

NBER WORKING PAPER SERIES

CONSUMER SURVEILLANCE AND FINANCIAL FRAUD

Bo Bian
Michaela Pagel
Huan Tang
Devesh Raval

Working Paper 31692
<http://www.nber.org/papers/w31692>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2023, Revised May 2026

We thank Simona Abis, Pat Akey, Hunt Allcott, Tania Babina, Nathan Blascak, Mark Eichorn, Emma Fletcher, Lewis Kirvan, Samuel Kruger, Kai Li, Markus Mobius, Jordan Nickerson, David Rothschild, Dave Schmidt, James Thomas, Andreanne Tremblay-Simard, Guillaume Vuillemeys, Jonathan Wallen, Brett Wendling, Paul Witt, Constantine Yannelis, Xiaoyun Yu, and Anthony Lee Zhang for valuable comments as well as seminar and conference participants at the Chicago Conference on Empirical Finance, Chinese University of Hong Kong, CFPB Research Conference, Future of Financial Information Conference, Georgia Tech - Atlanta Fed Household Finance Conference, Goethe University, Household Finance Brownbag Series, IMF, Microsoft Research, NBER SI (Household Finance, Digital Economics and Artificial Intelligence), NFA, Peking University, Renmin University, UBC Winter Finance Conference, University College London, UW-Madison Junior Finance Conference, and University of Innsbruck. Yile Chen, Rohan Joseph, Yiming Ma, Niels Wagner, and Andy Xin provided outstanding research assistance. The views expressed in this article are those of the authors and do not necessarily represent those of the Federal Trade Commission, any of its Commissioners, or the National Bureau of Economic Research. Bo Bian acknowledges financial support from the Social Sciences and Humanities Research Council of Canada (Grant Number: 430-2023-00553). We have nothing else to disclose. First Version: November 1, 2022.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Bo Bian, Michaela Pagel, Huan Tang, and Devesh Raval. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Consumer Surveillance and Financial Fraud
Bo Bian, Michaela Pagel, Huan Tang, and Devesh Raval
NBER Working Paper No. 31692
September 2023, Revised May 2026
JEL No. G5

ABSTRACT

In today's digital economy, firms continuously collect, store, share, and sell personal data, exposing customers to risks of financial fraud. Leveraging Apple's App Tracking Transparency policy as a natural experiment, we show that restricting data tracking and sharing significantly reduces consumer fraud complaints, particularly those involving personal information misuse. Effects are stronger in areas dominated by firms with risky data practices and coincide with a decline in dark web discussions and higher prices for sensitive data. By tracing effects along the fraud supply chain, our findings suggest that data regulation can reduce a consumer harm that firms may not fully internalize.

Bo Bian
University of British Columbia
bo.bian@sauder.ubc.ca

Huan Tang
University of Pennsylvania
huan.ht.tang@gmail.com

Michaela Pagel
Washington University in St. Louis
Olin Business School
and NBER
mpagel@wustl.edu

Devesh Raval
Federal Trade Commission
devesh.raval@gmail.com

1 Introduction

We live in an era of *surveillance capitalism*, in which firms collect, process, and combine personal data from multiple sources, turning it into a core input of production (Zuboff, 2019). A defining feature of this input is that it is non-rival and shareable: once collected, data can be reused across applications without being depleted (Jones and Tonetti, 2020). This shareability also enables data to be combined across contexts, as fragments generated in different settings can be shared and merged into unified customer profiles. While these properties underpin value creation in targeted advertising (Rafieian and Yoganarasimhan, 2021; Wernerfelt et al., 2025; Bian et al., 2024), open banking (Babina et al., 2025), and AI model training (Beraja et al., 2023), they also bring a “dark side.” Unlike capital or labor, data can be copied and repurposed with little friction, allowing information collected for legitimate purposes to be exploited elsewhere and potentially facilitating identity theft, impersonation, and other forms of financial harm.

The scale of these harms can be large. The Federal Trade Commission (FTC) estimates that over 10% of US adults fall victim to financial fraud annually and reported losses reached nearly \$12.5 billion in 2024.¹ In the meantime, policymakers have expanded privacy protections, from the General Data Protection Regulation (GDPR) in Europe to the California Consumer Privacy Act (CCPA) in California, as well as ongoing federal initiatives in the United States.² Do these privacy regulations have the potential to reduce financial fraud? To the extent that modern fraud depends on real-time, cross-linked data generated through opaque data collection and sharing, often without consumers’ awareness or consent, privacy regulation may limit fraud by disrupting profiling and targeting. However, fraud has long relied on traditional identifiers, such as Social Security numbers and home addresses, and perpetrators may adapt through broader targeting or social engineering. Accordingly, whether privacy protections reduce fraud is ambiguous *ex ante*.

We study this question in the context of Apple Inc.’s App Tracking Transparency (ATT) policy, which sharply limited the tracking and sharing of personal information across mobile applications (apps), websites, companies, and platforms. Introduced on April 26, 2021, ATT requires iOS apps to obtain explicit user permission before accessing mobile identifiers that enable cross-app tracking and third-party data sharing, reversing the default from *opt-out* (tracking allowed unless disabled) to *opt-in* (tracking allowed only with explicit consent). As of March 2022, only 17% of iOS users had opted in (Kraft et al., 2023), substantially

¹See FTC press release: “New FTC Data Show a Big Jump in Reported Losses to Fraud to \$12.5 Billion in 2024.” Other estimates suggest even higher prevalence, with surveys indicating up to half of Americans have experienced fraud (Anderson, 2019; Huff et al., 2010; DeLiema et al., 2017).

²For the FTC, see press release “FTC Explores Rules Cracking Down on Commercial Surveillance and Lax Data Security Practices.” For the CFPB, see “Required Rulemaking on Personal Financial Data Rights.” For the US Congress, see “Senate’s summary of the American Privacy Rights Act of 2024.”

limiting the volume and scope of personal data collected and shared across contexts.

As a major shock to the mobile data ecosystem, ATT provides an ideal setting to quantify a key consumer-side benefit of privacy regulation for several reasons. First, ATT creates sharp variation in exposure by affecting iOS but not Android users, enabling a more credible identification strategy. By contrast, broad regulations such as the GDPR and CCPA generate effects that are difficult to isolate, as they operate through firm-level responses while firms simultaneously serve many markets and user segments. Second, the policy directly targets real-time, cross-app data collection and linkage, isolating a channel that extends beyond the static, fragmented data used in traditional fraud schemes and thereby shedding light on the mechanics of modern fraud. Third, granular app-level data allow us to construct novel measures of data collection and sharing intensity by key data intermediaries and exploit this heterogeneity to trace mechanisms along the data supply chain.

We first estimate the effect of ATT on consumer fraud, focusing on consumer victimization as the most direct and policy-relevant measure of harm. Using granular SafeGraph foot traffic data, we construct zip-code-level measures of iOS usage and examine how fraud outcomes evolve across areas with different exposure to ATT, controlling for time-varying trends across counties. We use three different datasets of consumer fraud complaints in the US, starting with the CFPB’s public complaint database, which is especially relevant for studying financial fraud because it centers on disputes involving personal information and financial institutions.³

Our difference-in-differences (DiD) analysis indicates that limiting the tracking and sharing of personal information substantially reduces consumer complaints. We find that a 10 percentage point ($\sim$ one standard deviation) increase in the share of iOS users in a zip code leads to a 6.1% reduction in the number of CFPB complaints post-ATT. Given that 83% of iOS users opted out of tracking after ATT, a 10% increase in consumers opting out of tracking translates to a 7.3% reduction in CFPB complaints. Our identification assumes that, absent ATT, complaint trends would have evolved similarly across areas with high and low iOS penetration. While areas with more iOS users may differ in baseline complaint rates or consumer behavior, our dynamic DiD estimates show no evidence of differential pre-trends.

The dynamic estimates show a decline that emerges within a quarter after ATT and strengthens over time. This rapid response provides novel evidence on the short horizon over which modern, data-driven fraud operates. Much like legitimate digital advertising, fraud can depend on real-time data for rapid targeting and delivery, which helps explain why

³We refer to consumers’ voluntary submissions of information about fraud and other scams as “complaints.” Although the FTC and other institutions have long described this information as “complaints,” the FTC now describes this information as “reports” to emphasize the problems that consumers may observe as opposed to whether consumers were directly affected or lost money as a result.

disruptions to these data flows could generate effects quickly. Detection and reporting appear quick as well. Descriptive evidence points to a short horizon for both. On the reporting side, CFPB consumer complaints surged within a month of the 2017 Equifax breach, suggesting that data leakages translate into reported fraud almost immediately. On the detection side, the 2021 NCVS reports that 44% of victims discovered the fraud within a single day and 77% of identity-theft cases within one month.

We corroborate our main findings using proprietary data from the FTC, which provides uncensored coverage of all CFPB complaints, a complete census of fraud and identity theft cases from the Consumer Sentinel Network and Identity Theft databases, and information on reported dollar losses from fraud. The uncensored data confirms that censoring in the public CFPB database does not drive our results. The identity theft records show significant post-ATT declines consistent with reduced misuse of personal data. Broader complaints to the Consumer Sentinel Network exhibit smaller effects, as they include complaints on non-finance products and services and cases not triggered by data breaches. Importantly, the dollar loss data from Consumer Sentinel complaints suggests that ATT reduced consumer fraud losses by roughly \$274 million annually.

To sharpen identification, we exploit an alternative strategy based on the 2017 repeal of the Federal Communications Commission’s (FCC) broadband privacy rules. This reverse shock lifted anticipated constraints on Internet Service Providers’ (ISPs) collection and sharing of user data. The results are consistent with our main findings: following the repeal, financial fraud rose more in areas served by ISPs with more pervasive data practices. Beyond operating in the opposite direction, this design exploits a source of variation distinct from the iOS-Android comparison, namely, differences in local exposure to ISPs with more invasive data practices, thereby providing further support for identification.

The second part of our study explores the mechanisms behind how greater privacy protections reduce fraud. Providing direct evidence on how ATT leads to reduced data-driven fraud is exceptionally challenging. The fraud economy operates through opaque, unregulated channels where neither the flow of illicit data nor the actors involved are directly observable. Personal and financial information is collected covertly, trafficked through hidden online markets, and monetized in ways that rarely leave transparent traces. Even when consumers or firms are victimized, the mechanism from compromised data to fraudulent transactions is typically invisible in public or administrative records. We introduce several new databases to piece together complementary evidence that trace the effects of ATT from downstream victims, through data-collecting firms targeted by ATT, to the illicit markets supplying stolen or leaked data.

First, we dig deeper into the most visible point in the chain of events: downstream

consumers and their complaints. Using keyword searches and machine-learning classifications of complaint narratives, we separate fraud types that depend on compromised personal or financial data—such as account takeovers and identity theft—from those that do not. We find that ATT-induced declines are concentrated in data-dependent categories while unrelated categories covering, e.g., mortgages or student loans show little change.

We then move to the intermediaries that connect data collection to consumer fraud. Financial institutions and ISPs routinely collect, store, and share large volumes of individual-level behavioral and financial data. In doing so, they can serve both as gateways through which data reaches malicious actors and as targets whose data repositories attract cyber-criminals. To study this link, we combine local bank deposit and ISP market shares with measures of their data collection and sharing intensity, including the number of unique data items collected and the presence of third-party data-sharing platforms in their mobile apps. A triple-difference design shows that the decline in fraud complaints is significantly larger in markets where high-risk banks and ISPs are more prevalent.

Consistent with this mechanism, we also show that financial institutions with consumer-facing mobile apps experience significant declines in cyber incidents after ATT, especially those caused by data breaches and those resulting in violations of the Fair Debt Collection Practices Act or the Fair Credit Reporting Act, which are the types of incidents most likely to trigger consumer complaints. This pattern is matched by corresponding declines in relevant consumer complaints about these institutions, with no comparable changes in unrelated complaints. Together, these results connect downstream fraud reductions to intermediaries targeted by malicious actors and provide a window into the otherwise opaque activities of fraudsters, which we turn to next.

Finally, to examine the fraud supply chain at its most opaque point, we turn to illicit markets where compromised personal and financial data is traded, focusing on both dark web forum discussions (the conversation layer) and marketplace listings and prices (the transaction layer). Using novel data on 37 dark web forums and tens of millions of posts, we find sharp post-ATT declines in iOS/Apple-related discussions and smaller but still significant declines in broader topics on personal or financial data. These effects are estimated relative to a control group of posts about illicit trade categories (e.g., drugs, human trafficking, weapons) featuring actors and infrastructure that are unlikely to overlap with the data-driven fraud economy. Complementary analysis of dark web marketplace listings shows that prices for mobile-sourced and financial data rise after ATT, consistent with reduced supply. While these upstream markets are inherently opaque and our evidence is not definitive, the results align with ATT constraining the availability of compromised data, thereby limiting the supply chain that feeds downstream fraud.

A potential concern is that ATT may have displaced, rather than reduced, fraud by shifting targeting from iOS to Android users. Several pieces of evidence argue against such reallocation as the primary driver of our results. First, we find substantial declines in platform-agnostic fraud categories (e.g., phishing, identity theft, and credential-based attacks) that rely on data obtainable from both iOS and Android. If fraudsters had fully reallocated toward Android, these categories should be largely unaffected. Second, we observe increases in prices in illicit data markets following ATT, consistent with a contraction in the overall supply of usable data rather than substitution across platforms. Third, we find little evidence of spillover effects from iOS to Android users within zip codes or within target firms. Conceptually, such reallocation is limited because iOS users tend to be higher-value targets with greater financial stakes and are not easily substituted, while the return to investing in Android-based data collection was uncertain given anticipated privacy restrictions from Google.

We discuss several other empirical issues. First, we only observe reported complaints rather than true fraud incidences. Reporting propensities may vary across demographic groups, and ATT could itself affect awareness and reporting behavior. To address this, we adjust for differences in complaint propensities using demographic-based weights and find similar effects. Moreover, if ATT increased awareness and reporting, this would bias against finding declines. To the extent that ATT reduces consumer digital engagement or their willingness to share data, this aligns with our mechanism that reduced data flows curb fraud. Second, one may argue that iOS users differ systematically from Android users, as they tend to be wealthier and more educated. Our identification relies on similar trends rather than identical levels, and we find no evidence of differential pre-trends; results are also robust to interacting ATT with a rich set of zip code-level demographic controls. Third, our findings could be confounded by concurrent shocks, most notably COVID-19 and related stimulus programs that may have affected fraud patterns. We directly control for exposure to Economic Impact Payments and conduct placebo tests around major disbursement dates, finding no evidence that these factors drive our results. Finally, we conduct a large set of additional robustness checks—including adjustments for measurement error, alternative aggregation levels, and different fraud definitions—and obtain consistent results.

Taken together, our analysis provides a rare account of how privacy regulation reshapes the entire fraud supply chain. ATT reduced downstream consumer fraud in categories dependent on compromised data; these reductions were strongest for firms and localities with greater pre-ATT exposure to risky data practices; and ATT disrupted the illicit upstream markets that supply and utilize this data. By linking visible consumer outcomes to the hidden supply chain of illegally obtained information, we provide the first systematic evidence

tracing how privacy initiatives reduce fraud. Our findings therefore quantify a key benefit of stronger privacy protections and suggest that data practices, especially among firms engaged in extensive consumer surveillance, generate downstream harms to consumers that may not be fully internalized, warranting closer regulatory attention.

Literature Review. We first contribute to the literature on the economic implications of data privacy regulations. Prior work has largely examined the GDPR, connecting it to changes in web traffic (Goldberg et al., 2024), app entry and exit (Janssen et al., 2021), VC financing (Jia et al., 2021), firms’ abilities to collect, monetize, store, and use consumer data (Aridor et al., 2023; Bessen et al., 2020; Demirer et al., 2024; Peukert et al., 2022), as well as the visibility and quality of firms’ data collection disclosures (Ramadorai et al., 2025).⁴ Other studies have looked at open banking (Babina et al., 2025) and the effect of the CCPA on mortgage lending (Doerr et al., 2023) and firm risk (Wu, 2023). A small but growing set of papers focuses on Apple’s privacy initiatives, such as the privacy label policy (Bian et al., 2021) and the ATT framework (Kesler, 2022; Cheyre et al., 2023; Bian et al., 2024; Abis et al., 2025). We further extend this literature by providing new causal evidence on the *consumer-side* benefits of privacy regulation. We show that a standardized and uniform consent framework that led most users to decline tracking can curb financial fraud, in contrast to prior work that has primarily documented *firm-side* costs.

Our empirical evidence speaks to one important microfoundation of data privacy concerns. While surveys show that identity theft is the most cited reason for privacy worries among US households (Armantier et al., 2021), field evidence is rare.⁵ We complement the existing survey-based evidence by showing that weak data collection and sharing safeguards can cause higher rates of identity theft and financial fraud. This shifts the discussion from stated concerns to measurable harm, and adds to both the theoretical and empirical work on the welfare consequences of excess data collection, including price discrimination (Taylor, 2004; Acquisti and Varian, 2005; Bergemann et al., 2015; Bonatti and Cisternas, 2020), behavioral manipulation (Liu et al., 2023; Acemoglu et al., 2025), and surveillance (Tirole, 2021; Beraja and Yuchtman, 2025).⁶

More broadly, our paper relates to a growing literature that views data as a shareable input into economic activity with important externalities. One set of externalities arises from spillovers within and across firms: Jones and Tonetti (2020) emphasize the non-rival nature of data, Beraja et al. (2023) show that data obtained in one context are shareable

⁴See Johnson (2024) for a review of the literature examining GDPR.

⁵Armantier et al. (2021) document that around 90% of respondents in each demographic group report identity theft as an important concern, followed by abuse of data, personal safety, and reputation.

⁶See also Acquisti et al. (2016); Bergemann and Bonatti (2019); Agrawal et al. (2022) for excellent literature surveys.

across applications, generating spillovers to commercial AI innovation, and [Bian et al. \(2024\)](#) characterize data as a networked asset whose cross-firm linkages drive comovement in valuation and investment. Another set arises among consumers: [Choi et al. \(2019\)](#), [Acemoglu et al. \(2022\)](#), and [Ichihashi \(2021\)](#) argue that data generated by one consumer can reveal information about others, leading to excessive data collection in equilibrium. Our findings imply a unique negative externality: when firms collect, share, and inadequately protect consumer data, fraudsters can acquire and repurpose them to harm consumers. Because firms do not fully bear these losses, they may collect data at socially excessive levels.

Finally, our study contributes to the large literature on fraud and forensic finance. Prior research has examined various types of fraud, including cases involving elderly victims (e.g., [Alves and Wilson, 2008](#); [DeLiema et al., 2012](#); [Lichtenberg et al., 2013](#); [James et al., 2014](#); [DeLiema, 2018](#); [Carlin et al., 2023](#)), misconduct by financial professionals and investment advisers (e.g., [Egan et al., 2019, 2025](#); [Griffin and Kruger, 2024](#)), and patterns of victimization across demographic and geographic groups using consumer complaint data (e.g., [Raval, 2020a,b](#); [Sweeting et al., 2020](#)). These studies emphasize how individual characteristics, social factors, and institutional oversight shape who becomes a victim and how fraud propagates. We extend this literature by highlighting the role of digital surveillance and cross-linked data flows as a key determinant of financial fraud, providing the first evidence on the emergence of modern, data-driven fraud that differs from traditional forms relying on static, siloed information. In particular, we show that the data practices of ISPs and financial institutions—firms positioned at the core of consumer tracking or custodians of sensitive personal information—play a pivotal role in today’s fraud supply chain. This yields clear, actionable policy implications: regulating the data practices of these data intermediaries may produce spillover effects that reduce fraud.

2 Institutional Background

In this section, we first provide background information on how collecting and sharing of personal data can lead to financial fraud. We then describe an industry-led privacy initiative implemented by Apple and its relevance for data-driven fraud.

2.1 Characteristics of Data-Driven Fraud and Enforcement Activities

Contemporary financial fraud is enabled by two relevant inputs: first, real-time personal and behavioral data, and second, the ability to link information across multiple sources; both enhancing the scale and rapid execution of fraud. These inputs have been targeted by recent regulatory enforcement and are directly disrupted by our policy shock.

First, fraud schemes today often depend on access to live, constantly refreshed data. Unlike static identifiers such as Social Security Numbers (SSNs), dynamic information, such

as current addresses, recent transactions, active credentials, and live location, enables fraudsters to act quickly and exploit vulnerabilities before detection.⁷ Stolen credit card numbers can be used immediately for unauthorized purchases; newly breached login credentials can be tested in credential-stuffing attacks before passwords are reset; and geolocation data can help time account takeovers when victims are away from home. Such real-time behavioral data have become even more critical over time because banks and FinTech apps increasingly rely on contextual verification—matching logins against device fingerprints, IP addresses, and usage histories. This allows fraudsters to mimic the legitimate device environment (same city, device model, or time of day) and bypass these automated checks.⁸ Because consumer behavior changes quickly and institutions update security controls, stale data rapidly loses value, making real-time information especially important for successful fraud.

Second, and more important for the purpose of our study, fraudsters increasingly exploit the interconnected nature of digital data. Persistent identifiers, such as Apple’s Identifier for Advertisers (IDFA), have historically allowed data brokers, advertisers, and malicious actors to merge browsing, app usage, and transaction data into unified behavioral profiles. These linkages transform otherwise harmless fragments—for example, combining an address from one breach with a device ID or app activity from another—into actionable intelligence about a person’s financial habits, institutions, and online routines. Leveraging the linked data, fraudsters can micro-target deceptive ads to consumers most likely to fall victim to financial scams, e.g., promoting predatory loans or fake investment opportunities to financially distressed individuals. [Figure 1](#) provides an illustrative schematic of this process, showing how financial, behavioral, location, and social data collected across mobile apps can be linked via persistent identifiers into unified user profiles that facilitate several forms of fraud.

