Working Paper: Feb 2025

SMU
SINGAPORE MANAGEMENT
UNIVERSITY

Grand Theft Identity: The Privacy Costs of Digitalization

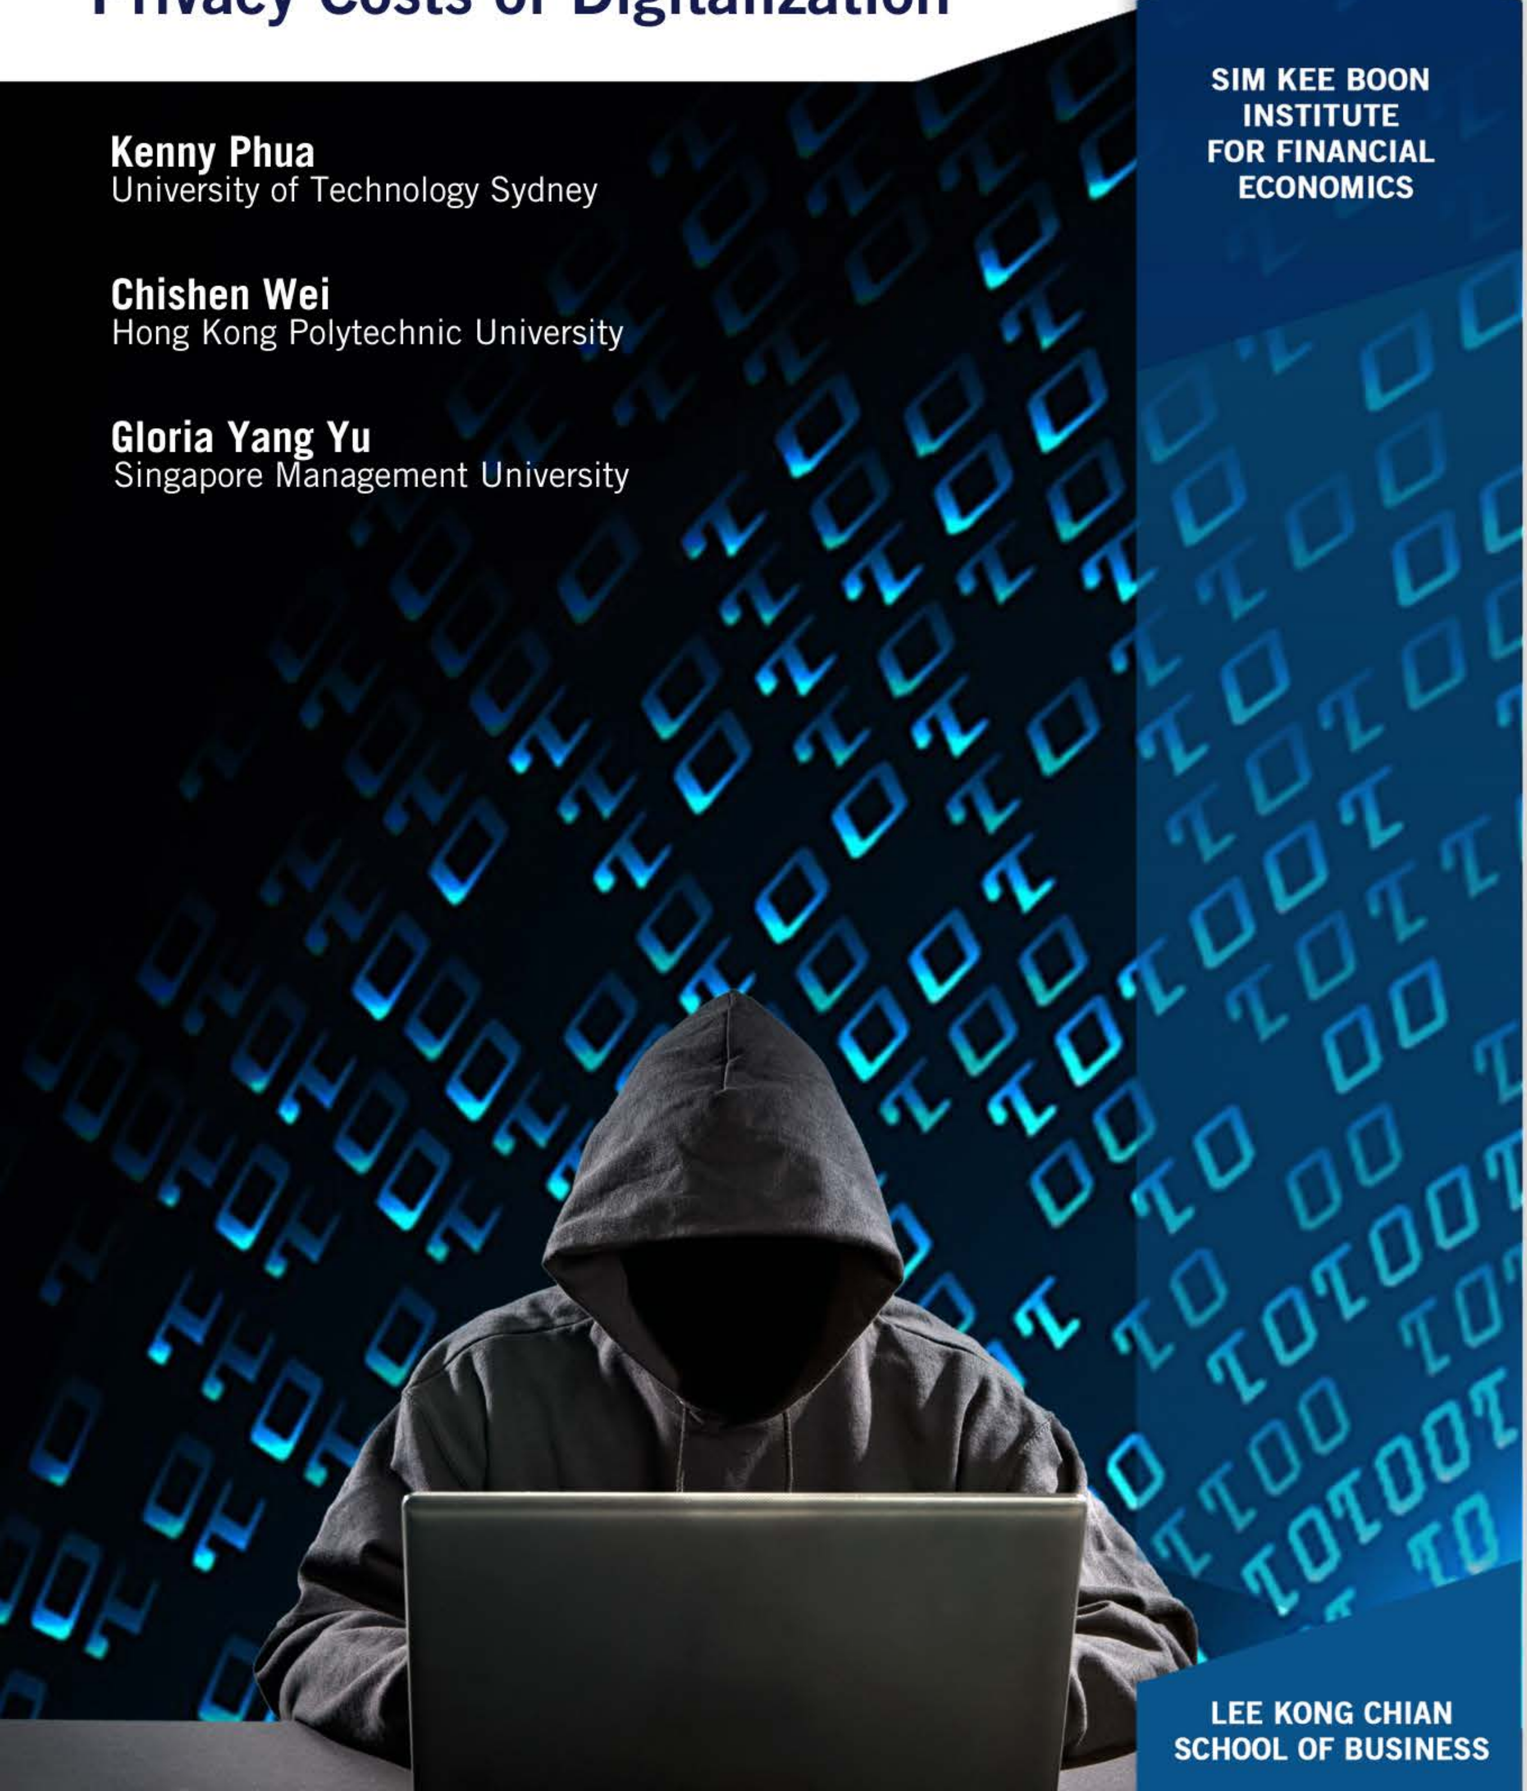

Kenny Phua
University of Technology Sydney

Chishen Wei
Hong Kong Polytechnic University

Gloria Yang Yu
Singapore Management University

SIM KEE BOON
INSTITUTE
FOR FINANCIAL
ECONOMICS

LEE KONG CHIAN
SCHOOL OF BUSINESS

Grand Theft Identity: The Privacy Costs of Digitalization *

Kenny Phua Chishen Wei Gloria Yang Yu

Abstract

We study whether greater digital engagement increases the risk of identity theft by exploiting bank branch closures as a shock that shifts economic activity online. Using a quasi-natural experiment, we find causal evidence that branch closures increase identity theft, particularly in more vulnerable communities. Exposed consumers spend more time on mobile apps and shift their expenditures from offline to online channels. Adversarial activities associated with identity theft, such as unwanted calls and phishing attempts, increase after branch closures. Overall, our evidence suggests that digitalization offers consumer benefits, but also imposes hidden privacy costs.

*Yu (corresponding author): gloriayu@smu.edu.sg, Singapore Management University. Phua: kenny.phua@uts.edu.au, University of Technology Sydney and Digital Finance Cooperative Research Center. Wei: chishen.wei@polyu.edu.hk, Hong Kong Polytechnic University. We are grateful for comments from Bo Becker, William Cong, Marcin Kacperczyk, Robert Marquez, Christine Parlour, Manju Puri, Adriano Rampini, Alberto Rossi. We also thank seminar participants at the Hong Kong Polytechnic University and the University of Technology Sydney, as well as participants at the Monash Winter Finance Conference and the ANU-FIRN Banking and Financial Stability Meeting.

1 Introduction

Privacy concerns have become a significant consumer issue in the digital age. Despite government regulations, platform privacy measures, and firm-level data security policies, adversarial actors continue to harvest personal identifying information (PII). Our paper focuses upstream of these institutional safeguards—we pinpoint consumers' increasing engagement with the digital economy as a structural source of exposure to privacy costs. As economic transactions increasingly shift online, digital engagement has become a modern necessity but it imposes privacy risks that are hidden and difficult to quantify. For example, 16 billion login credentials, including those linked to Apple, Facebook, and Google accounts, have been leaked online (Forbes Magazine, 2025). We hypothesize that digital engagement exposes consumers to greater privacy costs and can disproportionately affect vulnerable communities.

The digitalization of banking services provides an opportunity to test our hypothesis. Banks are uniquely positioned to shape consumer behavior in the digital economy because they facilitate many economic activities that involve payments and transfers. Traditionally, banks had extensive physical branch networks, but are now closing branches as they offer more digital services. Consumers affected by branch closures must learn to adopt digital banking and payment tools. In turn, these tools can reduce the marginal costs of using other digital services, creating network effects that reinforce digital adoption and drive greater digital engagement.

