website's content[2]. Such cookie walls are not uncommon, for example, in Europe, the guidance from authorities differs regarding the use of cookie walls[3]. Cookie walls are used in the Netherlands, France, UK, and are allowed under CCPA[4]. In addition, the broader choice environment is similar to the data-consent choices on Facebook when using its third-party plug-ins, or cookie banners that appeared on our first visits to different websites (Figure B1).

I measure the state dependence effect as a 15% decline in the probability of rejecting the next request, and this decline in the probability of rejection lasts for at least 3 choices after the use of the battery app. My finding of significant, but temporary, structural state dependence suggests an externality across app providers within the platform. This is a similar externality to the one identified in Nosko and Tadelis (2015), which documented that bad eBay buyer experiences with one seller had a negative impact on the platform overall.

I then explore potential explanations of the observed state dependence. I investigate two key factors influencing state dependence. By altering the magnitude of these two factors, I demonstrate how measured state dependence changes. First, I examine the role of consumer memory, drawing on the accessibility-diagnosticity framework from Feldman and Lynch (1988). Second, I consider the role of the strength of preference for one choice or another (Webb et al. (2024); Roy et al. (1996)).

I find evidence consistent with memory and preference strength playing important roles, but not the alternative explanations that would be elaborated in the later paragraphs. Specifically, in terms of memory, I demonstrate that the estimated state dependence decreases with each choice after the initial shock and that estimated state dependence decreases when the temporal gap between a past and current decision is larger.

I also find suggestive evidence that state dependence primarily occurs when preferences for accepting or rejecting the terms are relatively weak. In other words, the data consent decision is relatively harder for consumers as the value-cost tradeoff is ambiguous in the weak preference context, so they rely more on the last decision. To explore the strength of preference, I first use

[2]definition and overview of cookie wall: https://cookie-script.com/blog/cookie-wall

[3]from Deloitte Cookie Benchmark study: https://www2.deloitte.com/content/dam/Deloitte/nl/Documents/risk/deloitte-nl-risk-cookie-benchmark-study.pdf

[4]examples of the interface could be found here: https://www.cookieyes.com/blog/cookie-wall/.

the application program interface (API) in a large language model (LLM) to categorize each app in terms of whether the app is only useful if the consumer is physically present at the location where the services are provided and whether there is not another easy option for consuming the services besides using the app. In those cases, I categorize this as a context of stronger preference which suggests a high urgency for accepting the data request. Otherwise, I categorize the app as weaker preference context and a low urgency for accepting the request. Measured structured state dependence is only statistically and economically significant for low-urgency (i.e. weak preference context) apps, suggesting that past choices are more impactful on the more deliberative current choices. I performed a robustness check by categorizing preference strength based on app acceptance rate instead of use case. The results did not reveal significant differences between the two contexts, primarily because the acceptance rate is influenced not only by an app's use case but also by its reputation and quality, which involve endogenous consumer and firm choices. Consequently, the contexts annotated by the LLM provide a more suitable approach to measuring the strength of preference for approving data requests, conditional on downloading an app.

To further explore the role of the strength of preference, I distinguish between consumers who have read a privacy agreement before data consenting and consumers who do not. I find that consumers who have read privacy agreements do not exhibit structural state dependence in their choices. I interpret this to suggest that consumers with stronger privacy preferences exhibit less state dependence.

There are alternative explanations that serve as potential alternative reasons for the measured state dependence. For alternative explanations, I analyze the role of data permanence. Data could be measured as property rights, which can affect their use in the economy (Bonatti (2023); Jones and Tonetti (2020)). Once revealed, data is subjected to repeated, potentially unending secondary use (Acquisti (2023)). Furthermore, I explore the role of risk perception, where consumers are heterogeneous in the perception of privacy invasion (Acquisti (2023)).

I explore but do not find evidence consistent with a role for these two alternative explanations– data permanence and risk perception. In terms of data permanence, I explore whether consumers may give up on protecting the data they gave out previously. In other words, the cost of giving up

or exchanging a type of data becomes much lower if a consumer has already shared that data in previous periods. This would suggest that state dependence should persist over time, particularly when past choices involve the same types of data as future choices. In contrast, I find that state dependence is short-lived and that it still holds when new types of data are requested in the latter choice, with the effects being at least as strong when different data types are sought compared to when the same data type is requested.

In terms of risk perception, the state dependence could be due to consumers adjusting their risk perceptions after experiencing no negative consequences of accepting the data request. To explore this hypothesis, I compare users who are more likely to reject data requests prior to having used the battery app to users who are less likely to reject data requests. I find no significant difference in terms of state dependence for these groups, suggesting that users who might anticipate a high risk from sharing data do not adjust more than those who do not anticipate such a risk.

Beyond this suggestive evidence on why structural state dependence exists in privacy choices, recognizing and quantifying state dependence in consumer privacy decisions has important implications for digital platform strategy. Specifically, this structural state dependence suggests an externality from the data requests of some apps to the ability of other apps to collect data. By explicitly recognizing the connection between privacy choices over time, platforms should recognize the positive benefits of consumer-friendly privacy terms and the negative effects of terms that consumers are likely to reject. These externalities are particularly strong for apps and from privacy settings where the consumer is unlikely to have a strong preference. This leads to the managerial implication that the platform could boost the data consent rate of some apps, on which users are not sure about the value-cost tradeoff, leveraging users' state dependency on similar privacy choices.

Unlike prior research focused on state dependence for consumer-packaged goods and the implications for pricing and promotion strategy, the state dependence documented here suggests a connection between state dependence and externalities when considered in the context of privacy decisions on a digital platform. This finding aligns with the growing emphasis in privacy literature on the significant role of externalities in shaping privacy strategies and policies (Goldfarb and Que

(2023)).

The remainder of the paper is organized as follows. In section 2, I connect my results to research on privacy and on structural state dependence. I next combine insights from these two streams of research to develop a conceptual framework (section 3) for state dependent consumer privacy choices. In section 4, I describe the empirical strategy for the identification of state dependence and the institutional background. With this understanding of the institutional background in place, subsection 4.3 describes the construction of the dataset and suggestive evidence of state dependence and heterogeneity in privacy choices. section 5 presents the model estimates, documenting significant structural state dependence. section 6 examines the nature of this structural state dependence, and provides suggestive evidence that memory and preference strength play an important role in explaining the magnitude of structural state dependence. I conclude with a discussion of limitations and the managerial implications for the platform.

2 Literature

This paper relates to several strands of literature. The findings of this study contribute to our understanding of individual consumer privacy choices, addressing a gap where empirical evidence from the field has been limited. This study extends the concept of state dependence from traditional grocery decisions, as seen in scanner data, to digital privacy choices. It highlights that considering state dependence in consumer choices is crucial not only for brand promotion strategies in retailing but also for platform strategies aimed at retaining active users.

I. Consumer Privacy – First, my study relates closely to the literature on establishing an empirical framework to understand consumer privacy choices. Empirical evidence shows that consumers do care about their privacy. For example, given the price and functionality, consumers present lower demand for apps with sensitive data requests (Kummer and Schulte (2019)). When Apple's privacy label was released, consumers immediately decreased downloads on the iOS system (Bian et al. (2021)). In the Ad-Choices program, users opt out of behavioral advertising leaving the costs to the publishers and the exchange (Johnson et al. (2020)). In addition to the documented

evidence of the demand shift in response to privacy-related changes, there are endeavors to measure consumer privacy valuation. One of the latest comprehensive frameworks, a dual-preference framework that incorporates both intrinsic and instrumental components, measures privacy preferences from psychological and economic perspectives (Lin (2022)). In a related context, Tang (2019) quantifies the monetary value of privacy in an online borrowing context. She explicitly measures the value of several types of data, such as social network ID, and employer contact information. However, all the above-mentioned work measures consumers' privacy preferences in a static setting, or the changes in privacy choices in an aggregated fashion. There remains a dearth of systematic understanding regarding how a single user's privacy choices evolve over time. My research addresses this gap by examining consumers' sequential decisions to share data across time using field data.

Moreover, the study sheds light on unpacking a puzzle in current privacy debates: why are people giving out their information easily when they claim they care about their privacy? This question directly relates to the phenomenon of digital privacy paradox. The digital privacy paradox is a phenomenon characterized by a gap between consumers' stated privacy concerns and their actual behaviors (Norberg et al. (2007)). A body of empirical research seeks to understand this puzzle. Athey et al. (2017) established the digital privacy paradox empirically in a lab study where individuals give out private information for small incentives—such as a slice of pizza—even though they claim they care about privacy. Solove (2021) illustrates a methodological explanation– that the observed privacy behaviors are specific to particular contexts, whereas self-reported privacy concerns often derive from general surveys. This explanation suggests that privacy concerns are inherently context-dependent. Indeed, empirical evidence indicates that individuals' valuation of privacy is context-dependent and evolves over time (Goldfarb and Tucker (2012); Acquisti et al. (2015)). This perspective is consistent with Nissenbaum (2004)'s concept of contextual integrity.

Context serves as a key element in explaining the digital privacy paradox. Factors such as consumer ignorance and lack of attention affect privacy choices (Burtch et al. (2015); Adjerid et al. (2016, 2019)). Perceived control over personal data significantly influences user behavior, such that implementing a policy enhances the probability of users clicking on Facebook's advertisements (Tucker (2014)). Additionally, there is evidence suggesting a correlation between the benefits

derived from data flows and the desire to protect privacy, such that users with intense digital engagement could develop a stronger preference for privacy (Chen et al. (2021)). However, whether the change of contexts would affect how people make sequential digital privacy decisions has not been studied. Earlier field experiment work has shown that people's privacy evaluations are sensitive to contextual factors (Acquisti et al. (2013)). In recent advancements of data consent scenarios in mobile settings, contexts and the change of contexts become a more critical factor. By the definition nature of the mobile phone, the context changes when people move around. Marketers have found it more efficient to adopt contextual mobile advertising (Andrews et al. (2016)). The contextual information could not only improve the effectiveness of targeting but also could generate additional revenue by nudging customers to change future behavioral patterns (Ghose et al. (2019)). My results on the contextual state dependence effects have shown similar implications, where the platform could boost the data consent rate of some apps that users are not sure about the value-cost tradeoff, leveraging users' state dependency in similar privacy choices.

Choice architecture in privacy is an ongoing research stream, with lab studies showing that banner designs could impact users' future consent choices. This practice is called "dark patterns" by Bielova et al. (2024) and others. There is a related literature that seeks to predict privacy choices based on user characteristics and context in information systems and computer science (Wang and Li (2024); Serramia et al. (2023); Brandão et al. (2022)). All of this work highlights the discrepancy between consumers' true privacy preferences and their exhibited choices which, related to my results on consumer privacy choices, sheds light on the importance of understanding behavioral patterns. These papers often conclude with focus on implications for the design of privacy policies by firms and governments.

Nosko and Tadelis (2015) show that reputational externalities emerge when eBay buyers extrapolate their experiences with individual sellers to form impressions about the overall eBay market, and those who have a negative experience with a single seller show a decreased likelihood of returning to the platform. Jaffe et al. (2017) find similar results on Airbnb. The externalities in the privacy literature currently fall into three main themes[5]. First, an individual's choice to share data can directly impact others within their network (Acemoglu et al. (2022); Bonatti

[5]For a more comprehensive overview of the positive and negative externalities, refer to Goldfarb and Que (2023).

(2023)). Second, information about one individual provides probabilistic information about others since preferences and behaviors are correlated (Choi et al. (2019)). Third, consumers' decisions to withhold information can also expose their types through market interactions (Bergemann and Bonatti (2015)). This paper finds a new form of externality in the platform, where consumers' data consent choices to a specific app can positively or negatively impact the ability of other apps to collect data.

My research demonstrates that users are more likely to disclose their data if they have done so in previous decisions. Context plays a crucial role not only in individual privacy choices but also in the strength of state dependence between these choices. The identified state dependence in data-consent decisions in my work provides an answer to the public debate on why individuals claim to value data privacy yet appear to disregard it when sharing information across various digital platforms.

II. State Dependence – This work is also rooted in and significantly contributes to another stream of literature: the state-dependence literature in marketing and economics. The empirical regularity of state dependence is explained in two ways: structural state dependence, where experiencing an event alters future preferences, and spurious state dependence, which arises from unobserved consumer preferences. Since the 1980s, a consensus in the quantitative marketing literature has been growing that there is substantial evidence of persistent choices of consumers in the consumer packaged goods (CPG) industry (Keane (1997); Seetharaman et al. (1999); Dubé et al. (2008, 2010); Simonov et al. (2020)). Recently, Levine and Seiler (2022) found no evidence of such persistence when leveraging an exogenous shock that disrupted a single purchase occasion. Drawing on the Heckman (1981) framework, these works provide models and toolboxes to prove such persistence as a causal effect of past choices on current purchase behavior, which is structural state dependence, instead of the time-invariant preference heterogeneity, which is spurious state dependence.

In addition to the econometric models that quantify the state dependence, there is extensive research in consumer psychology on memory-based choices connected to the concept, particularly the accessibility-diagnosticity framework by Feldman and Lynch (1988). This framework posits that

memory input—in my setting, a prior data-consent choice—affects subsequent related behavior based on two factors: accessibility and diagnosticity. The observed diminishing effect of state dependence in my research aligns with the theory that past choices, recent and more accessible, exert a stronger influence. Furthermore, the heterogeneous results regarding context (familiar vs. unfamiliar use cases) and decision-making style (read vs. do not read privacy terms) indicate that individuals with lower diagnosticity of current information are more influenced by past choices.

Recognizing and quantifying state dependence holds significant implications. State dependence in consumer choices in the CPG industry directly impacts optimal firm policies, particularly pricing strategies. Specifically, the price of high-quality products declines relative to the low-quality products in the presence of consumer loyalty (Dubé et al. (2008)). Kong et al. (2022) discovered that structural state dependence accounts for at least 20% of choice persistence in the CPG market. Seetharaman (2004) looked into the sources of state dependence and underlined the lagged promotions to have carryover effects on habit persistence. Ignoring one or more sources of state dependence would underestimate the total incremental impact of a sales promotion. Similarly, Goldfarb (2006b) demonstrates the loyalty generated by switching costs on the Internet accounts for a large fraction of portal visits and revenues and suggests that random coefficient models overestimate true state dependence. Identifying structural state dependence is crucial, as short-term shocks or policy interventions can have significant, lasting effects on consumer behavior.

Moving from the CPG industry to digital platforms, particularly to the realm of digital privacy, highlights the importance of this dynamic. State-dependent consumer data-consent decisions indicate that even minor shocks can impact privacy choices considerably. Recognizing this, digital platforms can develop strategies, such as designing the ads display sequence, that comply with privacy regulations while promoting sustainable user engagement within their ecosystems. Simultaneously, policymakers can devise plans to anticipate and mitigate the extended consequences of temporary events. Furthermore, Levine and Seiler (2022) suggest that persistent behavioral changes develop over time and are influenced by a sequence of past choices. My findings similarly demonstrate that shocks temporarily affect privacy choices, with their impact largely dependent on the context. The extent to which privacy preferences are altered in the long term remains an

open question, warranting extensive future research.

3 Conceptual Framework

To provide intuition for my empirical analysis and identification strategy, I consider a simple model[6] of consumer privacy choices that allow for individual privacy preference heterogeneity as well as structural state dependence.

A consumer can choose to accept or reject the privacy terms of an app. I write $d \in \{0, 1\}$ to represent this decision, where $d = 1$ is acceptance and $d = 0$ is rejection. A consumer can only access the app's service by agreeing to the data-consent request. Upon downloading the app, consumer i decides to accept ($d = 1$) or reject ($d = 0$) the data-consent request from the app provider. Acceptance leads to the app being installed and opened, while rejection prevents the app from opening. The flow utility consumer i obtains from their decision d_{it} at choice t is given by:

$$u_{it}(d_{it})^j = \begin{cases} \underbrace{\delta_{it}^j - c_{it}^j}_{\text{value-cost tradeoff}} + \underbrace{\gamma_A \times \mathbf{1}(d_{i,t-1} = 1)}_{\text{state dependence}} + \quad \varepsilon_{it}^1, & d_{it} = 1 \\ \underbrace{\gamma_R \times \mathbf{1}(d_{i,t-1} = 0)}_{\text{state dependence}} + \quad \varepsilon_{it}^0, & d_{it} = 0 \end{cases} \tag{1}$$

When accepting at the current choice ($d_{it} = 1$), consumer i incurs a privacy cost of data consent, which is specified as c_{it}^j. This cost of data included in the model consists of a consumer-specific privacy preference, which is persistent individual heterogeneity, and a specific cost on observing the app-datatype pair. Therefore, in my setting, I am not able to separately identify the intrinsic preference from the instrumental as in Lin (2022). δ_{it}^j captures the value of the app j when consumer i accepts the data consent and adopts it at choice t. Together, consumer i faces the value-cost tradeoff of $\delta_{it}^j - c_{it}^j$ at every choice t, that is app-value and cost-of-data specific. The second term captures state dependence by allowing the utility to differ when the acceptance choice

[6]Note that I do not estimate a model of consumer choice. Instead, my analysis is based on the consequences of an underlying model of consumer choices with the pattern of persistence in acceptance or rejection.

was also made in the last choice $(t-1)$. $\mathbf{1}(d_{i,t-1}=1)$ is a dummy variable indicating whether or not the consumer accepted the data-consent request in the last choice. γ_A reveals the degree of state dependence of accepting if accepted last time, and γ_R reveals the degree of state dependence of accepting if rejecting last time. In the main results section below, the degree of state dependence, measured by γ_A and γ_R, is constant. In discussing potential explanations, I allow the degree of state dependence to vary across user groups and across contexts. Finally, ε_{it}^1 and ε_{it}^0 are logit taste shocks that are independent across consumers and choices. When rejecting $(d_{it}=0)$ at the current choice, consumer i does not enjoy the app value η_i^j or incur any cost of data consent. However, the individual receives the utility of state dependence if they rejected it in their last choice $(t-1)$. This is captured by γ_R.

Following Heckman (1981), structural state dependence and heterogeneity could be confounded. A core challenge in identifying structural state dependence lies in the fact that observed patterns of choice persistence could be produced by both preference heterogeneity and structural state dependence, making it difficult to distinguish between the two solely based on choice persistence data. In the data-consenting context of this study, the key challenge is to separately identify state dependence from consumers' underlying privacy preferences. Therefore, the key idea of the identification strategy that I detail below is that the privacy choice behavior in reaction to a shock is different if state dependence exists. An important aspect of comparing scenarios with and without state dependence is that behavior is different in the short run after an external shock. However, the impact of such a shock could wane over time, even when state dependence exists.

Revisiting the core idea behind my identification strategy, it posits that in the absence of structural state dependence, consumers' privacy decisions should be independent from one choice to the next. Consequently, the distribution of rejection share $(Pr(d_{it}=0))$ for each consumer is the same for each choice. This equality of rejection persistence should hold for each consumer and therefore, hold for the average rate of rejection persistence across the population. From this perspective, if the average persistence in choices $(Pr(d_{it}=0))$ does not immediately return to its baseline following a shock that forces users into acceptance or rejection, it suggests that the null hypothesis (that decisions across different time periods are independent) is not rejected. Im-

portantly, this characteristic of decision-making behavior is expected to persist irrespective of the distribution of preference heterogeneity in the population. In the following empirical analysis, I specifically focus on the external shock that forces users into acceptance. In this regard, I derive the probability of rejection ($Pr(d_{it} = 0)$) at choice t with respect to the condition where the consumer accepts the data request in their last choice ($d_{i,t-1} = 1$).