These targeted scams are prevalent and have become a major source of consumer harm. Internal documents obtained by Reuters revealed that Meta projected that about 10% of its 2024 advertising revenue—roughly \$13 billion—came from fraudulent promotions, especially deceptive e-commerce and investment schemes.⁹ A UK regulator reported that Meta’s platforms were linked to 54% of all payment-related scam losses in 2023.¹⁰ In the US, consumers reported more than three billion dollars in losses to the FTC’s Consumer Sentinel Network from fraud initiated via social media, websites and apps, or online ads, representing about half of all reported fraud losses with an identified point of contact.¹¹

These two characteristics are central to recent enforcement actions by US regulators,

⁷See ACFE: [“Emerging Trends in Fraud Law,”](#) November 2023.

⁸See Palo Alto Networks: [What Is Credential Stuffing?](#)

⁹See Reuters: [“Meta is earning a Fortune on a deluge of Fraudulent Ads, Documents show,”](#) November, 2025.

¹⁰See PSR: [“How Fraudsters Exploit Major Platforms to Scam Consumers,”](#) December, 2024.

¹¹See FTC: [“FTC Data shows Consumers report losing 2.7 billion in Social Media Scams,”](#) October, 2023.

as illustrated by the following two examples. First, in the case against Ideal Financial Solutions Inc., the FTC charged the company with purchasing Social Security numbers and bank account information from payday-loan applicants and using the merged data to debit consumer accounts without consent. The FTC also settled criminal charges with the data brokers who had sold the information.¹² Second, the Department of Justice settled criminal charges against Epsilon Data Management LLC, a major data broker that sold consolidated lists of financially vulnerable consumers to fraud rings promoting sweepstakes, astrology, and psychic scams. These schemes depended on timely and detailed data to identify and target susceptible individuals.¹³

As another example, the FTC has recently turned its attention to mobile location data, which can be used to profile consumers and pursue time-sensitive, targeted fraud opportunities. The agency settled with data brokers X-Mode and Outlogic and is currently litigating against Kochava for selling precise geolocation data.¹⁴ When combined with mobile device IDs and home addresses, this data can reveal where people are in real time, how they move, and what behaviors they engage in, and so has the potential to enable highly personalized and well-timed fraud. Kochava’s widely promoted “Household Mapping” is one example of such data. More broadly, to curb fraud via targeted deception, the Securities and Exchange Commission has launched an investigation into Meta for running ads for financial scams, according to Reuters. The FTC has also initiated a major study examining how large social media platforms detect and prevent deceptive advertising.¹⁵

2.2 Apple’s App Tracking Transparency Policy

In this paper, we study an industry-led privacy initiative. We use the implementation of the ATT policy as an exogenous shock to the gathering, sharing, or selling of detailed data of iOS users, which reduced the availability of high-quality data for fraudsters.

With the release of iOS 14.5 on April 26, 2021, Apple introduced a new privacy feature that required all apps to ask for explicit user permission before obtaining users’ mobile identifiers. These mobile identifiers were the primary way that allowed apps and data brokers to combine user data from different sources. This feature, dubbed “App Tracking Transparency (ATT),” grants users both greater and easier control over their data. An example

¹²See FTC: “FTC Action Leads Court Orders Against Scheme Charged Millions of Dollars to Consumers’ Bank and Credit Card Accounts,” March, 2016 and “Data Broker Defendants Settle FTC Charges They Sold Sensitive Personal Information to Scammers,” February, 2016.

¹³See DOJ: “Marketing Company Agrees to Pay \$150 Million for Facilitating Elder Fraud Schemes,” January, 2021.

¹⁴See FTC: “FTC Order Prohibits Data Broker X-Mode Social and Outlogic from Selling Sensitive Location Data,” January, 2024 and “FTC Sues Kochava for Selling Data that Tracks People,” August, 2022.

¹⁵See FTC: “FTC Issues Orders to Social Media and Video Streaming Platforms Regarding Efforts to Address Surge in Advertising for Fraudulent Products and Scams,” March, 2023.

of the prompt notification is provided in Panel a of [Figure 2](#). By default, a user is opted out of tracking. That is, Apple would no longer provide apps and websites with the user’s mobile identifier. Importantly, the opt-in design and the uniform consent prompt apply to all firms that serve iOS users. In addition, companies are forbidden from displaying the consent prompt to users who have already declined the request.

Industry reports suggest that the vast majority of users did not opt in for tracking upon seeing the notification ([Kraft et al., 2023](#)).¹⁶ [Bian et al. \(2021\)](#) document a sharp and negative stock market reaction for firms owning an active iOS app around the implementation of ATT, corroborating its substantial impact on the data economy (see [Appendix Figure A.1](#)).

ATT breaks down a core mechanism through which personal data is collected and circulated in the mobile ecosystem. Before ATT, tracking was enabled by default, allowing data collected by an app to flow far beyond the developer to a dense network of third-party firms—including ad networks, data brokers, and analytics providers—many of which had no direct relationship with the user and operated with limited oversight.¹⁷ This structure created multiple downstream copies of sensitive data, increasing the risk of resale, leakage, or misuse. ATT interrupts this chain by requiring user permission for cross-app tracking, thereby reducing both the volume of personal data collected and its broad distribution across the digital supply chain.

This disruption also directly targets the two characteristics of personal data that make it valuable for financial fraud: its linkability and real-time nature. By blocking access to Apple’s device-level identifier, ATT undermines the ability to stitch together data across apps and time, breaking the coherence needed to build detailed user profiles. At the same time, ATT limits the availability of up-to-date behavioral, transactional, and location data by curbing passive background tracking, reducing the chances that fraudsters can act before detection. In doing so, the policy weakens both the speed and precision of fraud operations built on mobile-sourced data.

3 Data

3.1 Exposure to ATT: Share of iPhone and Android Users

Because the ATT policy only affects iOS users, we measure treatment intensity using the share of iPhone users at the zip-code level. We construct this variable using data from SafeGraph, a company that tracks foot traffic using GPS location data from mobile devices. This data has information on daily visits to 6 million points of interest across the US. For

¹⁶For example, Flurry, a mobile app analytics platform, shows that only 18% of iOS users allowed tracking among those who were asked for permissions. For details, see source: <https://www.flurry.com/blog/att-opt-in-rate-monthly-updates/>.

¹⁷Panels b and c of [Figure 2](#) provide examples of such data collection, showing privacy labels for Mint and TikTok that disclose collection of financial information and identifiers used for cross-app tracking.

each point of interest, SafeGraph reports a rich set of information, including time-invariant information such as the POI’s operating company (e.g., a certain bank or retail store), NAICS code, postal code, and time-varying information, such as monthly visit or visitor counts. Crucial for our study, SafeGraph also reports whether visitors use Android vs. iOS devices. SafeGraph aims to provide a representative sample of US consumers. [Li et al. \(2024\)](#) document a near-perfect correlation (>0.97) between the number of sampled devices and census population, for both urban and rural areas, and minor sampling biases among a number of demographic categories such as age, gender, and moderate income, with less than 5% under- and over-representation.

For the purpose of our analysis, we aggregate all visits to retail and grocery stores (identified by the two-digit NAICS code 44) and financial institutions (identified by the two-digit NAICS code 52) based on the device operating system (iOS or Android) and zip codes. The data covers the period from January 2019 to June 2022, providing a comprehensive view of foot traffic trends over time. We specifically focus on foot traffic to retail locations as they represent the majority of visits, and any potential operating-system-specific bias is relatively limited compared to other types of locations such as workplaces or hospitals. We expect that the share of iOS users at these general-purpose retail locations is representative of the iOS share within the corresponding zip code. Although our primary focus is on retail locations, we also include banks and other financial institutions in our analysis due to our interest in understanding financial fraud patterns. However, it is important to note that the foot traffic to these financial institutions is relatively small compared to retail locations. Consequently, excluding these institutions has little effect on our measurement.

3.2 Financial Fraud Data: Consumer Complaints

Our principal measure of financial fraud is derived from publicly available data on complaint submissions to the Consumer Financial Protection Bureau (CFPB), the federal agency with a mandate to address consumer grievances involving financial products and services.¹⁸ The CFPB collects complaints through its website and phone line, as well as through forwarding from the FTC’s Report Fraud portal, and publishes them in the Consumer Complaint Database after relaying them to the firms involved.¹⁹ Existing work uses CFPB data to study the service quality and discrimination of financial institutions, especially mortgage providers ([Begley and Purnanandam, 2021](#); [Huang et al., 2024](#); [Dou and Roh, 2024](#); [Li, 2023](#); [Mazur, 2024](#); [Jou et al., 2024](#); [Haendler and Heimer, 2025](#)); we instead exploit the complaint narra-

¹⁸Consumers can also direct complaints to state regulators or private entities such as the Better Business Bureau (BBB); we focus on the CFPB because of its centralized, federal scope.

¹⁹The CFPB does not publish complaints referred to other regulators (e.g., complaints about depository institutions with less than \$10 billion in assets), and suppresses 5-digit zip codes with population below 20,000.

tives and product categories to identify fraud originating from digital surveillance, and link complaints to firm-level measures of data collection and sharing.

When filing a complaint, individuals select a “product” from a list of 18 pre-defined categories (e.g., “credit reporting,” “debt collection,” “mortgage”) and an “issue” from a list of 165, with optional “subproduct” and “subissue” refinements.²⁰ Several of these issues are direct fraud indicators—for instance, “Incorrect information on your report” could indicate that a fraudster opened a credit card under the consumer’s name—while others may not be related to fraud. Individuals can also provide a narrative statement, which the CFPB publishes after removing personal information. We analyze the product, issue, subproduct, subissue, and narrative fields to identify complaints related to financial fraud originating from tracking and sharing of personal data.

The CFPB also reports the company being complained about, but consumers and firms are often unable to identify the source of a data-related incident. A consumer with accounts at multiple banks may discover a fraudulent credit line only when reviewing their credit report, and consequently file a complaint against the credit bureau rather than the bank where the breach occurred. Roughly half of all complaints are directed at the three major credit bureaus for this reason. We therefore use aggregate complaint counts at a granular locality level, rather than at the firm level, as our primary outcome variable.

We supplement the CFPB data with the Consumer Sentinel database, which combines complaints from several sources including the CFPB, and the FTC’s Identity Theft database. We describe these datasets further in [Section 4.4](#).

3.3 Other Data Sources and Measures

Data Collection and Sharing Practices. We create novel measures of the data collection and sharing practices of financial institutions and ISPs. Information on data collection practices is from Apple’s App Store privacy labels, which provide standardized disclosures of firms’ data collection practices. We scraped and parsed these labels to count the number of unique data items collected (details in [Appendix Section A.2](#)). We measure data sharing activities using data on the third-party Software Development Kits (SDKs) usage from App-topia. SDKs are pre-built software components that enable specific functionalities, including data analytics and targeted advertising, within an app (details in [Appendix Section A.3](#)). We identify all SDKs that facilitate data sharing across firms following the approach in [Bian et al. \(2024\)](#). We obtain ISP market-share data from the FCC’s Broadband Data Collection, which reports provider coverage and market shares at the local level (details in [Appendix C](#)).

²⁰Common issues include “Incorrect information on your report,” “Problem with a credit reporting company’s investigation into an existing problem,” “Improper use of your report,” “Attempts to collect debt not owed,” and “Fraud or scam.”

We construct equivalent measures for financial institutions appearing in the CFPB complaint database, using the same scraped privacy label and SDK installation data, and obtain their local market shares from the FDIC’s Summary of Deposits (SOD). [Appendix Section A.2](#) provides details on how we identify the mobile presence of banks.²¹

Cyber Events and Dark Net Activities. We turn to Advisen to identify data breaches and other cyber security events at the firm level (details in [Appendix F](#)). Finally, we use novel data on dark web forum discussions and listings from the Cambridge Cybercrime Centre (CCC, details in [Appendix G](#)) and Top10VPN (details in [Appendix H](#)).

3.4 Summary Statistics

The main regression sample consists of a balanced panel of the public CFPB complaints at the zip code level, spanning from January 2019 to June 2022. We impute zeros for zip codes with no reported complaints and exclude outlier zip codes using a multi-criterion procedure detailed in [Appendix B](#), which flags persistent irregularities in complaint patterns that are more likely to reflect concentrated legal-aid activity, or localized service disruptions than ordinary variation in fraud complaints. As shown in [Appendix B](#), these zip codes exhibit abnormal and recurring volatility, such as frequent spikes far beyond typical levels, that standard winsorization techniques will not remove.²²

[Table 1](#) presents summary statistics. For the balanced sample around ATT, approximately 26% of zip codes have at least one complaint in any given month. The mean number of complaints per 1,000 residents in a zip code per month is 0.05, indicating that around 5 residents out of every 100,000 file a complaint in a given month. [Figure 3a](#) displays the number of CFPB complaints per 1,000 residents for each zip code in the US, providing a visual representation of the rich spatial variation in complaint rates. For the balanced sample around the FCC Broadband rule (see discussions on the institutional background in [Section 4.3](#) and [Appendix C](#)), both the probability and total number of complaints are lower. Examining local ISPs’ data collection and sharing intensities, we also observe substantial variation (see [Appendix Figure A.3](#) for the distributions at ISP level).

Using foot traffic data, we find that the average iOS share across US zip codes is 46%, close to Statista’s estimate that the national iOS share fluctuated around 50% during 2019–2022.²³

²¹Recent studies have explored the effects of mobile banking penetration on local competition ([Haendler, 2022](#); [Koont, 2023](#); [Jiang et al., 2025](#)), financial inclusion ([Jiang et al., 2025](#)), bank franchise value ([Koont et al., 2024](#)), as well as local small business lending and economic growth ([Haendler, 2022](#)). We contribute by showing that, conditional on mobile penetration, lax data practices by banks increase financial fraud.

²²For example, outlier zip codes in North Carolina exhibit multiple spikes ranging from 200 to 800 complaints per 1,000 residents.

²³See <https://www.statista.com/statistics/266572>, the 4% gap may reflect that our measure better captures usage frequency, and iOS and Android users may spend different amounts of time on their devices, making our measure a usage-weighted iOS share.

This share varies substantially across zip codes, with a share of 39% for the 25th percentile zip code and 52% for the 75th percentile. Figure 3b further illustrates the geographical variation across the US. Consistent with DeviceAtlas estimates, iOS adoption is higher in the Southern states and along the Northeast Corridor.²⁴

4 Effect of Privacy Rules on Financial Fraud Complaints

4.1 Regression Specification

We begin by examining the extensive margin using a linear probability model:

$$\mathbb{1}\{Complaints\}_{z,t} = \alpha_z + \alpha_{state/county(z),t} + \beta iOS\ Share_{z,pre-ATT} \times Post_t + \varepsilon_{z,t} \quad (1)$$

where the dependent variable $\mathbb{1}\{Complaints\}_{z,t}$ equals one if there is at least one complaint in zip code z during month t , and zero otherwise. To move beyond the extensive margin, we then estimate the average treatment effect on the total number of complaints using a Poisson model. Specifically, the regression specification is:

$$Complaints_{z,t} = \exp\left(\alpha_z + \alpha_{state/county(z),t} + \beta iOS\ Share_{z,pre-ATT} \times Post_t\right) \varepsilon_{z,t} \quad (2)$$

where $Complaints_{z,t}$ is the number of complaints in zip code z and month t . We then scale this variable by population of the zip code, using the 2020 Census data, to obtain complaints per 1,000 residents, and winsorize at the top and bottom 1%. We estimate Equation (2) using Poisson pseudo-maximum likelihood (PPML). This approach is well-suited for proportional count data, as it naturally accommodates the non-negative and skewed distribution of the outcome while allowing for high-dimensional fixed effects.

In both regression equations, we use the variable $iOS\ Share_{z,pre-ATT}$ to capture the variation in exposure to ATT. The $iOS\ Share_{z,pre-ATT}$ represents the average *pre*-ATT iOS share of users at the zip code level, calculated over the nine quarters preceding the policy’s introduction.²⁵ This variable remains constant for each zip code since it is based on pre-treatment data. The treatment event indicator, $Post_t$, takes a value of one starting from May 2021, the first month after the ATT policy took effect on April 26, 2021. The coefficient β on the DiD term, $iOS\ Share_{z,pre-ATT} \times Post_t$, captures the differential change in consumer complaints between zip codes with higher versus lower iOS device shares.

To account for time-invariant characteristics that contribute to fraud, we include zip code fixed effects. Additionally, we incorporate state-by-year-month or county-by-year-month fixed effects, denoted as $\alpha_{state(z),t}$ or $\alpha_{county(z),t}$, to control for time-varying confounders at the state or county level. These confounders may include region-specific data regulations,

²⁴See <https://deviceatlas.com/blog/mobile-os-popularity-by-us-state>.

²⁵We use the *pre*-ATT measure because SafeGraph’s foot traffic data rely on location tracking from mobile devices. Since ATT may reduce the precision of these measurements, using the pre-policy period avoids measurement error caused by the policy itself.

local fraud news, or local economic developments. To be conservative, we cluster the standard errors by state. In our robustness checks, we aggregate the data to larger units, such as the quarter-month level or the county-month level, and conduct similar DiD analysis.

4.2 Baseline Results

[Table 2](#) presents the regression results using the CFPB complaints database. At the extensive margin, Columns 1 and 2 show a significant, negative coefficient on the interaction term $iOS\ Share_{z,pre-ATT} \times Post_t$. This estimate implies that, within a given county and month, zip codes with a higher proportion of iOS users experienced a larger decline in the probability of consumer complaints following the implementation of the ATT policy, relative to zip codes with lower iOS user shares. In Column 2, the coefficient is approximately -0.059 when controlling for county-specific shocks, indicating that a zip code with a 10% ($\approx$ one standard deviation) higher iOS share experienced a 0.59 percentage point reduction in the probability of a complaint, equivalent to about 2.4% of the mean.

The extensive margin analysis, however, does not capture changes in complaint intensity. Moving to Columns 3 and 4, we examine the number of complaints per 1,000 residents. We again find negative and statistically significant coefficients at the 1% level. The Poisson estimates imply that a one standard deviation increase in iOS share translates into a 5.4%–6.1% decline in complaints per 1,000 residents.²⁶

The magnitude is larger using the total number of complaints because it captures both whether any complaints occur and the number of complaints filed, conditional on any complaints filed. Given the observed opt-out rate of 83%, our estimates imply that complaints to the CFPB could decline by roughly 6.5%–7.3% if 10% of mobile app users were to disallow data tracking (calculated as $5.4\%/0.83$ or $6.1\%/0.83$). While the implied magnitude appears large, it reflects a relative drop within an overall upward trend in CFPB complaints. Broader forces, such as growth in digital transactions, expanding attack surfaces, and increasingly sophisticated criminal tools, continue to push fraud upward, so even sizable relative effects may not be reflected in the aggregate trends.

Dynamics. While we compare zip codes within the same county-month, high iOS-share zip codes may differ from low iOS-share zip codes in various dimensions that could affect the trends of consumer complaints. For example, ownership of Apple products predicts higher income and better education ([Bertrand and Kamenica, 2023](#)), which can lead to changes in consumer complaints among iOS users over time if higher-income or better-educated consumers are hit with different shocks than other consumers. To rule out alternative expla-

²⁶In Poisson models, coefficients can be interpreted as semi-elasticities. The percentage change in the dependent variable for a change Δx in the regressor is given by $(e^{\beta \cdot \Delta x} - 1) \times 100$. In our case, the coefficient in Column 3 is 0.548 and we evaluate the effect of a one standard deviation increase in *iOS* share, yielding $(e^{-0.548 \times 0.108} - 1) \approx -5.4\%$.

nations, we examine pre-trends in consumer complaints and plot dynamic DiD coefficients.

In [Figure 4](#), we present the results. To reduce estimation error, we group all three months within a corresponding quarter. The analysis covers a total of nine quarters before the introduction of the ATT policy and four quarters after its implementation.²⁷ We define quarter -1 as the quarter immediately preceding the implementation (2021Q1), which serves as the benchmark quarter. We then plot the coefficients of quarters -2 to -4 as well as a coefficient for all quarters prior to -4 combined.

The dynamic DiD coefficients confirm that the reduction in complaints manifests after ATT’s implementation. Examining the probability of complaints in Panel a of [Figure 4](#), we observe that, prior to the introduction of ATT, the coefficients for all quarters are not statistically significantly different from zero. However, there is a clear negative post-trend, suggesting a decline following the policy’s implementation in zip codes with larger shares of iOS users. Examining complaints per 1,000 capita in Panel b of [Figure 4](#) reveals a similar pattern. We also estimate the monthly dynamic treatment effects and report the results in [Appendix Figure I.1](#). While the estimates are noisier at the monthly frequency, the overall pattern is consistent with the quarterly figures.

Immediacy of ATT’s Effects. The dynamic estimates show that the effects of ATT emerge within two months of implementation. This rapid response provides novel evidence on the short horizon over which modern, data-driven fraud operates. While one might expect a longer lag from increased data protection to reduced fraud and subsequent reporting, one plausible channel is targeted deceptive advertising, which relies on real-time data and can be disrupted almost immediately when data access is curtailed. Today’s surveillance economy also enables near-immediate monetization of stolen credentials and personal information through well-organized fraud-as-a-service markets and automated attack tools.²⁸

On the detection side, advances in real-time transaction monitoring, mobile banking alerts, and identity-theft protection services enable consumers to identify and report suspicious activity quickly. The widespread adoption of push notifications and credit-monitoring tools (e.g., LifeLock, Mint, Credit Karma) has further compressed the time between fraud occurrence and reporting.²⁹

To validate this short timeline, we examine CFPB complaints following the 2017 Equifax

²⁷As discussed in [Section 4.5](#), we also extend the sample through the end of 2023 as well as 2024 and find larger, statistically significant DiD coefficients ([Appendix Table I.6](#) Columns 5 and 6). We use an earlier cutoff in our main analysis because the pre-ATT iOS share may no longer accurately capture treatment intensity as the post-ATT period extends further out.