Bank branch closures may push consumers towards greater digital engagement and expose them to privacy risks through two pathways. First, branch closures nudge consumers to use online services and conduct more transactions digitally. Conducting digital activities typically requires the transfer of sensitive PII, which could leak and be compromised. Digital banking in itself is unlikely to incur privacy costs because banks invest heavily in cybersecurity infrastructure. However, third parties who facilitate many digital transactions may have weaker data security. They may also act as data brokers, who harvest digital footprints for sale and have been prosecuted for selling PII to scammers (Federal Trade Commission, 2015). Even sophisticated consumers can be affected because these data breaches often occur outside of one's control.

Second, branch closures eliminate a vital alternative for some vulnerable consumers who prefer in-person services. Physical branch services can limit PII exposure, allow face-to-face verification, and provide personalized security advice. For example, a financial advice columnist for New York Magazine recounts her experience with a bank teller, who warned that her $50,000 cash withdrawal was likely related to a scam (Cowles, 2024). Without these physical touchpoints, some consumers may become more susceptible to adversarial tactics such as phishing attacks using phone calls, text messaging, emails, SIM card swaps, and even generative artificial intelligence (AI) tools. These attacks can be especially harmful during the transition to digital services, where consumers are less familiar with security practices or warning signs of fraud.

Privacy costs are multifaceted. We focus on identity theft because it is the most common data privacy concern (Armantier, Doerr, Frost, Fuster, et al., 2024) and arises from the abuse of user data. Under 18 U.S.C. §1028(a)(7), identity theft is any crime that misuses personal information for financial gain, including credit card fraud, fraudulent loan applications, and unauthorized access to bank accounts.[1] Identity theft has serious long-lasting consequences, and the recovery of stolen identities is difficult. Victims may suffer severe emotional distress (Harrell and Langton, 2013) and face significant financial repercussions including reduced credit access and greater bankruptcy risk (Hamdi, Kalda, Sovich, 2024). The problem is so severe that specialized insurance products now exist to protect against losses related to identity theft.

Before turning to our main analysis, we examine changes in consumer banking behavior after branch closures. First, we estimate the elasticity of substitution to neighboring bank branches after the closure of a focal branch. Granular footfall data from the pass_by database show that, somewhat surprisingly, only 18% of foot traffic redirects to nearby branches—82% of footfall disappears entirely. This pattern suggests that most consumers cease physical banking after branch closures. It is also consistent with evidence that physical banking is hyper-localized, as minor inconveniences cause customers to adopt online payments and transactions (Choi and Loh, 2024). Second, using microdata on mobile application (app) usage from Global Wireless Solutions, we find that consumers spend 24.2% more time on the mobile app of a bank that closes a branch in that area. These preliminary

[1] The U.S. Congress passed the IDENTITY THEFT AND ASSUMPTION DETERRENCE ACT in 1998 to criminalize identity theft. In 2024, the U.S. House of Representatives passed another act to provide additional assistance to victims of identity theft.

findings validate our premise that branch closures can change consumer behavior.

Using MSA-level data from the FTC Consumer Sentinel Network database, we examine whether branch closures trigger more local cases of identity theft. The main threat to identification is that branch closures may be endogenous to local conditions that affect identity theft. For example, banks may close branches in areas with higher rates of digital adoption, which increases the baseline risk of identity theft. Thus, even in the absence of branch closures, we might observe more identity theft in these areas. To address this concern, we instrument for branch closures with staggered exposures to post-merger consolidation between large regional/national banks (Garmaise and Moskowitz, 2006; Nguyen, 2019). The idea is that an MSA with *both* acquirer and target branches is likely to have redundant branches, which the merged bank will close post-merger due to duplication in branch services.

For the instrument to be valid, it must satisfy the exclusion restriction, which requires that merger exposures be as good as randomly assigned with respect to local conditions. There are three reasons why the exclusion restriction holds. First, we only consider mergers where both acquirer and target banks have at least U.S. $1 billion in premerger assets. At these scales, the merging decision is plausibly removed from the local conditions of any particular MSA.[2] Second, we show that these mergers are unlikely to be solely motivated by local factors because (i) the acquirers and targets are large and geographically diversified, and (ii) exposed MSAs account for only a small share ($\simeq 2\%$) of their overall deposits. Third, exposed and unexposed MSAs do not significantly differ in measures of digital adoption and demographic characteristics in the pre-merger period, suggesting that merger exposures are uncorrelated with pre-existing digitalization trends.

In our setting, merger exposures are staggered across MSAs. Therefore, we estimate two Callaway and Sant'Anna (2021) difference-in-differences models—a first-stage and a reduced-form. As these exposures satisfy the exclusion restriction, the first-stage can isolate variation in branch closures that is not driven by local conditions. With the relevance condition met, the reduced-form identifies the effect of merger exposures on identity theft that comes through plausibly exogenous branch closures. Thus, the ratio of the reduced-form and first-stage estimates yields the causal effect of a branch closure on identity theft.

[2] Press releases indicate that these mergers are often driven by broad strategic objectives such as a strategic expansion into new markets.

Our results from the Callaway and Sant'Anna (2021) difference-in-differences estimator indicate that branch closures lead to more identity theft. On average, MSAs exposed to large bank mergers have 2.79 ($t = 2.90$) more branch closures and 455.56 ($t = 2.72$) more reports of identity theft. The ratio of these estimates is the local average treatment effect (LATE)—each branch closure leads to an increase of +163.28 (= 455.56/2.79) identity theft reports.[3] This effect is economically meaningful as a one standard deviation shock to branch closures causes an increase of 2,318 identity theft cases, representing 1.6 times of its unconditional sample mean. Although merger exposures are plausibly exogenous to local conditions, we further test whether our results are robust to more refined counterfactuals based on MSA characteristics and outcomes. Using the Arkhangelsky, Athey, Hirshberg, Imbens, et al. (2021) synthetic difference-in-differences estimator, we continue to find significant causal effects of branch closures on identity theft.

A possible concern is that merged banks ultimately select which branches to close in the exposed MSAs, and these decisions may reflect local conditions such as digital adoption rates. However, we find no differences in measures of digital adoption in the form of internet subscription rates and computing device ownership between exposed MSAs that close branches and those that do not. There are also no significant differences in local demographics. Collectively, our evidence indicates that the branch closure decisions ultimately made by the merged banks are orthogonal to observable MSA characteristics that may affect identity theft.

This selection concern is generally known as "noncompliance", which is a common feature in IV designs.[4] Our evidence suggests that merger exposures are as good as randomly assigned with respect to local conditions, so they can exert a plausibly *exogenous* encouragement that shifts branch closures. Even if a characteristic such as digital adoption rates affects compliance, the LATE remains causal—but it would be identified only for the subgroup of MSAs that comply with the IV (Imbens and Angrist, 1994). The economically relevant question is whether

[3] Comparing the LATE against the equivalent ordinary least squares (OLS) estimate, we find a close but larger (+213.42) OLS point estimate, which we speculate arises from uncorrected endogeneity (e.g., local conditions) that drive both branch closures and identity theft.

[4] Noncompliance often occurs in randomized control trials when assigned subjects fail to take treatment or improperly self-administer dosages. The assignment in a randomized control trial is analogous to our merger exposures, which are as good as randomly assigned across MSAs. Some exposed MSAs fail to comply because the merged bank close no branches there. In a corporate finance context, see Bernstein (2015) for an example of causal research design with noncompliance.

the LATE estimate can reasonably be applied to other MSAs. Our analysis indicates that the LATE can generalize to the broader population because complier MSAs, which provide the identifying variation in our IV design, resemble the average MSA across most characteristics.

To establish a more direct link between branch closures and the focal bank's customers, we turn to the Consumer Complaint Database (CCD) administered by the Consumer Financial Protection Bureau (CFPB). Although the CCD covers only the largest banks, a key advantage lies in its highly granular data, which comprise individual geotagged consumer complaints against specific banks. This granularity allows us to precisely trace how recent branch closures of a particular bank affect its own local customers. Following a branch closure, we find a significant increase in identity theft complaints from local customers of that particular bank, relative to other consumers and banks in that county.

We expect that the transition to digital platforms and services can affect consumers in unequal ways because technological shocks impact bank customers differently (Fuster, Goldsmith-Pinkham, Ramadorai, and Walther, 2022; Jiang, Yu, Zhang, 2025). Although some sophisticated consumers adapt seamlessly, vulnerable communities may face serious challenges, particularly from the loss of physical touchpoints they have previously relied on. First, we test whether consumers with limited digital capabilities are more vulnerable to the adverse effects of forced digitalization. We find that branch closures lead to more identity theft cases when (i) consumers are more reliant on bank branches, (ii) banks are less digitally focused to start with, and (iii) the local area has lower internet penetration rates.

If branch closures increase identity theft through our proposed pathway of digital engagement, we expect to observe shifts in online and offline consumption patterns. Consistent with this prediction, we estimate that branch closures lead the average consumer to spend an additional 10.76 hours per month on all mobile apps, excluding banking ones.[5] Furthermore, using data from Safegraph, we estimate that a branch closure tilts consumer expenditures and transaction volumes from offline to online channels by 3.6 pp and 3.0 pp, respectively. These findings indicate that branch closures lead consumers to increase overall digital engagement, which structurally exposes them to identity theft risks.

[5] This implies an increment of 0.36 hours per consumer per day. As a benchmark, Americans spend an average of 4.65 hours per day on their mobile phones (Statista, 2025).

Identity theft poses a significant challenge to law enforcement because adversaries often operate outside of legal jurisdictions and employ an ever-evolving suite of tactics. To shed light on their otherwise opaque activities, we examine two common adversarial tactics. Phone calls are often the first attack vector used to target potential victims. Adversaries can impersonate trusted institutions on these calls to extract PII such as Social Security numbers and bank account details. To compound the problem, advances in telecommunications technology have enabled the rise of "robocalls", which are programmatically automated calls that can reach large numbers of consumers at low cost. Using the FTC Do-Not-Call (DNC) Reported Calls database, we estimate that a branch closure leads to a 23.7% increase over the unconditional mean number of unwanted calls received by a county per month. These calls also impose additional social costs by wasting time and disrupting the provision of legitimate services.

Another adversarial tactic is phishing attacks. According to the FTC, a significant portion of identity theft cases stem from phishing attacks, where adversaries deceive consumers into providing personal information through deceptive emails or websites (Federal Trade Commission, 2024a). To identify phishing attacks, we exploit the idea that adversaries often clone the legitimate website as a template for malicious use. Using the search engine Shodan, we sweep the internet for suspicious websites that share features of their legitimate counterparts but lack digital security certificates. Using a two-stage estimation approach, we find that branch closures in an MSA can increase the exposure of its residents to identity theft through phishing attacks.

A limitation of our study is that we do not evaluate the *overall* welfare effects of digitalization. Identity theft is only one of many facets of privacy costs, and we do not address any benefits of digitalization. Moreover, we have not considered the nondigital counterfactual. Cash transactions create security and logistic costs, and the maintenance of bank branches requires significant resources. A welfare analysis that balances these costs against the convenience, cost savings, and efficiency brought about by the digital economy is beyond the scope of our study. Nevertheless, our findings suggest that bank branches can serve as a vital social good in the digital age, offering security and consumer protection.

Our study contributes to the literature on the data economy (Farboodi and Veldkamp, 2023; Abis and Veldkamp, 2023), particularly the economics of data privacy (Goldfarb and Tucker, 2012; Acquisti, Taylor, Wagman, 2016; Tirole, 2023;

Bian, Ma, Tang, 2023). Users bear privacy costs from digital businesses that over-collect data and under-invest in consumer data protection (Cong, Xie, Zhang, 2021; Fainmesser, Galeotti, Momot, 2023; Chen, Huang, Ouyang, and Xiong, 2025). However, large-scale evidence on privacy costs remains elusive (Johnson, 2022).[6] Ramadorai, Uettwiller, Walther (2025) examine the privacy policies of individual US firms. Bian, Pagel, Raval, and Tang (2024) assess privacy costs by analyzing restrictions on personal data collection from the Apple App Tracking Policy. Our contribution is upstream of data protection policies as we show that digital engagement is a structural source of exposure to identity theft risks.