Specifically, a consumer chooses to reject ($d_{it} = 0$) if and only if $u^j_{it}(0) \geq u^j_{it}(1)$. The probability of consumer i choosing rejection ($d_{it} = 0$) is

$$
\begin{aligned}
Pr(d_{it} = 0) &= Pr\Big(u^j_{it}(0) \geq u^j_{it}(1)\Big) \\
&= Pr\Big(\gamma_R \mathbf{1}(d_{i,t-1} = 0) + \varepsilon^0_{it} \geq -\delta_i + \gamma_A \mathbf{1}(d_{i,t-1} = 1) + \eta^j_i + \varepsilon^1_{it}\Big) \\
&= Pr\Big(\gamma_R \mathbf{1}(d_{i,t-1} = 0) - \gamma_A \mathbf{1}(d_{i,t-1} = 1) + \delta_i - \eta^j_i + \varepsilon^0_{it} - \varepsilon^1_{it} \geq 0\Big) \\
&= Pr\Big(\gamma_R \big(1 - \mathbf{1}(d_{i,t-1} = 1)\big) - \gamma_A \mathbf{1}(d_{i,t-1} = 1) + \delta_i - \eta^j_i + \varepsilon_{it} \geq 0\Big) \\
&= Pr\Big(-\big(\gamma_R + \gamma_A\big)\mathbf{1}(d_{i,t-1} = 1) + \gamma_R + \delta_i - \eta^j_i + \varepsilon_{it} \geq 0\Big)
\end{aligned}
$$

Conceptually, without structural state dependence, the rejection probability ($Pr(d_{it} = 0)$) at choice t would be independent, and solely depends on $\eta^j_i - \delta_i$. However, if structural state dependence exists, the distribution of rejection share ($Pr(d_{it} = 0)$) for each consumer would change. The state-dependence γ will be estimated in an analysis in section 4 using the change that low-battery mandatory-acceptance shock imposes on the probability of rejection ($Pr(d_{it} = 0)$). Note that in this study, I am not able to disentangle γ_A and γ_R from γ by the nature of only "mandatory-acceptance" shock is leveraged in the reduced-form analysis.

4 Empirical Strategy

In this section, I introduce the institutional context and describe my empirical approach to identify the structural state dependence. I detail the construction of the dataset used for the primary analysis and how the key variables are defined and provide summary statistics alongside motivating

evidence. Subsequently, I introduce a staggered difference-in-differences design to assess the causal impact of the shock. The empirical framework is centered on Alipay, a global leader with an extensive in-platform app ecosystem and a vast user base. Their data-consent request before app usage creates a useful setting to understand consumers' privacy choices over time.

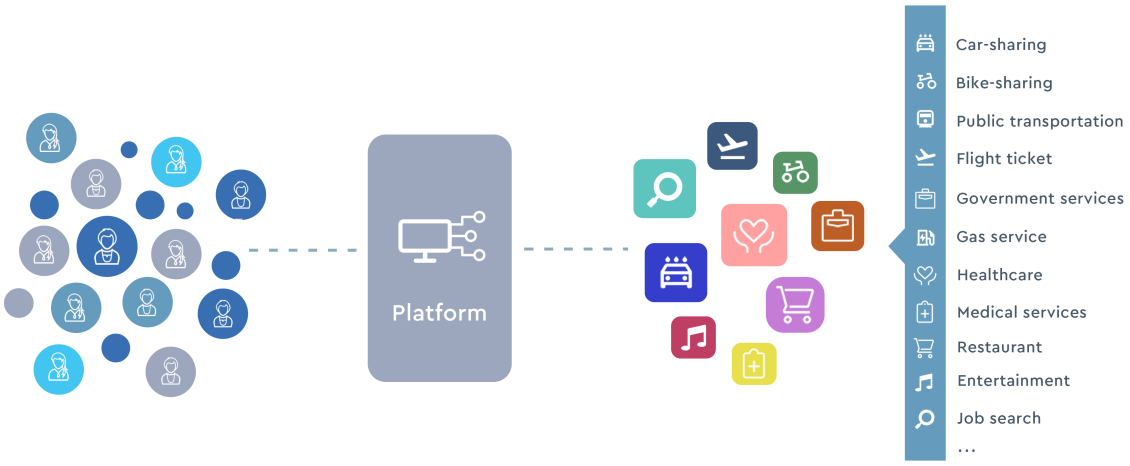
4.1 Institutional Background

Obtaining a suitable dataset to examine sequential consumer choices on data-consent decisions has been challenging. Assessing the central research question on state dependence requires individual-level data over a range of choice decisions. It requires a granular dataset that captures not only consumers' choices but also contexts applicable to other platforms worldwide, such as data-consent decisions on Facebook's third-party plug-ins or on cookies in the initial visits to various websites. I addressed these challenges by leveraging proprietary panel data from Alipay, which is at the individual-choice level, providing detailed information on user demographics, as well as abundant supply-side characteristics. With millions of data points, this data set allows for rich analysis of user data consent choices.

In empirical privacy research, privacy sensitivity is commonly measured by individuals' refusal to disclose personal information such as income, age, or zip code in surveys. A significant portion of previous research has relied on surveys and experiments to gather information on users' privacy preferences. These data sources provide useful information about consumer static preferences for privacy but they have limited the ability of researchers to study choices over time, particularly outside a laboratory setting.

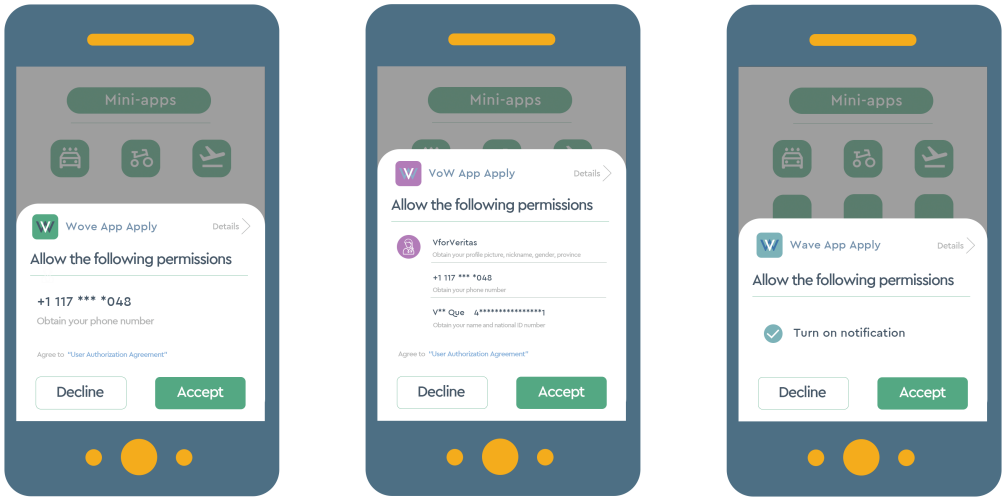
Using data from Alipay, I can study choices over time in an important real world setting. Alipay is a Chinese digital payment and lifestyle platform with more than a hundred million monthly active users. It is one of the largest digital platforms in the world. The platform supports hundreds of thousands of third-party apps for extensive digital products and services, such as tickets, food orders, health applications, personal tools, coupons and offers, life services, rent and resell, and entertainment activities (Figure 2). As the number of consumers has grown over time, the platform has attracted more and more individual apps with various uses, which in turn provide consumers

Figure 2: Empirical setting: a platform with consumers and third-party apps on both ends



with higher value through a diverse range of services offered by these separate nested apps. By using Alipay for apps, consumers avoid the friction of separately downloading, installing, and managing each app on their mobile phones.

Figure 3: Interface: binary choices & different data combinations

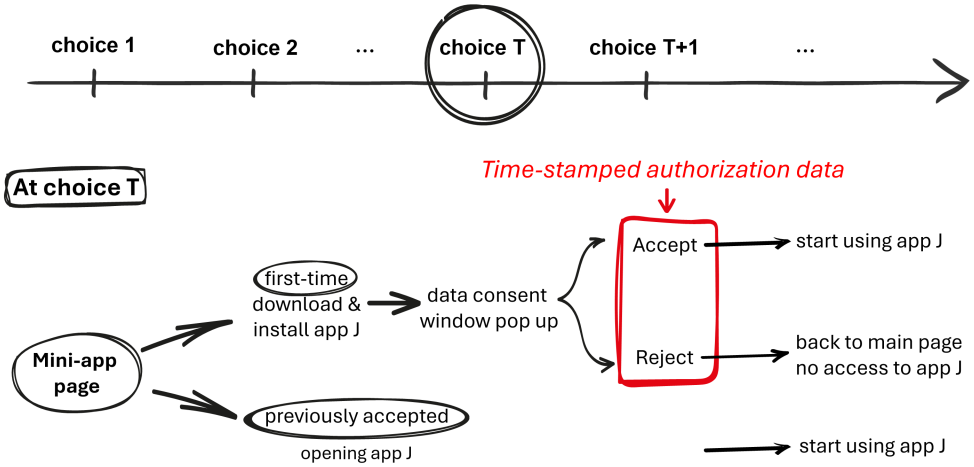


This concept is similar to third-party plug-ins found on Facebook or extensions available on the Chrome browser.

Through Alipay, in order to get access to the products and services provided by the apps, users need to authorize certain types of data when they first download and install an app. When prompted with a bundle request of data types, consumers have two options– to accept the data-sharing request by the app or to reject it (Figure 3). Only by accepting the whole bundle of data types can the consumers use the app. There are no partial acceptance or data customization options. This data-sharing setting makes the trade of consumer data for digital services explicit. Figure 4 illustrates the consumer journey, highlighting how privacy choices are integrated into the overall timeline of a user's app usage on the platform[7].

[7]At choice T on the timeline, if it is the initial download and installation of app J for consumer I, the individual must make data-consent decisions. If the individual rejects the data request, access to app J is denied; alternatively, if this is not the consumer's first interaction with app J, meaning they have previously engaged with the data request, consumer I can continue using app J without making further privacy decisions.

Figure 4: The Consumer Journey in the App Ecosystem



The bundle of data types requested varies across apps. It is the app owner's individual decision on what data types to request upfront, and there is considerable variation in the information requested. For example, some apps ask for general profile information only, others ask for credit information, and others request access to the step count.

Alipay standardized these mandatory upfront data requests in 2019, even though the payment platform itself launched in 2004 and the large scale third-party app ecosystem began in 2015. Therefore, this form of data-service consent decision was still new to consumers during the time period I study. The panel data underlying my analysis begins on March 1, 2020, after which comprehensive information covering the entire population was available, and my period of analysis continues until September 30, 2022.

4.2 Identification: A Quasi-Experiment

In the absence of structural state dependence, consumers' privacy choices should remain independent from one choice to the next. Consequently, the distribution of rejection shares for each consumer should be the same for each choice, irrespective of the previous rejection choice

To identify structural state dependence described in section 3, I leverage a quasi-experiment. Specifically, to separate the user's state dependence from heterogeneity, I use acceptance of data

requests from the portable power bank apps as an exogenous shock. The purpose is to identify whether the "mandatory" acceptance brought by the shock impacts a user's data-consent decision in the choices that follow.

Some prior work uses quasi-experiment variation to identify structural state dependence. Goldfarb (2006a) leverages internet denial of service (DoS) attacks by hackers to identify a medium-term lock-in effect on consumers' future choices of website visits. Sudhir and Yang (2014) study structural state dependence in the rental car setting where consumers get a different car from what they originally ordered in an upgrade event. Recently, Levine and Seiler (2022) show that households' purchases of bottled water do not exhibit structural state dependence when confronted with stockouts caused by hurricanes. In a related manner, Figueroa et al. (2020) show that the stockouts caused by an earthquake in Chile lower the market shares of the stock-out brands (top brands) post-shock. They interpret the effects as evidence that changes in the consideration set raised consumers' awareness of the competing products and thus changed their purchase behaviors persistently.

4.2.1 Identification Assumptions

The core identification assumption we need for this identification strategy to be valid is that the out-of-battery shock affects the subsequent choice only through its impact on the currently treated choice. In other words, the treatment is not correlated with any unobserved demand shocks to the next choice t.

For this assumption to be violated, either the low-battery situation should directly reduce the rejection probability on the subsequent choice or some omitted factor is leading to low rejection probabilities on consecutive choices. I argue that both of these are unlikely in the context of portable power bank apps.

The direct causal impact is unlikely because once the battery is charged then there is no longer an urgent need to rush decisions. The situation would be resolved. The omitted factor is unlikely for two reasons. First, consecutive decisions rarely occur over a short time period and are no driven by time-of-day or day-of-week effects. For example, only 5% of the observations occur

on the same day as the previous choice. I further test time-of-day and day-of-week effects in the robustness section below. Second, I observe no substantive changes in the variety or frequency of downloaded app categories before and after the treatment, suggesting treated users do not use the portable power bank apps because of a sudden change in preferences (Table C1)[8].

4.2.2 Industry Background of Portable Power Bank

In China, individuals rely heavily on their mobile device. Digital wallets enable people to go about their daily activities without the need for physical wallets or cards (Klein (2020))[9]. QR codes are ubiquitous, facilitating transactions in coffee shops, subway systems, malls, and health code verification. The average time spent on mobile devices per person has also surged, rising from an estimated 6.57 hours per day in May 2020 (the beginning of my data period) to 7.69 hours by May 2022[10].

The infrastructure of portable power banks is extensively developed across China, reaching urban and rural areas alike. By 2021, more than 200 million people in China had used shared portable power banks, averaging at least once per week[11]. Charging stations are deployed in various public facilities including but not limited to shopping malls, bars, restaurants, hotels, and tourist attractions. Figure 5 illustrates how the portable power bank works. When a user experiences low battery on their phone, they can turn to the portable power bank station, scan the QR code, and install the app. By granting data consent, the consumer can get the power bank out of the station and carry it around while charging. When the consumer finishes charging, they can return the power bank back to any station of the same brand and the consumer will be charged by the time of usage.

There were approximately 280 brands existing in China's market by 2022, 263 of which are

[8] for the treated and not-yet-tread users in the analysis, I calculate the downloading percentage of each category before and after the shock. The percentage numbers are calculated by the number of apps in this category over the total number of apps downloaded for each group. For a balancing purpose, I focus on three choices before and three after the shock for the treated users; the 2nd, 3rd, 4rd choices (comparable to the pre-shock period), and the 6th, 7th, 8th (comparable to the post-shock period) for the not-yet-treated ones.

[9] Details on the infrastructure sophistication can be found in this Brookings report: https://www.brookings.edu/wp-content/uploads/2020/04/FP_20200427_china_digital_payments_klein.pdf

[10] Source: FastData Industry Report on Sharing Power Bank 2023

[11] "Portable Charger Sharing Industry Appears to Have a Better Future Than Bike Sharing – But Still Holds Risks", chinamktginsights.com, January 7, 2021

Figure 5: An Illustration of Portable Power Bank Use Case



Note: A low-battery shock occurs when a user's phone battery depletes at public facilities (p1).To address this, users can rent a portable power bank by scanning a QR code at any station and installing the associated app. Upon accepting the privacy terms, users can remove a power bank (p2) and use it on the go (p3).Once charging is complete, they can return the power bank to any station of the same brand (p4).

covered in my dataset. The top three brands command nearly half of the market share (Table D8). While portable power banks have experienced significant growth in China, their market is not confined to the country. According to a recent report by Spherical Insights & Consulting, the global market for power bank rental services was valued at USD 7.1 billion in 2021, with projections suggesting it will reach USD 15.9 billion by 2030 [12].

4.2.3 Portable Power Bank Apps as the Instrument

Absent state dependence in a user's sequence of data-consent decisions, the "mandatory" acceptance of the power bank app during a low-battery emergency should not influence subsequent choices. Essentially, a consumer's likelihood of rejection is expected to return to baseline after such an event. The treatment is an increase in the likelihood that a consumer says "yes" in the middle of a digital privacy choice sequence. I argue that this allows me to separate structural state dependence in choices from persistent heterogeneity. I assume the shock happens to the treated users randomly in the position of their choice sequence, thereby enabling separate measurement of the state dependence from the underlying the sources of heterogeneity.

[12] https://www.globenewswire.com/en/news-release/2022/11/11/2554371/0/en/Global-Power-Bank-Rental-Services-Market-Size-to-reach-USD-15-9-Billion-by-2030-CAGR-of-13-5.html

As illustrated in Figure B5Panel A, the average rejection rate for portable power banks is markedly lower than for any other app type. This low rate of rejection holds across brands, with no brand having a rejection rate higher than 2 percent. The different portable power bank apps, however, do have substantial variation in the data requested, with some asking for more than one type of data and others only requesting basic profile information, or phone number. This exceptionally low average rejection rate regardless of the specific data requested supports the identifying assumption that the use of portable power bank apps is indeed urgent.

4.3 Data and Descriptive Evidence

The main data set is constructed from a panel of user data-consent decisions across all app categories on the platform. Data collection on individual data-consent choices began in September 2019; however, due to significant gaps in the records over the following few months, the data available for panel analysis starts March 1, 2020 and ends September 30, 2022, covering a total of 31 months. From the active user base during this period, I randomly sampled 800,000 users.

I described the shock I will use for the difference-in-differences research design above. To identify the structural state dependence in a sequence of consumer data-consent choices, the analysis includes only those who made at least eight data-consent decisions throughout the sample period. This selection criterion ensures adequate data for a burn-in period and for validating parallel trends pre-treatment, ultimately yielding a dataset of 475,139 active users. The primary analysis targets a subset of users who have both rejected and accepted a data request at least once, ensuring that their probabilities of rejection and acceptance are non-zero. This subset comprises 358,004 users. All individuals are de-identified before entering into the data analysis process.

The data includes three main categories of variables: user behavior (privacy decisions and reading privacy legal term decisions), user demographics, and app characteristics. Below are the main components of my data:

Data-consent decisions. An individual can choose to "accept" or "reject" the data request the first time they interact with an app. This is the focal decision I look into across time, which is observed on the user-app-choice level.

Reading privacy agreements decision. When consumers encounter a pop-up window requesting data consent for an app download, they have the option to review the privacy legal terms before selecting either "accept" or "reject". I observe the timestamp and which app's request they are responding to when they opt to read the privacy legal terms. However, the content of the legal terms and the duration of reading are not captured. Reading the agreement is optional. Consumers can make data-consent decisions without engaging with legal terms. Only 5% of the population ever read the agreements.

Demographic data. I observe the consumer's demographic information such as age, gender, education level, residential area (eg. province, city, city tier[13]), and the average balance on the account.

App data. I observe an app's merchant category code (MCC), which is a four-digit number assigned to a business by credit card companies (like Visa and MasterCard) to classify the business by the type of goods or services it provides. Additionally, I assess an app's popularity by aggregating the number of user downloads. I also observed the datatype combination an app requested, which may include any of the following, individually or in combination: user profile information (such as profile picture, nickname, province, and gender), phone number, address, credit history on the payment platform, national ID, student information, and daily step counts.