²⁸Dark web marketplaces facilitate near real-time resale of compromised data, while credential-stuffing and account-takeover attacks can be launched within hours of a breach. Emerging AI-driven tools further accelerate these processes. See the discussion on fraud-as-a-service [here](#) and AI tools [here](#).

²⁹See McKinsey’s discussion [here](#) and Thomson Reuters’ report [here](#).

data breach and document a sharp increase within three months ([Appendix Figure I.2](#)). Historical evidence from the Federal Trade Commission shows that even two decades ago, a large share of victims detected misuse within weeks, and more recent data from the National Crime Victimization Survey indicate that 77% of identity-theft cases are identified within a month and 44% within a day. These patterns are consistent with a compressed timeline in modern fraud—from data exposure, to exploitation, to detection and reporting—reflecting both the rapid use of stolen data by fraudsters and faster detection by consumers.

4.3 A Reverse Shock: Repeal of the FCC’s Broadband Privacy Rules

To provide further evidence, we examine a second policy shock that also affects the collection and sharing of consumer data, but in the opposite direction by weakening data protections. Specifically, we study the 2017 repeal of the FCC’s broadband privacy rules, which granted ISPs greater latitude in consumer surveillance. This setting offers two advantages. First, it provides a reverse shock relative to ATT, allowing us to test whether weakening privacy protections has effects that mirror our main results. Second, it introduces a distinct source of identifying variation beyond the iOS-Android comparison.

ISPs occupy a uniquely central position in the data economy due to their ability to observe user activity across websites, apps, and devices, unlike platforms constrained to logged-in or embedded environments. Their data collection is persistently tied to households and is difficult for consumers to avoid, given limited broadband competition. These characteristics make ISP data especially valuable for profiling and monetization in the data-sharing ecosystem ([Bian et al., 2024](#)).

In 2016, the Federal Communications Commission (FCC) adopted the Broadband Privacy Rules, which were set to require ISPs to obtain explicit customer consent before using or sharing sensitive information such as web browsing history, app usage data, and precise geolocation for advertising or other non-service purposes. The rule-making process started in 2015, when the FCC reclassified broadband as a telecommunications service under Title II, shifting privacy oversight from the FTC to the FCC and signaling that stronger privacy restrictions were forthcoming. From 2015 until the repeal in 2017, ISPs operated under the shadow of these pending rules, even though the rules themselves never fully took effect, constraining their ability to collect and exploit data.

In March 2017, Congress repealed the FCC’s privacy rules under the Congressional Review Act (henceforth “the Repeal”), and President Trump signed the repeal into law. This action prevented the FCC from ever enacting “substantially similar” rules, which not only removed the immediate threat of stringent privacy requirements but also credibly signaled

to ISPs a significantly lower likelihood of facing similar restrictions in the future.³⁰ Contemporaneous policymakers and privacy advocates described the repeal as creating a “rule-free zone” for ISPs.³¹ The repeal thus granted ISPs greater flexibility to collect, share, and monetize customer data under the FTC’s more general consumer protection framework, increasing the returns to existing data infrastructures and enabling firms with more invasive data practices to scale up data exploitation. Consistent with this shift, major ISPs simultaneously expanded their advertising and data monetization businesses: AT&T acquired AppNexus in 2018 to build a large-scale targeted advertising platform using subscriber data, while Verizon Communications integrated its broadband user data with its AOL and Yahoo ad-tech businesses to enable cross-platform targeting. A detailed description of the background and timeline of the Repeal is provided in [Appendix C](#).

This reverse shock complements the main ATT setting by allowing us to examine how an expansion in consumer surveillance, rather than a contraction, affects financial fraud. To measure a locality’s exposure to the Repeal, we construct two metrics that combine the local market shares of ISPs with the invasiveness of their data practices. Specifically, we consider (1) the number of unique data items collected (# Data Types Collected) and (2) the number of third-party platforms with which ISPs share or transmit data (# Data SDKs).³² To capture exposure to ISPs with invasive data practices, we aggregate the two metrics across ISPs within each zip code. Specifically, we first weight ISPs by their local market shares, measured at the census block level and available at semi-annual frequency. We then aggregate from census blocks to zip-code level using the population shares of blocks within each zip code, as defined in [Equation \(3\)](#) and [Equation \(4\)](#).³³

³⁰The “substantially similar” clause (5 U.S.C. §801(b)(2)) prohibits agencies from reissuing any future rules that closely resemble the repealed regulation, unless explicitly authorized by subsequent legislation.

³¹Senator Brian Schatz warned that overturning the FCC rule would leave “neither the FCC nor the FTC [with] clear authority” to protect consumer privacy, allowing ISPs to operate without meaningful oversight. Similarly, Eric Null, Policy Counsel at the Open Technology Institute, argued that the repeal substituted the FCC’s stronger *ex ante* privacy safeguards with the FTC’s weaker *ex post* “lowest common denominator” framework, which primarily targets deception rather than ensuring fairness in data practices.

³²Since ISPs often operate multiple apps, we compute the average of relevant metrics across all apps owned by a given ISP. Our results remain robust when using the aggregate or unique counts of data types collected and SDKs installed across all apps owned by the same ISP. These robustness checks are provided in [Appendix Table I.1](#) and [Table I.2](#).

³³Given the high concentration in ISP markets, we manually identify the 12 app-owning ISPs, listed in [Appendix Section A.3](#), which together represent 85% (92%) of the market as of June 2016 (June 2021). Two factors may lead to an underestimation of local exposure. First, we assume that the data collection and sharing intensity of the substantially smaller remaining ISPs is zero. This assumption is consistent with evidence that larger companies engage in heavier data extraction activities, as documented in [Bian et al. \(2021\)](#). Second, our measures of ISP data practices are based on their mobile apps, which may understate the true extent of surveillance, as ISPs can also track consumers through web-based activity.

$$\text{ISP Exposure}_z^{(m)} = \sum_{b \in \mathcal{B}(z)} \underbrace{\frac{p_b}{P_z}}_{\text{block population weight}} \left(\sum_i s_{ib} I_i^{(m)} \right), \quad (3)$$

where i indexes ISPs, b indexes census blocks, and $\mathcal{B}(z)$ denotes the set of census blocks in zip code z . Data collection and sharing intensities, $I_i^{(m)}$, are defined as

$$\begin{aligned} I_i^{(\text{Collection})} &= \# \text{ Data Types Collected}_i, \\ I_i^{(\text{Sharing})} &= \# \text{ Data SDKs}_i. \end{aligned} \quad (4)$$

We then compute the average of this exposure measure over the pre-Repeal period (January 2015 to October 2016). Higher exposure reflects ISPs with more developed data infrastructures prior to the Repeal, which allows them to expand data collection and monetization more aggressively once constraints are lifted, increasing exposure to data misuse and fraud. We employ a DiD design around the Repeal, between January 2015 and December 2018, and report the results in Table 3.³⁴ Panel a uses the exposure measure based on ISP data collection intensity, while Panel b focuses on data sharing intensity. We find that, following the Repeal, areas served by ISPs with more invasive data protection practices experienced an increase in complaints. For example, in Panel a, a one-standard-deviation increase in local exposure based on data collection practices is associated with a 2.4%-2.8% increase in the probability of complaints and a 3.9%-5.7% increase in the number of complaints per 1,000 residents. Results based on data sharing practices in Panel b are quantitatively similar. The dynamic DiD estimates in Appendix Figure I.4 show that the effects arise abruptly following the Repeal, with no evidence of pre-trends in the preceding period.

Like our analysis based on ATT, this setting links changes in data practices to shifts in fraud, strengthening our identification by showing that the narrative operates in both directions.

4.4 Analysis Using Consumer Sentinel and Identity Theft Complaints

We complement our main analysis using proprietary complaint-level data from the FTC, which offers three key advantages over the public CFPB database. First, the Consumer Sentinel data is uncensored: it includes all complaints made public by the CFPB as well as those withheld from the public release for privacy reasons (e.g., when disclosure could identify a consumer in a small zip code) and those excluded when firms fail to acknowledge a consumer relationship within 15 days.³⁵ Second, it includes two complementary non-public datasets: the Consumer Sentinel Network (CSN), which aggregates fraud and consumer

³⁴We restrict the sample to observations before 2019 to avoid overlap with the ATT analysis.

³⁵The Consumer Sentinel database has about 34% more CFPB complaints per month than the public CFPB database used in our previous analyses. See <https://www.consumerfinance.gov/data-research/consumer-complaints/> for more details on which CFPB complaints are made public.

complaints from a wide range of federal, state, and local agencies as well as private-sector partners, and the Identity Theft (IDT) database, which focuses specifically on identity theft cases.

Table 4 presents the results. Column 1 shows that using uncensored CFPB complaints from Consumer Sentinel yields an effect size of 5.6%, very close to the baseline 6.1% reduction in public CFPB complaints for a one-standard-deviation increase in iOS share. Thus, public-data censoring is unlikely to explain our results, and rules out the alternative explanation that firms might have been less able to acknowledge such relationships if ATT impaired firms’ ability to identify consumers, leading to filed complaints being omitted from the public CFPB database. Columns 2 and 3 turn to the Consumer Sentinel and Identity Theft datasets. For the Identity Theft database, we find a statistically significant decline in complaints in high-iOS-share areas after ATT, indicating that identity theft, as one of the most common forms of fraud leveraging illicit personal data, is affected. The broader CSN shows a smaller decline of 1.1%. This decline in magnitude is not surprising, as the database covers a much broader range of complaints that concern non-finance products and services and cases not triggered by data breaches.

Finally, Columns 4 and 5 exploit the FTC’s monetary loss data. The similar magnitudes in Columns 3 and 4 of a 1.1% decline suggest that both fraud cases involving a reported loss and those without a reported loss decline at similar rates after ATT. In contrast, the larger magnitude in Column 5 (1.8%) relative to Column 4 suggests that ATT may disproportionately reduce high-loss cases, implying that the policy curtails not only the frequency of fraud but also its more severe financial impacts.

These estimates also enable a simple back-of-the-envelope calculation of ATT’s potential monetary benefit to consumers. The Consumer Sentinel Network received reports of \$6 billion in consumer fraud losses in 2021. With iOS accounting for roughly 50% of the US smartphone market, an 83% opt-out rate from tracking under ATT, and the estimated 1.1% reduction in CSN fraud complaints with dollar losses per 10% increase in iOS device share, the implied reduction in fraud losses associated with ATT is approximately $0.011 \times (50\%/10\%) \times 83\% \times 6 \text{ billion} \approx \274 million , which corresponds to 4.75% of total losses due to financial fraud in 2021, in the year after its implementation. These estimates are conservative given that reported fraud losses have doubled since 2021 (at \$12.5 billion in 2024), and fraud losses reported to Consumer Sentinel are known to substantially understate total fraud losses.³⁶

³⁶See the discussion on underreporting in the FTC’s 2024 Older Adults Report <https://www.ftc.gov/reports/protecting-older-consumers-2023-2024-report-federal-trade-commission>.

4.5 Robustness

iOS vs. Android Users: Concurrent Shocks and Demographics. ATT was implemented in April 2021, which was in the aftermath of the COVID-19 pandemic hitting one year before in March 2020. It is possible that fraud complaints related to the pandemic increased for Android users relative to iOS users and this was still ongoing at the time of the implementation of ATT. The pandemic led to changes that could foster fraud, such as increases in online shopping and a surge in demand plus shortages for COVID-19 related items. We address this concern in the following ways.

We focus first on the Treasury’s Economic Impact Payments (EIP). To examine whether these payments affect our results, we directly control for the interaction between three variables related to EIP payments and the post-ATT indicator. The EIP variables are constructed using IRS zip code-level tax return data and include the total amount of EIPs received, the average household income, and the average number of children and other dependents in a zip code. The amount of EIPs received by a zip code captures the actual payments, while household income and the number of children determine eligibility for EIPs. If our baseline results are driven by rising EIP-related fraud in low-iOS-share zip codes, we should see a substantial reduction in the point estimates after controlling for these factors. However, [Table 5](#) Panel a shows that our DiD estimates remain similar in both economic and statistical significance, regardless of whether the interaction terms are based on actual EIP payments (Columns 1 and 3) or on eligibility (Columns 2 and 4).

To further address concerns that our ATT estimates could be confounded by COVID-19 relief fraud disproportionately affecting low-iOS-share areas, we conduct placebo tests using two alternative event dates: April 2020 and December 2020. These dates correspond to the disbursement of the first and second major waves of pandemic-related financial support under the CARES Act and subsequent legislation.³⁷ By shifting the event time to these earlier points, we can test whether our baseline results also appear around the timing of major COVID-19 relief disbursements. As shown in Panel b of [Table 5](#), after the COVID-19 relief disbursements, high-iOS-share areas do not experience declines in complaints compared to low-iOS-share areas, providing further assurance that our main results are not driven by fraud related to COVID-19 relief programs.

Beyond EIP-related controls, we further address potential confounding factors by interacting zip code-level socio-economic and demographic characteristics with the post-ATT indicator. The rationale is that such factors jointly influence COVID-19 and EIP exposure as

³⁷The first round of stimulus payments and expanded unemployment benefits began reaching individuals in late April and early May 2020, while the second round was distributed in late December 2020. Both periods saw sharp increases in fraud complaints, particularly related to identity theft and benefit scams.

well as the propensity to use iOS devices. [Haendler and Heimer \(2025\)](#) show that the propensity to claim fraud and request refunds in the complaint process correlates with education and income. We consider a broad set of variables, including age, gender, education, income, and unemployment. Unemployment is included to capture the potential effects of pandemic-related job support, such as generous unemployment insurance, as well as the increased rates of unemployment during COVID-19. Adding these interactions leaves the direction of our DiD estimates unchanged ([Appendix Table I.7](#)), though the magnitude is smaller (at most halved) as some of the identifying variation is absorbed by these socioeconomic factors.

Another contemporaneous change is Google’s “zero-out” policy, introduced between late 2021 and April 2022, which replaces the Android advertising ID with a string of zeros only for users who have already chosen to opt out of ad personalization. Because opt-out rates on Android were extremely low,³⁸ this policy affected only a small share of devices and represents a far more limited shift than Apple’s ATT. Moreover, any Android privacy initiatives like this would bias our estimates toward zero.

Propensity to File Complaints. The propensity to complain after victimization can vary across different communities, so declines in consumer complaints may not immediately translate to declines in fraud victimization. [Raval \(2020b\)](#) examines how several zip code-level demographic variables affect the likelihood of complaining by comparing complaints and victims for several consumer protection cases and found much lower complaint rates, conditional on victimization, in heavily Black and Hispanic areas. Using these estimates, [Raval \(2020b\)](#) develops zip code-level weights designed to be the inverse of the predicted complaint-to-victim ratio based on those demographics in order to “correct” complaint data for differences in the likelihood of complaining across demographic groups. The weight for the median zip code was normalized to 1, with Black zip codes averaging about 2 because residents there were roughly half as likely to file complaints.

In Column 1 of [Appendix Table I.3](#), we multiply our complaint counts by these weights to examine changes in fraud victimization. This exercise relies on the assumption that the differences between the propensity to complain across locations found in [Raval \(2020b\)](#) extrapolate to CFPB complaints and that the adoption of ATT did not change the propensity to complain.³⁹ We find that the estimated post-ATT decline in fraud victimization is 6.9% per one standard deviation increase in iOS share, which is larger than the unadjusted 6.1% from Column 4 of [Table 2](#). This pattern suggests that individuals most vulnerable to data-

³⁸According to a report by Singular based on over 176 million Android smartphones worldwide, only 2.08% of users opt out. See DOJ: [“Zeroing out the Android Limit Ad Tracking \(LAT\) will impact only 2% of devices globally,”](#) June, 2021.

³⁹ATT could raise awareness of data-driven fraud among iOS users and induce them to complain more, but that would lead to an increase in fraud complaints after ATT rather than a decrease.

driven fraud are, on average, less likely to file a complaint, implying that the unadjusted estimates modestly understate ATT’s true effects.

Measurement Error and Other Robustness Checks. In [Appendix D](#), we further examine the robustness of our estimates to potential measurement error in iOS shares, alternative aggregation, and other specification changes. To reduce possible measurement error in our iOS-share measure, we exclude zip codes with low foot traffic or few grocery stores and place greater weight on more populous zip codes, where the measure is more likely to reflect local residents rather than visitors. Our conclusions remain unchanged. The results are also robust to aggregating the data to alternative geographic and temporal levels, excluding complaints related to the three major credit bureaus, which account for about half of all complaints, and varying several specification choices, including extending the sample through 2024, winsorization, outcome scaling, outlier treatment, and standard-error clustering. The estimates remain close to the baseline and are in some cases slightly larger. Full results are reported and discussed in [Appendix D](#).

5 Mechanism

5.1 Downstream: Data-Related Fraud Complaints

The downstream analysis follows our earlier analysis but zooms in on consumer fraud complaints as the most visible consequences of compromised data. While many forms of fraud exist, only some rely directly on access to compromised personal or financial information. By singling out the more relevant categories, we can test whether ATT primarily affects fraud mechanisms that plausibly depend on compromised upstream data suppliers.

5.1.1 Identifying Relevant Complaints

Not all consumer complaints are directly linked to the collection and misuse of personal information by thieves or hackers. The “issues” or “sub-issues” fields in the CFPB database do not explicitly distinguish between more or less relevant categories for fraud arising from lax data privacy regulations. This is even more true in the Consumer Sentinel data, as it amalgamates complaints from many different sources.⁴⁰ To determine the relevance of complaints for data privacy issues, we employ two approaches using the consumer narrative field. We have narratives for 40% of the complaints in the public CFPB dataset, as well as for all of the complaints in the Consumer Sentinel and Identity Theft databases.

First, we conduct keyword searches based on the sub-product, issue, sub-issue, and narrative complaint fields for the CFPB database and based on the narrative field for the other

⁴⁰A consumer with the same underlying issue complaining to a specific data contributor like the CFPB could potentially classify the same complaint into different categories, and each data contributor also has its own way of classifying complaints into different categories, which then have to be translated into the Consumer Sentinel categorization.

databases. We compile a list of keywords related to fraud and data privacy, such as “incorrect,” “fraud,” “theft,” “identity,” and “data breach.”⁴¹ If any of these keywords appear in the relevant fields, we assign an indicator variable with a value of one. This approach allows us to identify complaints that might be affected by changes in data privacy standards.

Second, we use a language-model-based zero-shot learning (ZSL) technique to assess whether a narrative is related to fraud arising from data privacy issues. The procedure assigns each complaint narrative a continuous score reflecting its relevance to data-related fraud relative to other complaint types, without requiring a task-specific hand-labeled training set. Details of the procedure are provided in [Appendix E](#).

Since the narrative-based relevance score is only available for a subset of complaints with consumer narratives, we extrapolate the scores at the product-category level in the CFPB data to classify categories with higher average scores as more relevant. [Appendix Table E.1](#) presents the mean and standard deviation of complaint-level scores by product category.

Two patterns are worth noting. First, both the keyword search method and the machine learning approach generate meaningful variation in the average scores at the product level, allowing us to distinguish between more and less relevant complaints. For example, the highest and lowest product-level scores generated by the keyword search method are 0.82 and 0.30, respectively, while the highest and lowest scores generated by the ZSL method are 0.53 and 0.16, respectively. Second and more importantly, the scores generated by these two methods exhibit a high correlation at the tails, indicating a consensus on the most relevant and irrelevant complaints. Both methods consistently rank “Credit reporting” and “Debt collection” as the most relevant categories, while “Student loans” and “Mortgages” receive the lowest scores, suggesting lower relevance for data security regulations.

5.1.2 Heterogeneity by Complaint Relevance

We first estimate our main regression specification separately for the two CFPB categories with the highest relevance—Credit Reporting and Credit Repair and Third Party Debt Collection—and the category with the least relevance—Student Loan and Mortgage—and report the results in Panel a of [Table 6](#).

Consistent with the hypothesis that ATT reduces financial fraud enhanced by lax data privacy, we find negative and statistically significant effects on complaints within the top two fraud categories (Panels a and b). The magnitude of the effects is comparable to that observed in the full sample of complaints. Following the implementation of ATT, a 10% increase in the share of iOS users in a zip code is associated with a 5.5% decrease in the

⁴¹The full list of keywords used is “incorrect,” “improper,” “false,” “wrong,” “missing,” “fraud,” “scam,” “theft,” “embezzlement,” “imposter,” “unauthorized,” “unsolicited,” “identity,” “sharing,” “advertising,” “marketing,” “security,” “data breach,” “not owed.”

number of complaints related to credit reporting per 1,000 residents (Column 1) and a 5.7% decline in the number of complaints related to debt collection (Column 2). In contrast, ATT has an insignificant effect on complaints related to student loans or mortgages (Columns 3 and 4). Applying the same method to the Consumer Sentinel complaint database, we find similar results in [Appendix Table I.8](#).

We additionally leverage the full coverage of narratives in the Identity Theft and Consumer Sentinel complaint databases to directly split complaints into those including one of the keywords described above and those not including one of those keywords. We estimate the main regression specification separately for these two sets of complaints and report these estimates in Panel b of [Table 6](#). Once again, we find negative and statistically significant results for complaints that are more related to data privacy issues. Following the implementation of ATT, a 10% increase in the share of iOS users in a zip code is associated with a 7.2% decline for complaints about identity theft (Column 1) and a 2.7% decline in complaints to Consumer Sentinel (Column 2). We find null effects (0.0%, -0.4%) for complaints from both sources that do not include one of the relevant keywords. The lack of an effect of ATT on complaints that are not related to the misuse of personal information serves as an additional placebo test, suggesting that our results are unlikely to be driven by concurrent shocks or differential time trends for iOS and Android users.