We also add to the growing evidence that digitalization and technological disruption affect consumers unequally. Jiang, Yu, and Zhang (2025) show that bank digitalization tends to unbank elderly consumers. Koont (2024) finds that the surplus from the digital banking revolution is mostly captured by the wealthier consumers. Fuster, Goldsmith-Pinkham, Ramadorai, and Walther (2022) show that digital financial technologies may have different effects on racial minorities. Financial innovation may deliver participation benefits for some consumers, but it may also generate fraud that externalizes costs across the broader society (Cong, Feng, Liu, and Lu, 2025). We examine the privacy costs of digitalization and find that bank branch closures disproportionately harm communities with poorer digital-savviness.

Finally, our findings show that bank branches provide social benefits to the local community. Banks are vital to economic activity and are shaped by regulations (Kroszner and Strahan, 1996), political incentives (Liu and Ngo, 2014), and market forces (Bonfim, Nogueira, Ongena, 2020). Branches anchor banks' market presence, stabilize deposit funding, and facilitate the flow of soft information (Drechsler, Savov, Schnabl, 2017; Benmelech, Yang, Zator, 2023; Amberg and Becker, 2024; Keil and Ongena, 2024; Qi, De Haas, Ongena, Straetmans, et al., 2024). They also benefit local communities by supporting income, entrepreneurship, employment, financial inclusion, and health outcomes.[7] Using a pre-2000 sample, Garmaise and Moskowitz (2006) find that branch closures are accompanied by more property crime. We study the later era of bank digitalization and

[6] In the context of online lending, Tang (2019) quantifies the monetary value required for borrowers to share their personal data.

[7] See, for example, Jayaratne and Strahan (1996), Célerier and Matray (2019), Nguyen (2019), Martín-Oliver, Toldrà-Simats, Vicente (2020), Ji, Teng, Townsend (2023), Sakong and Zentefis (2024), Cramer (2024), and Fonseca and Matray (2024).

show that bank branches continue to provide a social benefit by mitigating the risk of identity theft.

2 Identity theft in the digital economy

This section motivates our empirical analysis by developing the hypothesis that bank branch closures increase the risk of identity theft. We begin by defining identity theft and documenting its prevalence. Next, we describe how branch closures push economic activity toward digital platforms. Finally, we discuss how this shift exposes consumers to greater risks of identity theft.

2.1 What is identity theft?

Identity theft is any illegal activity involving the unlawful acquisition of an individual's personal information, typically for financial gain. The Identity Theft and Assumption Deterrence Act of 1998 established identity theft as a distinct federal crime. Under 18 U.S.C. §1028(a)(7), identity theft is defined as any act of

> "knowingly transfer[ring] or us[ing], without lawful authority, a means of identification of another person with the intent to commit, or to aid or abet, any unlawful activity [. . .]"

Identity theft is widespread in the United States. According to FTC data, there were more than 1.1 million reports of identity theft in 2023 alone. The top three types, credit card fraud (33%), loan or lease fraud (13%), and bank account fraud (11%), comprise over half of all cases. FTC data also highlights that this crime affects a broad demographic. Individuals aged 30 to 49 are the most frequently targeted group, accounting for 470,663 reports or 43% of all identity theft cases in 2023. Identity theft cases steadily decline beyond these prime years of workforce participation. We tabulate the breakdowns of identity theft reports by types and age groups in the Internet Appendix.

Identity thieves use both online and offline strategies to acquire personal data. Common online tactics include impersonation scams, social media data mining, cyberattacks exploiting platform vulnerabilities, and intercepting data through unsecured public Wi-Fi networks. Offline methods include physical theft such as retrieving sensitive information from discarded documents and installing skimming devices at ATMs and point-of-sale terminals.

2.2 Branch closures push consumers to increase digital engagement

Consumers affected by branch closures must learn to adopt digital banking and payment tools. Once adopted, these tools can reduce the marginal costs of using other digital services. A 2021 report by fintech firm Plaid suggests that such adoption boosts users' financial confidence with technology, encouraging broader use of online services. Digital payments are also increasingly embedded in everyday platforms, such as e-commerce, food delivery, and subscription services. Thus, branch closures can expand the use of digital tools, serving as a gateway to deeper, habitual engagement with the digital economy.

The two-sided nature of digital platforms accelerates the shift of economic activity from the physical realm to the digital economy. As more consumers engage with digital tools for payments, shopping, and services, businesses have stronger incentives to expand their digital presence.[8] In turn, the wider availability of digital services reinforces consumer reliance on these platforms. Consumers and businesses reflect this mutually reinforcing dynamic in their behavior and strategy. The Federal Reserve's Diary of Consumer Payment Choice shows that the share of remote purchases more than doubled between 2016 and 2023 (Cubides and O'Brien, 2023). Major retailers such as Nike and restaurants such as Domino's Pizza have also expanded their digital platforms to meet consumer preferences.

2.3 Branch closures, digital engagement, and identity theft

Bank branch closures push consumers toward greater digital engagement, increasing their exposure to privacy risks through two pathways. The first pathway concerns the background risk of exposure inherent in digitalization. As consumers engage in more digital activities, their personal identifying information (PII) is transmitted and stored across a fragmented and less secure network of retailers, mobile apps, and service providers. Thus, their expanding digital footprints structurally increase the odds that their PII is leaked, stolen, or abused. Third parties

[8] Branch closures may also force local businesses to travel further for cash management services, increasing logistical complexity and security risks. Debt covenants and insurance contracts often prohibit businesses from storing cash in stores. These added burdens make accepting and handling physical cash less viable, particularly for small businesses. In response, businesses affected by branch closures may have greater incentives to adopt digital payment systems and even shift their operations online, where cash handling is no longer required.

that process digital transactions may have inadequate data infrastructure and poor security practices. For example, the Federal Trade Commission has prosecuted data brokers for selling PII to scammers. Even sophisticated consumers can be affected because these breaches or data leaks often occur outside of one's control. A 2019 Pew Research survey reflects these concerns, finding that over 80% of Americans feel they have little or no control over their personal data collected online.

The second pathway reflects the transition risks that emerge as branch closures remove a crucial alternative for some vulnerable consumers who prefer in-person services. Bank branches provide face-to-face verification, secure document storage, controlled access, and continuous surveillance to deter and detect fraud in real time. Human interaction also helps prevent fraud, as bank staff are trained to spot suspicious behavior and advise customers on security best practices. With the loss of these physical touchpoints, some consumers may become more susceptible to adversarial tactics such as phishing attacks using phone calls, emails, SIM card swaps, and even generative artificial intelligence (AI) tools. These tactics can be especially harmful during the transition period, when some consumers are less familiar with security practices or warning signs of fraud. Despite the widespread rollout of digital financial services, a 2016 survey by the National Telecommunications and Information Administration finds that 45% of U.S. households avoid online financial transactions due to privacy and security concerns.

As branch closures drive a digital shift in economic activity and eliminate physical touchpoints, we hypothesize that they lead to an overall increase in identity theft. These risks are unlikely to be evenly distributed across the population. The transition may be relatively seamless for tech-savvy consumers, but it poses a significant behavioral adjustment for many others, particularly those with limited digital literacy or a long-standing reliance on offline transactions.

3 Data and validation

We describe the primary datasets used in our main empirical analysis and validate that branch closures change the banking patterns of consumers.

3.1 Data and descriptive statistics

We obtain identity theft reports (*ID theft reports*) from the Consumer Sentinel Network (CSN), which is maintained by the Federal Trade Commission (FTC).[9] The CSN database is a collection of consumer fraud reports, a subset of which includes identity theft. Identity theft is categorized into the following seven types: credit cards, loans or leases, bank accounts, government documents or benefits, employment or tax, phone or utilities, and others. We provide a detailed summary of CSN identity theft reports by types in the Internet Appendix. Our sample covers 2009 to 2022, providing a panel with variation at the Metropolitan Statistical Area (MSA) and year level.

We collect bank branch closures from the Federal Deposit Insurance Corporation's (FDIC) Summary of Deposits (SOD) files, which provides annual branch-level data for all U.S. depository institutions. In our main analysis, we compute *net branch closures* at the MSA-year level by subtracting the branch count in the current year from that in the previous year. Our demographic variables at the MSA-year level are sourced from the U.S. Census Bureau's American Community Survey (ACS) database. These variables include population size, proportion of people aged 60+, unemployment rate, median household income, gender and racial compositions, and educational attainment (over high school).

- Figure 1 here -

Figure 1 shows consumer fraud complaints from 2009 to 2022, with the subcategory of identity theft presented in blue. At the start of our sample period in 2009, consumer fraud complaints total 1.33 million reports and increases steadily over the subsequent decade. By the end of our sample period in 2022, the overall number of fraud reports exceeded 5.1 million, with identity theft accounting for 21.4% of all incidents. Notably, consumer fraud and identity theft reports surged after the COVID-19 pandemic.

[9] The CSN compiles complaints from a wide range of data contributors including law enforcement agencies (e.g., F.B.I. and U.S. Postal Inspection Service), state attorneys general and consumer protection offices (e.g., New York State Attorney General and Ohio Attorney General), federal government agencies (e.g., U.S. Department of Justice and U.S. Social Security Administration), consumer advocacy organizations (e.g., AARP Fraud Watch Network and National Consumers League), financial institutions (e.g., Mastercard International and JPMorgan Chase), and technology and telecommunications companies (e.g., Apple and Verizon Wireless). The Internet Appendix contains a full list of these data contributors.

In terms of financial losses, Figure 1 shows a steady upward trend in consumer fraud over the period. In 2009, total losses were approximately $1.38 billion, with relatively minor fluctuations in the following years. However, losses spike in recent years, reaching $6.47 billion in 2022. This sharp rise suggests an escalation in the financial impact of consumer fraud, which may be driven by more sophisticated fraud schemes and the increasing stakes of digital fraud, particularly identity theft.

- Figure 2 here -

Figure 2 shows the geographic distribution of identity theft reports across MSAs in 2022. Larger circles indicate a greater number of reports. MSAs that are hit hard by identity theft are typically found in urban and coastal areas, like New York City and Los Angeles. These affected regions usually have higher population densities, better access to digital services, and rely heavily on online transactions. They also tend to have higher income levels, making them attractive to adversaries. Given the substantial variation across regions, it is crucial to control for these demographic traits in our analysis to estimate the relation between bank branch closures and identity theft.

- Table 1 here -

Table 1 summarizes the variables and pairwise correlations. Panel A shows that MSAs average 1,420 identity theft reports annually with a standard deviation of 5,320. The 90$^{\text{th}}$ percentile of MSAs reports 2,470 reports, highlighting concentration in specific areas. MSAs experience an average of 2.9 net branch closures per year with a substantial standard deviation of 14.2. Decomposing this statistic, the average MSA has 3.9 branch closures and 1.0 branch opening. The average MSA has 18.4% of its population over 60, 50.1% male, 6.6% unemployed, 76.6% White, and 63.4% with at least a high school education. Average household income is $56,100 with significant dispersion (σ = $13,000). MSA population size varies widely, averaging 730,000 with a large standard deviation of 1,580,000.

Panel B of Table 1 shows that identity theft correlates positively with net branch closures (57%) and strongly with population size (76%), indicating a link between urbanization and identity theft. A 28% correlation with household income suggests higher-income areas may be more prone to fraud. In contrast, identity theft shows a negative association with the proportions of white residents and individuals who finished high school, suggesting possible distributional effects.