Firm and service description. I complement the panel data from Alipay with service description data from Qichacha[14]. Qichacha is a leading Chinese corporate database platform that provides comprehensive information on Chinese companies. It provides detailed information on companies, including product offering details, registration, financial records, and legal proceedings. The platform aggregates data from various official sources to ensure the accuracy and reliability of the information provided[15]. Qichacha is widely used by more than 400 million users for due

[13]In China, cities are often categorized into five tiers based on a variety of criteria such as economic performance, population size, and infrastructure development. First-tier cities are the largest and most developed cities with significant international influence and advanced infrastructure. The typical first-tier cities include Beijing, Shanghai, Guangzhou, and Shenzhen. Second-tier cities are regional economic hubs with rapidly growing economies and improving infrastructure. They include provincial capitals and major coastal cities such as Chengdu, Hangzhou, Wuhan, and Nanjing. Third-tier cities include smaller provincial capitals and other significant cities that have a moderate pace of economic development, such as Qingdao, Dalian, and Xiamen. Fourth-tier cities generally have smaller populations and less developed infrastructure. Fifth-tier cities are often referred to as county-level cities or small towns. Life in fifth-tier cities is much slower, and the focus is often on local markets and agriculture.

[14]https://www.qcc.com/

[15]Including government databases such as National Enterprise Credit Information Publicity System

diligence, market research, competitive analysis, and investment decision-making[16]. This textual data is manually collected and used for the heterogeneity analysis of contextual differences. I gather detailed descriptions of all firms operating on least one app on the Alipay platform. These descriptions provide the operational contexts and use cases, and technological capabilities, if any.

4.3.1 Sample Characteristics

The summary statistics of the users are shown in Table C2Panel A. Specifically, in the main sample, 55.15% are male. The average age is 35. The annual average asset balance on the platform is $1,251.67 in US dollars (10,966 in local currency). The sample of the treated users exhibits similar distributions in all demographic dimensions as the main sample. 57.28% of the user group are male, with an average age of 32. Their annual average asset balance on the platform is $1,229.26 in US dollars, very close to the main-sample average of $1,251.67. Table C2Panel A demonstrates that one-third of the population has ever been treated.

Table C2Panel B shows the summary statistics of the main characteristics of the 7321 apps that appeared in the main sample data. In the main sample comprising 358,000 users, the average number of users per app is 6,116, while the median stands at 202. The user distribution shows a significant range, with the largest app amassing over 333,000 users, and the smallest registering a single user. The user base of the apps exhibits a long-tailed distribution, indicating that a small number of apps have a very high number of users, while the majority have relatively few. During the sample period of 31 months, the average time between entry and exit of an app is 15 months, and the median is 13 months. In terms of the aggregated monthly activities on each app, the mean of monthly average users is 567 while the median is 16.

4.3.2 Motivating Analysis

This section documents three aspects of the data that motivate the main analysis. First, that the rejection rate for the current choice is lower when the preceding choice was acceptance rather than

(www.gsxt.gov.cn), China Judgements Online (wenshu.court.gov.cn), China's Enforcement Information Disclosure Website (zxgk.court.gov.cn), China National Intellectual Property Administration (www.cnipa.gov.cn), Trademark Office of China National Intellectual Property Administration (sbj.cnipa.gov.cn), National Copyright Administration (www.ncac.gov.cn)

[16]https://finance.yahoo.com/news/chinese-corporate-data-provider-qichacha-093000829.html

Figure 6: Rejection Rates by Previous Choice



rejection. Second, that there is considerable heterogeneity across demographic groups in terms of data-consent choices. Third, that there is considerable heterogeneity across app categories in terms of the rejection rate, with the lowest rate for the portable power bank category,.

I. Rejection Rates in Sequence

The starting point for a study of structural state dependence is an assessment of the observed correlation between current and past choices. In my data sample, on average, users generally exhibit lower rejection rates if they accepted the last choice (as illustrated in Figure 6). This correlation suggests a potential state dependence effect.

The core research question of this paper is whether state dependence exists in consumer data-consent decisions in an online sequential setting. Figure 6 indicates consumer's that a consumer's current choice may be influenced by their previous choice or there may be underlying preference heterogeneity. I describe my strategy to identify structural state dependence in subsection 4.2. Below, I show the results of staggered difference-in-differences analysis in section 5.

II. Heterogeneous Consumer Privacy Choices

Consumers do reject requests for data use permission, with an average rejection rate of approximately 13% in the data sample. This indicates that consumers do not always consent to share their data. In the institutional setting, in order to get access to the apps, consumers need to authorize certain types of data during the initial download and installation. Partial acceptance or customization of data permissions is not available. Consequently, if consumers encounter a privacy-intrusive situation that makes them uncomfortable sharing specific data types, they will reject the data request and forgo using the app.

This tendency to reject data requests from an app is heterogeneous across consumers, as shown in Figure B2Panel A. The sample panel data also reveals heterogeneity in the number of app downloads (Figure B2Panel B). A closer examination of privacy preferences across demographic groups highlights significant variation by age (Figure B3Panel A). This finding aligns with existing literature, such as studies by Goldfarb and Tucker (2012) and Tang (2019), which indicate that older individuals are more concerned about privacy issues. My data corroborates this, showing that the average rejection rate increases with age.

Education level also presents notable differences (Figure B3Panel B), particularly between those with only a high school diploma and those with higher education. Other demographic factors, such as gender, city tiers, and asset balance (serving as a proxy for income level), do not show significant systematic differences in our sample (Figure B4).

III. Heterogeneous Rejection Rates Across Apps and Contexts

Figure B5Panel A displays the rejection rate of each category in ascending order. Apps related to services like food delivery, restaurants, ride-sharing, bike-sharing, parking, and vending machines tend to have relatively low rejection rates. Notably, the "Portable Power Bank" category stands out with an exceptionally low rejection rate, with less than 1% of scenarios involving user rejection.

Rejection rates also vary by the types in the requests (Figure B5Panel B). The data that an app requests may include any of the following, individually or in combination: user profile information (such as profile picture, nickname, province, and gender), phone number, address, credit history on the payment platform, national ID, student information, and daily step counts[17].

[17]Here, type "other" means any request that is not related to a consumer's personal information. For example, the

Since the data requests are endogenously selected by the apps, there is no particular reason for the average rejection rate to be monotonically increasing with scope of the data request. The descriptive statistics demonstrate this. For instance, the data type "step counts" alone has a higher rejection rate than the combination of step counts, profile information, and others. Therefore, the particular context likely plays a role as consumers decide whether the benefits of accepting the data request outweigh the costs. This will be discussed in detail in subsection 6.2.

Table D8 provides a summary of the average rejection rate and market size for each app category, as well as the same information for the top apps (top 1 and top 3 respectively) in that category. The top 1 app in this table is defined as the number 1 app with the largest market share and user base within its category. Similarly, the top 3 apps are the first three holding the highest market shares in their respective categories. The categories are listed according to their alphabetical order. Overall, there is a great deal of variability in terms of market share and market size across app categories.

4.4 Staggered Difference-in-Differences Design

Difference-in-Differences (DiD) is a widely adopted research design to evaluate causal effects after an external shock. This section lays out the common specifications available, subsequently discussing the rationale behind the selection of suitable estimators. These choices are informed by the need to address issues highlighted in recent DiD literature. Common specifications can be broadly categorized as either "static" or "dynamic". I define $D_{it}^{l} := \mathbf{1}\{t - E_i = l\}$ to be an indicator for user i being l periods away from initial treatment at their choice t. For never-treated users, $E_i = \infty$ and $D_{it}^{l} = 0$ for all l and all t. The conventional two-way fixed effect (TWFE) framework adopts a static specification, estimating a single treatment effect that is time invariant:

$$Pr(d_{it} = 0) = \alpha_i + \lambda_t + \gamma \sum_{l \geq 0} \cdot D_{it}^{l} + \nu_{it} \tag{2}$$

requests of "turning on notification of this app", or "allow transitioning to the payment system" are both considered in type "other".

However, this specification is problematic in its interpretability under staggered rollout of the treatment as pointed out by Callaway and Sant'Anna (2021), Goodman-Bacon (2021), Sun and Abraham (2021), Borusyak et al. (2021), and De Chaisemartin and d'Haultfoeuille (2020). Specifically, γ as an aggregated parameter, recovers a weighted average of some underlying treatment effect parameters but could put negative weights on some of these parameters. Even in the case where weights are not negative, they are sensitive to the size of each group, the timing of treatment, and the total number of time periods (see Goodman-Bacon (2021) for an explanation of the Decomposition Theorem underlying this concern).

4.4.1 Callaway Sant'Anna's Estimator

To address the issue of heterogeneous treatment effects, the above-mentioned papers propose alternative DiD estimators to avoid the pitfalls associated with TWFE. I adopt the approach from Callaway and Sant'Anna (2021), focusing on a disaggregated causal parameter called "group-time average treatment effect", where a "group" is defined by the time period when units are first treated. The Callaway-Sant'Anna (CS) estimator is most suitable in my setting since it has weaker assumptions on a) whether the parallel trends assumption holds after conditioning on covariates; and b) whether the possible aggregations of group-time average treatment effects are flexible, which allows treatment effect heterogeneity both on cohort groups and across time.

Callaway and Sant'Anna (2021) describe several aggregation schemes, among which I focus on a summary parameter that circumvents the issue of changing the composition of groups by balancing the sample with respect to event time. Specifically, in my analysis, this treatment effect parameter is local to groups that participated in the treatment for at least 3 periods before and after. In this sense, the treated users in the main analysis are those who have their first choice for the burn-in period and at least 3 choices before and after the treatment. The unit of observation in my estimation sample is at the user(i)–choice(t) level. From section 3, the state-dependence effect on rejection probability on choices could be assessed through reduced-form analysis after treatment.

Under the assumptions of conditional parallel trends and limited anticipation, I measure the

model to capture the effect of state dependence at each choice after the shock as:

$$Pr(d_{it} = 0) = \alpha_i + \lambda_t + \left(\sum_{\tau \in \{-3,\ldots,3\}} \gamma_\tau \cdot D_{i,\tau} \right) + \epsilon_{it} \tag{3}$$

where $Pr(d_{it} = 0)$ is the probability of user i rejecting at choice t, which is derived in the section 3. $D_{i,\tau}$ is the indicator that equals to one if user i was at τ periods from the user i's initial treatment choice. α_i and λ_t are user and choice fixed effects. The user-fixed effect α_i captures time- and choice-invariant individual characteristics that would affect the privacy choices, while the choice-fixed effect λ_t captures user-invariant, choice-order-dependent factors that would affect the decisions. γ_τ measures the causal effect of the exogenous shock on privacy choices, which is the state dependence of the last choice on the current choice. Last, ϵ_{it} represents unobserved idiosyncratic shocks for user i at choice t with mean zero.

5 Results

In this section, I present the empirical findings. The primary analysis focuses on a sub-sample of users who have at least rejected the data request once in their lifetime (not necessary in the regression), ensuring their likelihood of rejection is not zero. For instance, if a user's 10th choice in the dataset is a rejection, but the treatment (in which the consumer used a battery app) occurred at the 5th choice, this particular rejection record is not included in the regression. However, this user is still being treated as a qualified user in the analysis since his rejection likelihood is not 0. Importantly, for the primary analysis, qualified users must have made at least three choices before and three choices after the treatment to test for parallel trends and to assess the main effect of state dependence, respectively. Next, I also present robustness checks for the primary analysis using the full sample where all users are included. The significance of the effect holds. In another robustness check, I further examine a subset of users who registered on the platform midway through the sample period to show that the state dependence holds.

5.1 Main Effect

The main results are reported in Table 1. Columns (1) and (2) capture the difference-in-differences estimates under standard two-way fixed-effect (TWFE) regressions with user and choice fixed effects. Columns (3)-(5) in Table 1 report the Callaway and Sant'Anna (CS) estimator of the regression. I used the choice order as the time fixed effect. [18]

Table 1: TWFE & Callaway and Sant'Anna Estimates (not-yet-treated as control group)

	(1) TWFE	(2) TWFE	(3) CS	(4) CS	(5) CS
Treated × 1st Choice Post	-0.019***	-0.012***	-0.022***	-0.021***	-0.019***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Treated × 2nd Choice Post	-0.014***	-0.009***	-0.017***	-0.017***	-0.016***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Treated × 3rd Choice Post	-0.010***	-0.007***	-0.015***	-0.014***	-0.013***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
Observations	603,631	603,631	603,631	603,631	603,631
Individual FE	Yes	Yes	Yes	Yes	Yes
Choice-order FE	Yes	Yes	Yes	Yes	Yes
Category Control	No	Yes	No	Yes	Yes
Datatype Control	No	Yes	No	No	Yes
Clustered error	user	user	user	user	user
Note:	baseline rejection rate before treatment is 14.5%				

Using the TWFE, column (1) shows that the magnitude of the average rejection rate dropping is 1.9 percentage points at first choice after the treatment. Since the baseline average rejection rate before treatment is 14.5%, the rejection rate declines by 13% after the treatment. Column (2) adds the category and datatype controls into the model and yields an estimated effect of 1.2 percentage points.

The Callaway and Sant'Anna (2021) estimates show effects of similar magnitude. Column (3) shows the basic CS estimator with user and choice fixed effects. Compared to the baseline, the

[18]The rationale is that, unlike traditional consumer shopping behaviors where seasonal effects are salient, calendar time in privacy preferences would probably not be the first priority. Instead, due to the hypothesis that there is state dependence across choices, then the main relevant time dimension is likely to be the choice order. As shown in Figure B6Panel A, the rejection rate is going downward smoothly for all the users in the sample and the qualified users with more than 8 choices.

rejection rate declines 2.2 percentage points or 15% after the treatment, as suggested by the CS estimator.[19] Column (4) and (5) add the category and datatype controls to the model.

The Callaway and Sant'Anna (2021) model provides additional controls and flexibility. Variation by calendar month (as shown in Figure B6Panel B) is addressed in Table C3. Furthermore, the Callaway and Sant'Anna (2021) eliminates the use of already treated users as a control. This means it is possible to reliably examine different control groups such as the not-yet-treated. Table 1 uses the not-yet-treated users as the control group. The CS estimates in this table do not suffer from the concern that users who are treated by the out-of-battery shock are inherently different from the users who never use the portable power bank apps. By using the not-yet-treated users as the control group, I show that users in the control groups have the same likelihood of running into the out-of-battery situation. As a robustness check, I also display Table C3, which uses both the never-treated and the not-yet-treated users as the control group to verify the main results. Adding never-treated users into the control group does not change the estimates of the CS model.[20]

Figure 7 displays the Callaway and Sant'Anna (CS) estimators for each choice order pre-and post-treatment. The parallel trend assumption appears to hold pre-treatment. There is also an apparent decline in the rejection rate for at least 3 choices after the treatment. The magnitude of the average rejection rate dropping is 2.2 percentage point at first choice after the treatment, which represents 15% decline from the pre-treatment rejection rate.

5.2 Robustness Checks

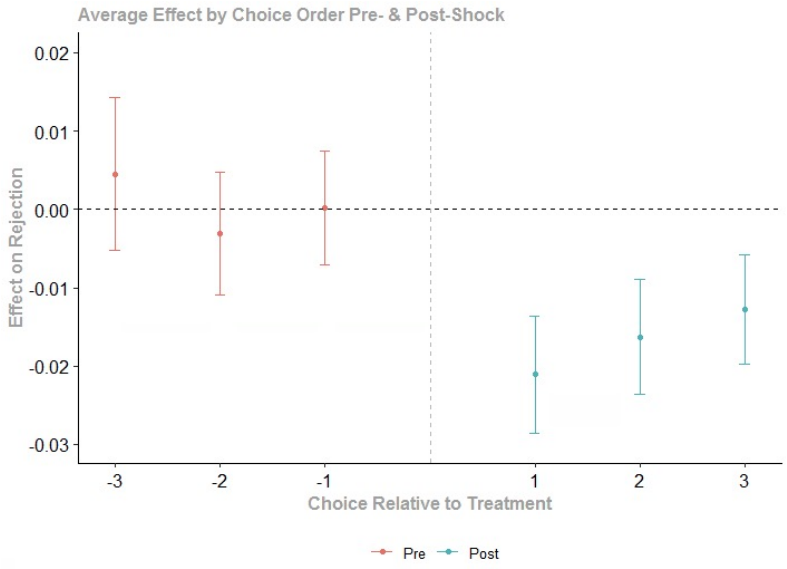
5.2.1 Time of day and day of the week

As noted earlier, one assumption underlying the use of portable power bank apps is that there is not an omitted variable that causes sequential decisions to have a low rejection rate. Several such potential omitted variables relate to to the time of day or day of the week. For example, users may be more likely to run out of battery on their phone at a later time of the day, say in the evening. If there are time effects on the rejection rate, for instance, that people are more likely to accept in

[19]The pre-treatment rejection rate is 14.5% and the percent change (ratio of 2.2 to 14.5) is 15%.

[20]The TWFE estimates in Table 1 and Table C3 are not reliable because they include the already treated in the control. I include them for completeness.

Figure 7: Event-Study Plot of Callaway and Sant'Anna Estimates



the evening, then treated users are more likely to have a lower rejection rate on their next choice. To test this possibility, I present the average rejection rates by hour of a day in Figure B7Panel A. Users demonstrate a consistently low rate of rejection during working hours, spanning from 9 am to 7 pm. This rate experiences a slight rise during the nighttime hours, particularly from 10 pm to midnight, and a significant surge is observed from 4 am to 7 am, attributed to the scarcity of data points in this early morning timeframe. This pattern suggests that, if anything time of day effects, if anything, bias the measured effects downward. Furthermore, only 5% of the observations with two consecutive choices happen on the same calendar day, further suggesting that the decreasing probability of rejection on the choice after the treatment is unlikely due to the time of day effect. In addition, Figure B7Panel B shows that the probability of rejection is constant across the day of a week.

5.2.2 Using not-yet-treated as the control group

I present robustness checks for the primary analysis using the full sample where all users are included[21]. The CS estimates using not-yet-treated as the control group are summarized in Ta-

[21]Recall that in the main analysis in Table 1,I include the users who have at least rejected the data request once in their lifetime. Therefore in this robustness check, I include the users that never rejected.

ble C4Panel A, while the estimates, in the model where the control group consists of users both never-treated and not-yet-treated, are summarized in Table C4Panel B.[22] The baseline rejection rates before treatment for both models decrease since we include the users who have zero likelihood of rejecting the request from the full panel. Therefore, even though the magnitude of the estimates in Table C4Panel A declines compared to the main results in Table 1, the relative drop from the baseline rejection rate stays unchanged. For instance, in the CS model Table 1, when I exclude the zero-likelihood of rejection users, the rejection rate declines 15% relative to the baseline. When including these tail users, the relative decrease is 14%.

5.2.3 Using a subsample of new users on the platform

The start of my analysis period, in March 2020, marks an early phase in the evolution of this data-consent format. However, many users will nevertheless have experience with data-consent decisions in the preceding months. Therefore, I examine robustness to completely new users who have yet to encounter this data-consent format on this platform.

There are 10,032 users who registered midway through the sample period or later, and made 7 or more choices (3 before and 3 after the treatment for the treated group). [23] This analysis is somewhat limited because 1) newly registered users are rare since the platform has its prevalent establishment in China for more than 15 years; and 2) just 887 users are in the treatment group, while 9,145 are the ones in never-treated or not-yet-treated groups. The results of the analysis on new-registered users are presented in Table C5 and show coefficients of a similar magnitude to the main analysis, although will less statistical precision.

This section documented that there exists state dependence after the use of a portable power bank app. Next, I explore possible explanations for this state dependence.

[22] As noted above, the TWFE estimates in Table C3, Table C4Panel A, and Table C4Panel B are not reliable because they include the already treated in the control. I include them for completeness.