5.2 Midstream: Financial Institutions and ISPs as Data Gateways

Midstream actors—banks, and ISPs—form a key link between fraudsters and consumers. They routinely collect large volumes of individual-level behavioral and financial data, which, if stolen or misused, can fuel downstream fraud. These institutions are thus both gateways through which data can reach malicious actors and targets whose data repositories attract cybercriminals. Understanding how their data-collection and sharing practices shape fraud risk is essential to connecting consumer fraud outcomes to upstream data markets. In addition, unlike more purely digital platforms such as Google or Meta, they have observable local market shares, allowing us to exploit geographic variation in both their presence and their data practices.

5.2.1 Heterogeneity by Local Exposure to High-Risk Data Collectors

We begin by extending our baseline zip code-level design, which exploits geographic variation in iOS share, to incorporate local variation in the prevalence and data practices of two key data collectors, financial institutions and ISPs, and test whether ATT’s effects are stronger in places where these intermediaries collect and share more consumer data. Financial institutions handle high-value identity and transaction data directly relevant to financial fraud, while ISPs sit at the center of online activity and capture broad swaths of consumer

behavior. Both can act as gateways through which data reaches upstream fraudsters.

We measure these institutions’ data-handling practices by analyzing their mobile apps, parsing Apple’s iOS privacy-label disclosures, and identifying third-party data-sharing SDKs installed in those apps. We then aggregate these measures to the local market level using bank deposit shares and ISP market shares to capture local exposure to aggressive data collectors. Interacting these exposure measures with zip code-level iOS share in a triple-difference regression allows us to test whether ATT’s effect on fraud is concentrated in areas where high-risk banks and ISPs are more prevalent.

Panel a of [Table 7](#) reports results exploiting local exposure to problematic data practices in the banking sector. We focus on institutions active in the mobile market and construct two exposure measures, each interacted with $Post \times iOS\ Share$. The first measure captures local exposure to data collection. For each bank, we multiply the number of unique data items collected by its app by the app’s popularity, proxied by download counts, and then sum across all banks.⁴² The second measure, local exposure to data sharing, is constructed analogously, replacing the number of data items collected with the number of third-party data-sharing SDKs integrated into the app. Both measures are constructed at the county-quarter level and then averaged across all quarters in the pre-ATT period. To facilitate comparison across specifications, we standardize the measures using z-scores. To approximate the share of downloads originating from local users, we weight exposure by the within-bank, across-county deposit share, as formally defined in [Equation \(5\)](#).

$$\text{Bank Exposure}_c^{(m)} = \sum_{j \in \mathcal{B}(c)} \left(I_j^{(m)} \times \text{Downloads}_j \times \% \text{Deposit}_{jc} \right), \quad (5)$$

where c indexes counties, j indexes banks, and $\mathcal{B}(c)$ denotes the set of banks with branches in county c . The data collection and sharing intensities, $I_j^{(m)}$, are defined analogously to [Equation \(4\)](#). Downloads_j is the total number of app downloads for bank j , and $\% \text{Deposit}_{jc}$ is the within-bank share of deposits held by bank j in county c .

Across all specifications, the coefficients on the triple interaction terms $Post \times iOS \times Exposure$ are negative and statistically significant, indicating that ATT’s fraud-reducing effect is strongest in markets where local banks collect or share more consumer data through their mobile apps. Based on Column 4 in Panel a, a one-standard-deviation increase in the interaction between iOS share and exposure based on banks’ data-sharing intensity is associated with a 1.9% additional decrease in complaints per 1,000 residents. This pattern suggests that banks act as information intermediaries, often inadvertently playing a key role

⁴²This product reflects both how many users are using the app and how much data is collected per user, providing an estimate of the total volume of data extracted in a given area. If a bank owns multiple mobile apps, we consider its flagship app to be the one with the highest number of downloads.

in collecting and transmitting personal data that is later exploited by fraudsters.

Panel b of [Table 7](#) repeats the analysis for ISPs, using county-level exposure to their data-collection and data-sharing intensities, defined analogously to [Equation \(3\)](#). The coefficients are again negative and statistically significant across specifications, indicating that ATT’s fraud-reducing effect is larger in markets where ISPs are more active in gathering and transmitting consumer data.

Overall, these results show that ATT’s fraud-reducing effects are most pronounced in markets where key local intermediaries collect and share more consumer data. This evidence links our baseline consumer-side results to the midstream institutions through which data may flow into the broader ecosystem.

5.2.2 Firm-Level Evidence on Cyber Incidents

We next study firm-level cyber incidents, which provide a complementary lens on the same mechanism while bringing us one step closer to fraudsters’ activities. Because financial institutions with consumer-facing mobile apps tend to collect and share rich consumer data, they should be especially exposed to ATT’s privacy restrictions. If ATT limits fraud by reducing the volume and quality of exploitable personal data, we should also observe declines in cyber incidents and related consumer complaints at these firms.

We examine this question using Advisen’s cyber-incident database, which covers over 90,000 events from publicly verifiable sources between 2000 and 2023.⁴³ For each firm-month in our sample, we identify whether the firm experienced any cyber incident, whether it stemmed from malicious breaches or privacy violations, whether it resulted in violations of the Fair Debt Collection Practices Act or Fair Credit Reporting Act—the two regulations most commonly linked to fraud in CFPB complaint narratives—or whether it was unrelated to lax data standards (e.g., lost devices or accidental disclosure).

Panel a of [Table 8](#) shows that ATT significantly reduces cyber-incident exposure for app-owning firms. The likelihood of any incident falls by 4.9 percentage points relative to firms without an app, roughly half of the baseline. The effects are strongest for privacy-related incidents and for those involving violations of the two key fraud-related regulations (Columns 1 and 3), with no significant effect for unrelated causes (Column 4).

We next examine whether the reduction in firm exposure to breaches and privacy violations translates into fewer consumer complaints about these firms. In Panel b of [Table 8](#), we find that app-owning firms are significantly less likely to receive fraud-related complaints after ATT than non-app-owning firms. Column 1 shows that the monthly probability of any complaint falls by 2.1 percentage points. The effect is concentrated in the most fraud-relevant categories—credit reporting, credit repair services, and debt collection—with no

⁴³More information on Advisen’s data sources can be found [here](#).

detectable changes in unrelated categories such as student loans or mortgages.

Together, these results connect downstream fraud reductions to the intermediaries targeted by malicious actors and provide a window into the otherwise hard-to-observe activities of fraudsters. This motivates our next step: examining illicit upstream markets where compromised personal and financial data is exchanged.

5.3 Upstream: Illicit Data Markets

Our upstream analysis digs into the supply side of the fraud economy—the illicit markets where compromised personal and financial data is traded. As the FTC notes, the dark web functions as a critical conduit for this trade, enabling cybercriminals to buy and sell compromised information that fuels a variety of fraudulent activities.⁴⁴ These markets operate in hidden, unregulated environments and are only partially observable through dark web forum discussions and marketplace listings, which together capture both the conversation layer (coordination and exchange of know-how) and the transaction layer (actual sales and pricing). We track these two layers and investigate discussions, availability, and prices of compromised data on the dark web. Importantly, we compare trends in categories directly exposed to ATT with those in unaffected control groups to assess whether privacy restrictions disrupted the supply chain feeding downstream fraud.

In contrast to downstream complaints and the data-collection practices of midstream entities, which can be readily measured in structured administrative or standardized datasets, upstream activity is inherently opaque. Dark web market prices and volumes can be noisy, influenced by strategic behavior, enforcement actions, or migration to less visible channels. As such, our evidence based on the dark web should be viewed as indicative rather than conclusive. Even so, these data offer a rare opportunity for forensic finance—direct observations of fraudsters’ actions rather than inferences pieced together from more distant outcomes. To our knowledge, they have not been used in finance research, and we hope our work encourages further studies that leverage such rare opportunities to shed light on the economy of fraudsters.

5.3.1 Posts on Dark Web Forums

We begin with dark web forum discussions—the conversation layer of the illicit data economy—where actors share techniques, coordinate activities, and exchange market intelligence. Changes in discussion patterns can signal shifts in supply conditions or attacker behavior.

Our data comes from CrimeBB, a structured collection of textual data scraped and maintained by the Cambridge Cybercrime Centre (CCC). CrimeBB contains approximately 124 million posts from 7 million members across 37 cybercrime and extremist forums, regularly

⁴⁴See FTC’s discussion on dark web [here](#).

updated and stored on CCC servers. We use posts from all forums from 2015 onward.

We apply a textual classification procedure to label each post according to the labels and keywords listed in [Appendix Table G.2](#) to [Table G.6](#). We then aggregate monthly post counts by label and forum to construct our regression panel.⁴⁵ The most exposed category includes posts explicitly related to Apple devices or the iOS ecosystem. While direct exposure to ATT is difficult to measure, the policy applies only to iOS, making Apple related discussions most likely to be affected. The partially treated category covers broader illicit activities that plausibly depend on personal or financial data—such as tracking, data breaches, identity theft, and financial data theft. Many of these activities are known to rely heavily on data originating from mobile devices, particularly in financial fraud schemes, so changes in these discussion volumes can provide an early signal of shifts in the supply of mobile-sourced and financially sensitive data.

Our control group consists of illicit-trade categories that operate in a fundamentally different, segmented market. These markets, such as drugs, human trafficking, and weapons, are characterized by distinct participants, supply chains, and transaction methods, with minimal overlap in actors or infrastructure with the data-driven fraud economy. As such, they provide a benchmark plausibly insulated from shifts in digital data availability and attacker strategies that could result from ATT.

This setup allows us to estimate DiD regressions comparing changes in post volume for the exposed and partially treated groups against the control group around ATT’s implementation. The results are displayed in Panel a of [Table 9](#). We find that post volumes in the most exposed group fell by over 30% following ATT (Column 1). The decline is smaller but still statistically significant for the partially treated group (Column 2), consistent with these illicit activities being affected through reduced access to mobile-sourced and financially sensitive data. Although some actors discuss potential workarounds immediately after ATT, overall activity declines, likely due to a combination of reduced supply and higher barriers to circumvention, which limit participation in these markets.

We show in [Appendix Table I.9](#) that the results are robust to using alternative, more restrictive keyword lists and to alternative counting methods (fractional assignment versus double counting ambiguous posts). The richness of forum topics also enables a placebo test: we examine unrelated digital technology categories such as zero-days, malware, and botnets, which do not rely on personal data for their operation. We find little or no change in these categories, helping to rule out the possibility that our results are driven by broader trends

⁴⁵Our classification of labels follows [Avarikioti et al. \(2018\)](#), which use supervised machine learning to develop a comprehensive categorization of the dark web content.

in digital-related discussions rather than ATT (Column 3 of [Table 9](#) Panel a).⁴⁶

5.3.2 Dark Web Listings and Prices

To connect conversations on dark web forums to actual transactions, we analyze the transaction layer through dark web marketplace listings and prices. Using data on products listed on the dark web, we provide suggestive evidence that ATT has driven up the price of data relevant for financial fraud and identity theft, potentially by reducing the supply of stolen or hacked data. The listing data, assembled by a research organization focused on fraud, is described in detail in [Appendix H](#).

We use two snapshots of dark web listings from 2020 and 2023 to compare the prices of data generated from mobile app user activities with those of other listings, and of financial information with other categories, before and after ATT.⁴⁷ Applying a DiD design, we find that ATT is associated with higher prices for data likely generated from consumers’ mobile activities and for financial information, relative to other products (Panel b of [Table 9](#)). This pattern is consistent with ATT reducing the availability of such data on the dark web.

5.3.3 Adaptation by Fraudsters

An important question is whether fraudsters adapt to ATT by reallocating their activities across platforms or toward alternative forms of fraud. For example, fraudsters can shift their targeting from iOS to Android users. This matters for both estimation and welfare. From an identification perspective, such reallocation could lead to overstating the treatment effect if reductions on iOS are offset by increases on Android. From a welfare perspective, it could imply a redistribution of harm toward potentially more vulnerable populations.

Our evidence suggests that such reallocation is not the primary force behind our findings. First, Column 2 of [Table 9](#) Panel a shows large declines in platform-agnostic fraud categories. These categories rely on data available from both iOS and Android and should therefore be unaffected if fraud were simply reallocated across platforms. Second, Panel b shows an increase in mobile footprint data prices, consistent with a contraction in the overall supply of usable data rather than substitution across platforms. Third, although directly estimating spillover effects is challenging, we present evidence that is inconsistent with both within-zip-code and within-firm spillovers. If within-zip-code spillovers were driving ATT’s effects, we would expect the largest effects in Android-heavy areas, where fraudulent activity could be

⁴⁶As with other work using dark web posts, some categories inevitably fall into a gray area. Even in the computer science literature, where more sophisticated classification algorithms are employed, researchers continue to face challenges in drawing sharp boundaries. In our case, the results are not sensitive to reassignments of these fuzzy categories (Column 5 of [Table I.9](#)).

⁴⁷Because only two cross-sectional snapshots are available, we cannot analyze changes in the number of listings over time and instead focus solely on the price. The detailed regression specification and results are provided in [Appendix H](#).

redirected; instead, as shown in [Appendix Figure I.3](#), the effects are strongest in iOS-heavy regions. At the firm level, [Section 5.2](#) shows reductions in overall complaints for firms with a mobile presence, providing no evidence of substitution from iOS- to Android-using customers within firms. Finally, our alternative identification strategy in [Section 4.3](#), based on variation in ISP data practices, does not rely on iOS–Android differences and yields consistent results.

Conceptually, such reallocation is likely limited. iOS users tend to be higher-value targets with larger financial accounts and are not easily substituted, while shifting toward Android requires rebuilding cross-platform data linkages and targeting infrastructure. Moreover, with Google already signaling tighter privacy restrictions around 2021—such as limits on advertising IDs and the rollout of its Privacy Sandbox on Android—the return to investing in Android-based data collection was uncertain, reducing incentives for fraudsters to adapt.

We also find little evidence of substitution into other technology-driven fraud categories, such as zero-day exploits, malware, or botnets: as shown in Column (3) of [Table 9](#) Panel a, these categories exhibit near-zero effects. While it is possible that fraudsters shift toward other illicit markets (e.g., drugs, human trafficking, or weapons), these operate in distinct ecosystems and serve as our control group, making direct testing infeasible. We argue that such substitution is unlikely, as these activities require substantially different expertise and networks than data-driven fraud.

While our listings and forum data provide an important glimpse into upstream markets, we caution that they likely represent only the tip of the iceberg. More sophisticated actors operate in closed or invitation-only markets, conduct transactions via encrypted communication channels, or route sales through intermediaries to avoid detection.⁴⁸ There are also gray areas where fraudsters engage in transactions of compromised data through seemingly legitimate channels, creating the appearance of lawful business activity while still trafficking compromised information.⁴⁹

6 Conclusion

In this paper, we show that Apple’s App Tracking Transparency (ATT) policy, that greatly limited the tracking of personal data across apps and websites, reduced financial fraud in areas where more iOS users are present. Our findings shed light on the full fraud supply chain and highlight how data privacy regulation benefits consumers by constraining illicit data flows. We interpret our estimates as capturing the effects of tighter constraints on data collection and sharing practices. Reductions in fraud arise not from suppressing online

⁴⁸For example, on many dark web platforms, user-to-user transactions reveal only minimal metadata, such as the timestamps and user-submitted ratings, while concealing the actual content or price of the transaction.

⁴⁹The FTC charged data brokers (Gen X) for selling sensitive financial information to scammers who used it to drain consumer accounts, under the cover of routine data market activity. The scammers operated through a seemingly legitimate business, Ideal Financial Solutions Inc. See the FTC newsletter [here](#).

engagement — neither practical nor desirable — but from altering how data is collected, processed, shared, and sold. These findings suggest that improving data practices and their regulation, rather than restricting digital consumption, have the potential to reduce fraud and generate spillovers by disrupting the broader illicit data economy that sustains financial fraud. We caveat that, although the policy is not designed to curb digital activity, such activity could nonetheless be affected—for instance, if restrictions on data use reduce app monetization or, conversely, if greater consumer trust encourages higher engagement (Armantier et al., 2024). Assessing the broader equilibrium consequences of improving privacy standards for online activity and welfare remains an important task for future research.

References

- Abis, S., H. Tang, and B. Bian (2025). Breaking the Data Chain: The Ripple Effect of Data Sharing Restrictions on Financial Markets. *Available at SSRN 5334566*.
- Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics* 14(4), 218–256.
- Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2025). When Big Data Enables Behavioral Manipulation. *American Economic Review: Insights* 7(1), 19–38.
- Acquisti, A., C. Taylor, and L. Wagman (2016). The economics of privacy. *Journal of economic Literature* 54(2), 442–492.
- Acquisti, A. and H. R. Varian (2005). Conditioning prices on purchase history. *Marketing Science* 24(3), 367–381.
- Agrawal, A., J. Gans, and A. Goldfarb (2022). *Prediction machines, updated and expanded: The simple economics of artificial intelligence*. Harvard Business Press.
- Alves, L. M. and S. R. Wilson (2008). The effects of loneliness on telemarketing fraud vulnerability among older adults. *Journal of elder abuse & neglect* 20(1), 63–85.
- Anderson, K. B. (2019). Mass-Market Consumer Fraud in the United States: A 2017 Update. *Federal Trade Commission. Washington, DC*.
- Aridor, G., Y.-K. Che, and T. Salz (2023). The Effect of Privacy Regulation on the Data Industry: Empirical Evidence from GDPR. *RAND Journal of Economics* 54(4).
- Armanter, O., S. Doerr, J. Frost, A. Fuster, and K. Shue (2021). Whom do Consumers Trust with their Data? US Survey Evidence. Technical report, Bank for International Settlements.
- Armanter, O., S. Doerr, J. Frost, A. Fuster, and K. Shue (2024). Nothing to hide? gender and age differences in willingness to share data. *Gender and age differences in willingness to share data (April 26, 2024). Swiss Finance Institute Research Paper* (24-99).
- Avarikioti, G., R. Brunner, A. Kiayias, R. Wattenhofer, and D. Zindros (2018). Structure and Content of the Visible Darknet. *arXiv preprint arXiv:1811.01348*.
- Babina, T., S. Bahaj, G. Buchak, F. De Marco, A. Foulis, W. Gornall, F. Mazzola, and T. Yu (2025). Customer Data Access and Fintech Entry: Early Evidence from Open Banking. *Journal of Financial Economics* 169, 103950.
- Begley, T. A. and A. Purnanandam (2021). Color and Credit: Race, Regulation, and the

- Quality of Financial Services. *Journal of Financial Economics* 141(1), 48–65.
- Beraja, M., D. Y. Yang, and N. Yuchtman (2023). Data-intensive innovation and the state: Evidence from AI firms in china. *The Review of Economic Studies* 90(4), 1701–1723.
- Beraja, M. and N. Yuchtman (2025). Generalized disruption: Society, work, and property rights in the age of AI. *NBER Chapters*.
- Bergemann, D. and A. Bonatti (2019). Markets for information: An introduction. *Annual Review of Economics* 11(1), 85–107.
- Bergemann, D., B. Brooks, and S. Morris (2015). The limits of price discrimination. *American Economic Review* 105(3), 921–957.
- Bertrand, M. and E. Kamenica (2023). Coming Apart? Cultural Distances in the United States over Time. *American Economic Journal: Applied Economics* 15(4), 100–141.
- Bessen, J. E., S. M. Impink, L. Reichensperger, and R. Seamans (2020). GDPR and the Importance of Data to AI Startups. Working paper, New York University, Boston University.
- Bian, B., Q. Huang, Y. Li, and H. Tang (2024). Data as a Networked Asset. *Available at SSRN 5183829*.
- Bian, B., X. Ma, and H. Tang (2021). The Supply and Demand for Data Privacy: Evidence from Mobile Apps. *Available at SSRN 3987541*.
- Bonatti, A. and G. Cisternas (2020). Consumer scores and price discrimination. *The Review of Economic Studies* 87(2), 750–791.
- Carlin, B., T. Umar, and H. Yi (2023). Deputizing financial institutions to fight elder abuse. *Journal of Financial Economics* 149(3), 557–577.
- Cheyre, C., B. T. Leyden, S. Baviskar, and A. Acquisti (2023). The Impact of Apple’s App Tracking Transparency Framework on the App Ecosystem. *Available at SSRN 4453463*.
- Choi, J. P., D.-S. Jeon, and B.-C. Kim (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics* 173, 113–124.
- DeLiema, M. (2018). Elder fraud and financial exploitation: Application of routine activity theory. *The Gerontologist* 58(4), 706–718.
- DeLiema, M., Z. D. Gassoumis, D. C. Homeier, and K. H. Wilber (2012). Determining prevalence and correlates of elder abuse using promotores: Low-income immigrant latinos report high rates of abuse and neglect. *Journal of the American Geriatrics Society* 60(7), 1333–1339.
- DeLiema, M., G. R. Mottola, and M. Deevy (2017). Findings from a Pilot Study to Measure Financial Fraud in the United States. *Available at SSRN 2914560*.
- Demirer, M., D. J. J. Hernández, D. Li, and S. Peng (2024). Data, Privacy Laws and Firm Production: Evidence from the GDPR. Technical report, National Bureau of Economic Research.
- Doerr, S., L. Gambacorta, L. Guiso, and M. Sanchez del Villar (2023). Privacy Regulation and Fintech Lending. *Available at SSRN 4353798*.
- Dou, Y. and Y. Roh (2024). Public Disclosure and Consumer Financial Protection. *Journal of Financial and Quantitative Analysis* 59(5), 2164–2198.
- Egan, M., G. Matvos, and A. Seru (2019). The Market for Financial Adviser Misconduct. *Journal of Political Economy* 127(1), 233–295.
- Egan, M., G. Matvos, and A. Seru (2025). Arbitration with Uninformed Consumers. *Review of Economic Studies*, Forthcoming.