3.2 Validation: Consumer response to branch closures

Before presenting our main results, we first validate the premise that consumers affected by branch closures change their banking behavior through two tests. First, we investigate whether these consumers substitute visits to nearby branches. Second, we examine whether consumers use more digital banking services after branch closures in the local area.

3.2.1 Do affected consumers switch to nearby branches?

In principle, consumers who are affected by branch closures could travel to another branch, hence avoiding the need to engage more with the digital economy. On the other hand, if physical banking is hyper-localized with high switching costs (Choi and Loh, 2024), consumers may not readily switch to other nearby branches.

To estimate the elasticity of branch visits between the closed branch and other nearby branches, we collect weekly branch-level footfall data from the pass_by database. Using the Google Geocoding API, we map the street address, zipcode, city, and state of every pass_by branch to its latitude and longitude coordinates to precisely compute the distances between pass_by branches and closed branches.[10] Then, for every pass_by branch, we merge in the total *net branch closures* that occur (i) within the past 180 days and (ii) within a 1-mile distance band. Finally, for branch i in week w and for branch closures/openings i' within the $(\omega, \omega+1)$ mile distance band, we estimate equation (1):

$$\begin{gathered} \textit{bank branch visits}_{i,w} = \alpha + \beta \sum_{t=w-180\text{ days}}^{w-1\text{ day}} \sum_{i'} \textit{net branch closures}_{i',t} + \mathbf{X}_{z,w}^{\top}\boldsymbol{\lambda} + \varepsilon_{i,w}, \\ \text{s.t.} \quad \omega \text{ mi} \le \text{distance}(i,i') \le (\omega+1) \text{ mi}. \end{gathered} \tag{1}$$

The vectors $\mathbf{X}_{i,w}$ and $\boldsymbol{\lambda}$ represent vectors of control variables at the zipcode-year level and their corresponding coefficients, respectively.

- Figure 3 here -

Figure 3 shows that most consumers affected by branch closures do not readily switch to other nearby branches. A branch closure is associated with just 14.77

[10] We focus on savings banks by filtering for "stores" tagged with the NAICS code for savings institutions (522120). We obtain the latitude and longitude coordinates of the closed branches from the FDIC Summary of Deposits dataset.

($t = 2.61$) more visits per week to another branch within 0–1 miles. To put this estimate into context, the average branch closure/opening is surrounded by 8.14 other branches within 1 mile, and the average bank branch receives 665.89 visits per week. Somewhat surprisingly, only 18% (= 14.77 × 8.14/665.89) of the foot traffic from a closed branch redirects to nearby branches within a distance band of 0–1 mile. We find a weaker and statistically insignificant foot traffic redirection from closed branches to more distant locations. For example, the implied elasticity of branch visits drops sharply from 18% to 5.1% between the 0–1 mile and 1–2 mile distance bands.

Overall, our findings support the view that physical banking patterns are hyper-localized. Consumers inconvenienced by branch closures largely choose not to switch to other branches even if the alternatives are in close proximity. This pattern could reflect the high costs of switching banks or the shift toward digital banking methods away from branch-based services.

3.2.2 Consumers spend more time on banking mobile apps

Next, we test whether customers spend more time on the mobile app of a bank after it closes branches in the local area. We obtain consumer-level mobile app consumption from the Global Wireless Solutions (GWS) Magnify database. GWS curates a panel of opted-in Android smartphone users that is demographically representative of the United States population. Within the panel, GWS continuously tracks every consumer's mobile activity 24 hours a day, seven days a week. When a consumer uses her smartphone, we can observe among many statistics the (i) name of app used, (ii) time and duration of app usage, and (iii) her current latitude and longitude coordinates. Crucially, we can infer the zipcode of a consumer's primary residence by her most frequented location between 9.00 PM and 6.00 AM in her local time zone. The GWS Magnify database covers 194,530 consumers between 2019 and 2022.[11] We aggregate the bank mobile app usage

[11] We hand-match banks to their mobile apps on the Google Play Store using the following procedure. First, we collect a bank's website URL from the FDIC BankFind Suite. Next, we carefully search the website for the URL to the bank's mobile app. This URL contains a unique identifier of the mobile app on the Google Play Store. We manage to match 2,493 (out of 3,823) banks to their respective mobile apps. Most of the shortfall arises from banks that have ceased operations as of June 2025. Such banks typically delist their apps from the Google Play Store, but we can only observe apps that are currently hosted on it. A small number of banks do not offer mobile banking services.

data to the bank-county-month level.

The granularity of the GWS Magnify database allows us to examine the relation between bank mobile app usage patterns and branch closures. For all branch closures of bank b within the past 180 days in county c, we estimate in equation (2) whether there is a greater usage of bank b's mobile app among consumers in that county in a given month m:

$$\begin{aligned}\log(\textit{bank mobile app usage}_{c,b,m}) = \alpha + \beta \cdot \sum_{t=m-180\text{ days}}^{m-1\text{ day}} \textit{net branch closures}_{c,b,t} \\ + \mathbf{X}_{c,m}^{\top}\boldsymbol{\lambda} + \varepsilon_{c,m}.\end{aligned} \tag{2}$$

The vectors $\mathbf{X}_{c,m}$ and $\boldsymbol{\lambda}$ represent vectors of county-level control variables and their corresponding coefficients, respectively. We saturate the model with state × month and bank × month fixed effects, and cluster standard errors at the state, bank, and month levels.

- Figure 4 here -

Figure 4 presents the binned scatterplot corresponding to equation (2).[12] On average, consumers spend 24.2% ($t = 8.37$) more time on the mobile app of a bank that closes a branch in the local area. Our estimates suggest that branch closures push consumers to use significantly more digital banking. Although our results are grounded in granular mobile activity data, two caveats remain. First, our analysis necessarily focuses only on surviving banks. Second, we observe mobile app usage only among Android users, so our results may not generalize to behavioral shifts among iOS users. Nevertheless, we find evidence that local branch closures spur greater usage of digital banking, which can facilitate the shift of economic activities to the digital economy.

[12] To construct the binned scatterplot, we first regress separately log(*bank app usage*) and *net branch closures* on the control variables, as well as state × month and bank × month fixed effects. Next, we sort the residualized *net branch closures* into bins and compute the averages of residualized log(*bank app usage*) within these bins. Finally, we plot these residuals and the OLS best-fit line through them. By the Frisch-Waugh-Lovell theorem, the slope of this best-fit line equals β in equation (2).

4 Main analysis

In this section, we estimate the effect of branch closures on identity theft.

4.1 Branch closures and identity theft

We examine whether MSAs with bank branch closures report more cases of identity theft. To test this hypothesis, we estimate specification (3) for MSA i in year t:

$$\textit{ID theft reports}_{i,t} = \alpha + \beta \cdot \textit{net branch closures}_{i,t} + \mathbf{X}_{i,t}^{\top}\boldsymbol{\lambda} + \varepsilon_{i,t}. \tag{3}$$

The vectors $\mathbf{X}_{i,t}$ and $\boldsymbol{\lambda}$ represent vectors of control variables at the MSA-year level and their corresponding coefficients, respectively.

- Table 2 here -

Our estimates in Table 2 show that MSAs with more bank branch closures experience more identity theft. In column 1, we find that an additional bank branch closure is associated with 213.42 ($t = 5.91$) more *ID theft reports*. Our estimated effect is economically significant. A standard deviation shock to *net branch closures* leads to an increase of 3,030.6 (= 14.2 × 213.42) identity theft cases, representing 2.13 times of its unconditional sample mean. In column 2, our findings remain unchanged as we control for various MSA demographic characteristics and saturate our models with MSA fixed effects and year fixed effects.

It is plausible that banks invest in consumer-facing cybersecurity technologies as they close branches. Thus, MSAs with more branch closures may also see greater cybersecurity adoption, which could mitigate or confound the observed increase in identity theft. To better isolate the effect of branch closures on identity theft, we further control for cybersecurity technologies deployed by banks. We query every bank against the BuiltWith database to collect the list of technologies it deploys on its website over time.[13] Then, we count the technologies that are related to "bot and fraud detection" (e.g., ThreatMetrix, Fireblade, Fortinet). Finally, we aggregate these bank-level counts to the MSA-year level, defining *cybersecurity*

[13] BuiltWith is a commercial web-technology database that identifies and tracks the technologies used by websites worldwide. It provides structured data on detected software stacks (e.g., CMS, analytics, hosting, cybersecurity, and e-commerce platforms) and historical trends in technology adoption across domains.

as the average number of such technologies deployed by banks operating branches within an MSA in a given year.

Our findings are robust to controls for cybersecurity investments made by banks in the local area. In Columns 3 and 4, we restrict the sample to 2014–2022, as the BuiltWith data are sparse before 2014. Column 3 reproduces our baseline specification in this subsample and confirms that MSAs with more branch closures continue to report significantly higher identity theft ($+145.22, t = 5.48$). Column 4 shows that the estimated coefficient on *cybersecurity* is negative and statistically significant ($-903.51, t = 1.89$). This finding is consistent with the idea that cybersecurity technologies can mitigate fraud and reduce identity theft cases. However, the magnitude of the estimated coefficient on *net branch closures* remain largely unchanged ($+145.02, t = 5.46$). Thus, digital safeguards are useful, but they appear largely orthogonal to the effects of branch closures.

Overall, we find that identity theft increases with branch closures in the local area, and this finding holds even after accounting for banks' cybersecurity investments. Furthermore, MSA and year fixed effects ensure that our findings cannot be explained by persistent, unobserved local factors that affect both bank branching decisions and identity theft.

4.2 Identification strategy

The main threat to identification is that branch closures may coincide with unobserved, time-varying local factors that also affect identity theft. For example, banks may close more branches in areas where digital adoption rates are higher. Consumers in these areas may be wealthier or use more digital services, making them more attractive targets for cybercrime in the first place. Thus, even absent branch closures, we would expect to observe more identity theft in these areas.

To make causal inferences, we need variation in the incidence of branch closures that is plausibly exogenous to time-varying local factors. Our main identification strategy in this section is an instrumental variable (IV) approach within a staggered difference-in-differences framework.

4.2.1 Instrument for bank branch closures

Our instrument for branch closures is staggered exposures to postmerger consolidation of large banks following the approach used in Nguyen (2019). Bank mergers create operational redundancy such that a merged institution often closes branches in areas where the two previously separate banking networks overlap. Therefore, an MSA is first exposed in the year of a merger if it has branches of both the acquirer and target banks.

For the instrument to be valid, it must satisfy the exclusion restriction that MSA-level exposures to bank mergers are as good as randomly assigned with respect to local conditions. Below, we provide three reasons why this key identifying assumption holds for our merger exposure IV.

Bank mergers could be endogenous to local conditions if, for example, economic distress in an area forces banks to merge for cost-cutting purposes. Therefore, the first step in our identification strategy is to focus on mergers where both acquirer and target banks have at least U.S. $1 billion in premerger assets. At such a scale, the merging decision is plausibly removed from local economic conditions. Indeed, press releases suggest that these mergers are often motivated by broader goals, such as strategic expansion into new markets.

Next, we show that each MSA represents only a small share of the merging banks' overall operations. By construction, these banks are large institutions with extensive branch networks. The median acquirer (target) in our sample holds $8.6 billion ($2.2 billion) in assets and controls 81 (24) branches across 10 (4) MSAs. The median bank in the U.S. in comparison holds only $0.19 billion in assets and controls only 3 branches in a single MSA. Moreover, the median percentage of the acquirer (target) banks' deposits held in exposed MSAs before the mergers is only 1.03% (2.12%). Thus, the merger decision is unlikely driven by the local conditions of any specific MSAs.