[23] Figure B8 shows the overall the distribution of the timing of the first choice for each user.

6 Potential Explanations

The main results show state dependence exists after the shock. That is, in consumers' online sequential data-consent decisions, their current choices are affected by their past choices. After being treated by the need to use a portable power bank app, consumers reduce their rejection rate on subsequent data requests and are more willing to give out data in the following privacy decisions.

There is not a literature focused on state dependence in privacy per se. There are, however, a number of papers that provide possible mechanisms for the state dependence that I document.

Decaying Effect– Lynch and Srull (1982) introduce the concept of "stimulus-based choice" to describe situations where all relevant information is available and processed at the time an individual makes decisions. On the contrary, a pure "memory-based choice" happens when individuals make decisions without any present information, relying solely on memory. However, most real-world choices are a blend of both. That is, the customer uses a combination of remembered information and options, and information available at the moment of decision-making. A rich stream of literature has shown that past attitudes and elicitation of preferences, as well as choices, can influence subsequent behavior. A particularly relevant theoretical model for understanding these dynamics is the Accessibility-Diagnosticity framework (Feldman and Lynch (1988); Lynch et al. (1988)), which underlines that past choices, when highly accessible, can serve as indicators of preferences and thus be repeated (Biehal and Chakravarti (1983)). Higher accessibility of past choices occurs when they are more readily recalled from memory; diminished memory recall over time reduces the likelihood of repeating these choices, and so state dependence, due to decreased accessibility (Alba et al. (1992)). For accessibility, I demonstrate that state dependence intensifies when the temporal gap between a past choice[24] and the current decision is shorter.

Strength of Preference– The role of preferences is a simple reflection of the marginal impact of state dependence on choice. If preferences are strong, then a one-time shock will have little impact on behavior. This is consistent with the discussion in Webb et al. (2024), that persistent preferences and behaviors may take time to form, and that a short-term interruption in consumption did not sufficiently disrupt established habits. In choice settings, Roy et al. (1996) lay out a similar

[24]Here, the past choice for treated users is their choice under the shock (treatment)

framework where state dependence incentivizes trials, while strong habit persistence (i.e. strong preferences) hinders consumers from switching brands. Similarly, Levine and Seiler (2022) found that strong preferences persisted despite an exogenous shock disrupting a single purchase. The focus on marginal effects remains whether the strong preferences are driven by the benefits of data sharing or the benefits of privacy. Thus, if consumers appear to have strong habits in their data consent decisions or to be deeply engaged in privacy deliberation then measured state dependence is likely to be small because a one-time shock should have little effect on their choices.

Data Permanence– This relates to the idea that once a consumer discloses a particular type of data on the platform, the consumer may no longer see a benefit to protecting that data further. As Bonatti (2023) notes, data is non-rival and so secondary usage can be repeated without degrading the quality of the initial data disclosure. Therefore, if consumers perceive that disclosing data to one app effectively implies disclosure to all apps, then consumers will only be privacy-protective the first time they disclose a particular type of data.

Risk Perception– The lowered rejection rates could emerge from consumers observing that data sharing carries no adverse effects, thus lowering their risk perception after the shock. Lin (2022) shows consumers' beliefs on the consequences of sharing data are of first-order importance in understanding context-dependent privacy preferences. Risk perception refers to the user's perception of potential negative consequences after data disclosure (Tucker (2023); Acquisti (2023)).

6.1 The Decaying Effect

The Accessibility-Diagnosticity framework (Feldman and Lynch (1988); Lynch et al. (1988)) predicts diminished memory recall over time reduces the likelihood of repeating these choices, and thus, the magnitude of the state dependence effect. In the main result (Figure 7 and Table 1), we see the magnitudes of the estimators are decreasing with the number of choices post-treatment. This is consistent with what the framework predicts.

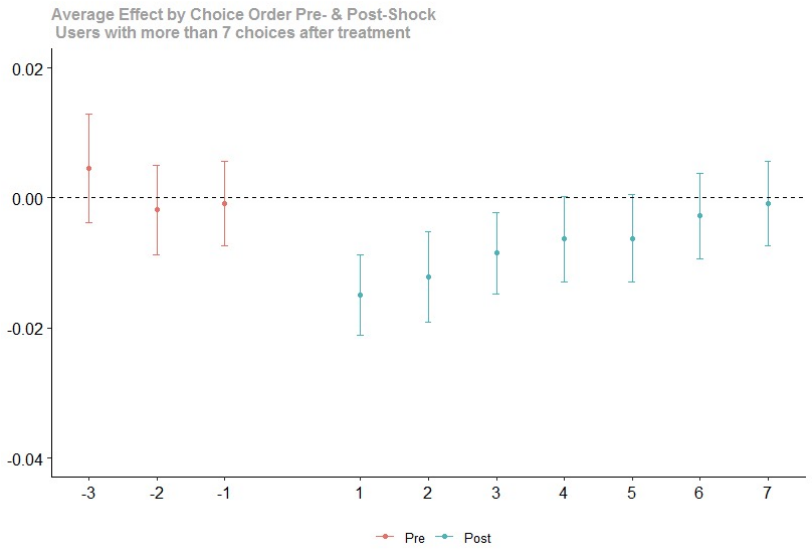
To further investigate this decreasing trend and the length of this effect, I sub-sampled users with more choices after the treatment in the data period. Specifically, the subsample with users making more downloads, and therefore the corresponding data consent decisions, follow the same

analysis using Callaway and Sant'Anna (CS) estimation:

$$Pr(d_{it} = 0) = \alpha_i + \lambda_t + \left(\sum_{\tau \in \{-3,\ldots,7\}} \gamma_\tau \cdot D_{i,\tau} \right) + \epsilon_{it} \tag{4}$$

Figure 8 shows that the average effect size decreases with each choice after the treatment. In this sub-sample of users having more than 12 choices[25], the average effect goes to null after the 7th choice post-treatment. Therefore, consistent with a role for memory, the effect of the treatment decreases over time.

Figure 8: Decaying estimate with more choices



I also look into the role of memory by examining the duration between choices. I classify users into "short-duration" versus "long-duration" groups, depending on whether their first choice post-treatment occurs within a week[26].

Assigning control groups to these "short" and "long" groups presents challenges, as it is difficult to pinpoint the first post-treatment choice for control individuals. To address this, I narrowed the focus to treated users who received their treatment at their fifth choice. Consequently,

[25] I do not include any users with fewer than 12 choices in this sub-sample.

[26] The rationale for defining "short" versus "long" duration as one week stems from the distribution of intervals across all observations. The instances where the duration that is shorter than one week account for approximately 25% of the observations. The median duration is three weeks.

control users can be divided into two duration categories based on the interval between their fifth and sixth choices.

Figure B10 demonstrates that the main effect is significant and much larger for the users who have a short duration between the shock and their first choice following the shock. More specifically, if the next choice following the shock happened within a week, then the magnitude of the effect (-0.0382) is more than twice the magnitude of the group with duration over a week (-0.0185).

As the choice of one week as a threshold is necessarily arbitrary, I present the coefficients of CS estimators in Table C7 to account for the decaying effect if we define the "long" group by various durations between the treatment and the first post-treatment choice. Column 1 shows three days or more as the threshold (comprising 82% of the sample). The relative drop in rejection rate from the baseline is 20% (i.e. a 2.2% drop from 11%). Column (2) uses 7 days or more, encompassing 75% of the sample, with a reduction in the baseline rejection rate of 17.6%. Following the same rationale, column (3) accounts for 50% of the users in the sample, capturing those who make their subsequent choice three weeks post-treatment. Finally, in column 4 the treatment effects become smaller and statistically insignificant for decisions made beyond five weeks from the treatment.

Thus, the magnitude of the state dependence effect diminishes with each subsequent choice following the initial shock and the effect size decreases as the time interval between a previous and current decision increases. These findings, which highlight the decaying nature of the effect over time, suggest the role of memory in state dependence and that the effect is temporary.

6.2 Strength of Preference

The strength of preferences directly influences the marginal effect of state dependence on choice. Webb et al. (2024) suggests that a single, isolated shock will have minimal impact on switching behavior when preferences are habitual and enduring. The Accessibility-Diagnosticity framework also implies that individuals are more likely to retrieve past choices to help current decision when they do not have a strong or clear preference in the current context. To evaluate the role of preference strength, I examine the heterogeneous effects across both app usage contexts and user

attributes.

6.2.1 State Dependence is Stronger When the Next Choice is in Weak Preference Contexts

The existing descriptive statistics show varying rejection rates across app categories (Table D8), suggesting that individuals could follow different decision strategies depending on the context. To further explore the mechanism of this observed state dependence between choices, I next examine whether and how the magnitude of the estimated state dependence shifts when the context differs, specifically, based on a proxy for the strength of preferences.

The context of the decision may affect the estimated state dependence. In scenarios involving data-consent decisions in app downloads and installations, relevant context factors may include situational factors such as location, use case, and urgency. In my analysis, I use large textual data of firms' service descriptions to create a proxy for the likely strength of preference for the particular app conditional on being sufficiently interested in the app to download it.

In order to do this large-scale classification task, I leverage recent advances in large language models (LLMs). With LLMs, it is now possible to classify large quantities of text through clearly articulated prompts.

In the fields of economics and marketing, machine learning techniques like the Random Forest classifier are often used to categorize mobile apps by transforming app descriptions into vectors (Ershov (2024)). These methods assign weights to words in the descriptions based on frequency, so that the descriptions could be mapped to the classified apps. Embedding techniques such as BERT (Devlin et al. (2018)), Sentence-BERT (Reimers and Gurevych (2019)), and SimCSE (Gao et al. (2021)) have been implemented to measure the similarity between patents and app descriptions, or between patents and occupation descriptions, to represent exposure to technology (Sun and Trefler (2023)). These pre-trained network models are nevertheless limited in their ability to classify text descriptions that vary widely in wording and style, as is the case with the app descriptions that I analyze.

In contrast, the most recent advancements in GPT-3 (Brown et al. (2020)) and GPT-4

(Achiam et al. (2023)) introduce a more suitable approach for classifying situational factors related to app use cases. These models leverage in-context learning, or few-shot learning, where the model learns a natural language instruction and a few demonstrations of the task. This new learning paradigm allows the model to adapt to further instances of the task. The primary advantage of GPT-3 and GPT-4 is that the models are trained to follow instructions presented in prompts, shifting the researcher's task to instruction tuning (Ouyang et al. (2022)).

This flexibility in LLMs allows me to develop a nuanced definition of situational differences in privacy decisions in interactions with apps. Specifically, current embedding models have their own well-defined reward policies to measure bag of words and to provide similarity measures. However, keyword similarity does not always correspond to the same level of user preference strength.

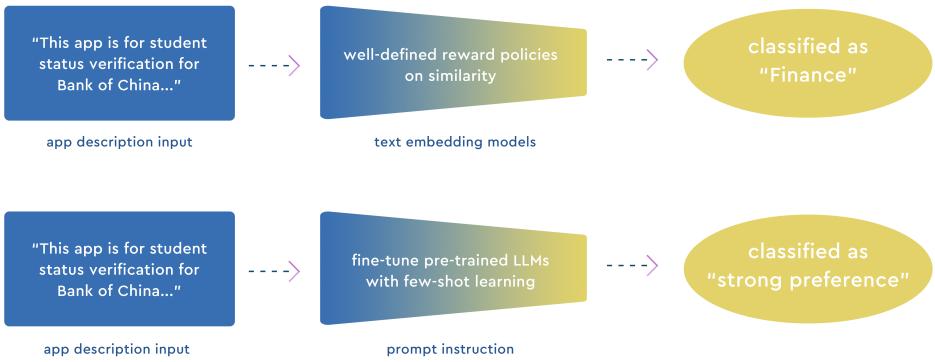
For instance, a restaurant's app for menu exploration and payment might show high textual similarity with a restaurant discovery app, yet these contexts are categorized as different levels of preference strengths in my rubric (detailed below). A restaurant discovery app might exhibit low textual similarity with a job search app, but my rubric groups them together as same group of preference strengths due to their functional and situational similarities. Figure 9 illustrates an example of how the app for student status verification purposes could be classified as Finance, where most apps in this category fall in weak preference context in my definition. On the contrary, with detailed prompt instructions and examples, the LLMs classified the app as strong preference context.

At the time of labeling, my analysis included models in the GPT-3.5, GPT-4, and GPT-4o families. The difference between the models lies in the contextual handling. GPT-4 and GPT-4o family handles subtle distinctions and complex decision trees more effectively, which is essential for my task where multi-step classification rules need to be applied accurately. Generally, in-context learning models provide better tools for understanding the nuances between contexts. I follow similar steps in Eloundou et al. (2023) for instruction tuning and data labeling[27].

To construct my primary classification to proxy for preference strength conditional on downloading, I collected both human annotations and GPT classifications, using a prompt tuned for

[27]The rubric is derived from the general manual (Brand et al. (2023)) of using ChatGPT as a research tool.

Figure 9: An Illustration of the Advantages of LLMs on Capturing Situational Factors



agreement with a sample of labels from three individuals (including the author). I observe similar agreement levels in GPT responses and between human and machine evaluations.

- Human Ratings: The author and two other human annotators (undergraduate and graduate research assistants) labeled an extensive set of service descriptions to the context of preference strength[28]. Three human annotations are mapped to a single classification standards.
- GPT-3.5 Ratings: Agreement rate is 72.8%.
- GPT-4 Ratings: Agreement rate is 84.1%.

I propose a new rubric for service contexts and the summary of the rubric is listed below, while the complete rubric can be found in Appendix A. I divide apps into strong preferences, weak preferences, and lottery. I include lottery as a separate category because I anticipated that they would not fit well into the descriptions I gave for either strong or weak preferences.

Summary of Service Rubric

- **Strong Preferences (S0) if:**
  - services under strong preferences are often on-site necessities services, which cover consumers' standardized tasks and are largely repetitive;

[28]Human annotations are performed on a sample of 800 apps. This sample is labeled to check the agreement rates. The full sample contains 7,321 apps in total.

  - the exclusive justification for downloading the designated app arises when a user is physically present at the specified facility, as the services provided by the app are confined to this particular location and are not applicable elsewhere;
  - if a consumer is at the facility, there is no space for outside options – meaning the main activities can not be achieved without using the app.

- **Weak Preferences (S1) if:**
  - services under weak preferences are often value-added services, which provide additional benefits and supplementary value but are not critical for routine life;
  - there is no need to be physically present at a specified facility to justify the app download;
  - if a consumer is at the facility, there is relative space for outside options – meaning the main activities can still be achieved without using the app.

- **Lottery (S2) if:**
  - apps under this context are the ones for marketing campaigns, coupon, cashback or lottery purposes;
  - this category is specialized in a lottery and gambling nature

I present the results in Table 2, which is based on the service rubric previously defined, where context refers to the preference strength of the use scenario to the consumer[29]. Users are grouped based on the app context of their first choice post-shock. Similar to the accessibility analysis in the previous section, identifying the "first post-treatment" choice for control individuals is challenging. To address this, I focus on treated users who encountered their treatment at their fifth choice. Consequently, control users are grouped into the context classifications based on their sixth choice. Column (1) shows the state dependence effect is both statistically and economically significant. Conditioning on interacting with apps of weak preference strength in the sixth choice, users who were exposed to the low-battery acceptance shock significantly reduced their rejection probability by an average of 3.7%, which translates to a 31.4% reduction from the baseline rate. Column (2) reveals that the state dependence effect is absent when comparing the treatment and control groups within the routine app context. I include column (3) for completeness. I urge caution in

[29]Figure B9 presents an illustration of the scenarios of the classification where most apps in a category fall in. Pictures are generated by GenAI.

interpreting the result in Column (3) due to the high variance in rejection rates within this category, which may contribute to the statistical insignificance.

This difference in rejection rates is about the apps rather than the users. There are no systematic differences in the pre-shock baseline rejection rate between individuals in the weak and strong preference groups at their sixth choice.

Table 2: Heterogeneity in Context – Weak versus Strong Preferences

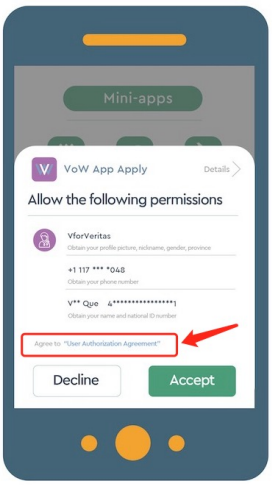
	(1) Weak Preference	(2) Strong Preference	(3) Lottery
Treated × 1st Choice Post	**-0.037*** **	**0.008**	**0.014**
	(0.007)	(0.005)	(0.016)
Observations	147,349	299,586	54,114
Individual FE	Yes	Yes	Yes
Choice-order FE	Yes	Yes	Yes
Clustered error	user	user	user
*baseline pre-shock rej_rate	11.80%	10.98%	14.31%

Note: Preference strength (weak vs. strong) was based on categorization by a large language model (LLM) on service description. This categorization focused on whether the app provided an on-site service that required physical presence without a readily available substitute.

Column (1) and (2) in Table 2 reveal that users maintain their established preferences in strong preference contexts, showing no significant difference in rejection probabilities between those who have and have not recently experienced the shock. Conversely, in weak preference contexts, users are more likely to base their decisions on their most recent choices, resulting in a notable state dependence effect.

Overall, strong habitual preferences hinder consumers from switching, where we would not expect state dependence to happen after one shock (Roy et al. (1996)). Consequently, if consumers exhibit strong habits in their data consent decisions, the measured state dependence effect is expected to be negligible. The impact of state dependence is most pronounced among consumers who are on the margin. This heterogeneity in contextual decision-making aligns with psychological

Figure 10: An Illustration of the Privacy Agreement that Attentive User Clicks On



theories, which suggest that when decision problems require deliberation or are complex (i.e., weak preferences in the current setting), individuals are more likely to rely on heuristics derived from past choices (Payne et al. (1993)).

6.2.2 State Dependence is Stronger for Users Who Do Not Have Strong Preference on Privacy

This section explores the heterogeneous effect of the observed state dependence by whether the user appears to have strong privacy preferences. I proxy for privacy preferences with information on whether the user ever clicked on the "user authorization agreement" (illustrated in Figure 10) before any data consent decision. The "user authorization agreement" specifies the legal terms that regulate the apps on the data they collect. These users that examined the privacy information consist of 5% of the population. Because I use whether a user examined privacy information as a proxy for strong privacy preferences, I treat it as a fixed trait across individuals.

Table 3 shows that the main effect is insignificant for the users who examined privacy information (column 1), while the significance and magnitude of the effects hold for the rest of the users. Moreover, the average pre-treatment baseline rate of rejection of the users who examined

privacy information is slightly but not systematically higher than the average rate of the other users– 19.5% compared to 14.3%. According to the distribution of the rejection rate across the whole sample Figure B2Panel A, the rate of rejection for users that examined privacy information is not systematically located on the right tail of the distribution.

Table 3: State dependence by whether user read privacy terms

	(1) Users who read privacy terms	(2) Users who did not read privacy terms
Treated × 1st Choice Post	-0.007	-0.022***
	(0.016)	(0.003)
Treated × 2nd Choice Post	-0.011	-0.016***
	(0.015)	(0.003)
Treated × 3rd Choice Post	-0.005	-0.013***
	(0.017)	(0.003)
Observations	31,241	734,974
Individual FE	Yes	Yes
Choice-order FE	Yes	Yes
Clustered error	user	user
Note:	baseline rate pre-shock: 19.5% (attentive) and 14.3% (inattentive)	

In summary, the heterogeneity in users' level of engagement with privacy deliberation, serves as a critical indicator of strong or weak preferences, respectively. In line with Webb et al. (2024), consumers with stickier and stronger preferences are expected to exhibit negligible state dependence effects.