- Goldberg, S. G., G. A. Johnson, and S. K. Shriver (2024). Regulating Privacy Online: An Economic Evaluation of the GDPR. *American Economic Journal: Economic Policy* 16(1), 325–358.
- Griffin, J. M. and S. Kruger (2024). What is forensic finance? *Foundations and Trends® in Accounting* 14(3), 137–243.
- Haendler, C. (2022). Keeping Up in the Digital Era: How Mobile Technology is Reshaping the Banking Sector. *Available at SSRN 4287985*.
- Haendler, C. and R. Heimer (2025). The Hidden Costs of Financial Services: Consumer Complaints and Financial Restitution. *Available at SSRN*.
- Huang, R., J. S. Linck, E. J. Mayer, and C. Parsons (2024). Can Human Capital Explain Income-Based Disparities in Financial Services. *Review of Financial Studies*, 22–16.
- Huff, R., C. Desilets, and J. Kane (2010). The 2010 National Public Survey on White Collar Crime. *National White Collar Crime Center* 44.
- Ichihashi, S. (2021). The economics of data externalities. *Journal of Economic Theory* 196, 105316.
- James, B. D., P. A. Boyle, and D. A. Bennett (2014). Correlates of susceptibility to scams in older adults without dementia. *Journal of elder abuse & neglect* 26(2), 107–122.
- Janssen, R., R. Kesler, M. Kummer, and J. Waldfogel (2021). GDPR and the Lost Generation of Innovative Apps. NBER Working Paper 146409, University of Zurich, University of Minnesota, University of East Anglia, Georgia Institute of Technology.
- Jia, J., G. Z. Jin, and L. Wagman (2021). The Short-Run Effects of the General Data Protection Regulation on Technology Venture Investment. *Marketing Science* 40(4), 661–684.
- Jiang, E. X., Y. G. Yu, and J. Zhang (2025). Bank Competition amid Digital Disruption: Implications for Financial Inclusion. *Journal of Finance*, Forthcoming.
- Johnson, G. A. (2024). 4. economic research on privacy regulation: Lessons from the gdpr and beyond. In *The economics of privacy*, pp. 97–126. University of Chicago Press.
- Jones, C. I. and C. Tonetti (2020). Nonrivalry and the Economics of Data. *American Economic Review* 110(9), 2819–58.
- Jou, J., A. Kleymenova, A. Passalacqua, L. Sándor, and R. Vijayaraghavan (2024). Disciplining Banks through Disclosure: Evidence from CFPB Consumer Complaints. *Available at SSRN*.
- Kesler, R. (2022). The Impact of Apple’s App Tracking Transparency on App Monetization. *Available at SSRN 4090786*.
- Koont, N. (2023). The Digital Banking Revolution: Effects on Competition and Stability. *Available at SSRN 4624751*.
- Koont, N., T. Santos, and L. Zingales (2024). Destabilizing Digital “Bank Walks”. Technical report, National Bureau of Economic Research.
- Kraft, L., B. Skiera, and T. Koschella (2023). Economic Impact of Opt-In versus Opt-Out Requirements for Personal Data Usage: The Case of Apple’s App Tracking Transparency (ATT). *Available at SSRN 4598472*.
- Li, X. (2023). Does the Disclosure of Consumer Complaints Reduce Racial Disparities in the Mortgage Lending Market. *Available at SSRN 4741819*.
- Li, Z., H. Ning, F. Jing, and M. N. Lessani (2024). Understanding the Bias of Mobile Location Data across Spatial Scales and over Time: a Comprehensive Analysis of SafeGraph Data

- in the United States. *Plos one* 19(1).
- Lichtenberg, P. A., L. Stickney, and D. Paulson (2013). Is psychological vulnerability related to the experience of fraud in older adults? *Clinical Gerontologist* 36(2), 132–146.
- Liu, Z., M. Sockin, and W. Xiong (2023). Data privacy and algorithmic inequality. Technical report, National Bureau of Economic Research.
- Mazur, L. (2024). The Impact of the CFPB Complaint Database Disclosure on Banks’ Provision of Mortgage Credit to Low-and Moderate-Income Communities.
- Peukert, C., S. Bechtold, M. Batikas, and T. Kretschmer (2022). Regulatory Spillovers and Data Governance: Evidence from the GDPR. *Marketing Science* 41(4), 746–768.
- Rafieian, O. and H. Yoganarasimhan (2021). Targeting and privacy in mobile advertising. *Marketing Science* 40(2), 193–218.
- Ramadorai, T., A. Uettwiller, and A. Walther (2025). Privacy policies and consumer data extraction: Evidence from us firms. *Review of Finance*, Forthcoming.
- Raval, D. (2020a). Which Communities Complain to Policymakers? Evidence from Consumer Sentinel. *Economic Inquiry* 58(4), 1628–1642.
- Raval, D. (2020b). Whose Voice do We Hear in the Marketplace? Evidence from Consumer Complaining Behavior. *Marketing Science* 39(1), 168–187.
- Sweeting, A., D. J. Balan, N. Kreisler, M. T. Panhans, and D. Raval (2020). Economics at the ftc: fertilizer, consumer complaints, and private label cereal. *Review of Industrial Organization* 57, 751–781.
- Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics*, 631–650.
- Tirole, J. (2021). Digital dystopia. *American Economic Review* 111(6), 2007–2048.
- Wernerfelt, N., A. Tuchman, B. T. Shapiro, and R. Moakler (2025). Estimating the value of offsite tracking data to advertisers: Evidence from meta. *Marketing Science* 44(2), 268–286.
- Wu, X. (2023). Mobile App, Firm Risk, and Growth. Available at SSRN 4519061.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs.

Figure 1: An Illustration of Fraud Pipeline w/ Data Collected and Merged via Mobile Apps





Figure 2: Examples of App ATT Prompt and Privacy Nutrition Labels


a. Mint ATT Prompt b. Mint Privacy Labels c. Tiktok Privacy Labels




Apple’s App Tracking Transparency (ATT) policy was introduced on April 26, 2021. Panel a shows the ATT prompt through which apps (Intuit’s Mint in this example) could get user permission to obtain mobile identifiers that allow them to track, share, or sell consumer data with or to other apps and websites. Panel b shows Mint’s privacy label, which describes how mobile identifiers are used to track consumers and what data is collected and linked through data sharing. Panel c shows TikTok’s privacy label, which describes that TikTok uses mobile identifiers to obtain, e.g., financial data of consumers.

Figure 3: Number of Complaints Scaled by Population and iOS Share by Zip Code

a. # Complaints per Million Residents




b. iOS Share




This figure presents the number of CFPB complaints per million residents (Panel a) and the average iOS share prior to the ATT rollout (Panel b), across zip codes over the full sample period (January 2019 to June 2022). For readability, Panel a displays CFPB complaints per million residents; however, our empirical analyses use CFPB complaints per 1,000 residents.

Figure 4: Dynamic Effects of ATT on CFPB Complaints

a. Any Complaints



b. # Complaints per 1,000 Residents



This figure illustrates the dynamic effect of ATT on CFPB complaints around the implementation of ATT (April 26, 2021). Quarter -1 is the omitted category. All three months in a quarter are grouped. Panel a uses an indicator outcome (LPM), and Panel b uses complaints per 1,000 residents (PPML). Coefficients interact timing indicators with pre-ATT iOS share. Zip-code fixed effects and county $\times$ year-month fixed effects are included. Standard errors are clustered at the state level.

Table 1: Summary Statistics: Zip code-Month Level

variable	mean	sd	p25	p50	p75	count
<i>ATT sample (2019m1–2022m6)</i>						
Any complaints (0/1)	0.26	0.44	0.00	0.00	1.00	927,780
# Complaints per 1,000 residents	0.05	0.15	0.00	0.00	0.03	927,780
iOS share	0.46	0.11	0.39	0.45	0.52	927,780
<i>FCC Broadband Rule sample (2015m1–2018m12)</i>						
Any complaints (0/1)	0.21	0.40	0.00	0.00	0.00	1,166,496
# Complaints per 1,000 residents	0.03	0.09	0.00	0.00	0.00	1,166,496
Exposure (data collection)	0.93	0.78	0.04	0.85	1.51	1,166,496
Exposure (data sharing)	0.15	0.14	0.00	0.13	0.26	1,166,496

This table presents summary statistics for our key explanatory and outcome variables at the zip code-month levels. The top part shows the summary statistics for the main sample that spans January 2019 to June 2022, which we use to examine the effect of the ATT. The bottom part shows the summary statistics for the sample that spans January 2015 to December 2018, which we use to examine the effect of the 2017 Repeal of the FCC Broadband Rule. We primarily examine two outcome variables constructed from the CFPB complaints: an indicator for whether a zip code has at least one complaint (*Any complaints (0/1)*) and the total number of complaints per 1,000 residents (*# Complaints per 1,000 residents*). *iOS share* denotes the share of iOS devices in a given zip code. *Exposure (data sharing)* refers to the market-share-weighted number of data SDKs used by ISPs in their mobile applications in a given zip code. *Exposure (data collection)* represents the market-share-weighted number of data types collected by the top 12 fixed broadband ISPs, as disclosed in their privacy labels, in a given zip code.

Table 2: Effect of ATT on Consumer Complaints - CFPB

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post $\times$ iOS share	-0.036* (0.018)	-0.059*** (0.015)	-0.548*** (0.095)	-0.621*** (0.094)
Zip code FE	✓	✓	✓	✓
State $\times$ Year-month FE	✓		✓	
County $\times$ Year-month FE		✓		✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.264	0.264	0.053	0.053
Magnitude ($\uparrow$ 1 SD iOS share)	-1.4%	-2.4%	-5.4%	-6.1%
Observations	926,772	926,772	926,772	926,772
R ² / Pseudo R ²	0.441	0.432	0.094	0.134

This table presents the estimated effects of ATT on CFPB complaints. The unit of observation is at the zip-code-month level. Columns 1 and 2 report estimates from a linear probability model, where the outcome is a binary indicator equal to one if at least one complaint is filed in a zip code during a given month. Columns 3 and 4 report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 to June 2022. All columns include zip code fixed effects. The odd-numbered columns include state $\times$ year-month fixed effects, while the even-numbered columns include county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 3: Effect of 2017 Repeal of FCC Broadband Privacy Rule on CFPB Complaints**Panel a. Local Exposure Based on ISP Data Collection Intensity**

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post Repeal $\times$ Exposure (data collection)	0.008*** (0.002)	0.007*** (0.002)	0.074*** (0.014)	0.058*** (0.022)
Zip code FE	✓	✓	✓	✓
State $\times$ Year-month FE	✓		✓	
County $\times$ Year-month FE		✓		✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.207	0.207	0.034	0.034
Magnitude ($\uparrow$ 1 SD in exposure)	2.8%	2.4%	5.7%	3.9%
Observations	1,166,496	1,166,496	1,166,496	1,166,496
R ² / Pseudo R ²	0.406	0.460	0.061	0.109

Panel b. Local Exposure Based on ISP Data Sharing Intensity

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post Repeal $\times$ Exposure (data sharing)	0.044*** (0.008)	0.038*** (0.011)	0.414*** (0.059)	0.311*** (0.102)
Zip code FE	✓	✓	✓	✓
State $\times$ Year-month FE	✓		✓	
County $\times$ Year-month FE		✓		✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.207	0.207	0.034	0.034
Magnitude ($\uparrow$ 1 SD in exposure)	2.9%	2.5%	6.1%	3.9%
Observations	1,166,496	1,166,496	1,166,496	1,166,496
R ² / Pseudo R ²	0.406	0.460	0.061	0.109

This table presents the estimated effects of the 2017 FCC Broadband Privacy Rule repeal on CFPB complaints. The unit of observation is at the zip-code-month level. Columns 1 and 2 report estimates from a linear probability model, where the outcome is a binary indicator equal to one if at least one complaint is filed in a zip code during a given month. Columns 3 and 4 report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. Panel a uses the market-share-weighted number of data types collected by ISPs, as disclosed in their privacy labels, as the measure of exposure to the shock. Panel b uses the market-share-weighted number of data SDKs used by ISPs in their mobile applications. The sample period is January 2015 to December 2018. All columns include zip code fixed effects. The odd-numbered columns include state $\times$ year-month fixed effects, while the even-numbered columns include county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 4: Effect of ATT on Consumer Complaints -
Consumer Sentinel Network and FTC Identity Theft Databases

	# Complaints per 1,000 residents				
	(1) Uncensored CFPB	(2) Identity Theft	(3) CSN	(4) Any loss	(5) Loss $\geq$ \$1k
Post $\times$ iOS share	-0.575*** (0.085)	-0.255** (0.124)	-0.114*** (0.044)	-0.108* (0.064)	-0.183** (0.093)
Zip code FE	✓	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓	✓
Model	PPML	PPML	PPML	PPML	PPML
Mean outcome var.	0.125	0.223	0.733	0.165	0.062
Magnitude ($\uparrow$ 1 SD iOS share)	-5.6%	-2.5%	-1.1%	-1.1%	-1.8%
Observations	757,635	955,332	1,131,115	1,024,758	825,936
Pseudo R ²	0.209	0.240	0.143	0.099	0.096

This table presents the estimated effects of ATT on consumer complaints using proprietary internal data from the Consumer Sentinel Network (CSN) and FTC Identity Theft databases. The unit of observation is at the zip-code-month level. In all specifications, we estimate a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. Column 1 uses all uncensored CFPB complaints from the Consumer Sentinel Network. Column 2 uses all Identity Theft complaints. Column 3 uses all Consumer Sentinel Network complaints. Column 4 uses Consumer Sentinel Network complaints that report a positive dollar loss. Column 5 uses Consumer Sentinel Network complaints that report a loss greater than \$1,000. The sample period is January 2019 (February 2019 for Identity Theft) to June 2022. All columns include zip code and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 5: Effect of ATT on Consumer Complaints - CFPB
Addressing COVID-19 Related Concerns

Panel a. Controlling for Exposure to EIP

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post $\times$ iOS share	-0.073*** (0.017)	-0.056*** (0.016)	-0.684*** (0.098)	-0.483*** (0.090)
Post $\times$ EIP amount	-0.012*** (0.003)		-0.015 (0.011)	
Post $\times$ Total income		-0.006*** (0.002)		-0.019** (0.008)
Post $\times$ Child care credit		0.001 (0.002)		-0.022* (0.013)
Constant	0.291*** (0.003)	0.288*** (0.002)	-2.084*** (0.019)	-2.120*** (0.018)
Zip code FE	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.280	0.280	0.070	0.070
Magnitude ($\uparrow$ 1 SD iOS share)	-2.7%	-2.1%	-6.6%	-4.7%
Observations	876,414	876,414	876,288	876,288
R ² / Pseudo R ²	0.512	0.512	0.133	0.133

Panel b. Placebo Tests

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post-EIP 1st $\times$ iOS share	0.005 (0.014)		-0.170 (0.146)	
Post-EIP 2nd $\times$ iOS share		0.007 (0.019)		0.059 (0.182)
Zip code FE	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.250	0.250	0.065	0.065
Magnitude ($\uparrow$ 1 SD iOS share)	0.2%	0.3%	-1.7%	0.6%
Observations	617,848	617,848	617,848	617,848
R ² / Pseudo R ²	0.504	0.504	0.131	0.131

This table presents the estimated effects of ATT on CFPB complaints by controlling for COVID-19 stimulus payments (Panel a) and reports results from placebo tests (Panel b). The unit of observation is at the zip-code-month level. Columns 1 and 2 report estimates from a linear probability model, where the outcome is a binary indicator equal to one if at least one complaint is filed in a zip code during a given month. Columns 3 and 4 report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. Panel a controls for exposure to economic impact payments (EIPs) by adding the interaction terms between the post-ATT indicator and variables related to EIPs in the regression. The variables include total amount of EIPs received (Columns 1 and 3), the average household income, and the amount of childcare credit (Columns 2 and 4), all constructed using IRS data. Panel b considers placebo tests using two alternative event dates: April 2020 and December 2020. These dates correspond to the disbursement of the first and second major waves of pandemic-related financial support under the CARES Act and subsequent legislation. The sample period is January 2019 to June 2022 in Panel a. In Panel b, the sample is restricted to the pre-ATT period (January 2019-April 2021) to avoid contamination from ATT's actual treatment effect. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 6: Effect of ATT on Consumer Complaints
More vs. Less Relevant Categories of Complaints—Classified by Narratives

Panel a. CFPB

	# Complaints per 1,000 residents			
	Top 2 Fraud Categories		Bottom 2 Fraud Categories	
	(1) Credit Reporting and Repair	(2) Debt Collection	(3) Student Loan	(4) Mortgage
Post $\times$ iOS Share	-0.574*** (0.163)	-0.600*** (0.178)	0.019 (0.597)	-0.285 (0.209)
Zip code FE	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓
Model	PPML	PPML	PPML	PPML
Mean outcome var.	0.022	0.008	0.002	0.005
Magnitude ($\uparrow$ 1 SD iOS share)	-5.5%	-5.7%	0.2%	-2.8%
Observations	782,292	782,292	782,292	782,292
Pseudo R ²	0.138	0.088	0.059	0.072

Panel b. Identity Theft and Consumer Sentinel Network

	# Complaints per 1,000 residents			
	Narratives w/ Keywords		Narratives w/o Keywords	
	(1) Identity Theft	(2) CSN	(3) Identity Theft	(4) CSN
Post $\times$ iOS Share	-0.745*** (0.154)	-0.274*** (0.060)	-0.001 (0.142)	-0.042 (0.043)
Zip code FE	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓
Model	PPML	PPML	PPML	PPML
Mean outcome var.	0.106	0.276	0.139	0.472
Magnitude ($\uparrow$ 1 SD iOS share)	-7.2%	-2.7%	0.0%	-0.4%
Observations	800,386	1,034,586	880,678	1,089,791
Pseudo R ²	0.199	0.137	0.214	0.117

This table presents the estimated effects of ATT on consumer complaints after classifying complaints by their likelihood to be affected by ATT. The unit of observation is at the zip-code-month level. The outcome variable is the number of complaints per 1,000 residents. Panel a considers the top 2 relevant and bottom 2 relevant product categories in CFPB complaints: Credit Reporting and Credit Repair Services (Top 1), Debt Collection (Top 2), Student Loans (Bottom 2), and Mortgages (Bottom 1). Panel b leverages the narratives available in the Identity Theft and Consumer Sentinel Network complaints to classify complaints into those with and without any relevant words in the narrative. All columns report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 (February 2019 for Identity Theft) to June 2022. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 7: Effect of ATT on Consumer Complaints
Heterogeneity by Local Exposure to Data Practices of Banks and ISPs

Panel a. Data Practices of Banks

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post $\times$ iOS share	-0.059*** (0.015)	-0.060*** (0.015)	-0.597*** (0.091)	-0.603*** (0.094)
Post $\times$ iOS share $\times$ Exposure (data collection)	-0.042*** (0.012)		-0.167* (0.087)	
Post $\times$ iOS share $\times$ Exposure (data sharing)		-0.040*** (0.014)		-0.189** (0.086)
Zip code FE	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.267	0.267	0.068	0.068
Magnitude ($\uparrow$ 1 SD iOS share)	-2.4%	-2.4%	-5.8%	-5.9%
Magnitude ($\uparrow$ 1 SD iOS share $\times$ Exposure)	-1.7%	-1.6%	-1.7%	-1.9%
Observations	925,428	925,428	925,428	925,428
R ² /Pseudo R ²	0.514	0.514	0.131	0.131

Panel b. Data Practices of ISPs

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post $\times$ iOS share	-0.060*** (0.015)	-0.059*** (0.015)	-0.604*** (0.095)	-0.598*** (0.092)
Post $\times$ iOS share $\times$ Exposure (data collection)	-0.040*** (0.014)		-0.197** (0.087)	
Post $\times$ iOS share $\times$ Exposure (data sharing)		-0.042*** (0.012)		-0.174* (0.089)
Zip code FE	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.267	0.267	0.069	0.069
Magnitude ($\uparrow$ 1 SD iOS share)	-2.4%	-2.4%	-5.9%	-5.9%
Magnitude ($\uparrow$ 1 SD iOS share $\times$ Exposure)	-1.6%	-1.7%	-2.0%	-1.7%
Observations	925,428	925,428	925,428	925,428
R ² / Pseudo R ²	0.514	0.514	0.134	0.134

This table reports the heterogeneous effects of ATT on consumer complaints by local exposure to the data practices of banks and ISPs. The unit of observation is at the zip-code-month level. Columns 1 and 2 of both panels report estimates from a linear probability model, where the outcome is a binary indicator equal to one if at least one complaint is filed in a zip code during a given month. Columns 3 and 4 report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The variables used in the triple interaction terms capture heterogeneity in the data collection practices of local banks and ISPs across counties. In Panel a, *Exposure (data collection/sharing)* represents the download-weighted number of data types collected or data SDKs used by all banks that own a mobile app as of April 2021. The precise definitions are provided in Equation (5). In Panel b, *Exposure (data sharing)* refers to the market-share-weighted number of data SDKs used by ISPs in their mobile applications. *Exposure (data collection)* represents the market-share-weighted number of data types collected by the top 12 fixed broadband ISPs, as disclosed in their privacy labels. The sample period is January 2019 to June 2022. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 8: Effect of ATT on Firm-Level Cyber Incidents and Consumer Complaints**Panel a. Cyber Incidents**

	Any cyber incidents (0/1)			
	(1) All	(2) Breach/Data Misuse	(3) Debt Collection/Credit Reporting Regulation	(4) Other Causes
Has an app $\times$ Post	-0.049*** (0.010)	-0.046*** (0.009)	-0.009*** (0.003)	-0.003 (0.004)
Firm FE	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓
Mean outcome var.	0.097	0.072	0.015	0.021
Observations	306,933	306,933	306,933	306,933
R ²	0.183	0.185	0.198	0.123