Finally, we show in the next section that there are no significant demographic differences between exposed and unexposed MSAs in the pre-merger period. Importantly, they also do not differ significantly in measures of digital adoption. Thus, our IV is uncorrelated with digitalization trends or other observable characteristics that may drive identity theft.

4.2.2 Staggered difference-in-differences estimation

We outline our empirical design to identify the effect of branch closures on identity theft. Because merger exposures are staggered across years and MSAs, we estimate two Callaway and Sant'Anna (2021) difference-in-differences models—a first-stage and a reduced-form. Because the merger exposure IV satisfies the exclusion restriction, the first-stage isolates variation in branch closures that are orthogonal to local conditions. With the relevance condition met, the reduced-form identifies the effect of merger exposures on identity theft that comes through plausibly exogenous branch closures. Thus, the ratio of the reduced-form to first-stage estimates gives the causal effect of a branch closure (x) on identity theft (y), or the local average treatment effect (LATE):

$$\underset{\text{(ID theft per branch closure)}}{\text{LATE}} = \frac{\overbrace{E\left(y \mid exposed = 1\right) - E\left(y \mid exposed = 0\right)}^{\text{Reduced-form effect}}}{\underbrace{E\left(x \mid exposed = 1\right) - E\left(x \mid exposed = 0\right)}_{\text{First-stage effect}}}. \tag{4}$$

The average treatment effects on treated (ATTs) are the building blocks of this estimator and are defined as follows,

$$\text{ATT}(g,t) = E[Y_t - Y_{g-1} \mid G_g = 1] - E[Y_t - Y_{g-1} \mid C = 1]. \tag{5}$$

Among MSAs that are first exposed to merger shocks at year g (i.e., $G_g = 1$), the first expectation takes the difference in the outcomes of Y at year t and at year $g-1$. The second expectation computes the same difference among control MSAs that are never exposed to the merger shocks (i.e., $C = 1$). These ATT(g,t)'s form the building blocks of the estimator.

To understand dynamic effects, we can use these building blocks to construct an event-time plot. At every event-time τ in years, we aggregate the ATT(g,t)'s by averaging over all exposed MSAs that have been observed at that event-time:

$$\text{ATT}(\tau) = \frac{1}{\sum_{g,t} 1(t-g=\tau)} \sum_{g,t} 1(t-g=\tau) \cdot \text{ATT}(g,t). \tag{6}$$

We first support our key identifying assumption by plotting the event-time trends of standardized MSA characteristics in the preexposure period (i.e., $\tau < 0$).

Figure 5 shows that preexposure differences in demographics and measures of digital adoption between exposed and unexposed MSAs are statistically indistinguishable from zero. Furthermore, the point estimates of the ATT($\tau < 0$)'s are economically small, mostly well below 0.1 standard deviations of the characteristics. In every plot, we also report the average of the preexposure ATTs. Across all characteristics, these averages are small and lack statistical insignificance. These patterns show that our IV is uncorrelated with MSA characteristics, supporting the exclusion restriction.

- Figure 5 here -

Next, we examine how merger exposures affect branch closures and identity theft by plotting their event-time ATT(τ) for $\tau \in [-4, +4]$.

- Figure 6 here -

The top subfigure in Figure 6 shows no statistically significant differences in net branch closures in the (blue) preexposure period. However, we find a mostly increasing trend of net branch closures in the (red) postexposure period. These patterns validate the relevance condition of our IV design by showing exposed MSAs have significantly more branch closures after the shocks. The bottom subfigure corresponds to the reduced-form of our IV design. Here, we find a statistically significant and steady rise in identity theft reports in the postexposure period. The lack of pretrend differences supports our identifying assumption that the assignment of merger exposures is orthogonal to local factors that affect identity theft.

- Table 3 here -

We summarize the ATTs of merger exposures on branch closures and identity theft reports in Panel A of Table 3. The preexposure ATT averages the ATT(τ)'s for event-time $\tau \leq -1$. The postexposure ATT does likewise but for $\tau \geq 0$. Equation (7) defines the overall ATT, which is the simple average of all ATT(g,t)'s for all $t \geq g$:

$$\text{Overall ATT} = \frac{1}{\sum_{g,t} 1(t \geq g)} \sum_{g,t} 1(t \geq g) \cdot \text{ATT}(g,t). \tag{7}$$

Figure 6 shows that treatment effects obtain only after the merger exposures with large and statistically significant postexposure ATTs. Consistent with the assumption of parallel trends, the preexposure ATTs in branch closures and identity theft reports are small and statistically insignificant. The first-stage indicates that exposed MSAs have +2.79 ($t = 2.90$) more branch closures. In the reduced-form, we find that they have (+455.56, $t = 2.72$) more identity theft reports. The ratio of these estimates yield the local average treatment effect (LATE): each branch closure leads to an increase of +163.28 (= 455.56/2.79, $t = 4.75$) identity theft reports. Reassuringly, this estimate is smaller but relatively close in magnitude compared to its OLS counterpart (column 1 of Table 2) because we generally expect unobserved heterogeneity to induce an upward bias in the OLS estimate.

We perform robustness checks to strengthen the credibility of our causal estimates. Although we have shown that our counterfactuals (i.e., unexposed MSAs) are—on average—similar to exposed MSAs across observable characteristics and preexposure outcomes, this aggregate similarity could mask MSA-level heterogeneity. To create a more credible counterfactual trajectory, we estimate in Section IA.6 of the Internet Appendix synthetic difference-in-differences models (Arkhangelsky, Athey, Hirshberg, Imbens, et al., 2021) that optimally weight unexposed MSAs to create granular counterfactuals by matching on the pre-exposure outcomes and characteristics of each exposed MSA. In these models, a branch closure leads to an increase of similar magnitude of +127.45 ($t = 2.22$) identity-theft reports.

4.2.3 Do selection concerns affect our identification strategy?

In this section, we discuss the selection concern inherent in IV research designs. An instrument encourages units to change their treatment status, but typically not all units will respond accordingly (i.e., noncompliance). For example, the passage of state-level legislation that incentivizes firms to increase research & development (R&D) activities may not induce all firms to expand R&D efforts.[14] Similarly, in our setting, merger exposures exogenously encourage the closure of redundant branches, but the merged bank may not "comply" and ultimately selects the actual branches to close.

- Panel A of Table 4 here -

We evaluate the selection issue in two ways. Among exposed MSAs, we first compare the characteristics of those that close and those that do not. One concern is that consolidated banks systematically close redundant branches only in MSAs

[14] In randomized control trials, treated units can also exhibit noncompliance. For example, treated individuals in clinical trials often fail to take treatment or properly administer dosage (e.g., medication).

where consumers are digitally-ready. However, our data do not support this view. Panel A of Table 4 shows no statistically significant differences in internet penetration rates and ownership of computers and smartphones ($t = 0.11$–0.18). Another possibility is that banks close redundant branches in economically stagnant areas or areas where consumers are less sophisticated. We also find no support for this view. Exposed MSAs that close redundant branches are similar to those that do not close branches across demographic measures including income ($t = 0.08$), age ($t = 0.87$), education attainment ($t = 0.69$), and unemployment ($t = 1.56$). Overall, our evidence suggests that the decision to close redundant branches in exposed MSAs is not systematically related to observable local characteristics that could correlate with identity theft.

Next, we provide a simple example to illustrate the source of identification in our research design. Suppose there is noncompliance such that merged banks close redundant branches only in MSAs with high digital adoption rates (i.e., high digital). This would imply that branch closures (x) do not vary with exposure in low-digital MSAs because banks never close branches in those areas. Thus, exposed low-digital MSAs contribute no identifying variation to the first-stage estimate because:

$$E(x \mid exposed = 1, digital = L) = E(x \mid exposed = 0, digital = L) = 0. \tag{8}$$

Instead, the variation that identifies the first-stage originates only from a comparison of exposed and unexposed high-digital MSAs:

$$E(x \mid exposed = 1, digital = H) - E(x \mid exposed = 0, digital = H) > 0. \tag{9}$$

If the exclusion restriction is satisfied, merger exposures affect identity theft only through branch closures, so our example above shows that low-digital MSAs cannot contribute identifying variation in our IV design. Therefore, our research design identifies a causal LATE through the comparison of exposed and unexposed high-digital MSAs.

More generally, our identification strategy is valid as long as the exclusion restriction holds such that merger exposures exert an exogenous "encouragement" that shifts branch closures. Even if there is a characteristic, such as digital adoption rates, that affects compliance, the LATE is still causal but it would be identified only for the subgroup of MSAs that comply with the IV (Imbens and Angrist, 1994). However, this observation naturally raises questions of how well the LATE generalizes to other MSAs. Would the effect of branch closures on identity theft in these MSAs, if any, be smaller or larger?

To examine the generalizability of our LATE, we use the Marbach and Hangartner (2020) framework to estimate characteristic means of the following subgroups. Complier MSAs, for which the LATE is identified, respond to the IV by (not) having at least one branch closure in the (absence) presence of merger exposures. Always-taker (never-taker) MSAs do not respond to the IV because they always (never) close branches regardless of merger exposures. We defer technical details of this framework to the Internet Appendix. Similarity in characteristics between the compliers and the average MSA would increase confidence in the generalizability of our LATE.

- Panel B of Table 4 here -

Panel B of Table 4 presents the estimated characteristic means in 2010 for the full sample and these subgroups. We find that complier MSAs resemble the average MSA on measures of digital adoption. They have comparable internet subscription rates and ownership of internet-enabled computing devices (i.e., desktops, laptops, and smartphones). They also have similar demographics to the average MSA, apart from having slightly larger populations. Overall, these patterns suggest that our LATE has external validity to the average MSA in our sample.

For completeness, we also describe the always-taker and never-taker MSAs. Against the secular trend of branch closures in the U.S., always-taker MSAs make up a large proportion of our sample. Always-taker MSAs have smaller, older populations and fewer bank branches, suggesting that they are relatively unattractive banking markets. In contrast, never-taker MSAs have significantly larger populations, more branches, and higher household incomes. These MSAs are likely population centers where branch presence may be profitable or strategically important. Interestingly, the digital adoption rates in both subgroups are similar to the average MSA and the compliers.

4.2.4 Quantifying the effects of branch closures on identity theft

Given that the complier characteristics analysis supports the generalizability of our LATE, we attempt to quantify the effects of branch closures on the prevalence of and losses from identity theft. First, we compute the Callaway and

Sant'Anna (2021) calendar ATTs, defined as the average treatment effect for MSAs that are or are already exposed to large bank mergers in that year. For MSAs that were first exposed to merger shocks at year g, the calendar ATT in year t is the average over all ATT(g,t) with $t \geq g$:

$$\text{Calendar ATT}(t) = \frac{1}{\sum_g 1(t \geq g)} \sum_g 1(t \geq g) \cdot \text{ATT}(g,t) \quad \text{for } t \in [2011\,..\,2022]. \tag{10}$$

To obtain a measure of the increase in identity theft per instrumented branch closure, we then compute the (calendar) yearly Wald estimate ω_t as the following ratio,

$$\omega(t) = \frac{\text{Calendar ATT}^{\text{ID theft reports}}(t)}{\text{Calendar ATT}^{\text{net branch closures}}(t)}. \tag{11}$$

Next, we impute the yearly total increase in identity theft due to branch closures by multiplying $\omega(t)$ by the actual *net branch closures* in the MSAs and summing them up. This imputation, in spirit, applies the LATE to all MSAs, even though it is defined as the treatment effect for only complier MSAs. This extrapolation is reasonable because Panel B of Table 3 shows that complier MSAs closely resemble the average MSA across most observable dimensions. Subscripting MSAs by i, we compute the number of identity theft cases attributable to branch closures as follows,

$$\textit{total num. ID theft reports}(t) = \sum_i \omega(t) \times \textit{net branch closures}_{i,t}. \tag{12}$$

Finally, we impute the total annual losses stemming from the increase in identity theft reports by factoring in the average per-report dollar loss (*avg. loss*) and the proportion of reports with reported losses (*% reports with loss*) in the state-year:

$$\begin{aligned} \textit{total losses}(t) = \sum_i \omega(t) & \\ & \times \textit{net branch closures}_{i,t} \\ & \times \textit{avg. loss}_{\text{state}(i),t} \\ & \times \textit{\% reports with loss}_{\text{state}(i),t}. \end{aligned} \tag{13}$$

The FTC Consumer Sentinel Network provides these statistics only at the state level, so we match them with their constituent MSAs. Not all fraud reports involve financial losses because some victims may ultimately be able to recover or reverse their damages (e.g., banks waiving fraudulent transactions). The proportions of fraud reports accompanied by dollar loss amounts vary significantly across states and years.