6.3 Alternative Explanations for the Observed Effect

In this section, I explore potential alternative explanations for the observed effect where treated users lower their rejection rates post-treatment.

Due to the permanence of data, consumers might relinquish data protection efforts on data they have previously exposed to third-party apps. The perceived cost of surrendering or sharing a particular data point diminishes if the consumer has previously shared it. This relates to the literature on the non-rivalry of data (Bonatti (2023)). The non-rival nature of data enables its secondary usage without compromising the quality of the original disclosure. Consequently, if

consumers believe that sharing data with a single app essentially equates to revealing it to all apps, they will likely exhibit privacy-protective behavior only during their initial disclosure of a specific data type.

Consumers may also be more inclined to share data after being treated by the need to use a specific app because of an adjustment in risk perception when experiencing no negative consequences by sharing their data. Risk perception, which refers to users' beliefs about the potential negative consequences of data disclosure, plays a role in shaping consumers' context-dependent privacy preferences (Lin (2022); Acquisti (2023)), and so may affect choices in the context of data consent decisions on Alipay.

In the empirical evidence presented below, I do not find support for the arguments that estimated state dependence is related to data permanence and risk perception, contrasting with the evidence supporting state dependence based on memory and preference strength theories.

6.3.1 Data Permanence

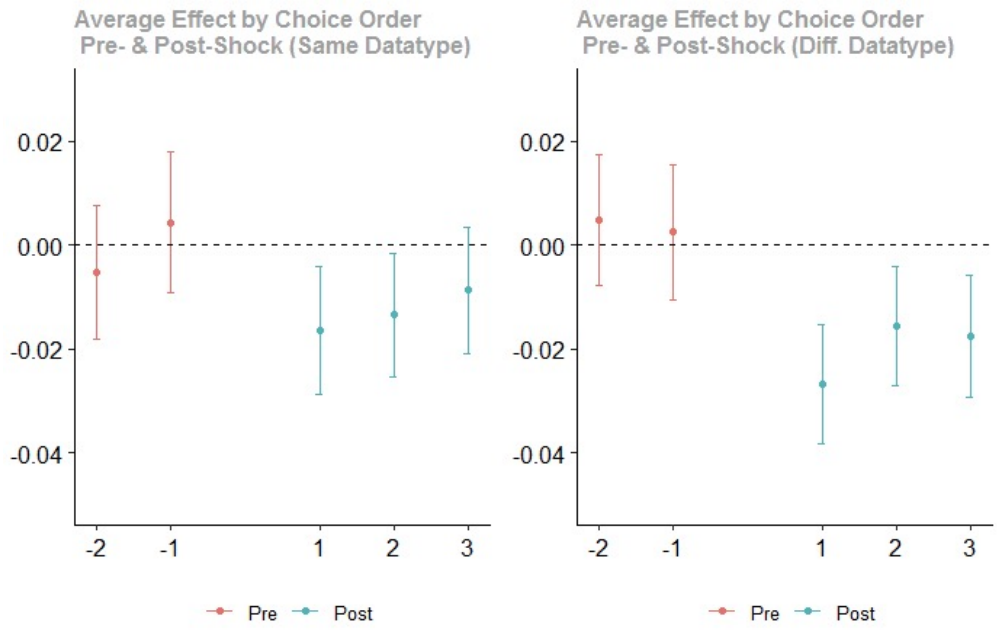
This potential explanation suggests that if consumers believe disclosing data to one app implies disclosure to all apps, they will likely only be privacy-protective during their first disclosure of a particular data type, due to the data's non-rival nature.

To examine whether state dependence may be driven by data permanence, I examine whether the subsequent choices involve different data requests than the portable power bank choice of the treatment. This analysis explores whether a consumer who shares a particular data type during the treatment is less likely to guard this information closely thereafter. The marginal cost of giving out this data decreases and the consumer has little reason to actively protecting it anymore. In this regard, I expect the state dependence effect would not exist on the following choices with different data types ($data_{-d_j}$) from the one asked in the treatment($data_{d_j}$). In other words, since consumers are protecting their privacy with respect to individual data items, being forced to give out a specific data type in the treatment should only have the state-dependence effect on the same data type in future decisions. I examine this possibility in two separate analyses.

First, I categorize users into two distinct groups: those encountering the same or fewer data

types (d_g where $g \in J$) on the subsequent choice as those presented during the treatment (d_J where $J = \{1, ..., J\}$), and those exposed to new or different data types (d_k where $\exists k_0$ s.t. $k_0 \notin J$) on the subsequent choice than those in the treatment. Again, to address the challenge for assigning control groups, the analysis is confined to users who received their treatment at their fifth choice. Consequently, the "next choice" for the control group is simply their sixth choice.

Figure 11: Test for data permanence: State dependence exists for both conditions



If the hypothesis is true that consumers give up protecting a particular data type subsequent to a compulsory data-consenting decision, one would anticipate a more pronounced treatment effect in instances with the same or fewer data types. Conversely, the effect is expected to be small or potentially negligible in scenarios where the choice contains different data types. Yet, as illustrated in Figure 11, significant treatment effects are observed across both circumstances (and, if anything, effects are smaller when the same data type is requested).

Second, I focus on consumers' decisions regarding specific types of data. There are three types of data requested in the portable power bank app treatment – profile information, phone number, and credit information, with different apps seeking different data types. I perform the staggered difference-in-differences estimation on the pre- and post-treatment choices with all but

the focal treated data type.

Figure 12: Test for data permanence: Treatment as another data type

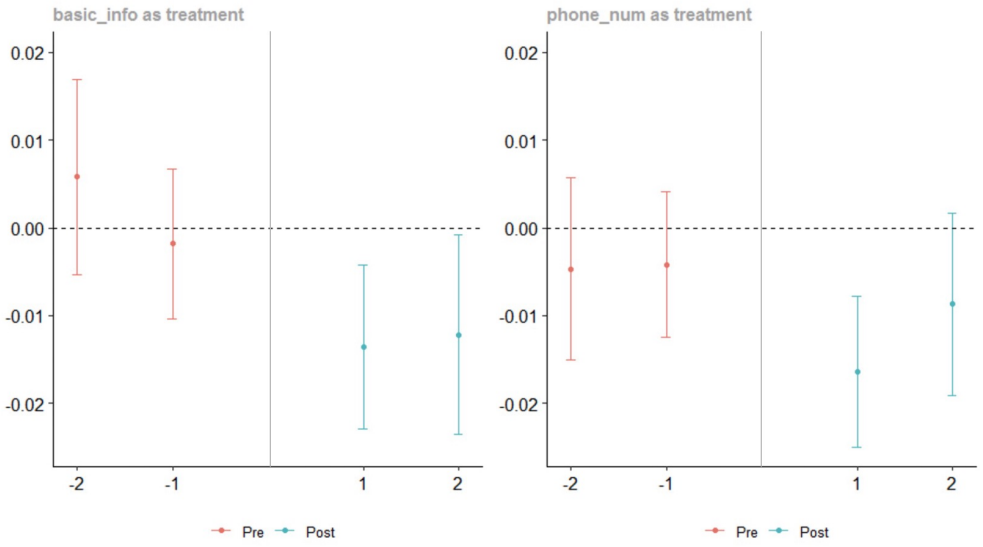


Figure 12 shows that when profile information or phone number is requested and shared during the treatment, subsequent choices involving other data types also exhibit a state dependence effect. The analysis results for the scenario where credit information is requested in the treatment are not shown due to sparse data in this estimation. Overall, the test suggests that the state dependence in the sequential privacy decisions exists regardless of data types.

Furthermore, if it is consumers giving up protecting their data, the size of the main effect should not vary due to the time duration between choices. Therefore, data permanence is not consistent with the evidence presented above that the magnitude of the effect is decaying across the choice order and that the size of the main effect decreases with the duration between the treatment and the next choice.

Overall, I do not find evidence to suggest that data permanence concerns underlie the state dependence.

6.3.2 Risk Perception Adjustment

This potential explanation for the estimated state dependence suggests that once consumers share their data with portable power bank apps and observe no adverse outcomes, their initial concerns dissipate. Consequently, their perceived risk associated with data-consent decisions decreases. If the theory holds, the treatment should exert a more substantial impact on those with higher risk perceptions. Specifically, I anticipate a more pronounced decrease in rejection rates among the high-risk perception group, as they are more likely to adjust their beliefs significantly.

To examine this explanation, I proxied for high and low risk perception using rejection rates from the historical period prior to my main panel (Sep 2019 to Feb 2020)[30]. Those who had previously ever rejected a data request during this historical timeframe were categorized as the high-risk perception group, constituting 18% of the sample. This measurement captures the initial consumer risk perception on data-consent interactions on the digital platform prior to the beginning of the analysis.

The event-study plots, comparing high and low risk perception groups, are presented in Figure B11. The estimated values for the high risk perception group are twice as large as those for the low- risk perception group. However, the baseline rejection rate pre-shock is 21% for the high risk perception group and 10% for the low risk perception group. The drop from the baseline for the high risk perception group is 17.2% and for the low risk perception group, it is 15.8%. This similarity between the high and low-risk perception groups indicates that the test does not support the hypothesis that the drop in rejection is driven by adjustments in consumers' risk perceptions.

Furthermore, the evidence presented above that there is a decaying effect over time further implies that the mechanism does not involve a permanent change in beliefs about the riskiness of providing data.

Overall, my analysis suggests the magnitude of adjustment does not differ significantly between users with high and low rejection rates in the period before the analysis. Therefore, the changes in risk perception are unlikely to be the primary factor underlying the measured state

[30]The first several months are the period when the firm starts recording consent decision data. Data in this period was not used for the main analysis because of missing data and low quality concerns. It is sufficient to assess the rate of rejection for those users who appear in the data.

dependence.

7 Conclusion

This paper provides empirical evidence of state dependence in consumers' digital privacy choices. To address the challenge where researchers rarely observe individual data-consent decisions from field data, I use data from Alipay, one of the largest digital platforms in the world, in which consumers consistently engage with apps from third-party providers. To identify state dependence, I leverage a quasi-experiment in which consumers are highly likely to accept data requests due to the need of immediate charging. The staggered difference-in-differences estimates suggest a 15% reduction in the rejection rate of subsequent data-consent choices, and this effect is significant for at least 3 choices after the treatment. The shock leads to a temporary change in data consent decisions. The rejection rate eventually goes back to baseline, suggesting that people will ultimately hold on to their well-calibrated privacy preferences.

I examine the roles of consumer memory and preference strength in driving the identified state dependence effect. First, I provide evidence that the magnitude of state dependence is greater when the treatment choice is temporally closer to the current choice. Second, I find that users in usage contexts with apparently weak preferences exhibit significant state dependence following the shock, while those in contexts with apparently stronger preferences do not show this effect. Furthermore, the role of preference strength is supported by the analysis of user traits, revealing that users who examine privacy terms do not demonstrate state dependence after the shock.

Alternative explanations could account for the declining rejection rates among treated individuals. One theory suggests that data permanence leads consumers to abandon data protection efforts for data types previously exposed to third-party apps. However, In addition to the phenomenon that the main effect decays over time, I find that state dependence in sequential privacy decisions exists even for data types not requested in the treatment. These results indicate that the treatment does not cause individuals to completely abandon data protection efforts. Another explanation assumes that consumers adjust their risk perception when experiencing no negative consequences after data disclosure. However, I find that both high- and low-risk perception groups

exhibit similar decreases in rejection probability, making it unlikely that the effect is driven by changes in consumers' risk perceptions. Overall, the decaying effect over time suggests that the state dependence is not a function of the abandoning data protection or changing risk perceptions.

The findings of this paper have important implications. Recognizing state dependence in consumers' privacy decisions suggests an important aspect of platform privacy strategy. Data consent decisions on one app may change the ability of other apps to collect data from the same user. App-specific data decisions have an externality similar to the reputation externality highlighted in Nosko and Tadelis (2015).

Any discussion of specific strategies to address this externality is necessarily speculative. For example, if the externality is negative, the platform could incentivize apps in consumers' strong preferences to mitigate the effect. Conversely, if the externality is positive, the platform could promote apps in consumers' weak preferences. This leverages the result on strong effects for apps that have apparently weaker preferences. While I did not test these specific data strategy designs, future theoretical and empirical research in this area is welcome. More generally, the state dependence documented in this paper suggests an incentive for platforms to encourage consumer-friendly data request designs.

To the extent that the state dependence across apps would also apply to data consent decisions within an app, the results may also affect scenarios where a single app engages the same consumer with multiple sequential data requests. For example, an app may benefit from asking for data in sequence rather than asking for data all in once.

Furthermore, the temporary nature of the estimated state dependence suggests that underlying privacy preferences may be relatively stable, at least for the Alipay users studied here. This contrasts with assertions by various policymakers and legal scholars that consumers do not a have well-established understanding of privacy, and therefore privacy preferences fluctuate rapidly (Acquisti (2023)).

This paper acknowledges several limitations. First, the main data lacks records of the initial choices for all consumers. Therefore the analysis does not build from true initial conditions. That is, the platform did not start collecting consumer decision data when it introduced the current

data-consent rules where consumers exchange their data for app services. Second, for the main analysis, the identified state dependence is observed only in the three subsequent choices following the treatment. Examining the long-term effect is therefore limited to a relatively low-powered analysis of a subsample of the main data. Furthermore, the implications for platform strategy with respect to externalities are inferred, rather than directly tested.

Overall, understanding consumer privacy choices is important because firms trade off the competing goals of data access and privacy protection. This paper demonstrated that past privacy choices affect current privacy choices. Such state dependence implies an externality in the data requests across apps. Ignoring such this externality may affect the ability of apps and platforms to collect data.

References

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4):218–256.

Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al. (2023). Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Acquisti, A. (2023). The economics of privacy at a crossroads. *NBER Book Chapters: The Economics of Privacy*.

Acquisti, A., Brandimarte, L., and Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221):509–514.

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

Acquisti, A., Taylor, C., and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2):442–92.

Adjerid, I., Acquisti, A., and Loewenstein, G. (2019). Choice architecture, framing, and cascaded privacy choices. *Management Science*, 65(5):2267–2290.

Adjerid, I., Acquisti, A., Telang, R., Padman, R., and Adler-Milstein, J. (2016). The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Science*, 62(4):1042–1063.

Alba, J. W., Marmorstein, H., and Chattopadhyay, A. (1992). Transitions in preference over time: The effects of memory on message persuasiveness. *Journal of Marketing Research*, 29(4):406–416.

Andrews, M., Luo, X., Fang, Z., and Ghose, A. (2016). Mobile ad effectiveness: Hyper-contextual targeting with crowdedness. *Marketing Science*, 35(2):218–233.

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

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

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

Biehal, G. and Chakravarti, D. (1983). Information accessibility as a moderator of consumer choice. *Journal of Consumer Research*, 10(1):1–14.

Bielova, N., Litvine, L., Nguyen, A., Chammat, M., Toubiana, V., and Harry, E. (2024). The effect of design patterns on (present and future) cookie consent decisions. In *USENIX Security Symposium. USENIX Association. Accepted for publication*.

Bonatti, A. (2023). The platform dimension of digital privacy. *NBER Chapters*.

Borusyak, K., Jaravel, X., and Spiess, J. (2021). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*.

Brand, J., Israeli, A., and Ngwe, D. (2023). Using gpt for market research. *Available at SSRN 4395751*.

Brandão, A., Mendes, R., and Vilela, J. P. (2022). Prediction of mobile app privacy preferences with user profiles via federated learning. In *Proceedings of the Twelfth ACM Conference on Data and Application Security and Privacy*, pages 89–100.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. (2020). Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.

Burtch, G., Ghose, A., and Wattal, S. (2015). The hidden cost of accommodating crowdfunder privacy preferences: A randomized field experiment. *Management Science*, 61(5):949–962.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of econometrics*, 225(2):200–230.

Chen, L., Huang, Y., Ouyang, S., and Xiong, W. (2021). The data privacy paradox and digital demand. Technical report, National Bureau of Economic Research.

Choi, J. P., Jeon, D.-S., and Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173:113–124.

De Chaisemartin, C. and d'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American economic review*, 110(9):2964–2996.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Dubé, J.-P., Hitsch, G. J., and Rossi, P. E. (2010). State dependence and alternative explanations for consumer inertia. *The RAND Journal of Economics*, 41(3):417–445.

Dubé, J.-P., Hitsch, G. J., Rossi, P. E., and Vitorino, M. A. (2008). Category pricing with state-dependent utility. *Marketing Science*, 27(3):417–429.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*.

Ershov, D. (2024). Variety-based congestion in online markets: evidence from mobile apps. *American Economic Journal: Microeconomics*, 16(2):180–203.

Farboodi, M. and Veldkamp, L. (2022). A model of the data economy. Technical report, Working Paper, Columbia University.

Feldman, J. M. and Lynch, J. G. (1988). Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *Journal of applied Psychology*, 73(3):421.

Figueroa, C., Musalem, A., and Noton, C. (2020). Earthquakes and brand loyalty: Beyond the short-term effects of stockouts. *Available at SSRN*.

Gao, T., Yao, X., and Chen, D. (2021). Simcse: Simple contrastive learning of sentence embeddings. *arXiv preprint arXiv:2104.08821*.

Ghose, A., Li, B., and Liu, S. (2019). Mobile targeting using customer trajectory patterns. *Management Science*, 65(11):5027–5049.

Goldfarb, A. (2006a). The medium-term effects of unavailability. *Quantitative Marketing and Economics*, 4:143–171.

Goldfarb, A. (2006b). State dependence at internet portals. *Journal of Economics & Management Strategy*, 15(2):317–352.

Goldfarb, A. and Que, V. F. (2023). The economics of digital privacy. *Annual Review of Economics*, 15.

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

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of econometrics*, 225(2):254–277.

Heckman, J. J. (1981). Heterogeneity and state dependence. In *Studies in labor markets*, pages 91–140. University of Chicago Press.

Jaffe, S., Coles, P., Levitt, S., and Popov, I. (2017). Quality externalities on platforms: The case of airbnb. Technical report, Working paper.

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

Jones, C. I. and Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–58.

Keane, M. P. (1997). Modeling heterogeneity and state dependence in consumer choice behavior. *Journal of Business & Economic Statistics*, 15(3):310–327.

Klein, A. (2020). China's digital payments revolution. *Brookings Institution, Washington*.

Kong, X., Dubé, J.-P., and Daljord, Ø. (2022). Nonparametric estimation of habitual brand loyalty. *Available at SSRN 4070747*.

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

Levine, J. and Seiler, S. (2022). Identifying state dependence in brand choice: Evidence from hurricanes. *Marketing Science*.

Lin, T. (2022). Valuing intrinsic and instrumental preferences for privacy. *Marketing Science*.

Lynch, J. G., Marmorstein, H., and Weigold, M. F. (1988). Choices from sets including remembered brands: Use of recalled attributes and prior overall evaluations. *Journal of Consumer Research*, 15(2):169–184.

Lynch, J. G. and Srull, T. K. (1982). Memory and attentional factors in consumer choice: Concepts and research methods. *Journal of consumer research*, 9(1):18–37.

Nissenbaum, H. (2004). Privacy as contextual integrity. *Wash. L. Rev.*, 79:119.