Panel b. Consumer Complaints

	Any complaints (0/1)				
	(1) All	(2) Credit Reporting and Repair	(3) Debt Collection	(4) Student Loan	(5) Mortgage
Has an app $\times$ Post	-0.021*** (0.006)	-0.031*** (0.005)	-0.024*** (0.005)	-0.001 (0.001)	-0.001 (0.003)
Firm FE	✓	✓	✓	✓	✓
Year-month FE	✓	✓	✓	✓	✓
App-size-specific linear trend	✓	✓	✓	✓	✓
Mean outcome var.	0.037	0.037	0.079	0.006	0.030
Observations	394,243	394,243	394,243	394,243	394,243
R-square	0.396	0.316	0.352	0.423	0.475

This table presents the estimated effects of ATT on cyber incidents (Panel a) and firm-level consumer complaints received via CFPB (Panel b). The unit of observation is at the firm-month level. In Panel a, the outcome variables are an indicator variable for whether the firm was exposed to (1) any cyber incident, (2) cyber incidents that were caused by data breach or data misuse, (3) cyber incidents that violated the Fair Debt Collection Practices Act or the Fair Credit Reporting Act, and (4) cyber incidents that were caused by other reasons unrelated to data breach. All columns report linear probability model estimates for the interaction term between the post-ATT indicator and an indicator for whether the firm owns an app. In Panel b Column 1, the outcome variable is an indicator variable for whether the firm was mentioned in a CFPB complaint. In Panel b Column 2-5, the outcome variable is an indicator for whether the firm was mentioned in the two most relevant and two least relevant product categories in CFPB complaints: Credit Reporting and Credit Repair Services (Top 1), Debt Collection (Top 2), Student Loan (Bottom 2), and Mortgage (Bottom 1). All columns include firm fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the firm level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 9: Effect of ATT on Dark Web Activities**Panel a. Posts on Dark Web Forums**

	# Posts		
	(1)	(2)	(3)
Post $\times$ Treated	-0.413*** (0.024)		
Post $\times$ Partially treated		-0.241* (0.126)	
Post $\times$ Placebo			-0.006 (0.173)
Label FE	✓	✓	✓
Forum $\times$ Year-month FE	✓	✓	✓
Model	PPML	PPML	PPML
Mean outcome var.	46.0	60.4	101.8
Magnitude (Treated vs. Control)	-33.8%	-21.5%	-0.6%
Observations	27,612	43,719	57,525
Pseudo R ²	0.871	0.863	0.879

Panel b. Listing Price of Data

	Listing price per unit (log)			
	Worldwide		US	
	(1)	(2)	(3)	(4)
Post $\times$ Mobile footprint data	0.474*** (0.090)		0.520*** (0.080)	
Post $\times$ Financial info data		0.470*** (0.145)		0.322* (0.173)
Firm FE	✓	✓	✓	✓
Year FE	✓	✓	✓	✓
Currency FE	✓	✓		
Mean outcome var.	3.025	3.025	2.345	2.345
Observations	3,938	3,938	2,703	2,703
R ²	0.862	0.863	0.636	0.636

This table presents the estimated effects of ATT on upstream dark web activities. Panel a focuses on posts on dark web forums and the unit of observation is at forum-label-month level. The full list of labels and their exposure to ATT can be found in [Appendix G](#). The outcome variable is the number of posts. Column 1 examines the explicitly treated group (Apple/iOS-related posts), Column 2 the partially treated group (categories reliant on personal or financial data such as tracking, breaches, identity theft, financial data theft), and Column 3 a placebo group (e.g., zero-days, malware, botnets). All groups are compared against a control group of illicit trade categories (e.g., drugs, human trafficking, weapons) operating in separate markets. Standard errors are double clustered at the forum and label level and reported in parentheses. The sample period is January 2015 to December 2022. Panel b reports the effect of ATT on the price of data sold in the Dark Web. The unit of observation is a listing on the Dark Web, with details described in [Appendix H](#). In Columns 1 and 3, we interact the post-ATT indicator with an indicator for whether the listing sells data that is likely generated from consumers' mobile activities. In Columns 2 and 4, we interact the post-ATT indicator with an indicator for whether the listing sells financial information. The outcome variable is the logarithm of the listing price per unit. One snapshot of listings is in 2020 and the other one is in 2023. Columns 1-2 use all the listings, while Columns 3 and 4 use only listings in which the price is quoted in USD. Company and year fixed effects are included in all columns. We additionally include currency fixed effects in Columns 1 and 2. Standard errors clustered at the company level (i.e., the company from which the data originate in each listing) are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Internet Appendix to “Consumer Surveillance and Financial Fraud”


This Appendix has 10 sections. Section [A](#) contains information on policy shocks in the mobile app economy and the construction of our data collection/sharing intensity variables. Section [B](#) discusses irregularities in the CFPB complaints data and methods to deal with outliers. Section [C](#) covers background information on the 2017 repeal of the FCC broadband privacy rule. Section [D](#) provides motivations and descriptions of additional robustness checks. Section [E](#) presents the classification methods of fraud-related complaints and example narratives on data-driven fraud incidents. Section [F](#) contains information on data on cyber incidents from Advisen. Section [G](#) covers information on the posts from dark web forums, including classification methods, summary statistics, and examples. Section [H](#) describes the data on dark net listings related to data and the analysis. Section [I](#) covers additional analysis, including tables and figures.

A	Background: Data Collection and Sharing in the Mobile Economy	A-2
A.1	The App Tracking Transparency and Privacy Nutrition Label Policies	A-2
A.2	Measuring Data Collection Intensity via Privacy Nutrition Labels	A-2
A.3	Measuring Data Sharing Intensity via Third-Party Data Sharing SDKs . . .	A-3
B	Dealing with Irregularities in CFPB Complaints Data	A-6
C	2017 FCC Privacy Rule Repeal—Timeline and Nature of the Shock	A-9
D	Measurement Error, Alternative Aggregation, and Additional Checks	A-11
E	Classification of Fraud-related Complaints	A-12
F	Description of Advisen Data	A-15
G	Dark Web Forum Posts: Classification, Statistics, and Examples	A-16
H	Dark Web Listings for Data	A-21
I	Additional Figures and Tables	A-22

A Background: Data Collection and Sharing in the Mobile Economy

A.1 The App Tracking Transparency and Privacy Nutrition Label Policies

Figure A.1: Stock Market Reactions around ATT



This figure is from [Bian et al. \(2021\)](#). This figure plots the average cumulative abnormal returns (CARs) around the implementation of the App Tracking Transparency Policy on April 26, 2021. The event window includes 30 days before and after the implementation date. CARs are computed using the Fama-French factor models.

A.2 Measuring Data Collection Intensity via Privacy Nutrition Labels

Following [Bian et al. \(2021\)](#), we use iOS app privacy nutrition labels to quantify the data collection intensity of firms that own these apps. Since December 14, 2020, Apple has required all app developers to disclose their data collection practices in a standardized and accessible format. An example of such a privacy label can be found at: <https://apps.apple.com/kg/app/facebook/id284882215>.

The standardized structure of these labels enables consistent comparison across firms. Specifically, we measure data collection intensity by counting the number of unique data types an app collects across all declared purposes (i.e., the second layer of the label). These data types correspond to the third-layer items illustrated in [Figure A.2](#), including categories such as identifiers, location, and financial information.

Identifying Mobile Presence of ISPs and Banks. For banks, we match the highest-download iOS app in SensorTower’s *Finance* category to FDIC-registered institutions using cleaned developer/official-website URLs (from SensorTower and the FDIC WEBADDR field)

and standardized institution names (stripping suffixes such as “THE”, “NATIONAL ASSOCIATION”, and “LTD”). This procedure identifies 2,618 banks that own an active mobile app as of the data collection date.


For ISPs, given the high concentration in ISP markets, we manually identify the 12 largest app-owning ISPs, listed in Panel B of [Figure A.3](#), which together accounted for 85% of the market as of June 2016. Since ISPs often operate multiple apps, we compute the average of relevant metrics across all apps owned by a given ISP. For example, Comcast owns 248 apps and Verizon 124, corroborating their central position in the digital surveillance economy.

A.3 Measuring Data Sharing Intensity via Third-Party Data Sharing SDKs

Following [Bian et al. \(2024\)](#), we construct a measure of data sharing intensity based on the use of third-party Software Development Kits (SDKs), using data from Apptopia. SDKs are pre-built software components that provide functionalities such as analytics and targeted advertising. We identify SDKs that facilitate data sharing across firms by linking user activity through unique identifiers, thereby enabling the creation of detailed consumer profiles. Firms integrate these SDKs into their apps to gain access to such profiles, which can inform marketing strategies, product development, and other business decisions.

We measure data sharing intensity as the number of data-sharing SDKs used in a firm’s iOS apps. When a firm owns multiple apps, we take the average number of such SDKs across its app portfolio. The distributions of data collection and sharing intensity across the top 12 ISPs and banks are provided in [Figure A.3](#).


Figure A.2: Apple Privacy Nutrition Label




This figure shows the structure of the mandatory privacy nutrition label from Apple for iOS apps. Apple privacy label has four layers. The first layer consists of three categories: *Data Used to Track You*, *Data Linked to You*, and *Data Not Linked to You*. If an app doesn't collect any data, it will have *Data Not Collected* as the only layer in its privacy label. For the second layer, only *Data Linked to You* and *Data Not Linked to You* have this layer, which shows 6 different purposes of data use. The third layer includes 14 different data types that the app collects; all data types can appear under each of the 6 purposes of data use in the second layer. The fourth layer reports 32 data items under the corresponding data types in the third layer. The first and the third layers are displayed on the main App Store page, while the second and the fourth layers are only displayed in a pop-up window when one clicks on the "See Details" button in the upper right corner of the App Privacy section.

Figure A.3: Distributions of Data Collection & Sharing Intensity


a. Banks



b. ISPs - Average Across Apps



c. ISPs - Total Across Apps



This figure presents the distribution of data collection intensity, measured by the number of data types collected (blue bars), and data sharing intensity, measured by the number of data-related SDKs used (orange bars). Panel (a) shows the distributions for banks, with the horizontal axis showing the number of data types collected or SDKs installed. Panels (b) and (c) display the average and total data collection and sharing intensity across app portfolios of the 12 ISPs with the largest market shares.

Identity Theft Complaints. The FTC also maintains a separate database of Identity Theft complaints that consumers file using different channels from the complaints in Consumer Sentinel (for example, by visiting identitytheft.gov instead of reportfraud.ftc.gov). The Identity Theft database contains complaints with broadly similar information to those in Consumer Sentinel. Like the Consumer Sentinel complaints, the Identity Theft complaints are non-public. Like the CFPB complaints, they are more focused on a specific issue (here, identity theft rather than financial problems) compared to Consumer Sentinel.

B Dealing with Irregularities in CFPB Complaints Data

This section documents our approach to identifying irregularities in complaint activity. Extreme complaint intensity in certain zip codes may arise from small population denominators, concentrated legal-aid activity, or localized service disruptions. Because such patterns often reflect structural or institutional factors rather than the broader variation of interest, we remove zip codes that exhibit persistently extreme levels or volatility in consumer complaints. These units can exert disproportionate influence on estimates and obscure broader patterns.

We show below that the excluded units differ markedly from the regression sample in both temporal behavior and socioeconomic composition, and tend to cluster in specific geographic areas with elevated complaint activity.

While one alternative would be to winsorize the outcome variable, we find that aggressive right-tail winsorization (e.g., at the 5% or 10% level) yields results similar to our main estimates. This suggests that the distortion is not driven by a few isolated extreme observations, but rather by a small subset of structurally atypical zip codes. Excluding these units therefore yields a cleaner and more interpretable estimation sample without materially affecting our substantive conclusions.


Procedure to Identify Outliers. To mitigate the influence of highly skewed complaint volumes, we identify and exclude zip codes based on four criteria, all constructed from population-scaled monthly complaint volumes and intended to capture persistent or structurally atypical patterns rather than transitory noise.

A zip code is flagged as an outlier if it satisfies at least two of the following: (i) the absolute month-over-month change exceeds the 90th percentile in more than 10 months; (ii) monthly complaint volume exceeds the 90th percentile in more than 10 months; (iii) the standard deviation of monthly complaints over the full sample exceeds the 90th percentile; and (iv) the average monthly complaint volume over the full sample exceeds the 90th percentile.

Applying this rule identifies 2,147 outlier zip codes. Our main regression sample, after excluding these units, contains 24,237 zip codes. The geographic distribution of the excluded zip codes is shown in [Figure B.1](#). These areas exhibit unusually high complaint intensity

and are concentrated in several metropolitan clusters, particularly in the Southeast, the mid-Atlantic, and parts of California and Texas. Notable concentrations appear around Los Angeles, Miami, Atlanta, Washington, D.C., and central and northern Florida.

Figure B.1: Number of complaints per Million Residents
Outlier Sample



This figure presents the number of CFPB complaints per million residents for the outlier zip codes.

Table B.1: The Number of Complaints per 1,000 Residents
Regression Sample vs. Outlier Sample


	mean	sd	p25	p50	p75	p90	p99	count
<i>Regression Sample</i>								
Any complaints (0/1)	0.26	0.44	0.00	0.00	1.00	1.00	1.00	927,780
# Complaints per 1,000 residents (raw)	0.07	4.65	0.00	0.00	0.03	0.14	0.74	927,780
<i>Outlier Sample</i>								
Any complaints (0/1)	0.83	0.38	1.00	1.00	1.00	1.00	1.00	90,174
# Complaints per 1,000 residents (raw)	0.35	5.99	0.05	0.17	0.35	0.63	1.95	90,174

This table reports summary statistics for our two main outcome variables: whether a zip code has at least one complaint and the total number of complaints per 1,000 residents, separately for the regression sample and the outlier sample.

Table B.1 summarizes the distribution of consumer complaints per 1,000 residents across the regression sample and the excluded outlier sample. Zip codes in the outlier group are substantially more likely to report complaints in a given month (mean of 0.83 vs. 0.26) and exhibit significantly higher complaint volumes, with a mean of 0.35 complaints per 1,000 residents compared to 0.07 in the main sample. The outlier sample also shows markedly greater dispersion, with the 99th percentile complaint rate reaching 1.95, nearly triple that of the main sample (0.74).

Time-Series of Complaints for Outlier Zip Codes. Figure B.2 compares population-scaled complaint volumes for the regression sample (Panel a) and the outlier sample (Panel b). The mean, 70th, and 90th percentiles remain stable in both, but the 99th percentile in the outlier sample is far more volatile and diverges sharply after mid-2020, indicating a growing concentration of complaints in a subset of zip codes. Figure B.3 zooms into twelve states most representative of these irregular patterns: P99 spikes exceed 800 complaints per 1,000 residents in states such as Indiana and North Carolina, and even low-average states (e.g., Arizona, Nevada) display repeated short-term surges. By contrast, the regression sample’s 99th percentile rarely exceeds 1. Outlier zip codes are thus characterized by both persistently elevated levels and abrupt, episodic deviations, and their exclusion prevents such behavior from distorting our estimates.

Figure B.2: Time Series of the Number of Complaints per 1,000 Residents




This figure compares various statistics of the number of complaints (scaled by population) between the regression sample and the outlier zip code sample. For each month in the sample period (January 2019 to June 2022), we compute the mean, 75th percentile, 90th percentile, and 99th percentile across zip codes.

Comparison of Zip Codes in Outlier Sample vs. Regression Sample. Finally, Table B.2 compares zip code-level socioeconomic characteristics between the regression and outlier samples. Outlier zip codes differ systematically from those in the main sample along several dimensions. They tend to have slightly lower iOS usage, higher shares of young adults (age 20-39), and a greater proportion of female residents. Outlier zip codes also have significantly higher unemployment rates and lower median income levels. These areas are less likely to have older populations (age 50+) and show a modestly higher share of residents with a bachelor’s degree.

Extreme complaint intensity in certain zip codes may arise from small population denominators, concentrated legal aid activity, or localized service disruptions. As such patterns

Figure B.3: Time Series of the Number of Complaints per 1,000 Residents — *Outlier Sample*



This figure presents various statistics for outlier zip codes by state. We display twelve states that are most representative of the irregular complaint patterns observed in the data. For each month in the sample period (January 2019 to June 2022), we compute the mean, 75th percentile, 90th percentile, and 99th percentile across zip codes.

often reflect structural or institutional dynamics rather than random variation, their exclusion provides a cleaner estimation sample and strengthens the interpretability of aggregate results.

C 2017 FCC Privacy Rule Repeal—Timeline and Nature of the Shock


This Appendix Section provides additional detail on the timeline and interpretation of the FCC privacy rule repeal as a regulatory shock to consumer surveillance capacity. We also elaborate on the distinctive role of ISPs in the modern data economy, motivating their centrality in our empirical design.

Timeline of Events.

Table B.2: Regression Sample vs. Outlier Sample

	(1)		(2)		(3)	
	Regression sample		Outlier sample		Outlier – Main	
	mean	sd	mean	sd	b	t
iOS share	0.46	0.11	0.45	0.10	-0.01***	(-6.18)
Female	0.50	0.04	0.52	0.03	0.02***	(23.11)
Age 10-19	0.13	0.04	0.13	0.03	-0.00	(-0.13)
Age 20-29	0.11	0.06	0.14	0.05	0.03***	(25.80)
Age 30-39	0.12	0.03	0.15	0.04	0.03***	(30.43)
Age 40-49	0.12	0.02	0.12	0.02	0.01***	(20.15)
Age 50+	0.41	0.10	0.34	0.08	-0.07***	(-39.33)
Unemployment rate	0.05	0.04	0.07	0.04	0.02***	(20.52)
Bachelor	0.27	0.17	0.31	0.17	0.04***	(9.86)
Median income	65,962	27,517	63,043	25,342	-2,919***	(-5.02)
Observations (# ZIP codes)	22,090		2,147		24,237	

This table reports two-sample t -tests comparing socioeconomic factors between the regression sample and the outlier sample of zip codes.



In 2015, the FCC reclassified broadband internet as a telecommunications service under Title II of the Communications Act, thereby transferring privacy oversight authority from the Federal Trade Commission (FTC) to the FCC. In October 2016, the FCC formally adopted the Broadband Privacy Rules, which would have required ISPs to obtain opt-in consent before using or sharing sensitive customer data. These rules were scheduled to take effect in phases starting in 2017, but were repealed before full implementation. The repeal was enacted in April 2017 through the Congressional Review Act (CRA), which also barred the FCC from issuing substantially similar rules in the future absent new congressional authorization.

Nature of the Shock: A Positive Shift in Surveillance Capacity. Although the FCC’s rules never fully took effect, ISPs faced regulatory overhang during 2015-2017, which limited their ability to expand surveillance practices. The 2017 repeal abruptly removed these constraints, restoring FTC jurisdiction and granting ISPs greater discretion under a less prescriptive enforcement regime. In contrast to the FCC’s ex-ante consent requirements, the FTC framework relies primarily on ex-post enforcement against unfair or deceptive practices.

The CRA’s prohibition on reissuing substantially similar regulations further signaled that stringent privacy rules were unlikely to return in the near term. This combination of

deregulatory policy and legal durability constitutes a clear and plausibly exogenous positive shock to the permissible scope of ISP surveillance.

The Central Role of ISPs in the Data Economy. ISPs occupy a privileged position in the digital ecosystem because they serve as the conduit through which consumers access the internet. Unlike individual platforms or apps, they can observe nearly all unencrypted traffic from a subscriber, including website visits, app usage, device identifiers, and in some cases precise geolocation.

ISP data is distinctive for three reasons. First, ISPs can track behavior across websites, apps, and devices, whereas platforms such as Google or Facebook observe only activity within their own services or embedded tools. Second, ISP data is linked to persistent customer accounts or IP allocations, enabling long-term tracking and linkage to household characteristics. Third, limited broadband competition in many U.S. markets constrains consumers' ability to avoid such surveillance.

These features make ISP data especially valuable for advertising, profiling, and data brokerage. ISPs are therefore central actors in the surveillance economy, especially after the 2017 repeal loosened restrictions on monetizing such data. [Bian et al. \(2024\)](#) further shows that ISPs have the highest level of data connectedness with other firms in the data economy.

D Measurement Error, Alternative Aggregation, and Additional Checks

Measurement Error in iOS Shares. We address potential measurement error in our foot-traffic-based iOS share measure by restricting the sample in several ways. First, we exclude zip codes in the bottom quartile or bottom half of foot traffic. Second, we drop zip codes with a small number of grocery stores (bottom quartile or bottom half), as residents in such areas may shop in other zip codes, making our measure less representative of local residents. Finally, we weight regressions by population, as larger populations are more likely to support multiple grocery and retail stores, ensuring that the measured iOS share better reflects local residents rather than visitors. These alternative sampling criteria improve the precision of our estimates by reducing potential noise in low-traffic areas. Columns 2 to 6 in [Appendix Table I.3](#) show that the estimated ATT effect ranges from -6.0% to -6.8% across specifications and remains negative and statistically significant with a slightly larger magnitude than our baseline estimates.

Alternative Aggregation. Our baseline analysis is conducted at the zip code-month level. While this fine-grained panel makes full use of the identifying variation in our data, it also raises concerns about sparsity and noise in the data, as the number of complaints in many zip code-month cells is small. To address this concern, we aggregate the data to both the quarterly level and to the county-month level to develop datasets at a coarser

level of aggregation. Estimates using the zip code-by-quarter panel are both qualitatively and quantitatively similar to our baseline estimates (Column 1 of [Appendix Table I.4](#)). When aggregating to the county level, the estimated effect is -3.7% (Column 2 of [Appendix Table I.4](#)), about one-third smaller than our baseline estimates (Column 3 of [Table 2](#)), but remains highly significant. This difference could reflect weighting effects from aggregating to larger geographic units, as the treatment effect likely differs across zip codes.

Excluding Complaints about Credit Bureaus. Another potential concern is that reports concerning the three major US credit reporting agencies make up about half of the complaints in the CFPB complaints database. If our estimated treatment effects were driven primarily by complaints about these firms, the results could reflect idiosyncratic dynamics in that segment rather than patterns of consumer fraud driven by ATT. To address this, we re-estimate our main specifications after excluding all complaints related to these three agencies. The results are similar in magnitude and significance (Column 1 of [Appendix Table I.5](#)), so our findings are not driven solely by complaints about the credit bureaus.