- Figure 7 here -

Figure 7 shows that the imputed financial losses to identity theft have risen steadily over the years, peaking at over U.S.$1.8 billion in 2021 during the COVID-19 pandemic. During the pandemic, consumers experienced an accelerated push onto the digital economy as many in-person economic activities were curtailed. Even as the pandemic began to ease in 2022, the losses remained high at almost U.S.$1.4 billion. The steady rise in identity theft cases is not only due to the multi-year decline in physical branch banking—our Wald estimates show that the impact per branch closure is also stronger in later years. This pattern may reflect the increasing risk of identity theft faced by consumers as more economic activities digitalize over time. We tabulate the annual statistics from this analysis in the Internet Appendix.

4.2.5 Linking banks to customer harm

Despite the strength of our IV design, it remains possible that the rise in identity theft is driven by a subgroup of local consumers who are *not* actually customers of banks that close branches. Given the exogenous nature of our IV, such misattribution is unlikely to be systematic across MSAs. Nevertheless, this possibility challenges us to establish a more direct link between branch closures and the focal bank's customers.

To this end, we turn to the Consumer Complaint Database (CCD) administered by the Consumer Financial Protection Bureau (CFPB). Although the CCD covers only the largest banks, its key advantage lies in its highly granular data, which record individual geotagged consumer complaints against specific banks.[15] This granularity allows us to observe both the exposure (i.e., branch closures) and

[15] Only banks with at least U.S.$ 10 billion in assets are subject to the supervisory authority of the CFPB. Identity theft in financial services is also a particularly serious problem. According to the FTC, there were 614,711 reports of identity theft related to credit card and bank fraud in 2022. These reports account for 44.2% of all identity theft cases, marking a 16.8% year-over-year increase.

the outcome (i.e., complaints) on a single group of consumers—the bank's own customers. Thus, we can more directly identify the harm from branch closures that falls on the bank's own customers, rather than being misattributed to other local consumers.

We process the CCD data by first hand-matching the names of banks on the CCD to their FDIC identifiers. These banks are inherently large because only depository institutions with over \$10 billion in assets are subject to CFPB oversight. Altogether, we match 386,549 complaints to 136 banks that operate branches in 2,533 counties. Although we observe a consumer's zipcode, we cannot know whether she used a specific bank branch in her residential zipcode. Thus, we aggregate at the bank-county-month level, (i) consumer complaints and (ii) net branch closures within the past 180 days.

We hypothesize that branch closures are followed by a higher incidence of complaints filed by consumers in the area. To test this hypothesis, we estimate equation (14) for county c, bank b, in month m:

$$1(complaint_{c,b,m} > 0) = \alpha + \beta \cdot \sum_{t=m-180\text{ days}}^{m-1\text{ day}} net\ branch\ closures_{c,b,t} + \mathbf{X}_{c,m}^{\top}\boldsymbol{\lambda} + \varepsilon_{c,m}. \tag{14}$$

The vectors $\mathbf{X}_{c,m}$ and $\boldsymbol{\lambda}$ represent vectors of control variables at the county-year level and their corresponding coefficients, respectively.

- Table 5 here -

The results in Table 5 indicate that branch closures are followed by a higher incidence of consumer complaints. Column 1 shows that a branch closure in the county is associated with a 326 bps (t = 9.78) higher probability of complaints lodged by customers of the bank in the area. As we have both county and month fixed effects, this finding is not explained by unobserved heterogeneity across counties or variation in complaint incidence across time. All complaints have the potential to be relevant to consumer fraud even when they are not explicitly labeled in the CCD as such (Bian, Ma, and Tang, 2023). For example, an "incorrect information on your report" label may indicate that an adversary has applied for a loan using a customer's stolen identity. Likewise, a complaint about a bank "closing your account" could stem from fraudulent activity if the closure was due to suspicious transactions or identity theft.

To sharpen our analysis, we distill consumer narratives from complaints to identify those most likely related to identity theft. We use a zero-shot-learning model, the bart-large-mnli model developed by Facebook, to classify the nature of every complaint. For every complaint narrative, we apply the hypothesis format "I am reporting a case of {label}" with the following labels: "fraud", "harassment", "inaccuracy", and "identity theft". The model produces a probabilistic score for each label, indicating the likelihood that the complaint matches each category. Then, we classify each complaint by assigning it the label with the highest score.

Having classified complaints that are related to identity theft, we reestimate equation (14). Column 2 shows that a bank branch closure in the county increases the probability of identity-theft-related complaints from local customers by 34 bps ($t = 3.71$). This estimate is economically meaningful—it nearly doubles the baseline probability of 18.8 bps per bank-county-month.

Finally, we saturate the model with bank × county, county × month, and bank × state × month fixed effects. This stringent specification help us address endogeneity concerns related to (i) banks matching to specific geographies and (ii) unobserved heterogeneity in consumers across counties in a given month. With this stringent specification in column 3, we continue to find a significantly positive relation (+17 bps, $t = 2.00$) between branch closures and identity theft complaints.

5 Vulnerable communities

The transition to digital platforms and services can affect consumers in unequal ways. Although sophisticated consumers may be more able to navigate the digital economy safely, some vulnerable communities can face significant challenges. Thus, branch closures may disproportionately affect consumers who are less digitally savvy. To test this idea, we condition our merger-exposure IV on the ability of bank customers to engage safely with the digital economy. We expect consumers who are less able or ready to make the digital transition to experience a stronger effect of branch closures on identity theft.

We first classify each large bank merger by the acquirer-bank customers' reliance on bank branches. Following Jiang, Yu, and Zhang (2025), we measure *branch reliance* of a bank as the ratio of the number of branches to its total deposits. A high *branch reliance* measure of a bank implies that its customers are likely more accustomed to physical banking and hence less able or prepared to go online. Among the set of large bank mergers, we classify a merger as having high (low) *branch reliance* if its value is above (below) the yearly median. We predict that branch closures has a stronger effect on identity theft when consumers rely more on bank branches in the first place.

- Table 6 here -

The results in Panel A of Table 6 support our prediction. The IV in columns 1 and 2 is an MSA's exposure to large bank mergers in which the acquirer has high *branch reliance*. On average, this exposure leads to 11.67 (t = 2.70) net branch closures and 2,362.36 (t = 2.92) ID theft reports. Thus, the Wald estimate reveals that a branch closure causes 202.43 (= 2,362.36/11.67, t = 8.76) more cases of identity theft. We find weaker results in columns 3 and 4 where the IV is an MSA's exposure to mergers characterized by low *branch reliance*. The second-stage estimate is less than one-fifth the size of the estimate in column 2. Moreover, the Wald estimate is noisy (t = 1.70) and uninformative because the first-stage estimate is statistically insignificant.[16]

In Panel B, we classify bank mergers by the *digital focus* of the acquirers. We measure the *digital focus* of a bank as the ratio of the all-time download volume of its mobile app on the Google Play Store to its number of branches. A digitally focused bank likely has fewer number of branches per mobile app user, compared to another bank that has a digital presence but also caters to the physical banking preferences of some customers. So, customers of a bank with lower (higher) *digital focus* are more (less) likely to be pushed online after branch closures. Columns 1 and 2 indicate that a branch closure leads to 145.96 (= 281.70/1.93, t = 2.52) more cases of identity theft in MSAs exposed to mergers with low *digital focus*.[17] We find a much weaker effect (86.11, t = 1.80) when we switch to bank mergers marked by high *digital focus*.

In Panel C, we use Census Bureau data to separate exposed MSAs by the proportion of people who have internet subscriptions (*consumer tech-savviness*). Columns 1 and 2 show that a branch closure leads to 208.54 (= 398.31/1.91, t = 10.56) more cases of identity theft in MSAs with low *consumer tech-savviness*. In

[16] The Wald estimate will be inflated as it is the second-stage estimate divided by a first-stage estimate that is statistically indistinguishable from zero (Jiang, 2017).

[17] We classify mergers by the *digital focus* of their acquirers. An acquirer is labelled as low (high) *digital focus* if its value is below (above) the yearly median.

contrast, this effect is 15.4% (176.29, $t = 3.03$) smaller among MSAs with high *consumer tech-savviness*.

Overall, we find that the treatment effect is stronger when (i) consumers are more reliant on bank branches, (ii) banks are less digitally focused to start with, and (iii) consumers are less tech-savvy. Therefore, our findings support the view that branch closures may disproportionately harm some consumers who are ill-equipped or unable to safely engage with the digital economy.

6 Pathways

We perform additional tests to investigate two pathways that underpin our main findings. First, we examine whether branch closures lead consumers to engage more deeply with the digital economy, increasing the margin over which identity theft can occur. Second, we demonstrate that consumers' exposure to adversarial activities makes the transition to the digital economy perilous.

6.1 Evidence of digital engagement

If branch closures increase identity theft through our digital engagement pathway, we expect consumer's consumption patterns to change as a result. To measure consumption, we track the mobile app usage of consumers and analyze their online and offline expenditures.

6.1.1 Mobile app usage

Mobile app usage is a behavioral measure of how consumers engage with the digital economy. Many economic activities, such as finance, commerce, and entertainment, are now conducted on mobile platforms. Thus, changes in time spent on mobile apps are likely to reflect shifts in consumption from offline to online channels. We return to the GWS Magnify database introduced in Section 3.2.2. Our structural relation of interest is β—the effect of local branch closures over the past 180 days in county c on a consumer's (i) time spent on all mobile apps (excluding banking apps) in month m:

$$\textit{mobile app usage}_{i,c,m} = \alpha + \beta \sum_{t=m-180\text{ days}}^{m-1\text{ day}} \textit{net branch closures}_{c,t} + \mathbf{X}_{i,c,m}^{\top}\boldsymbol{\lambda} + \varepsilon_{i,c,m}. \tag{15}$$

The vectors $\mathbf{X}_{i,c,m}$ and $\boldsymbol{\lambda}$ represent vectors of county and consumer characteristics and their corresponding coefficients, respectively.

To address endogeneity concerns, we instrument branch closures with consumers' staggered exposures to large bank mergers (Section 4.2). Although our panel of mobile app usage has a relatively short sample period, its high frequency (monthly) and granularity (consumer-level) ensure that we have sufficient statistical power to perform this test. We present the Callaway and Sant'Anna (2021) ATTs of merger exposures in Table 7.

- Table 7 here -

We find that branch closures in the local area lead consumers to increase mobile app usage. The first stage in column 1 shows that consumers exposed to large bank mergers encounter 0.098 ($t = 8.73$) more branch closures per month. In the second stage contained in column 2, we find that an exposed consumer increases her mobile app usage by 2.450 ($t = 2.11$) hours per month. Taken together, the Wald estimate implies that one branch closure causes the average consumer to spend an additional 24.91 ($= 2.450/0.098, t = 2.04$) hours on mobile apps per month. The average consumer in our dataset experiences 0.432 branch closures in the county per month. Thus, she can expect to spend an additional 10.76 ($= 24.91 \times 0.432$) hours per month, or 0.36 hours per day, on mobile apps due to branch closures. As a benchmark, an industry survey finds that Americans spend an average of 4.65 hours per day on their mobile phones in 2022 (Statista, 2025).