Norberg, P. A., Horne, D. R., and Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126.

Nosko, C. and Tadelis, S. (2015). The limits of reputation in platform markets: An empirical analysis and field experiment. Technical report, National Bureau of Economic Research.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. (2022). Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.

Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). Chapter 5: Constructive processes in decision making. *The adaptive decision maker*.

Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.

Roy, R., Chintagunta, P. K., and Haldar, S. (1996). A framework for investigating habits,"the hand of the past," and heterogeneity in dynamic brand choice. *Marketing science*, 15(3):280–299.

Seetharaman, P. (2004). Modeling multiple sources of state dependence in random utility models: A distributed lag approach. *Marketing Science*, 23(2):263–271.

Seetharaman, P., Ainslie, A., and Chintagunta, P. K. (1999). Investigating household state dependence effects across categories. *Journal of Marketing Research*, 36(4):488–500.

Serramia, M., Seymour, W., Criado, N., and Luck, M. (2023). Predicting privacy preferences for smart devices as norms. *arXiv preprint arXiv:2302.10650*.

Simonov, A., Dubé, J.-P., Hitsch, G., and Rossi, P. (2020). State-dependent demand estimation with initial conditions correction. *Journal of Marketing Research*, 57(5):789–809.

Solove, D. J. (2021). The myth of the privacy paradox. *Geo. Wash. L. Rev.*, 89:1.

Sudhir, K. and Yang, N. (2014). Exploiting the choice-consumption mismatch: A new approach to disentangle state dependence and heterogeneity.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of econometrics*, 225(2):175–199.

Sun, R. and Trefler, D. (2023). The impact of ai and cross-border data regulation on international trade in digital services: A large language model. Technical report, National Bureau of Economic Research.

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

Tucker, C. (2023). The economics of privacy: An agenda. *NBER Book Chapters: The Economics of Privacy*.

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

Wang, W. and Li, B. (2024). Learning personalized privacy preference from public data. *Information Systems Research*.

Webb, R., Fong, J., Mazar, A., Levine, J., Wellsjo, A. S., Natan, O., Zhao, C., Lally, P., de Wit, S., Odoherty, J., et al. (2024). Integrating neuro-psychological habit research into consumer choice models.

A Service Rubric

Service Taxonomy

As an analytical expert in service categorization, your task is to classify the main services provided by the given company descriptions into one of three categories (S0, S1, S2).

Use the following context and information for categorization:

Assume you are a consumer with an average income level and have been living in your city for more than 5 years. You have a strong preference for services (S0) that, conditional on your downloading, you are in an urgent situation to use them, such as public transportation, bike sharing, or car sharing; vending machines, utility payment channels, or electric car charging stations. However, for other services, here denoted as services with weak preference (S1), probably recommended by your friend, you may find them value-added, but not with strong preferences at the moment downloading– meaning you will find it valuable or time-saving to use, but it won't disturb your life if you finally decide do not use them.

To distinguish between services of strong preference (S0), services of weak preference (S1), and lottery-nature services (S2), use the following thought process:

- Step 1: Determine whether physically presenting at the facility of the focal app is the only way that makes sense to justify downloading the app.
- Step 2: If the consumer is already standing in front of the facility, determine how much outside option they have if they decide not to use the app.

In Step 1, if the user must be physically present to justify the download, move to Step 2. If not, and this is not a lottery-nature app, the service can be under weak preference (S1). In Step 2, if users have relative space for outside options, the app is labeled S1. If not, it's labeled S0.

In Step 1, there are two scenarios about you downloading an app: First, you might be standing in front of a service infrastructure. Without downloading and using the app, you are not able to use the service. Therefore, the use of the app directly leads to the viability of your next move. Second, you might be sitting at home downloading and exploring the app. The app does not directly lead to the viability of your next move.

For example, you are downloading an e-commerce app. You could be either downloading it at some random place outdoors or sitting at home exploring the brand. According to Step 1, this e-commerce app could be directly categorized as S1.

For instance, if you are downloading a car charger app. It only makes sense to download it when you need to use a specific brand of car charger at a random place. It does not make sense that you download it

at home since you won't predict your car is running out of battery. Therefore, the car charger app belongs to the first scenario only. According to Step 1, being physically present at the facility is the only way that justifies the download, we could move to Step 2. In Step 2, not downloading the car charger app leaves little space for outside options. Therefore, the car charger app should be categorized as S0. For another example, let us look at an app for a mall directory or the electronic record system (ERS) of a hospital, the apps can make information search or data storage more efficient for the user. According to Step 1, it fits the description and we move to Step 2. The user can still achieve the original purpose of the focal activities, such as shopping in the mall or seeing a doctor, without using the app. Then the mall directory app and the ERS app are S1.

The key distinction is that apps under strong preference are integrated into the user's regular habits and serve essential needs, while apps under weak preference offer supplementary benefits or services that enhance the overall digital experience when required, but are not critical for the routine life operation.

Here is the taxonomy:

- ## S0 – Strong Preferences
  - Services under strong preferences are often on-site ncessities services. Following the two-step process above, if a firm's description of its main service can be seen as 1) physically presenting at the facility of the focal app is the only way to justify the download; and 2) without using the app, the user cannot continue his next step which is the direct intention of downloading the app. By satisfying the 2 steps, then it should be classified as S0.
  - This includes services such as: Electric Car Charger, Coffee and Dessert stores (QR codes for orders), Restaurants (QR codes for orders), Lockers or Vending Machines, Food or Floral Delivery, Parking, Public Transportation, Sharing bikes, Sharing cars, Car Rent, Utility Payment Channels, COVID QR code for mandatory checking.
- ## S1 – Weak Preferences
  - Services under weak preferences are often on-demand value-added services. These services provide additional benefits and supplementary value but are not critical for routine life. Following the two-step process above, if a firm's description of its main service can be seen as 1) physically presenting at the facility of the focal app is not the only way to justify the download – meaning the user can be sitting at home exploring; OR 1) physically presenting at the facility of the focal app is the only way to justify the download but 2) the user has some space for outside options even without the app at the facility they attempt to download it.

  - This includes services such as: Brand online stores, Courier Service Tracking, E-Commerce, Financial Services, Gas Stations, Labor Service Search, Information Search, Productivity, Health Management, Hospital ERS, Luxury Rent, Mall Directory, Charity, Resell, Social, Tourist Spots Guide, SME Brands and Stores.

- ## S2 – Lottery-nature Services
  - Lottery-nature Services are the ones for advertising, marketing campaigns, or coupon purposes. People download the apps for lottery or play games to receive coupons or cashback. This category is less valuable in utility compared to S0 and S1, but is more hedonic.
  - This includes services such as: Entertainment, Marketing, Coupon, Games, and Lottery.

Please take into consideration the examples attached below:

Annotation Examples:

Example 1:

Description: zFrontier Equipment Frontline is a geek equipment purchase and interactive community platform. By directly cooperating with designers and manufacturers, the platform allows members to enjoy high-end products at exclusive preferential prices. Users can also discuss and share geek equipment through the platform. Questions and answers are available online.

Label (S0/S1/S2): S1

Explanation: This app is an online community connecting designers and manufacturers. Implementing the 2-step process, physically presenting at the facility of the focal app is not the only way to justify the download – meaning the user can be sitting at home exploring. Therefore, this could be categorized as S1. Moreover, according to the taxonomy of S1, this app is mostly related to "Labor Service Search".

Example 2:

Description: The name 85°C comes from "coffee tastes best at 85°C." Founded in 2003, "85°C" is a fashionable catering company in Taiwan, China, focusing on coffee, cakes, and baking.

Label (S0/S1/S2): S0

Explanation: The app is a coffee chain shop. Implementing the 2-step process, without a need to order in a coffee store, a consumer would never download a specific app for this coffee chain. Therefore, we move from Step 1 to Step 2. Since ordering coffee is the only purpose to download this app, then a consumer does not have much space for outside options. Therefore, this could be categorized as S0. Moreover, according to the taxonomy of S0, this app is mostly related to "Coffee and Dessert stores (QR codes for orders)".

Example 3:

Description: Shengdianhua APP is an online shopping APP released by Guangzhou Nuomici Information Technology Co., Ltd. Shengdianhua cooperates with mainstream e-commerce platforms such as Taobao, Tmall, and Pinduoduo. Users can not only receive coupons for shopping, but also enjoy shopping cashback, life privileges, and other rights after becoming members.

Label (S0/S1/S2): S2

Explanation: The app is a coupon/cashback app. Implementing the 2-step process, physically presented at the facility of the focal app is not the only way to justify the download, and the app is a lottery-dominated app. Therefore, this could be categorized as S2. Moreover, according to the taxonomy of S2, this app is mostly related to "Marketing" or "Coupon".

Example 4:

Description: Atour, a company with a comfortable smile. Atour is a new lifestyle hotel brand that advocates a humanistic, warm and interesting "on the road" fourth space lifestyle. The brand has four major products: Atour Hotel, ATHouse, THE DRAMA, and long-term rental apartments. Atour, providing guests with a high-quality life experience covering all areas such as accommodation, e-commerce, tourism, and finance. ATHouse is based in the urban core, leading the dynamic social space of urban travel. Each ATHouse is located in the core blocks of first-tier cities and provincial capitals. The public area of the hotel is named United Park (abbreviated as UP public area). It is not limited to the broad concept of socializing and urban living room. It will become a gathering place for young people to share the urban trend culture of travel life. This is a place that gathers trend culture, urban activities and active young groups. It provides coffee, wine, food and accommodation, shares and selects sales content, and provides young people with novel travel experiences and quality life on the road. THE DRAMA is the first drama-themed hotel in China, dedicated to building a Shanghai Renaissance distribution center. THE DRAMA is located in the Jing'an Temple business district, located in SMG Shangyan Valley, adjacent to the world-famous immersive drama "Sleep No More". As an epoch-making hotel product, THE DRAMA brings together art peripherals such as drama, vinyl, and movies, and redefines hotel accommodation space through art exhibitions, markets, performances, etc.

Label (S0/S1/S2): S1

Explanation: This app is an all-round service of the Atour facilities. Implementing the 2-step process, without physically being at the Atour facility, a consumer would not download the app. Therefore, we move

from Step 1 to Step 2. Different from Example 2, a consumer in this case does have some space for outside options. They could enjoy shopping and other activities even without this app. Therefore, this could be categorized as S1. Moreover, according to the taxonomy of S1, this app is mostly related to "Mall".

B Figures

Figure B1: Example screenshots of data-consent decisions

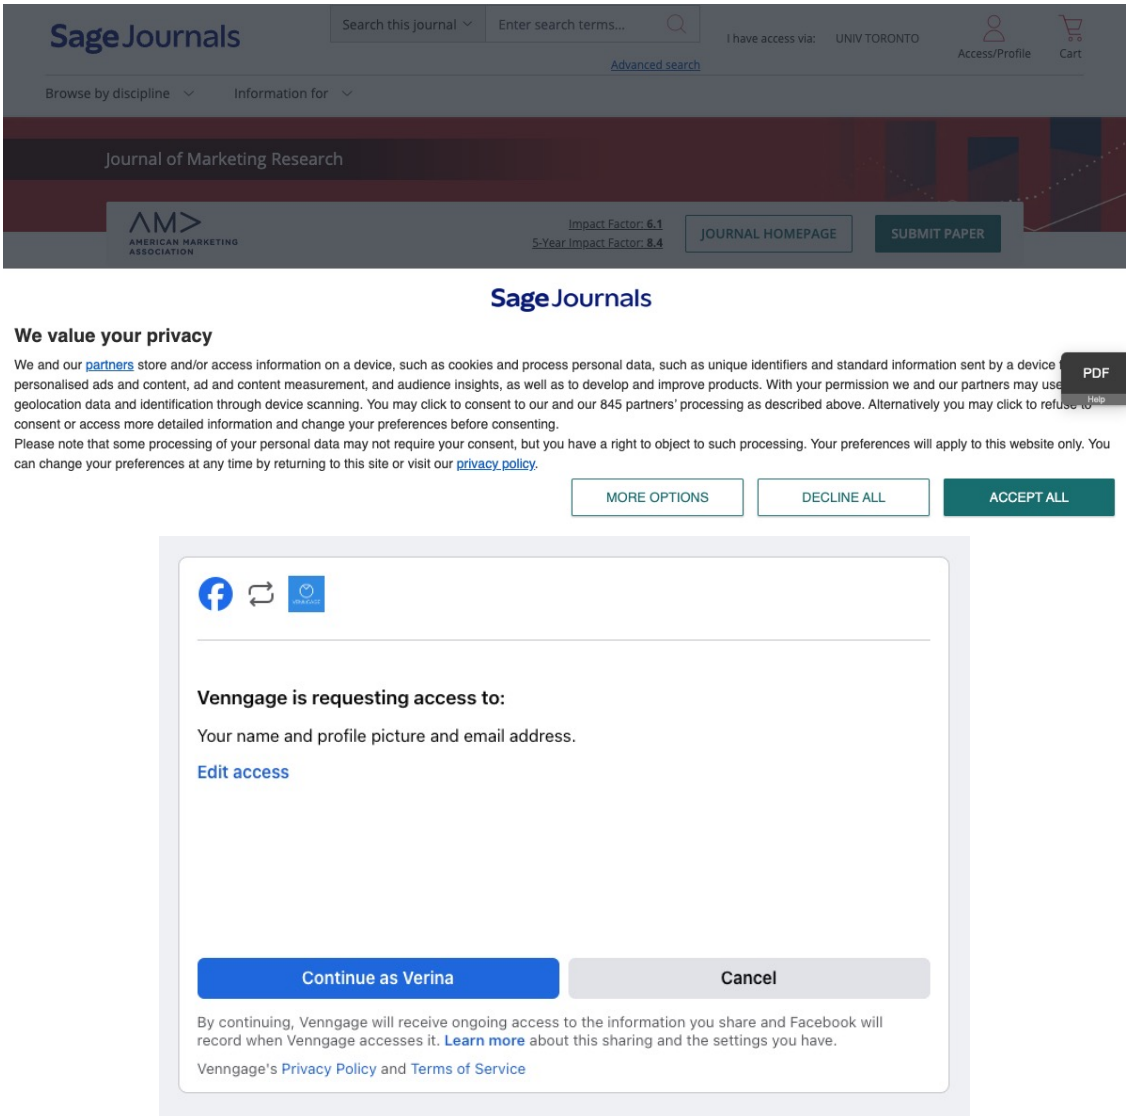
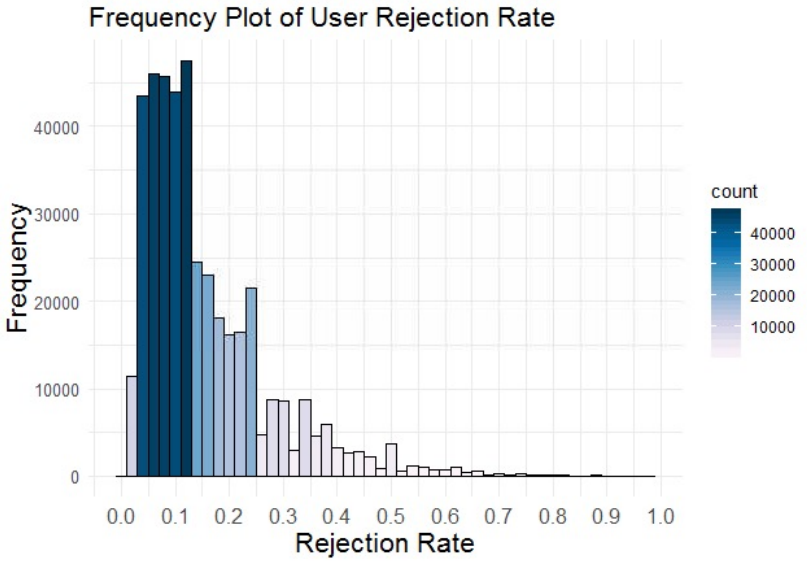


Figure B2: Distribution of User Behaviors – User Heterogeneity

Panel A: distribution of user rejection rates



Panel B: distribution of number of apps a user downloads

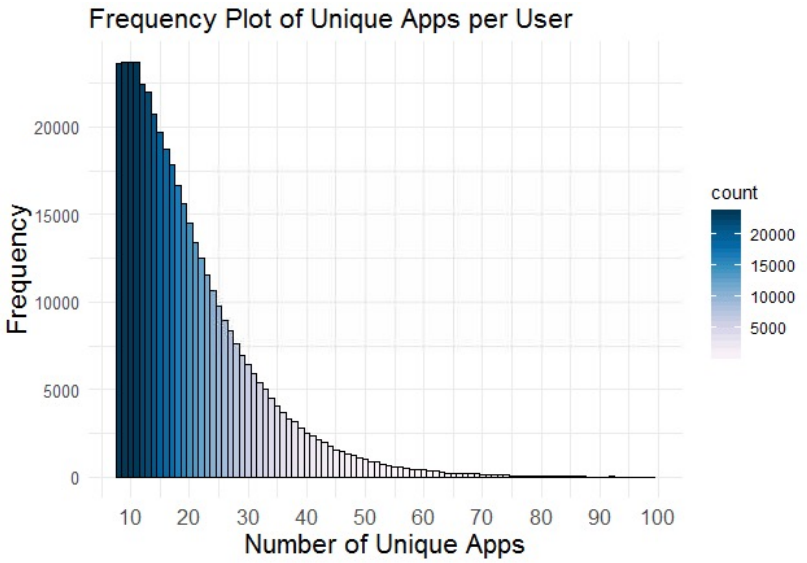
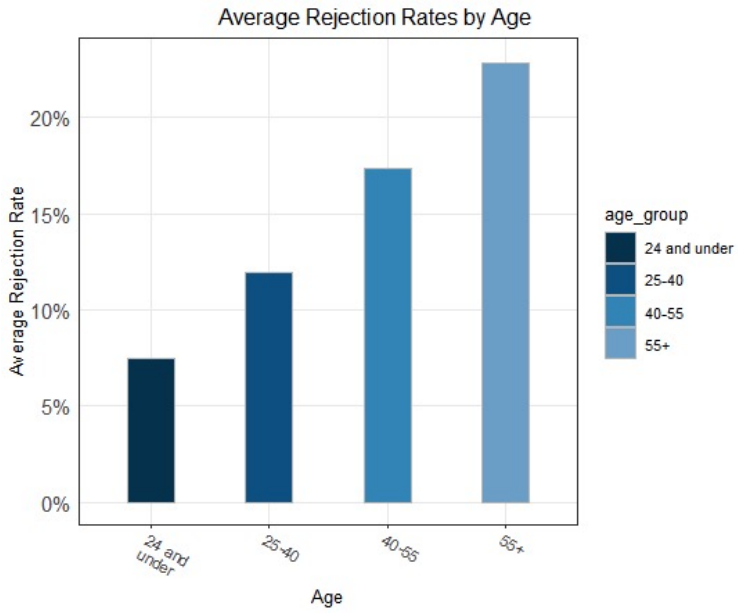