Other Robustness. In [Appendix Table I.6](#), we augment our regression specification with several additional robustness checks. First, we replace our baseline 1% two-sided winsorization with a more conservative threshold (0.5% on each side), and find almost no change in the DiD coefficient (Column 1). Second, we re-estimate the Poisson model without scaling the outcome variable by population, which yields results that are both economically and statistically similar to the baseline (Column 2). Third, we relax the criteria for identifying outlier zip codes and continue to find robust effects when we do not remove any outlier zip codes (Column 3). Next, we double-cluster standard errors by state and year-month. Column 4 of [Appendix Table I.6](#) shows that the results remain highly significant with a t-statistic almost identical to the baseline (6.21 vs. 6.60). Lastly, as mentioned previously, we extend the sample through the end of 2023 as well as 2024 and find larger, statistically significant DiD coefficients (Columns 5 and 6). This pattern indicates that the effect of ATT is not short-lived but rather strengthens over time, consistent with the dynamics shown in Panel b of [Figure 4](#).

E Classification of Fraud-related Complaints

[Table E.1](#) reports the likelihood of fraud cases by product category based on two approaches: keyword search and zero-shot learning (ZSL). The two methods deliver similar rankings, with “Credit reporting, credit repair services, or other personal consumer reports” and “Debt collection” being the top two relevant fraud categories, and “Student loan ” and “Mortgage” being the bottom category. Below we describe the details of both approaches.

Keyword Search. Our goal is to classify complaints into cases that are more versus

Table E.1: Classification of Complaints using Keyword Search and Zero Shot Learning

	Mean		St. Dev.		N
	keyword	ZSL	keyword	ZSL	
Credit reporting, credit repair services, or other personal consumer reports	0.822	0.526	0.383	0.284	349,106
Debt collection	0.727	0.517	0.445	0.238	99,484
Money transfer, virtual currency, or money service	0.656	0.426	0.475	0.239	18,792
Checking or savings account	0.510	0.436	0.500	0.251	36,263
Credit card or prepaid card	0.494	0.406	0.500	0.250	54,899
Vehicle loan or lease	0.392	0.289	0.488	0.180	13,023
Payday loan, title loan, or personal loan	0.345	0.315	0.475	0.207	8,472
Student loan	0.341	0.248	0.474	0.167	10,490
Mortgage	0.313	0.159	0.464	0.116	42,559

This table reports the likelihood of relevant fraud cases by product category based on two approaches: keyword search and zero-shot learning (ZSL). The keyword search method returns a binary outcome that is equal to one if any of the keywords is found in the issue, sub-issue, and consumer narrative fields. The zero-shot-learning method returns a continuous variable that represents the likelihood of a fraud-related complaint. Columns 1 and 2 report the mean of the fraud measure, the next two columns report the standard deviation, and the last column is the number of observations in each product category. The sample only includes complaints with narratives, and about 40% of complaints have narratives.

less likely to be triggered by data privacy-related issues. To do this, we search for certain keywords in the issue, sub-issue, and, more importantly, consumer narrative fields. The following keywords are included: “incorrect”, “improper”, “false”, “wrong”, “missing”, “fraud”, “scam”, “theft”, “embezzlement”, “imposter”, “unauthorized”, “unsolicited”, “identity”, “sharing”, “advertising”, “marketing”, “security”, “data breach”, “not owed”. The keyword search method returns a binary outcome that is equal to one if any of the keywords was found in the issue, sub-issue, and consumer narrative fields.

Zero-Shot Learning. We develop an alternative measure to classify complaints using the machine learning approach “zero-shot learning”. The method does not require manual annotations and is, therefore, a more robust approach when few labeled observations are available, as its understanding of language is rooted in a large, diverse sample of text. We use the BART-large-mnli model from Facebook, which uses the pre-trained BART-large language model and adds a task-specific head. Within this model structure, we consider the hypothesis format “I am reporting {label}” with the following 17 labels: “a data breach”, “a mistake”, “an inaccuracy”, “an oversight”, “an unauthorized action”, “an unrecognized action”, “card fraud”, “collection scam”, “debt collection scam”, “embezzlement”, “fraud”, “harassment”, “identity theft”, “mistreatment”, “mortgage scam”, “scam”, and “unresponsiveness”. Varying the hypothesis from “I am reporting a data breach” to “I am reporting unresponsiveness”, for example, while keeping the narrative constant will change the scores generated since the relationship between the narratives and the hypotheses changes.

The relationship between the premise and hypothesis can either be an entailment, neutral, or a contradiction. The model outputs a logit score for each case (e_i, n_i, c_i , respectively). An

example of a full query to determine if a specific narrative refers to identity theft includes a narrative (premise) such as *“I am the victim of identity theft. Please remove the fraudulent accounts from my credit report.”* and a hypothesis *“I am reporting identity theft.”* In this case, a good model outputs a high logit score for entailment and a low score for contradiction.

To combine these multiple logit scores into a single fraud probability, we run a lasso-penalized logistic regression on a manually annotated sample of 1,400 narratives whose product distribution mirrors the full sample’s. The outcome is the manual fraud label, and the regressors are the entailment, neutral, and contradiction scores for all labels (51 features in total). Relative to ridge, the lasso penalty shrinks unimportant coefficients to zero, allowing the model to retain only the most informative scores while tailoring the ZSL representation to our context. The estimated model assigns non-zero weight to 17 features—10 entailment, 3 neutral, and 4 contradiction—with the largest positive coefficient on “identity theft” entailment (1.05) and the most negative on “fraud” neutral (−0.57) and “mistreatment” entailment (−0.38). We then apply this model to combine all 51 features into a final fraud score for every narrative.

Using the manually annotated sample, we verify that the ZSL learning has a satisfactory performance. Setting the threshold score at 0.5 for data-driven fraud complaints, we obtain the following out-of-sample statistics: an F1-score of 0.66, an accuracy of 0.81, a precision of 0.61, and a recall of 0.71.

Example Narratives on Data-driven Fraud Incidents. Below, we list a few example complaint narratives from the public CFPB complaints that scored highly under both methods. We can see that these narratives clearly reveal that the reporting individuals have been a victim of data breach and identity theft and that the unverified inquiries/accounts/debts are typical consequences.

Complaint ID - 3758105 “I am a victim of identity theft. Due to the Corona Virus Pandemic, we are all facing which has me sitting still at home and I saw the recent news about the multiple XXXX Data breaches. I decided to look at my credit reports from the 3 major credit bureaus and found that someone had used my Identity. I have no idea how the theft took place. I also have no knowledge of any suspects. I did not receive any money, goods, or services as a result of identity theft. I contacted the Credit Bureau and told me to file an Identity Theft Report which I am doing. I appreciate your effort in getting this matter resolved. Thank you. Please let me know if you need any other information from me to block this information from my credit report. Thank you..”

Complaint ID - 1488173 “Today I was Contacted by XXXX from credit control at XXXX on XXXX/XXXX/15 for the purpose of a debt collection. She Previously called on XXXX/XXXX/15 XXXX and was unable to provide information substantiating a debt she

was attempting to collect from XXXX XXXX When we first spoke I informed her that There exist the possibility that I may be a victim of identity theft. To day when she called I informed I would not provide her with any verification information and to no longer contact me in regards to the matter or I would be forced to contact your agency and execute my rights under the law. I was very adamant and calm when I informed her of my wishes. XXXX informed me that the calls would continue despite my strict instructions that I do not want her to call my residence any more. To paraphrase her words, “it might not be me who calls but someone will call you”.

F Description of Advisen Data

This section describes the set of information that we use at the incident level. The cases in Advisen’s cyber dataset involve billions of unauthorized disclosures, thefts, or serious disruptions of customer and employee identities, corporate assets, and system capabilities. Over our sample period, 2019/01-2022/06, there are 51,105 incidents.

Causes of Incidents. We use the subcategory risk to identify cases that are more likely to be affected by ATT (as highlighted in bold). The share of each case type is reported in parentheses.

- Data – Unintentional Disclosure (23.56%)
- Data – Physically Lost or Stolen (4.93%)
- **Data – Malicious Breach (44.87%)**
- **Privacy – Unauthorized Data Collection (1.17%)**
- **Privacy – Unauthorized Contact or Disclosure (11.15%)**
- Identity – Fraudulent Use/Account Access (0.69%)
- Industrial Controls & Operations (0.05%)
- Network/Website Disruption (7.11%)
- **Phishing, Spoofing, Social Engineering (4.48%)**
- Skimming, Physical Tampering (0.29%)
- IT – Configuration/Implementation Errors (0.65%)
- IT – Processing Error (0.63%)

Regulations Violated. When specific laws or regulations are violated by the cyber event, Advisen reports the names of the laws and regulations. 5,566 or 10.89% of incidents are recorded as leading to violations of laws or regulations. Among those incidents, the three most frequently violated regulations are Telephone Consumer Protection Act (TCPA) (73%), Fair Debt Collection Practices Act (FDCPA) (29.6%), and General Data Protection Regulation (GDPR) (3.5%). Note that multiple regulations can be violated in a single event.

Conditional on violating the FDCPA, the companies that experienced the most incidents are Synchrony Bank (57 incidents), Midland Credit Management Inc (52 incidents), Capital One Bank (45 incidents), Bank of America National Association (44 incidents), and Portfolio Recovery Associates LLC (35 incidents). These observations suggest that incidents that resulted in violations of FDCPA are more likely to result in fraud that involves financial companies.

G Dark Web Forum Posts: Classification, Statistics, and Examples

We use data from CrimeBB, a structured collection of posts scraped from 37 dark web forums by the [Cambridge Cybercrime Centre \(CCC\)](#). CCC started their work on 1 October 2015. Access to their dataset is restricted to accredited academic researchers under the CCC’s legal and ethical framework, requiring formal applications and compliance with their strict data-handling agreements. We report the volume of posts by each forum in [Table G.1](#).

Table G.1: Total Posts and Average Monthly Posts by Forum (2015–2022)

Forum	Total Posts	Avg. Monthly Posts
Antichat	44,133	1,583
Blackhatworld	2,486,478	38,915
Breached	136,991	24,560
Cracked	174,050	10,168
Crackedto	50,149	6,310
Deutschland-Im-Deep-Web	13,568	1,621
Dread	148,875	48,224
Elhacker	63,044	1,508
Envoy-Forum	1,412	779
Forum-Team	12,301	447
Freehacks	2,398	54
Garage-For-Hackers	327	22
Greysec	5,599	403
Hack-Forums	2,998,744	109,157
Hackers-Armies	1,725	74
Ifud	5,090	170
Indetectables	8,279	441
Kernelmode	3,614	186
Lolzteam	4,270,894	441,309
Mmo4Me	733,294	13,831
Multiplayer-Game-Hacking	982,005	24,998
Nullid	818,866	28,121
Offensive-Community	71,281	32,319
Ogusers	1,322,813	123,882
Piratebay-Forum	13,566	835
Probiv	396,350	12,855
Raidforums	117,860	6,972
Runion	43,996	1,698
Safe-Sky-Hacks	15,893	5,113
Stresser-Forums	2,912	366
The-Hub	26,410	1,598
Torum	17,313	3,846
Underc0De	28,976	736
Unknowncheats	991,618	21,791
V3Rmillion	613,149	28,009
Xss-Forum	33,966	1,738
Zismo	509,137	20,238

Classification. We first present the classification of posts into different labels ([Avarikioti et al., 2018](#)) and groups, including treated ([Table G.2](#)), partially treated ([Table G.3](#)), control ([Table G.4](#)), placebo ([Table G.5](#)), and excluded ([Table G.6](#)). Both the full keyword list and the narrower, more precise keyword lists are presented in these tables. The treated group consists of posts explicitly related to Apple devices or the iOS ecosystem, which are directly exposed to ATT’s impact. The partially treated group includes posts covering illicit activities that plausibly depend on access to personal or financial data—such as tracking, data breaches, identity theft, and financial data theft—and could therefore be indirectly affected by ATT through reduced availability of such data.

Table G.2: Classification of Dark Web Forum Posts - Treated

Label	Full Keyword List	Precise Keyword List
iOS/Apple	ios, iphone, apple, apple id, icloud, idfa, app store, apple pay, jailbreak, mdm bypass, apple watch	ios, iphone, apple

Table G.3: Classification of Dark Web Forum Posts - Partial Treated

Label	Full Keyword List	Precise Keyword List
Tracking & AdTech & ATT	att, app tracking, tracking transparency, app tracking transparency, cross-app tracking, sdk, mobile attribution, spyware, device fingerprinting, tracking app, tracking id, advertising id, attribution	att, app tracking, tracking transparency
Data Breaches & Dumps	leak, data dump, database leak, db dump, combo list, combolist, logs, fullz, dox, cracked db, stealer logs, public dump, private dump	leak, data dump, database leak
PII & Identity Theft	ssn, dob, driver's license, passport, email+pass, identity, kyc, selfie with id, gov id, biometric, facial recognition, personal data	ssn, dob, driver's license
Phishing & Exploits	phishing, spoofing, keylogger, rat, stealer, scraper, clone app, malware, fake update, exploit, 0day, drive-by download	phishing, spoofing, keylogger
Financial Data Theft	bank login, paypal, credit card, cvv, bin, apple pay, crypto wallet, account takeover, 2fa, sms bypass, otp	bank login, paypal, credit card
Scam & Fraud	scam, fraud, rip, ripoff, fake, counterfeit, trick, deception, fake review, fake escrow	scam, fraud, rip
Personal Security & Privacy	opsec, operational security, privacy, anonymity, encryption, vpn, tor, secure messaging	opsec, operational security, privacy
Stolen Accounts & Credentials	account, login, credential, password, username, email+pass, netflix, spotify, social media account	account, login, credential

Table G.4: Classification of Dark Web Forum Posts - Control

Label	Full Keyword List	Precise Keyword List
Hardware & Electronics	hardware, electronics, device, component, circuit, mining rig, usb, phone, computer	hardware, electronics, device
Drugs	cannabis, weed, cocaine, heroin, lsd, mdma, ecstasy, fentanyl, oxy, xanax, stimulants, opiates, psychedelics, dark-net drugs	cannabis, weed, cocaine
Weapons & Explosives	gun, firearm, pistol, rifle, shotgun, ammo, ammunition, explosive, bomb, grenade, c4, tnt, ar-15, ak-47	gun, firearm, pistol
Counterfeit & Forged Items	fake id, counterfeit, forged, replica, passport, driver's license, currency, credit card, branded goods	fake id, counterfeit, forged
Child Exploitation Material (CSAM)	child porn, csam, child sexual abuse material, cp, illegal images, underage	child porn, csam, child sexual abuse material
Guides & Tutorials	guide, tutorial, how to, guide for, handbook, manual, wiki, method, tips, tricks	guide, tutorial, how to
Human Trafficking	human trafficking, slavery, forced labor, organ trafficking, sex trafficking, exploitation	human trafficking, slavery, forced labor
Leaks & Whistleblowing	wikileaks, leak, whistleblower, dump, disclosure, anonymous tip	wikileaks, leak, whistleblower
Vendors & Markets	vendor, market, shop, store, seller, buyer, product list, listing, escrow, review	vendor, market, shop
Tools & Utilities	tool, utility, software, script, generator, checker, builder, scanner, tester	tool, utility, software
Legal & Law Enforcement	law enforcement, fbi, police, customs, court, lawyer, seized, arrest, bust, investigation	law enforcement, fbi, police

Table G.5: Classification of Dark Web Forum Posts - Placebo

Label	Full Keyword List	Precise Keyword List
Android OS & Device Terms	android, apk, apk mod, apk crack, android device, android id, device fingerprint, imei, rooted, root access, custom rom, magisk, bootloader	android, apk, apk mod
Market Listings & Services	[sell], [leak], autoshop, market, marketplace, vendor, escrow, buyer, seller	[sell], [leak], autoshop
Hacking & Cracking Tools	exploit kit, bypass, crack, patch, keygen, license key, loader, ddos, botnet, vpn service, rdp, ssh, proxy, shell, web shell, admin panel, cms exploit, remote access tool, trojan, rootkit, ransomware	exploit kit, bypass, crack
Malware Viruses	malware, virus, worm, ransomware, trojan, rootkit, spyware, adware, backdoor, dropper, exploit, keylogger	malware, virus, worm
Cryptocurrencies & Finance	bitcoin, btc, monero, xmr, zcash, zec, ethereum, eth, crypto, blockchain, wallet, exchange, mixer, tumbler, fiat, usd, eur, gbp	bitcoin, btc, monero
Hosting & Servers	hosting, vps, server, dedicated server, dns, domain, ssl, ip address, bulletproof hosting	hosting, vps, server
VPN & Anonymity Tools	vpn, tor, socks, proxy, anonymity, privacy, secure communication, ip address mask	vpn, tor, socks
Botnets & DDoS	botnet, ddos, distributed denial of service, stresser, booter, attack service	botnet, ddos, distributed denial of service
Exploits & Zero-days	exploit, zero-day, 0day, vulnerability, patch, cve	exploit, zero-day, 0day
Hacking Services	hack for hire, ddos service, website deface, database hack, social media hack, email hack, phone hack	hack for hire, ddos service, website deface
Cryptography/Encryption	encryption, decryption, pgp, gpg, cipher, hash, private key, public key	encryption, decryption, pgp
Software Development & Programming	code, script, api, programming, dev, development, open source, repository, github, framework	code, script, api
Security & Anonymity Guides	guide, tutorial, opsec, vpn, tor, secure, anonymity, privacy, encryption, safety	guide, tutorial, opsec
Counter-Exploitation & Defense	honeypot, anti-malware, firewall, intrusion detection, security research, vulnerability assessment	honeypot, anti-malware, firewall

Table G.6: Classification of Dark Web Forum Posts - Excluded

Label	Full Keyword List	Precise Keyword List
Illegal Services	money laundering, carding, hacking for hire, assassination, hitman, fake documents, drug trafficking, weapon sales, cyber attack for hire	money laundering, carding, hacking for hire
Forums & Communities	forum, board, community, discussion, thread, post, member, user	forum, board, community
Darknet Operations & Community	darknet, onion, tor network, hidden service, deep web, community, forum, market	darknet, onion, tor network
General Discussion & News	news, discussion, current events, forum, thread, post, general chat, updates	news, discussion, current events

Table G.7: Summary Statistics by Group

Group	Mean	SD	p25	p50	p75	N
Number of Posts (Treated)	23	66	0.25	3.58	17.90	2,441
Number of Posts (Partially Treated)	84.6	469	0.111	1.75	15.00	19,528
Number of Posts (Placebo)	238	2,143	0.45	3.82	27.10	34,174
Number of Posts (Control)	45	192	0	0.75	9.27	26,851
Number of Posts (Excluded)	626	4,269	0.167	5.21	113.00	9,764

The table presents summary statistics for the number of posts for each group. The unit of observation is at the forum-label-month level. The sample includes all the posts over the period 2015–2022.

The control group contains illicit-trade categories, such as drugs, human trafficking, and weapons, that operate in segmented markets with minimal overlap in actors or infrastructure with the data-driven fraud economy, providing a benchmark insulated from ATT-related shocks. The placebo group includes unrelated digital technology categories such as zero-days, malware, and botnets, which do not rely on personal data collection for their operation and thus serve to detect whether observed effects reflect broader shifts in cyber-related discussion rather than ATT. All remaining categories, including those with ambiguous or mixed exposure, are assigned to the excluded group. We provide summary statistics for the number of posts for each group in each forum in [Table G.7](#).

Examples. We next provide a few examples, showcasing the discussions in both the (partially) treated and control groups.

A user named “sludgepuppy” posted on Dread in a thread about phishing Apple logs (post label: iOS/Apple, treated):

Ok so ive successfully phished some apple logs, they're full info including ssn, mmn, dob, live cvv and email access. through email access i gained license front and back as-well as license plate number and other tax info. The original goal was to use it with apple pay but the card unfortunately isnt linked to the same email, regardless i know theres much more i can do with this info. Ive worked with pros for a while so im not new but im wondering what you would do with all of this information. Hearing peoples ideas might spark something up.

A user named “ilovetobun69” posted on Multiplayer-Game-Hacking in a thread about fake IDs and PayPal, sharing advice on account setup and identity details (post label: Financial Data Theft, partially treated):

Thanks for understanding, i did some testing and that seems to be the problem (made a fake account under my real info and a different ISP.). It has the same problem. It could be the SSN but I am not sure about the SSN. It most likely was caused by the address, maybe a mix.

Think about it like this, you are PayPal and you see two accounts and one has a different ISP and is newer. The bot would flag it and won't let anyone send money to the account until they call and send in their ID (that is what they would have said if I called).

Also I really do not think he should sell accounts with real SSNs; fake ones work just as well and are not identity theft. That was misleading - I thought I was buying an account with a fake SSN attached. Most sellers sell fake SSNs, if they attach one at all.

A user named “vcpu” posted on Breached in a thread about credential theft and account compromise (post label: Phishing & Exploits, partially treated):

Essentially, most threat actors will try to find as much stuff on you - your email address, any compromised passwords, your phone number, and other PII. I'd say most people reuse passwords or at least variations of the same password. Most hackers just try to use "public"

knowledge to gain access. 2FA would require a SIM swap, or access to their email depending on verification method. Sometimes, hackers will run phishing campaigns or compromise websites and use harvested credentials from other sites.

A user named “afin” posted on Runion in a thread offering firearms for sale (post label: Weapons & Explosives, control):

You will NEVER have to pay for postage on delivery, and we make sure all customs forms are filled out correctly. Sending the parts at once or in separate packages is your choice. You do have to pay upfront for the item so most people choose to buy the first part, then the second, then the third. However, we are easily able to ship the firearm in multiple parts or even all at once. Note that the firearm will NOT be assembled in any event. This means we are going to have to disassemble the firearm (strip, not completely disassemble) and put it in the shielding (decoy packaging). We do ask that you pay for at least 40% of the firearm upfront.....