Overall, using microdata on mobile app usage, we find a causal link between branch closures and engagement with the digital economy.

6.1.2 Consumer expenditures

To complement our analysis of mobile app usage, we examine whether branch closures are accompanied by measurable changes in consumers' spending and transaction activity between online and offline channels. Changes in consumption patterns reflect how economic activity reallocates across transaction modes. If branch closures push consumers to engage with the digital economy, we expect to see a greater reliance on online transactions at the expense of offline ones.

To measure shifts in consumption patterns, we use SafeGraph data to construct two aggregate measures at the MSA-month level. The *online spending gap* is the difference in dollar values between online transactions and offline transactions, scaled by the total dollar value of all transactions. The *online transaction gap* is the corresponding measure for transaction volume. Both measures capture the relative tilt of consumer activity toward online channels and abstract away from changes in aggregate local economic activity.

- Table 8 here -

As before, we instrument for branch closures with consumers' staggered exposures to large bank mergers and estimate the effects using the Callaway and Sant'Anna (2021) doubly-robust estimator. Table 8 shows that an MSA exposed to these mergers have 0.311 ($t = 3.79$) additional branch closures per month. These exposures increase the monthly *online spending gap* and *online transaction gap* by 1.1 pp ($t = 2.39$) and 0.9 pp ($t = 2.01$) transactions, respectively. The Wald estimates indicate that a branch closure widens the monthly online spending gap by 3.6 pp ($= 0.011/0.311, t = 2.90$) and the online transaction gap by 3.0 pp ($= 0.009/0.311, t = 2.58$). These shifts to online channels are economically meaningful, representing 0.58 (0.53) times the standard deviation of *online spending gap* (*online transaction gap*).[18]

In summary, we find causal evidence that branch closures shift consumption activity toward online channels, reflecting a deeper engagement with the digital economy. Although this change can improve convenience and access, it also expands the surface over which adversaries can target consumers for identity theft.

6.2 Adversarial tactics

Law enforcement faces significant challenges in combating identity theft because adversaries continually evolve their tactics to exploit security vulnerabilities. To shed light on their otherwise opaque operations, we focus on two common

[18] The sample standard deviations of *online spending gap* and *online transaction gap* are 6.18% and 5.66%, respectively.

adversarial tactics used in identity theft—unwanted calls and phishing attacks.

6.2.1 Unwanted calls

In many identity theft schemes, phone calls serve as the initial attack vector used to target consumers. Adversaries are known to impersonate trusted institutions (e.g., banks and government agencies) on these calls to extract PII such as Social Security numbers and bank account details. Greater digital engagement after branch closures make consumers more vulnerable to these attacks for two reasons. First, consumers' phone contact details are often shared in the delivery of digital services. Security breaches on the provider's end may then expose these details to adversaries. Second, as digital services often involve phone-based interactions, such as SMS verification or customer support, adversaries can readily exploit consumers' trust placed on these systems.

To compound the problem, advances in telecommunications technology have enabled the rise of "robocalls", which are programmatically automated calls that can reach a large number of consumers at a low cost. This trend prompted the FTC to alert consumers to the use of robocalls in phone scams, especially in relation to identity theft (Federal Trade Commission, 2024b). In a notable enforcement action, the FTC fined VoIP (Voice-over-Internet-Protocol) provider XCast Labs U.S. $10 million in January 2024 for enabling billions of illegal robocalls, many of which impersonated government agencies to commit fraud.[19]

We collect data to test whether unwanted calls become more prevalent after branch closures in the local area. From the FTC Do-Not-Call (DNC) Reported Calls database, we first collect 1,015,744 complaints of unwanted calls and robocalls lodged between 2014 and 2022. For brevity, we refer to both categories collectively as "unwanted calls". Next, we aggregate these complaints at the county-month level and merge in net branch closures in the county over the past 180 days. Our structural relation of interest in equation (16) is β—the effect of local branch closures over the past 180 days in county c on the number of unwanted

[19] Under the FTC's Telemarketing Sales Rule, an illegal robocall is one that, among other violations, fails to obtain prior consent from the recipient, disregards the National Do-Not-Call (DNC) registry, or uses false or misleading information to induce sales or payments.

calls received by consumers in month m:

$$\textit{unwanted calls}_{c,m} = \alpha + \beta \sum_{t=m-180 \text{ days}}^{m-1 \text{ day}} \textit{net branch closures}_{c,t} + \mathbf{X}_{c,m}^{\top}\boldsymbol{\lambda} + \varepsilon_{c,m}. \quad (16)$$

The vectors $\mathbf{X}_{c,m}$ and $\boldsymbol{\lambda}$ represent vectors of control variables at the county-year level and their corresponding coefficients, respectively.

- Table 9 here -

Our Callaway and Sant'Anna (2021) difference-in-differences estimates in Table 9 indicate that local consumers receive more unwanted calls following branch closures in the area. Column 1 shows that exposure to large bank mergers can instrument for branch closures. Exposed counties have 1.135 ($t = 4.81$) more *net branch closures* per month after the occurrence of these mergers. In column 2, we find that exposed counties receive 2.105 ($t = 2.02$) more *unwanted calls*, implying a Wald estimate of 1.855 ($= 2.105/1.135, t = 1.86$) calls per branch closure. This effect is economically significant, representing 23.7% ($= 1.855/7.840$) of the average number of unwanted calls targeting a county per month.

To verify that our results are not driven by general shifts in calling activity unrelated to fraud, we examine telemarketing calls, which are separately reported to the FTC, as a placebo outcome. Column 3 shows that exposed counties do not receive significantly more telemarketing calls ($-1.271, t = 1.35$) after large bank mergers. Thus, the increase in *unwanted calls* following local branch closures cannot be attributed to benign shifts in telemarketing trends and instead likely reflects a rise in adversarial activity. As a robustness check, column 4 repeats our estimation on *unwanted calls (net)*, defined as *unwanted calls* minus *telemarketing calls*, to better isolate adversarial activity. We find that exposed counties report 2.690 ($t = 2.33$) more cases of *unwanted calls (net)*, or a Wald estimate of 2.370 ($= 2.690/1.135, t = 2.10$) calls per branch closure.

Overall, our evidence suggests that consumers are more targeted through unwanted phone calls after local branch closures. These patterns are consistent with the idea that the transition to the digital economy expands the exposure of PII, creating new pathways for identity theft.

6.2.2 Phishing attacks

According to the FTC, a significant portion of identity theft cases stem from phishing attacks, where adversaries deceive consumers into providing PII through deceptive emails or websites (Federal Trade Commission, 2024a). The adversaries then use the information to open bogus accounts or access existing ones. As branch closures push consumers toward digital banking platforms, they may become more exposed to such phishing attacks.

In phishing attacks, adversaries often buy unlicensed developer kits from black markets to clone legitimate websites for malicious use. Thus, many phishing websites have the same exact "favicons" (i.e., abbreviation for favorite icons) as their legitimate counterparts. Favicons are small logos shown on internet browser tabs to provide a recognizable brand identity to consumers.

By tracking the appearances of suspicious websites using the favicon of a bank's website, we can measure the intensity of phishing attacks on its customers. We begin by hand-collecting "favicons" from the websites of the top 100 U.S. banks by total assets in 2022. Next, we represent every favicon as an integer hash using the MMH3 (MurmurHash3) hashing algorithm. For example, the favicon of the Wells Fargo Bank website maps to the hash value "893468414".[20] We then query this hash value on the Shodan search engine to obtain the monthly number of web servers that host websites with the same favicon. Shodan indexes devices connected to the internet, including their open ports, running services, and potential vulnerabilities. To exclude legitimate websites from our query, we filter out those with secure-socket-layers (SSL) certificates, which authenticate website identities and enable secure data transmission. Some of the remaining websites might not be related to phishing attacks but are still suspicious from a cybersecurity standpoint. If anything, this measurement noise biases our estimates toward zero.

We create a novel measure of consumers' exposure to phishing attacks at the MSA-bank-month level. For bank b servicing the set S_b of MSAs indexed by i in month m, we construct a MSA-bank-month measure of *phishing exposure* as the product of (i) the month-over-month change in the number of phishing sites impersonating a bank, and (ii) the population weight of the MSA relative to all

[20] This favicon can be found at `https://www.wellsfargo.com/favicon.ico`.

MSAs serviced by the bank,

$$\textit{phishing exposure}_{i,b,m} := \Delta_{m-1,m}(\#\ \textit{phish sites})_b \times \frac{\textit{population}_{i,m}}{\sum_{i \in S_b} \textit{population}_{i,m}}. \tag{17}$$

Intuitively, phishing attacks on a bank have a greater effect on an MSA with a larger population weight. With more potential victims, consumers in these MSAs are exposed more to any given phishing attack. Thus, *phishing exposure* could vary over time due to the emergence of phishing attacks and changes in the bank's operational exposure to different MSAs.[21]

- Table 10 here -

To investigate the relation between *phishing exposure* and branch closures, we estimate the equation (18) at the MSA-bank-month level:

$$\textit{phishing exposure}_{i,b,m} = \alpha + \beta \cdot \textit{net branch closures}_{i,b,m} + \varepsilon_{i,b,m}. \tag{18}$$

Column 1 of Table 10 shows a positive and statistically significant relation between *net branch closures* and *phishing exposure*. Thus, when a bank closes branches in an MSA, its customers in that area are potentially exposed to more phishing attacks.

Next, we examine whether greater exposure to phishing attacks translates to more identity theft. We aggregate *phishing exposure* to the MSA-year level by computing the average of fitted values from equation (18), weighted by the number of branches each bank operates in the MSA:

$$\textit{MSA phishing}_{i,t} := \sum_{\substack{i, \\ \text{year}(m)=t}} \textit{phishing exposure}^{(\text{pred})}_{i,b,m} \times \frac{(\#\ \textit{branches})_{i,b,m}}{(\#\ \textit{branches})_{i,m}}. \tag{19}$$

Column 2 shows that *MSA phishing* has a positive and statistically significant relation with identity theft. A standard deviation increase in *MSA phishing* is associated with 375.46 (= 0.0245 × 14,590) more reports of identity theft, representing 26.4% of its unconditional sample mean. In columns 3 and 4, we repeat

[21] We focus on the change in counts of phishing sites because they tend to have very short lifetimes, often disappearing within days or weeks. Thus, the month-over-month change better captures consumers' incremental exposure to new phishing attacks. Whereas, the level would reflect transient noise rather than a persistent stock of risk.

our analysis using an out-of-sample variant of *MSA phishing*. Using predictive loadings from equation (18) estimated between 2018–2019, we continue to find that *MSA phishing* is strongly associated with identity theft in 2020–2022.

Overall, our results indicate that branch closures in an MSA could expose local consumers to more phishing attacks. In turn, higher exposures to these attacks are associated with more identity theft reports.

7 Conclusion

The digitalization of financial services has reshaped the conduct of economic activities. We examine the privacy costs borne by consumers in the shift from brick-and-mortar banking to online financial services. Following branch closures, we show that consumers increase their overall engagement with the digital economy and face greater exposure to adversarial attacks. These changes expand the margin over which consumers' PII is exposed, increasing the risks of identity theft. Digitalization often affects communities in unequal ways. We find that vulnerable communities, such as the less digitally savvy, are disproportionately affected by the transition to the digital economy.