Figure B3: Heterogeneous Privacy Choices by User Demographics

Panel A: Consumers are Heterogeneous in Privacy Choices across Age Groups



Panel B: Consumers are Heterogeneous in Privacy Choices across Education Groups

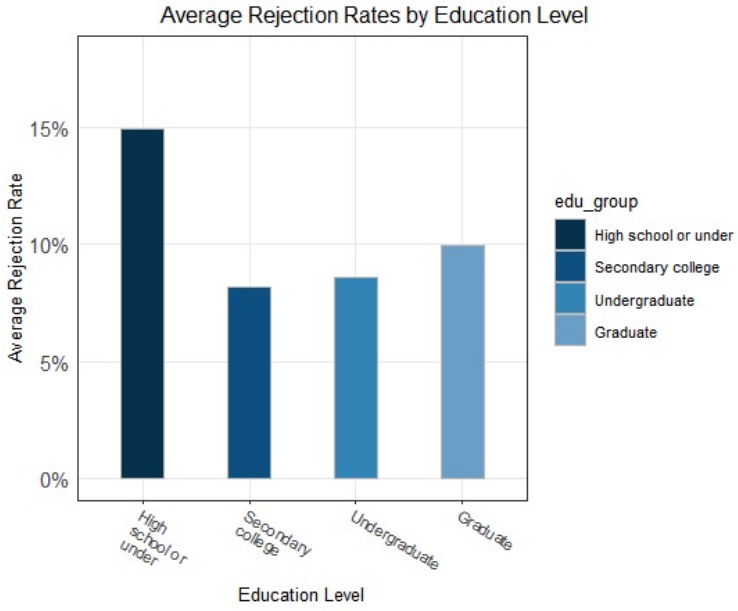
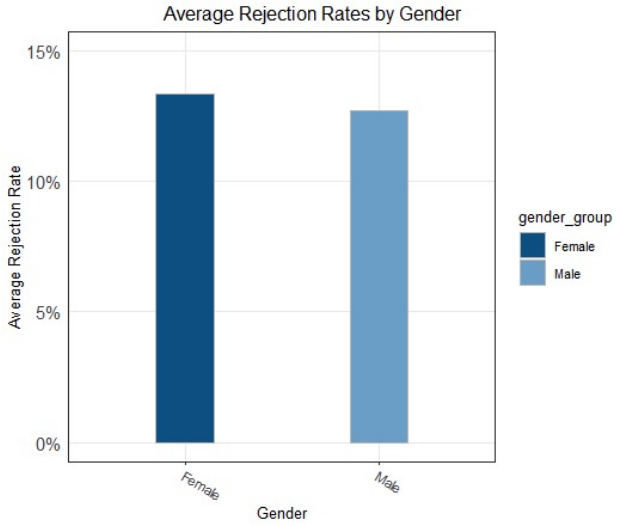
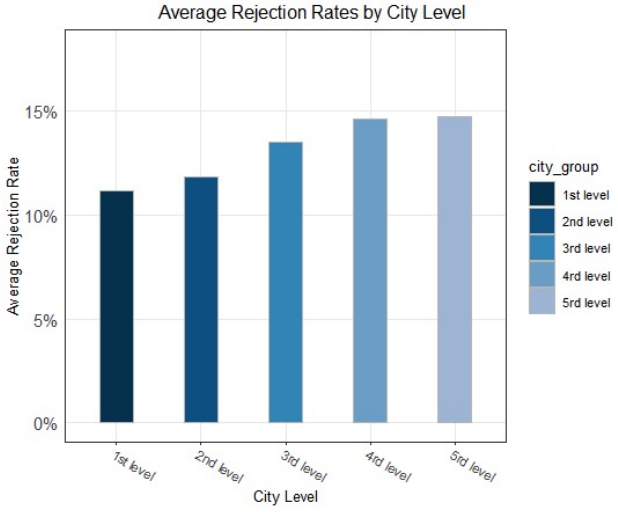


Figure B4: User Demographics that Lack Heterogeneity in Privacy Choices





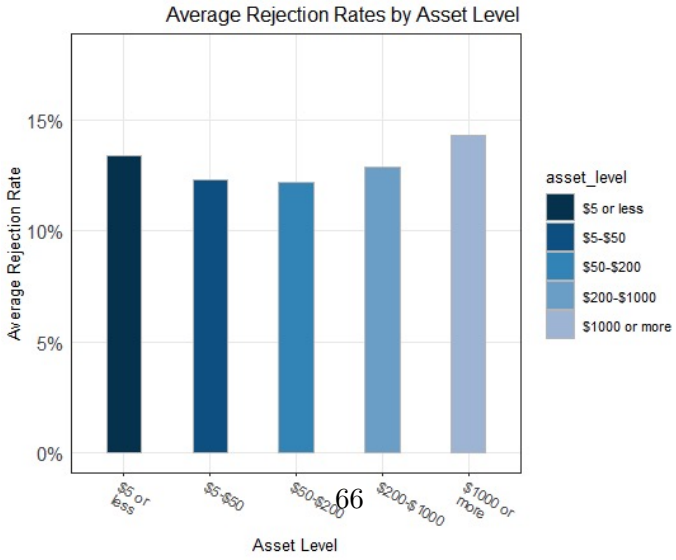
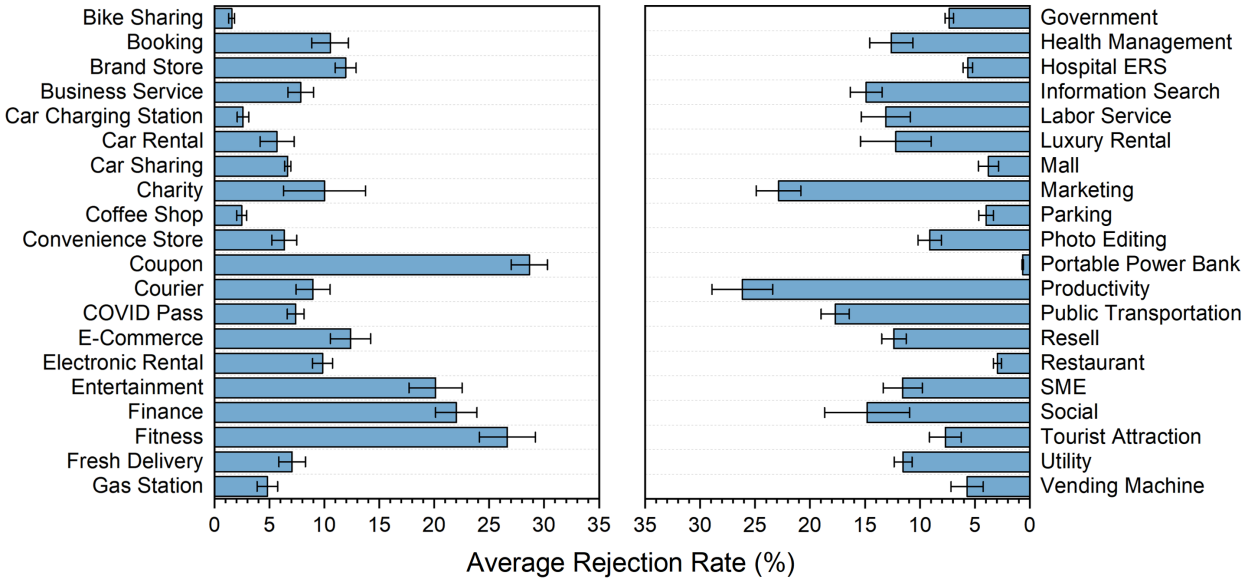


Figure B5: Consumer's Privacy Choices are Context-Dependent

Panel A: average rejection rate by app functionalities



Panel B: average rejection rate by datatype combinations

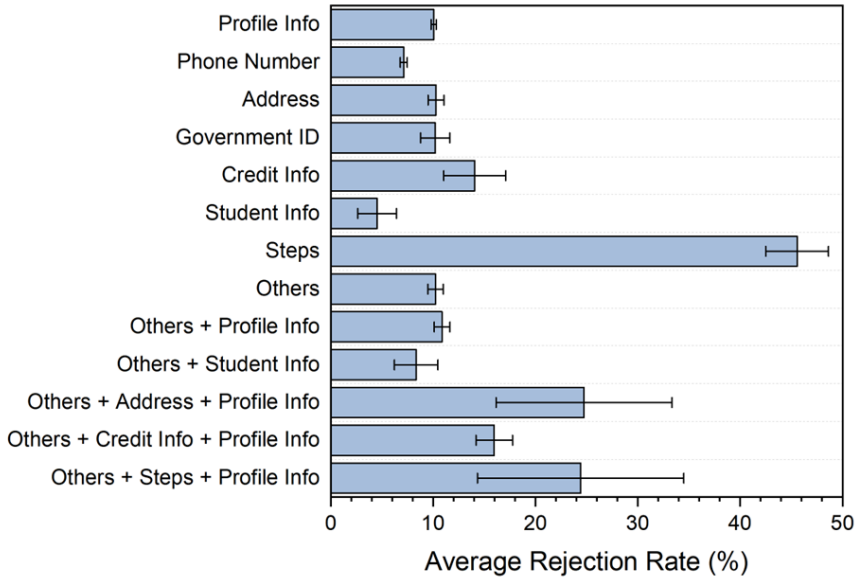
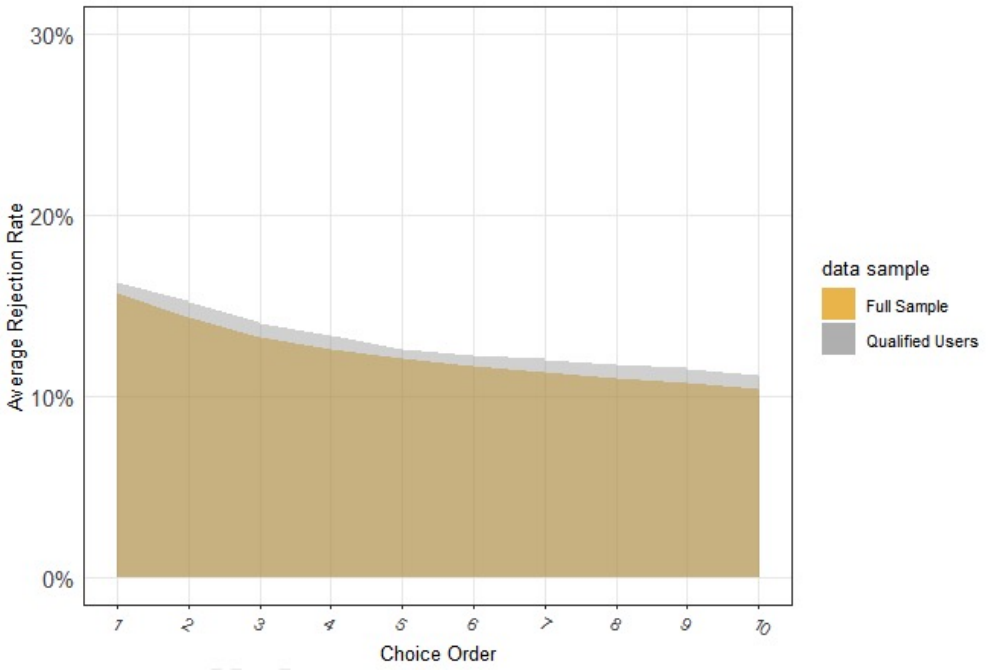


Figure B6: Average Rejection Rate Across Time and Choices

Panel A: average rejection rate across choice order



Panel B: average rejection rate across calendar time (month)

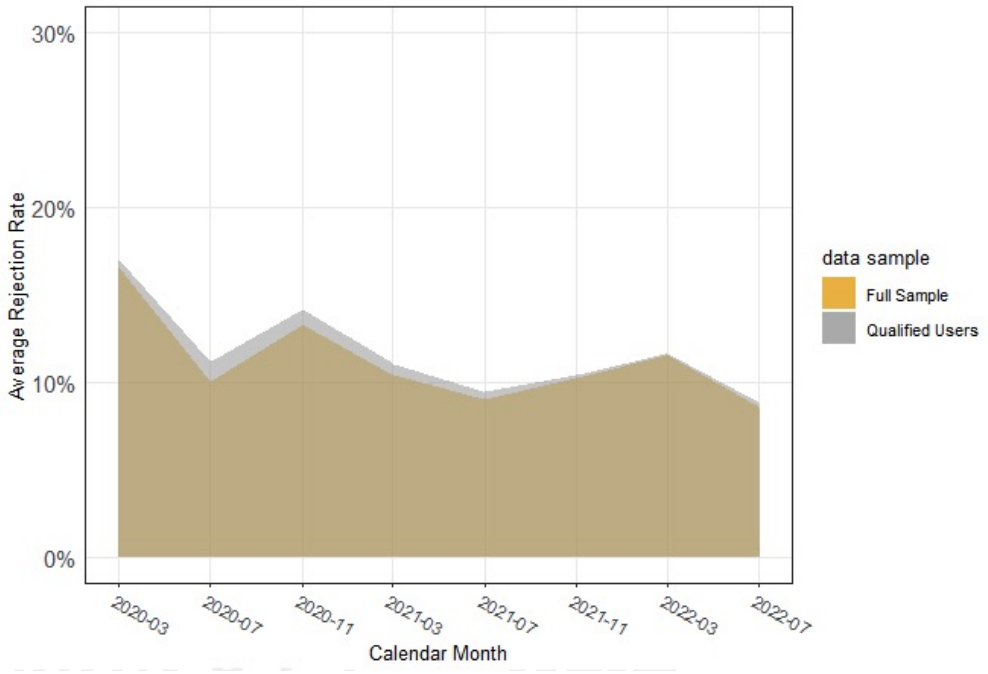
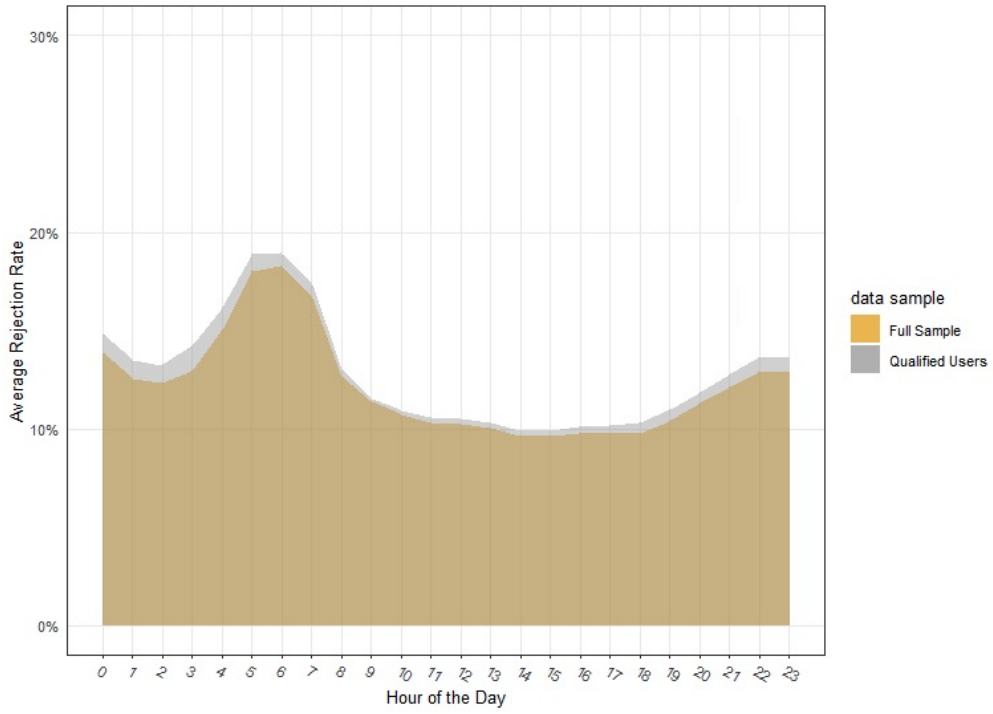


Figure B7: Average Rejection Rate – Time effects

Panel A: average rejection rate by hour in a day



Panel B: average rejection rate by day in a week

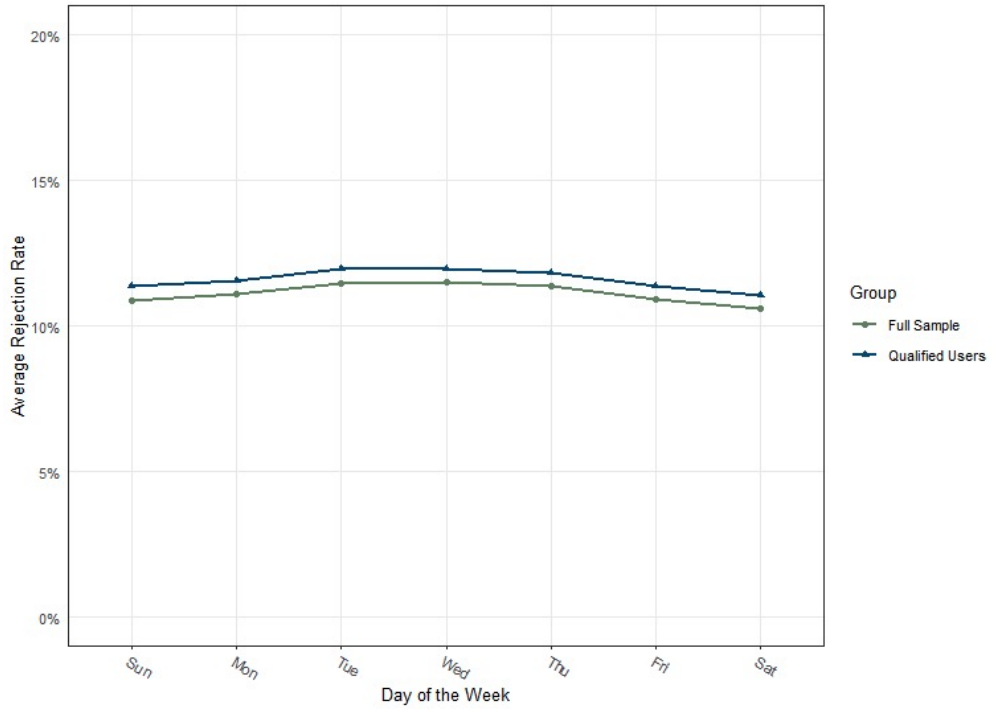


Figure B8: Distribution of time when the first choice happened

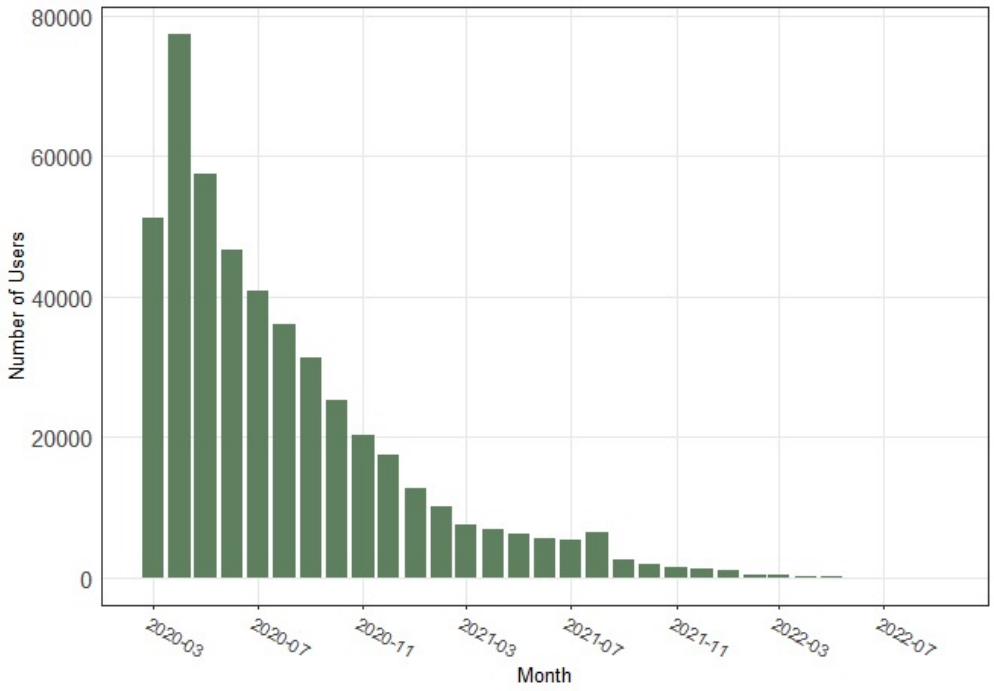


Figure B9: Context Categorization Under LLMs



Weak Preference

Examples of exploratory scenarios: e-commerce, job-search, financial services, productivity, ...