A user named “Jubs” posted on Multiplayer-Game-Hacking in a thread about drug consumption, sharing advice on how to pass a drug test (post label: Drugs, control):

There are many ways to pass a drug test.

The first and probably hardest one is to chug 2 liters of water every 8 hours. It dilutes your pee. Make sure to pee frequently.

Second would be taking supplements with garnicia or matcha extract and working out a lot.

Third would be borrowing urine from your friend. Buy a heat pack at Walmart or whatever, put the pee in a condom, tuck it in your waistband or pelvis area, and apply the heat pack 5 minutes before your test.

A user named “Diplopia” posted on Breached in a thread about his choice of VPN (post label: VPN & Anonymity Tools, placebo):

I recommend Mullvad vpn its a solid privacy respecting vpn and it also provides decent speeds for a good price its in my top 3 favourite vpns for sure.

A user named “Moody Exchanger” posted on Multiplayer-Game-Hacking in a thread advertising an online currency exchange and cash out service for various payment methods and digital assets (post label: Cryptography/Encryption, placebo):

We are happy to bring biggest online exchange service to this forum! Probably all of you people heard for us... You can google about Moody Exchange. We are biggest online vendors.

Buying/Selling:BTC, WMZ, PP, PM, WU, MG, PSC, GIFTCARDS, PREPAID CARDS, Game cards, Many more!

Please contact us on ICQ only: 679751619

If you don't have ICQ and you use email or Skype please PM ME your contact details and someone from our support team will contact you!

We can cashout your BTC with best rates!!!

Robustness. Table I.9 shows that our main results are robust to a range of alternative samples and measurement approaches. Column 1 restricts the analysis to forums active for at least five years, ensuring results are not driven by transient or short-lived communities. Column 2 changes the counting rule so that posts matching multiple labels are counted in each relevant category rather than split across them, capturing the full presence of multi-topic discussions. Column 3 instead allocates posts proportionally to each label based on keyword frequency, addressing concerns that some labels may dominate a post's content. Column 4 redefines the control group by excluding "Tools & Utilities" and "Leaks & Whistleblowing", which could plausibly interact with digital products and be weakly affected by ATT. Column 5 applies a narrower, more restrictive keyword list, following the definitions in Table G.2-Table G.5, to test whether results are robust to a more conservative labeling approach. Across all these specifications, the estimated ATT effect on Apple/iOS-related posts remains negative, statistically significant, and similar in magnitude to the baseline estimate in Column 1 of Table 9 Panel a, suggesting that the decline in related dark web discussions is not sensitive to these alternative definitions or sample restrictions.

H Dark Web Listings for Data

The research team at Top10VPN periodically scrapes fraud-related listings from active darknet markets, including Nemesis, Kingdom, Empire, Bohemia, and Kraken. We take two snapshots that capture listings in 2020 (July-August) and 2023 (February-March), respectively.¹ Each listing contains the following information: market, listing name, company, category, listing price, listing currency, units, unit price, and listing URL. Examples of listing names include "FRESH PAYPAL ACCOUNT WITH KNOWN BALANCE", "Fully Verified USA COINBASE + BANK LINKED + FULL ACCESS", and "PAXFULL DROP VERIFIED ACCOUNT + FULLZ + EMAIL AND MOBILE ACCESS". There are more than 20 categories of products sold on darknet markets, with the most popular categories being streaming, VPN, payment, shopping, entertainment, crypto, and learning. Popular companies in the darknet markets include NordVPN, Netflix, Paypal, Hulu, and Coinbase. Sometimes a certain quantity of accounts are bundled for sales ("PACK OF 5 CVV/CARDS DETAILS OF U.S WITH GOOD VALIDITY "), and a unit price is calculated for these listings.


We append the listings in 2020 and 2023 in one dataset and construct two variables. First, we determine whether the data being sold is likely generated from consumers' mobile activities. The following categories receive a value of zero: Internet Service Providers, Education (e.g., Masterclass Premium Account), Productivity (e.g., Microsoft Office), Reading, and Communication (e.g., phone

¹More details about these two snapshots can be found at <https://www.top10vpn.com/research/dark-web-prices/2020/> and <https://www.top10vpn.com/research/dark-web-prices/2023/>.

Verizon PIN). User activities concerning these categories are likely to take place via laptops or desktops as opposed to apps on mobile devices. Second, we determine whether the listing involves financial information. The finance-related categories include payment, crypto, personal finance, trading, and gambling.


I Additional Figures and Tables

Figure I.1: Dynamic Effects of ATT on CFPB Complaints - Monthly Frequency



This figure illustrates the dynamic effect of ATT on CFPB complaints around the implementation of ATT at the monthly frequency. Month -1 is the month before the implementation (March 2021) and is the omitted category. In Panel a, the outcome variable is an indicator for whether a zip code has at least one complaint and estimates are from a linear probability model. In Panel b, the outcome variable is the number of complaints per 1,000 residents and estimates are from a Poisson model using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. Coefficients on the interaction terms between indicators for the relative timing to ATT and the pre-ATT iOS device share are plotted. The sample period is January 2019 to June 2022. Zip code fixed effects and county $\times$ year-month fixed effects are included. Standard errors are clustered at the state level.

Figure I.2: Event Timeline from Data Breach to Reported Complaints




a. Changes in CFPB Complaints around Equifax Data Breach in 2017

b. Time to Discover Data Misuse from National Crime Victimization Survey


This figure illustrates the timeline from data breach or leakage to reported complaint. Panel a shows changes in CFPB consumer complaints related to Equifax relative to all other complaints around the July 2017 Equifax breach. Panel b presents the fraction of identity theft cases discovered within different timeframes, based on the most recent National Crime Victimization Survey (NCVS) in 2021. Survey respondents were asked: “How long was the information used before you found out?”

Figure I.3: Monotonicity—Categorizing iOS Share into Quintiles



This figure illustrates the effect of ATT on CFPB complaints, in which we interact the post-ATT indicator with indicators for each of the quintiles of the pre-ATT iOS share. The outcome variable is the number of complaints per 1,000 residents and estimates are from a Poisson model using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 to June 2022. Zip code fixed effects and county $\times$ year-month fixed effects are included. Standard errors are clustered at the state level.

Figure I.4: Dynamic Effects of 2017 FCC Privacy Rule Repeal on CFPB Complaints



This figure shows the dynamic effects of the 2017 FCC privacy rule repeal on CFPB complaints, estimated at a quarterly frequency using a Poisson model. Quarter $t = -1$ (2016Q4), the quarter immediately preceding the repeal, serves as the omitted reference category. In both panels, the outcome variable is the number of complaints per 1,000 residents. Panel a uses the market-share-weighted number of data types collected by ISPs, as disclosed in their privacy labels, as the measure of exposure to the shock. Panel b uses the market-share-weighted number of data SDKs used by ISPs in their mobile applications. Plotted coefficients correspond to the interaction terms between relative-quarter indicators and pre-Repeal exposure levels. Zip code fixed effects and state $\times$ year-month fixed effects are included. Standard errors are clustered at the state level.

Table I.1: Effect of 2017 Repeal of FCC Broadband Privacy Rule on CFPB Complaints
Robustness: Unique # Data Types Collected and SDKs Installed Across ISP-Owned Apps

Panel a. Local Exposure based on ISP Data Collection Intensity

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post Repeal $\times$ Exposure (data collection)	0.004*** (0.001)	0.004*** (0.001)	0.039*** (0.007)	0.025* (0.013)
Zip code FE	✓	✓	✓	✓
State $\times$ Year-month FE	✓		✓	
County $\times$ Year-month FE		✓		✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.207	0.207	0.034	0.034
Magnitude (1SD increase in iOS share)	3.0%	2.7%	6.4%	4.0%
Observations	1,166,496	1,166,496	1,166,496	1,166,496
R ² / Pseudo R ²	0.406	0.460	0.061	0.109

Panel b. Local Exposure based on ISP Data Sharing Intensity

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post Repeal $\times$ Exposure (data sharing)	0.002*** (0.000)	0.002*** (0.001)	0.018*** (0.003)	0.010* (0.006)
Zip code FE	✓	✓	✓	✓
State $\times$ Year-month FE	✓		✓	
County $\times$ Year-month FE		✓		✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.207	0.207	0.034	0.034
Magnitude (1SD increase in iOS share)	3.2%	2.9%	6.5%	3.6%
Observations	1,166,496	1,166,496	1,166,496	1,166,496
R ² / Pseudo R ²	0.406	0.460	0.061	0.109

This table presents the estimated effects of the 2017 FCC Broadband Privacy Rule repeal on CFPB complaints. The unit of observation is at the zip-code-month level. Columns 1–2 report estimates from a linear probability model, where the outcome is a binary indicator equal to one if at least one complaint is filed in a zip code during a given month. Columns 3–4 report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. Panel a (b) uses the count of unique data types collected (the count of unique data SDKs installed) across all apps owned by the same ISP. The sample period is January 2015 to December 2018. All columns include zip code fixed effects. The odd-numbered columns include state $\times$ year-month fixed effects, while the even-numbered columns include county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.2: Effect of 2017 Repeal of FCC Broadband Privacy Rule on CFPB Complaints
Robustness: Total # Data Types Collected and SDKs Installed Across ISP-Owned Apps

Panel a. Local Exposure Based on ISP Data Collection Intensity

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post Repeal $\times$ Exposure (data collection)	0.012*** (0.003)	0.008** (0.004)	0.109*** (0.021)	0.050 (0.034)
Zip code FE	✓	✓	✓	✓
State $\times$ Year-month FE	✓		✓	
County $\times$ Year-month FE		✓		✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.207	0.207	0.034	0.034
Magnitude (1SD increase in iOS share)	2.5%	1.7%	5.1%	2.4%
Observations	1,166,496	1,166,496	1,166,496	1,166,496
R ² / Pseudo R ²	0.406	0.460	0.061	0.109

Panel b. Local Exposure Based on ISP Data Sharing Intensity

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post Repeal $\times$ Exposure (data sharing)	0.005*** (0.001)	0.003** (0.002)	0.044*** (0.010)	0.020 (0.014)
Zip code FE	✓	✓	✓	✓
State $\times$ Year-month FE	✓		✓	
County $\times$ Year-month FE		✓		✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.207	0.207	0.034	0.034
Magnitude (1SD increase in iOS share)	2.3%	1.6%	4.6%	2.2%
Observations	1,166,496	1,166,496	1,166,496	1,166,496
R ² / Pseudo R ²	0.406	0.460	0.061	0.109

This table presents the estimated effects of the 2017 FCC Broadband Privacy Rule repeal on CFPB complaints. The unit of observation is at the zip-code-month level. Columns 1–2 report estimates from a linear probability model, where the outcome is a binary indicator equal to one if at least one complaint is filed in a zip code during a given month. Columns 3–4 report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. Panel a (b) uses the aggregate number of data types collected (the aggregate number of data SDKs installed) across all apps owned by the same ISP, allowing for double counting. The sample period is January 2015 to December 2018. All columns include zip code fixed effects. The odd-numbered columns include state $\times$ year-month fixed effects, while the even-numbered columns include county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.3: Effect of ATT on Consumer Complaints - CFPB
Robustness: Propensity to Complain and Measurement Error

	# Complaints per 1,000 residents					
	(1) victimization	(2) foot traffic>25 ^{pct}	(3) foot traffic>50 ^{pct}	(4) # grocery store>25 ^{pct}	(5) # grocery store>50 ^{pct}	(6) pop. weight
Post $\times$ iOS share	-0.719*** (0.111)	-0.742*** (0.103)	-0.657*** (0.099)	-0.658*** (0.093)	-0.701*** (0.087)	-0.723*** (0.094)
Zip code FE	✓	✓	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓	✓	✓
Model	PPML	PPML	PPML	PPML	PPML	PPML
Mean outcome var.	0.070	0.073	0.077	0.072	0.076	0.069
Magnitude ($\uparrow$ 1 SD iOS share)	-6.9%	-6.8%	-6.0%	-6.1%	-6.5%	-7.0%
Observations	878,346	698,601	469,885	701,673	472,122	926,772
Pseudo R ²	0.131	0.125	0.110	0.124	0.107	0.086

This table presents the estimated effects of ATT on CFPB complaints using alternative regression specifications. The unit of observation is at the zip-code-month level. Column 1 reports estimates adjusted for victimization, using outcome variables weighted following [Raval \(2020b\)](#) to account for differences in the propensity to complain across zip code level demographics. Column 2 reports estimates for observations with pre-ATT foot traffic above the 25th percentile and Column 3 restricts to those above the 50th percentile. Column 4 reports estimates for observations with a pre-ATT number of grocery stores above the 25th percentile, and Column 5 restricts to those above the 50th percentile. Column 6 reports estimates from a population weighted regression. All columns report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 to June 2022. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.4: Effect of ATT on Consumer Complaints - CFPB
Robustness: Alternative Aggregation by Zipcode-quarter and County-month

	# Complaints per 1,000 residents	
	(1) Zip-code-quarter	(2) County-month
Post $\times$ iOS share	-0.510*** (0.116)	-0.444*** (0.162)
Zip code FE	✓	
County FE		✓
County $\times$ Year-quarter FE	✓	
State $\times$ Year-month FE		✓
Model	PPML	PPML
Mean outcome var.	0.193	0.042
Magnitude ($\uparrow$ 1 SD iOS share)	-5.0%	-3.7%
Observations	309,161	134,177
Pseudo R ²	0.175	0.105

This table presents the estimated effects of ATT on CFPB complaints using alternative aggregation by zip code-quarter (Column 1) and county-month (Column 2). The unit of observation in Column 1 is at the zip code-quarter level. The unit of observation in Column 2 is at the county-month level. All columns report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 to June 2022. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.5: Effect of ATT on Consumer Complaints - CFPB
Robustness: Excluding and Including Complaints to the Three Credit Bureaus

	(1) # Complaints per 1,000 residents Excl. 3 credit bureaus	(2) # Complaints per 1,000 residents 3 credit bureau only
Post $\times$ iOS share	-0.420*** (0.077)	-0.670*** (0.150)
Zip code FE	✓	✓
County $\times$ Year-month FE	✓	✓
Model	PPML	PPML
Mean outcome var.	0.069	0.087
Magnitude ($\uparrow$ 1 SD iOS share)	-4.1%	-6.3%
Observations	908,292	908,292
Pseudo R ²	0.091	0.164

This table presents the estimated effects of ATT on CFPB complaints by separating complaints related to the three major credit bureaus from all others. The unit of observation is at the zip code-month level. Column 1 reports estimates using only complaints not concerning the three credit bureaus (Equifax, Experian, and TransUnion), while Column 2 reports estimates using only complaints concerning these three credit bureaus. All columns report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 to June 2022. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.6: Effect of ATT on Consumer Complaints - CFPB
Robustness: Alternative Regression Specifications

	# Complaints per 1,000 residents or # Complaints (unscaled)					
	(1) winsorize 0.5%	(2) unscaled count	(3) incl. outlier zip codes	(4) double cluster	(5) sample to 2023Q4	(6) sample to 2024Q4
Post $\times$ iOS share	-0.618*** (0.101)	-0.686*** (0.081)	-0.602*** (0.090)	-0.621*** (0.100)	-0.824*** (0.152)	-1.034*** (0.157)
Zip code FE	✓	✓	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓	✓	✓
Model	PPML	PPML	PPML	PPML	PPML	PPML
Mean outcome var.	0.069	0.069	0.093	0.069	0.073	0.100
Magnitude ($\uparrow$ 1 SD iOS share)	-6.0%	-6.8%	-5.9%	-6.1%	-7.9%	-10.0%
Observations	926,772	943,488	1,025,472	926,772	1,325,640	1,590,768
Pseudo R ²	0.149	0.445	0.172	0.134	0.187	0.241

This table presents the estimated effects of ATT on CFPB complaints using alternative regression specifications. The unit of observation is at the zip code-month level. Column 1 reports estimates with less aggressive winsorization at 0.5% on each side. Column 2 reports estimates using unscaled number of complaints as the outcome variable. Column 3 reports estimates including outlier zip codes in the regression. Column 4 reports estimates with standard errors double clustered at state and year-month level. Column 5 extends the sample to include CFPB complaints from January 2019 through the end of 2023. Column 6 extends the sample to include CFPB complaints from January 2019 through the end of 2024. All columns report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 to June 2022 in Columns 1 to 4. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level (except Column 4) are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.7: Effect of ATT on Consumer Complaints - CFPB
Robustness: Controlling for Socio-economic and Demographic Factors

	Any complaints (0/1)		# Complaints per 1,000 residents	
	(1)	(2)	(3)	(4)
Post $\times$ iOS share	-0.038*** (0.014)	-0.035*** (0.013)	-0.446*** (0.099)	-0.345** (0.139)
Post $\times$ %Female	0.010*** (0.002)		0.013 (0.019)	
Post $\times$ %Age 10-19	-0.010*** (0.003)		-0.001 (0.033)	
Post $\times$ %Age 20-29	-0.004 (0.003)		-0.051** (0.024)	
Post $\times$ %Age 30-39	0.002 (0.002)		0.005 (0.024)	
Post $\times$ %Age 40-49	-0.002 (0.002)		-0.061*** (0.021)	
Post $\times$ %Age 50+	-0.025*** (0.005)		-0.113*** (0.038)	
Post $\times$ Median income		-0.019*** (0.004)		-0.051*** (0.017)
Post $\times$ Unemployment rate		0.004*** (0.001)		0.041*** (0.014)
Post $\times$ %Bachelor		0.011*** (0.004)		0.021 (0.019)
Zip code FE	✓	✓	✓	✓
County $\times$ Year-month FE	✓	✓	✓	✓
Model	Linear	Linear	PPML	PPML
Mean outcome var.	0.267	0.273	0.069	0.070
Magnitude ($\uparrow$ 1 SD iOS share)	-1.5%	-1.3%	-4.4%	-3.4%
Observations	926,772	900,942	926,772	900,648
R ² / Pseudo R ²	0.514	0.513	0.134	0.133

This table presents the estimated effects of ATT on CFPB complaints by adding interactions between the Post-ATT indicator and various socioeconomic factors. The unit of observation is at the zip code-month level. Columns 1–2 report estimates from a linear probability model, where the outcome is a binary indicator equal to one if at least one complaint is filed in a zip code during a given month. Columns 3–4 report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. Columns 1 and 3 control for demographic factors (gender and age), whereas Columns 2 and 4 control for economic factors (income, unemployment, and education). The sample period is January 2019 to June 2022. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.8: Effect of ATT on Consumer Complaints - Consumer Sentinel Network
Robustness: Top vs. Bottom Fraud Categories by Relevance

	# Complaints per 1,000 residents	
	Top Fraud Category (1)	Bottom Fraud Category (2)
Post $\times$ iOS share	-0.223*** (0.079)	-0.024 (0.040)
Zip code FE	✓	✓
County $\times$ Year-month FE	✓	✓
Model	PPML	PPML
Mean outcome var.	0.215	0.556
Magnitude ($\uparrow$ 1 SD iOS share)	-2.2%	-0.2%
Observations	1,029,942	1,122,071
Pseudo R ²	0.161	0.126

This table presents the estimated effects of ATT after classifying complaints from Consumer Sentinel Network (CSN) by their relevance to ATT. The unit of observation is at the zip code-month level. The outcome variable is the number of complaints per 1,000 residents. Column 1 includes complaints from products for which at least 50% of complaints include at least one of the relevant words in the narrative. Column 2 includes complaints from products for which less than 25% of complaints include one of the relevant words in the narrative. All columns report estimates from a Poisson model where the dependent variable is the number of complaints per 1,000 residents, using the Poisson Pseudo-Maximum Likelihood (PPML) estimator. The sample period is January 2019 (February 2019 for Identity Theft) to June 2022. All columns include zip code fixed effects and county $\times$ year-month fixed effects. Standard errors clustered at the state level are reported in parentheses. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table I.9: Effect of ATT on Dark Web Activities—Posts on Dark Web Forums
Robustness: Alternative Samples and Measures

	# Posts				
	(1) forum age>60m	(2) double count	(3) keyword-freq. weighted	(4) alt. control	(5) alt. keywords
Post $\times$ Treated	-0.433*** (0.017)	-0.359*** (0.057)	-0.420*** (0.021)	-0.362*** (0.074)	-0.390*** (0.070)
Label FE	✓	✓	✓	✓	✓
Forum $\times$ Year-month FE	✓	✓	✓	✓	✓
Model	PPML	PPML	PPML	PPML	PPML
Mean outcome var.	44.5	46.0	46.0	42.8	20.8
Magnitude (Treated vs. Control)	-35.1%	-30.2%	-34.3%	-30.4%	-32.3%
Observations	27,600	27,060	27,612	27,612	23,010
Pseudo R ²	0.872	0.877	0.869	0.898	0.806

This table presents the estimated effects of ATT on number of posts on dark web forums using alternative samples and measures. The unit of observation is a forum-label-month. The full list of labels and their exposure to ATT can be found in Internet Appendix Section G. All columns compare the explicitly treated group (Apple/iOS-related posts) with a control group of illicit-trade categories (e.g., drugs, human trafficking, weapons) that operate in separate markets. Column 1 restricts the sample to forums active for at least five years. Column 2 counts a post in every applicable label rather than splitting it evenly across labels when it meets multiple label criteria. Column 3 allocates posts proportionally based on the frequency of matching keywords within each label. Column 4 uses an alternative control group that excludes “Tools & Utilities” and “Leaks & Whistleblowing”, which may interact with digital products and thus be tangentially treated. Column 5 applies a narrowly defined keyword list, as in Table G.2–Table G.5. Standard errors double clustered at the forum and label level are reported in parentheses. The sample period is January 2015 to December 2022. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.