Strong Preference

Examples of necessity scenarios: vending machines, public transportation, bike-sharing, car charger, ...

Figure B10: Decaying estimate over time

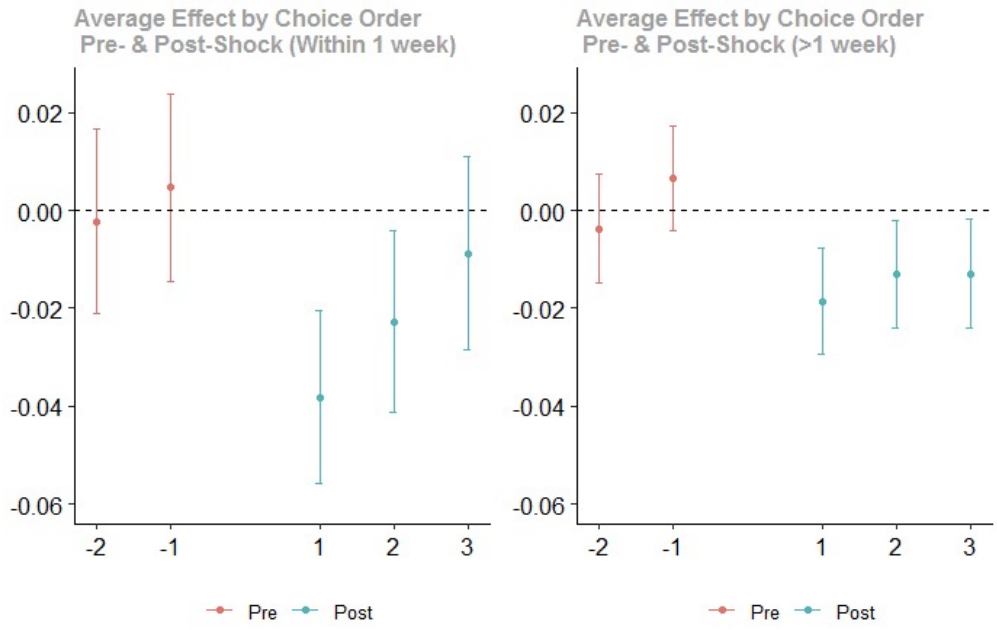
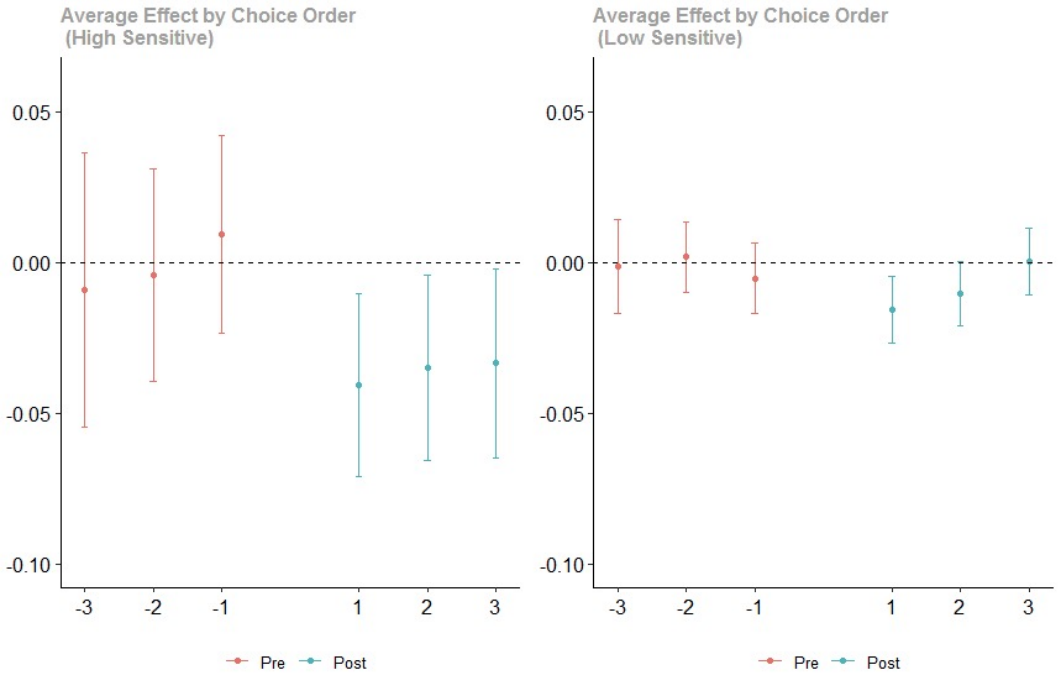


Figure B11: Similar magnitudes of effect across groups



C Tables

Table C1: Percent Apps Downloaded by Category Pre/Post Portable Power Bank App

	Treated		Control	
category	3 choices pre-shock	3 choices post-shock	choice 2,3,4	choice 6,7,8
Bike Sharing	4.1%	3.2%	2.6%	2.6%
Booking	4.1%	3.4%	3.5%	4.4%
Brand Store	1.2%	0.9%	1.6%	1.2%
Business Service	0.8%	0.7%	0.7%	0.7%
Car Charging Station	0.4%	0.4%	0.3%	0.3%
Car Rental	0.9%	0.7%	0.8%	1.0%
Car Sharing	5.3%	4.1%	4.5%	4.7%
Charity	1.2%	0.8%	1.8%	1.4%
Coffee Shop	3.3%	3.3%	3.4%	4.0%
Convenience Store	1.2%	1.5%	0.9%	1.4%
Coupon	6.0%	4.6%	7.6%	7.6%
Courier	3.8%	3.7%	3.7%	4.5%
COVID Pass	6.4%	6.9%	6.2%	4.6%
E-Commerce	2.0%	1.8%	2.2%	2.5%
Electronics Rental	1.9%	1.5%	2.1%	2.3%
Entertainment	1.2%	0.7%	1.6%	1.1%
Finance	1.1%	1.1%	1.2%	1.4%
Fitness	0.8%	0.8%	1.6%	1.7%
Fresh Delivery	5.5%	5.0%	4.3%	5.9%
Gas Station	1.1%	1.3%	0.8%	1.2%
Government	16.2%	15.0%	14.4%	14.2%
Health Management	1.2%	2.1%	0.9%	1.9%
Hospital ERS	0.4%	0.6%	0.3%	0.5%
Information Search	1.1%	1.0%	1.3%	1.4%
Labor Service	0.3%	0.3%	0.4%	0.4%
Luxury Rental	0.03%	0.02%	0.03%	0.03%
Mall	0.4%	0.2%	0.6%	0.4%
Marketing	0.3%	0.2%	0.6%	0.5%
Parking	0.9%	1.0%	0.7%	0.9%
Photo Editing	0.9%	0.9%	1.0%	1.2%
Productivity	3.4%	1.6%	4.6%	2.5%
Public Transportation	0.9%	0.6%	0.9%	0.8%
Resell	0.9%	0.8%	1.5%	1.5%
Restaurant	4.8%	5.0%	4.4%	5.6%
SME	0.3%	0.3%	0.2%	0.3%
Social	1.1%	0.1%	2.3%	0.2%
Tourist Attraction	0.1%	0.2%	0.1%	0.2%
Utility	8.2%	6.6%	8.3%	8.2%
Vending Machine	1.1%	1.0%	1.1%	1.1%

Table C2: Summary Statistics

Panel A: Summary Statistics for Users

	N	Mean	Std	Min	Median	Max
All Users						
age_i (in year)	358,000	34.58	10.00	16.00	34.00	79.00
$male_i$	358,000	56.02%	n.a	n.a	n.a	n.a
$higher_edu_i$	358,000	29.28%	n.a	n.a	n.a	n.a
$blue_collar_i$	358,000	26.72%	n.a	n.a	n.a	n.a
city $degree_i$ (1-5)	358,000	2.66	n.a	1.00	3.00	5.00
avg daily account $balance_i$ (in USD)	358,000	1251.67	4946.96	0.01	77.87	478,894.31
Treated Users						
age_i (in year)	119,864	31.93	8.14	16.00	31.00	78.00
$male_i$	119,864	58.29%	n.a	n.a	n.a	n.a
$higher_edu_i$	119,864	38.65%	n.a	n.a	n.a	n.a
$blue_collar_i$	119,864	21.19%	n.a	n.a	n.a	n.a
city $degree_i$ (1-5)	119,864	2.48	n.a	1.00	2.00	5.00
avg daily account $balance_i$ (in USD)	119,864	1229.26	4834.52	0.01	94.98	315,092.91

Panel B: Summary Statistics for Apps

	N	Mean	Std	Min	Median	Max
App Stats for Full Sample						
total user $count_j$	7321	6115.84	38017.51	1.00	202.00	333256.00
life $length_j$ (in months)	7321	14.82	9.90	3.00	13.00	31.00
monthly average $user_j$	7321	566.79	5548.56	1.00	15.65	228051.00
monthly average $duration_j$ (in hours)	7321	61.62	1812.56	0.006	1.00	108423.65
monthly average $launch_j$	7321	5162.72	63284.96	0.33	59.78	2853103

Table C3: TWFE & Callaway and Sant'Anna Estimates (never-treated & not-yet-treated as control group)

	(1) TWFE	(2) TWFE	(3) CS	(4) CS
Treated × 1st Choice Post	-0.018***	-0.020***	-0.021***	-0.024***
	(0.002)	(0.002)	(0.003)	(0.003)
Treated × 2nd Choice Post	-0.013***	-0.016***	-0.016***	-0.020***
	(0.002)	(0.002)	(0.003)	(0.003)
Treated × 3rd Choice Post	-0.010***	-0.012***	-0.010***	-0.017***
	(0.002)	(0.002)	(0.003)	(0.003)
Observations	2,943,504	2,943,504	2,943,504	2,943,504
Individual FE	Yes	Yes	Yes	Yes
Choice-order FE	Yes	Yes	Yes	Yes
Calendar month Control	No	Yes	No	Yes
Clustered error	user	user	user	user
Note:	baseline rejection rate before treatment is 18.5%			

Table C4: TWFE & Callaway and Sant'Anna Estimates with Full Sample

Panel A: not-yet-treated as control group

	(1) TWFE	(2) TWFE	(3) CS	(4) CS
Treated × 1st Choice Post	-0.013***	-0.012***	-0.015***	-0.024***
	(0.002)	(0.002)	(0.002)	(0.002)
Treated × 2nd Choice Post	-0.009***	-0.007***	-0.011***	-0.024***
	(0.002)	(0.002)	(0.002)	(0.002)
Treated × 3rd Choice Post	-0.006***	-0.004**	-0.009***	-0.024***
	(0.002)	(0.001)	(0.002)	(0.002)
Observations	787,177	787,177	787,177	787,177
Individual FE	Yes	Yes	Yes	Yes
Choice-order FE	Yes	Yes	Yes	Yes
Calendar month Control	No	Yes	No	Yes
Clustered error	user	user	user	user
Note:	baseline rejection rate before treatment is 10.8%			

Panel B: never-treated & not-yet-treated as control group

	(1) TWFE	(2) TWFE	(3) CS	(4) CS
Treated × 1st Choice Post	-0.012***	-0.014***	-0.015***	-0.018***
	(0.001)	(0.001)	(0.002)	(0.002)
Treated × 2nd Choice Post	-0.008***	-0.010***	-0.011***	-0.015***
	(0.001)	(0.001)	(0.002)	(0.002)
Treated × 3rd Choice Post	-0.006***	-0.007***	-0.008***	-0.013***
	(0.001)	(0.001)	(0.002)	(0.002)
Observations	3,888,286	3,888,286	3,888,286	3,888,286
Individual FE	Yes	Yes	Yes	Yes
Choice-order FE	Yes	Yes	Yes	Yes
Calendar month Control	No	Yes	No	Yes
Clustered error	user	user	user	user
Note:	baseline rejection rate before treatment is 13.6%			

Table C5: TWFE & Callaway and Sant'Anna Estimates For Newly Registered Users (sub-sample)

	(1) TWFE	(2) TWFE	(3) CS	(4) CS
Treated × 1st Choice Post	-0.024**	-0.024**	-0.023	-0.029*
	(0.009)	(0.009)	(0.012)	(0.012)
Treated × 2nd Choice Post	-0.020*	-0.019*	-0.016	-0.030*
	(0.009)	(0.009)	(0.012)	(0.013)
Treated × 3rd Choice Post	-0.022*	-0.021*	-0.021	-0.034**
	(0.009)	(0.009)	(0.012)	(0.013)
Observations	83,421	83,421	83,421	83,421
Individual FE	Yes	Yes	Yes	Yes
Choice-order FE	Yes	Yes	Yes	Yes
Calendar month Control	No	Yes	No	Yes
Clustered error	user	user	user	user
Note:	baseline rejection rate before treatment is 10%			

Table C6: Heterogeneity of Main Effects across Age & Educational groups

	Age				Education			
	<25	25-40	40-55	>55	high-school	secondary	university	graduate
	(1)	(2)	(3)	(4)	(1)	(2)	(3)	(4)
baseline rejection	7.9%	11%	15.1%	21.1%	12.6%	8.6%	8.9%	9.5%
Treated × 1st Choice Post	-0.013***	-0.013***	-0.010*	-0.015	-0.010***	-0.016***	-0.015***	-0.011
	(0.002)	(0.002)	(0.005)	(0.020)	(0.002)	(0.004)	(0.002)	(0.007)
Treated × 2nd Choice Post	-0.010***	-0.009***	-0.006	0.004	-0.008***	-0.001	-0.012***	-0.002
	(0.002)	(0.002)	(0.005)	(0.021)	(0.002)	(0.005)	(0.002)	(0.007)
Treated × 3rd Choice Post	-0.007**	-0.006**	-0.001	-0.001	-0.005**	-0.004	-0.006**	-0.003
	(0.002)	(0.002)	(0.005)	(0.020)	(0.002)	(0.004)	(0.002)	(0.007)
Observations	286,768	556,127	150,953	9,729	525,898	88,502	343,692	45,485
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Choice-order FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustered error	user	user	user	user	user	user	user	user
Note:					$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$			

Table C7: Decaying effect with time since use of Portable Power Bank App

	(1) > 3 days	(2) > 1 week	(3) > 3 weeks	(4) > 5 weeks
Treated × 1st Choice Post	-0.022***	-0.019***	-0.015***	-0.013
	(0.004)	(0.004)	(0.005)	(0.006)
Treated × 2nd Choice Post	-0.016***	-0.013***	-0.014***	-0.013
	(0.004)	(0.004)	(0.005)	(0.006)
Treated × 3rd Choice Post	-0.014***	-0.013***	-0.012	-0.011
	(0.004)	(0.004)	(0.005)	(0.006)
number of treated users	15,923	14,426	10,791	7,777
number of control users	27,945	24,318	16,152	10,283
baseline avg rej_rate for treated	10.75%	10.78%	10.98%	10.99%
baseline avg rej_rate for control	11.87%	11.77%	11.56%	11.64%
Note:				$^{*}p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$

D App Category

Table D8: Summary Statistics for App Categories

				Top1 App		Top3 Apps	
category	avg rejection	app #	user #	avg rejection	market share	avg rejection	market share
Bike Sharing	1.58%	131	235,614	5.62%	79.30%	5.34%	86.17%
Booking	10.53%	104	101,657	3.82%	14.57%	7.25%	39.42%
Brand Store	11.94%	307	95,373	8.55%	13.52%	8.96%	34.46%
Business Service	7.86%	162	47,482	2.45%	17.70%	2.36%	41.61%
Car Charging Station	2.60%	102	24,946	3.25%	26.66%	2.21%	53.45%
Car Rental	5.71%	48	28,679	1.66%	20.33%	2.09%	49.19%
Car Sharing	6.67%	36	210,080	2.21%	32.58%	7.60%	71.80%
Charity	10.01%	28	55,380	10.77%	54.53%	8.25%	77.10%
Coffee Shop	2.48%	177	143,259	13.49%	15.45%	7.96%	42.98%
Convenience Store	6.36%	188	101,676	16.67%	24.62%	10.21%	52.02%
Coupon	28.66%	307	241,660	22.60%	15.11%	27.12%	39.63%
Courier	8.98%	66	193,826	14.51%	34.84%	10.36%	60.62%
COVID Pass	7.39%	95	397,510	10.01%	30.08%	10.39%	43.95%
E-Commerce	12.29%	124	134,409	22.19%	20.37%	13.41%	47.39%
Electronics Rental	9.84%	290	95,552	18.23%	35.34%	15.01%	51.90%
Entertainment	20.12%	109	65,348	1.90%	22.17%	17.57%	56.90%
Finance	22.01%	196	69,716	4.93%	9.76%	14.78%	25.43%
Fitness	26.66%	82	72,603	47.87%	29.16%	33.35%	55.18%
Fresh Delivery	7.07%	142	235,211	14.09%	32.92%	15.79%	67.35%
Gas Station	4.81%	176	88,416	6.17%	27.67%	3.19%	66.97%
Government	7.33%	434	445,000	12.26%	15.04%	8.68%	30.25%
Health Management	12.59%	93	151,798	32.31%	26.59%	21.51%	53.70%
Hospital ERS	5.63%	700	78,366	3.97%	16.85%	7.03%	30.31%
Information Search	14.89%	194	417,541	13.42%	59.43%	13.35%	82.67%
Labor Service	13.10%	112	21,033	5.39%	33.42%	3.91%	59.82%
Luxury Rental	12.20%	14	1,210	25.67%	30.19%	27.93%	68.24%
Mall	3.77%	87	26,389	26.83%	39.67%	20.92%	71.64%
Marketing	22.87%	154	30,908	14.72%	20.19%	31.12%	35.31%
Parking	3.98%	112	50,768	1.54%	23.95%	3.82%	59.02%
Photo Editing	9.12%	189	84,505	8.05%	37.05%	9.53%	55.94%
Portable Power Bank	0.65%	241	128,971	1.97%	18.84%	1.70%	47.97%
Productivity	26.16%	137	154,996	11.11%	47.26%	23.40%	77.40%
Public Transportation	17.71%	158	50,818	2.27%	22.14%	12.99%	38.24%
Resell	12.37%	115	74,916	22.36%	23.28%	20.93%	46.22%
Restaurant	2.96%	896	192,431	2.53%	17.32%	7.95%	32.85%
SME	11.57%	159	12,984	0.95%	20.73%	3.90%	42.88%
Social	14.80%	22	35,141	54.68%	93.75%	53.84%	97.41%
Tourist Attraction	7.69%	154	16,945	3.70%	26.49%	2.59%	44.80%
Utility	11.53%	448	339,195	17.80%	13.72%	36.70%	31.33%
Vending Machine	5.70%	126	83,371	3.89%	47.16%	5.43%	68.89%