Our findings underscore the need for caution as societies transition towards a digital economy. This includes promoting financial literacy, providing support to consumers experiencing financial difficulties, and investing in cybersecurity measures to prevent cybercrime. In addition, our results highlight the complexity of navigating the digital landscape and the need for proactive policy responses. Policymakers now spend substantial resources to educate the public and have created cybercrime networks to quickly alert and shut down fraudulent transactions. More research is needed to understand the long-term impacts of digitalization and to develop strategies to mitigate the risks associated with it.

Finally, we caution that we do not evaluate the *overall* welfare effects of digitalization. We have not considered the additional logistical, security, and operational costs in a non-digital counterfactual (e.g., a cash-based economy). These costs must also be balanced against the convenience, cost savings, and efficiency brought about by the digital economy. This exercise is outside the scope of our study. However, our evidence indicates that consumers who are pushed into the digital economy may suffer significant privacy costs in the form of identity theft.

Thus, bank branches can remain a vital social good in the digital age, providing security and consumer protection.

References

Abis, S. and Veldkamp, L. (Aug. 2023). “The changing economics of knowledge production”. Review of Financial Studies 37, 89–118.

Acquisti, A., Taylor, C., and Wagman, L. (2016). “The economics of privacy”. Journal of economic Literature 54, 442–492.

Amberg, N. and Becker, B. (2024). “Banking without branches”. Working paper.

Angrist, J., Imbens, G., and Rubin, D. (1996). “Identification of causal effects using instrumental variables”. Journal of the American Statistical Association 91, 444–455.

Arkhangelsky, D., Athey, S., Hirshberg, D., Imbens, G., and Wager, S. (2021). “Synthetic difference-in-differences”. American Economic Review 111, 4088–4118.

Armantier, O., Doerr, S., Frost, J., Fuster, A., and Shue, K. (2024). “Nothing to hide? Gender and age differences in willingness to share data”. Working paper.

Benmelech, E., Yang, J., and Zator, M. (2023). “Bank branch density and bank runs”. Working paper.

Bernstein, S. (2015). “Does going public affect innovation?” The Journal of finance 70, 1365–1403.

Bian, B., Ma, X., and Tang, H. (2023). “The supply and demand for data privacy: Evidence from mobile apps”. Working paper.

Bian, B., Pagel, M., Raval, D., and Tang, H. (2024). “Consumer surveillance and financial fraud”. Working paper.

Bonfim, D., Nogueira, G., and Ongena, S. (Nov. 2020). ““Sorry, we’re closed” Bank branch closures, loan pricing, and information asymmetries”. Review of Finance 25, 1211–1259.

Callaway, B. and Sant’Anna, P. (2021). “Difference-in-differences with multiple time periods”. Journal of Econometrics 225, 200–230.

Célerier, C. and Matray, A. (2019). “Bank-branch supply, financial inclusion, and wealth accumulation”. The Review of Financial Studies 32, 4767–4809.

Chen, L., Huang, Y., Ouyang, S., and Xiong, W. (2025). “The data privacy paradox and digital demand”. Working paper.

Choi, H.-S. and Loh, R. K. (2024). “Physical frictions and digital banking adoption”. Management Science 70, 6597–6621.

Cong, L. W., Feng, T., Liu, Y., and Lu, F. (2025). “Crypto ATMs: Material effects of virtual currencies”. Working paper.

Cong, L. W., Xie, D., and Zhang, L. (2021). “Knowledge accumulation, privacy, and growth in a data economy”. Management Science 67, 6480–6492.

Cowles, C. (2024). “How I fell for an Amazon scam call and handed over $50,000”. New York Magazine. URL: https://www.thecut.com/article/amazon-scam-call-ftc-arrest-warrants.html.

Cramer, K. F. (2024). “Bank presence and health”. Working paper.

Cubides, E. and O’Brien, S. (2023). “2023 Findings from the Diary of Consumer Payment Choice”. Working paper.

Drechsler, I., Savov, A., and Schnabl, P. (2017). “The deposits channel of monetary policy”. Quarterly Journal of Economics 132, 1819–1876.

Fainmesser, I., Galeotti, A., and Momot, R. (2023). “Digital privacy”. Management Science 69, 3157–3173.

Farboodi, M. and Veldkamp, L. (2023). “Data and markets”. Annual Review of Economics 15, 23–40. ISSN: 1941-1391.

Federal Trade Commission (2015). *FTC Charges Data Brokers with Helping Scammer Take More Than* $7MillionfromConsumers'Accounts$. URL: https://www.ftc.gov/news-events/news/press-releases/2015/08/ftc-charges-data-brokers-helping-scammer-take-more-7-million-consumers-accounts (visited on 08/12/2015).

— (2024a). *Phishing Scams*. Accessed: 2024-07-11. URL: https://www.ftc.gov/news-events/topics/identity-theft/phishing-scams.

— (2024b). *Unwanted Calls*. Accessed: 2024-10-18. URL: https://consumer.ftc.gov/unwanted-calls-emails-and-texts/unwanted-calls.

Fonseca, J. and Matray, A. (2024). “Financial inclusion, economic development, and inequality: Evidence from Brazil”. Journal of Financial Economics 156, 103854.

Forbes Magazine (2025). *16 Billion Apple, Facebook, Google And Other Passwords Leaked*. URL: https://www.forbes.com/sites/daveywinder/2025/06/20/16-billion-apple-facebook-google-passwords-leaked---change-yours-now/ (visited on 06/20/2025).

Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2022). “Predictably unequal? The effects of machine learning on credit markets”. Journal of Finance 77, 5–47.

Garmaise, M. J. and Moskowitz, T. J. (2006). “Bank mergers and crime: The real and social effects of credit market competition”. Journal of Finance 61, 495–538.

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

Hamdi, N., Kalda, A., and Sovich, D. (2024). “The costs of financial fraud victimization”. Working paper.

Harrell, E. and Langton, L. (2013). *Victims of identity theft, 2012*. US Department of Justice, Office of Justice Programs, Bureau of Justice.

Imbens, G. and Angrist, J. (1994). “Identification and Estimation of Local Average Treatment Effects”. Econometrica 62, 467–475.

Jayaratne, J. and Strahan, P. (1996). “The finance-growth nexus: Evidence from bank branch deregulation”. Quarterly Journal of Economics, 639–670.

Ji, Y., Teng, S., and Townsend, R. M. (2023). “Dynamic bank expansion: spatial growth, financial access, and inequality”. Journal of Political Economy 131, 2209–2275.

Jiang, E. X., Yu, G. Y., and Zhang, J. (2025). “Bank competition amid digital disruption: Implications for financial inclusion”. Journal of Finance, Forthcoming.

Jiang, W. (2017). “Have instrumental variables brought us closer to the truth”. Review of Corporate Finance Studies 6, 127–140.

Johnson, G. (2022). “Economic research on privacy regulation: Lessons from the GDPR and beyond”. Working Paper.

Keil, J. and Ongena, S. (2024). “The demise of branch banking - Technology, consolidation, bank fragility”.

Koont, N. (2024). “The Digital Banking Revolution: Effects on Competition and Stability”. Working paper.

Kroszner, R. and Strahan, P. (1996). “Regulatory incentives and the thrift crisis: Dividends, mutual-to-stock conversions, and financial distress”. Journal of Finance 51, 1285–1319.

Liu, W.-M. and Ngo, P. (2014). “Elections, political competition and bank failure”. Journal of Financial Economics 112, 251–268.

Marbach, M. and Hangartner, D. (2020). “Profiling compliers and noncompliers for instrumental-variable analysis”. Political Analysis 28, 435–444.

Martín-Oliver, A., Toldrà-Simats, A., and Vicente, S. (2020). “The real effects of bank branch closings and restructurings”. Working paper.

Nguyen, H.-L. (2019). “Are credit markets still local? Evidence from bank branch closings”. American Economic Journal: Applied Economics 11, 1–32.

Qi, S., De Haas, R., Ongena, S., Straetmans, S., and Vadasz, T. (July 2024). “Move a little closer? Information sharing and the spatial clustering of bank branches”. Review of Finance 28, 1881–1918.

Ramadorai, T., Uettwiller, A., and Walther, A. (2025). “The market for data privacy”. Working paper.

Sakong, J. and Zentefis, A. K. (2024). “Bank branch access: Evidence from geolocation data”. Working paper.

Statista (2025). *Time spent with nonvoice activities on mobile phones every day in the United States from 2019 to 2024*. URL: https://www.statista.com/statistics/1045353/mobile-device-daily-usage-time-in-the-us/ (visited on 06/26/2015).

Tang, H. (2019). “The value of privacy: Evidence from online borrowers”. Available at SSRN 3880119.

Tirole, J. (2023). “Competition and the industrial challenge for the digital age”. Annual Review of Economics 15, 573–605.

Yin, W., Hay, J., and Roth, D. (2019). “Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach”. arXiv preprint arXiv:1909.00161.

Figure 1. Annual trends of consumer fraud in the U.S.

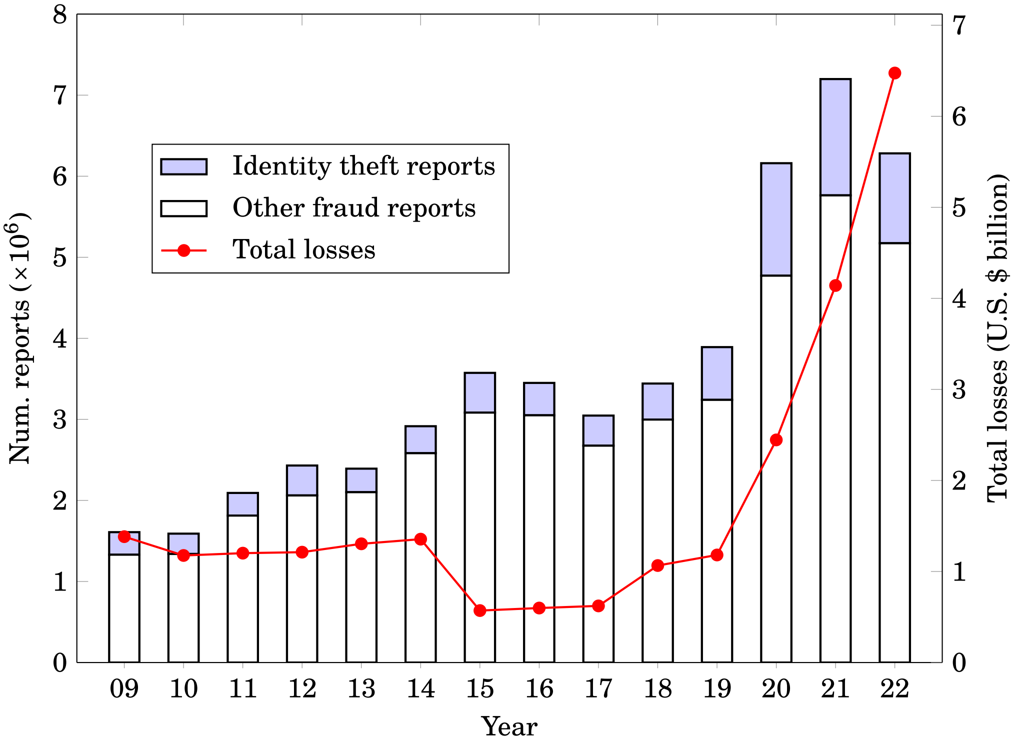

This figure plot the national annual trends of consumer fraud statistics. The left axis present the number of reports (in millions) of identity theft and all other consumer fraud plotted using stacked barcharts. The right axis reported total losses in U.S. billion dollars plotted using the red line. Data on consumer fraud are sourced from the Consumer Sentinel Network, administered by the U.S. Federal Trade Commission.

Figure 2. Geography of identity theft reports in the U.S.

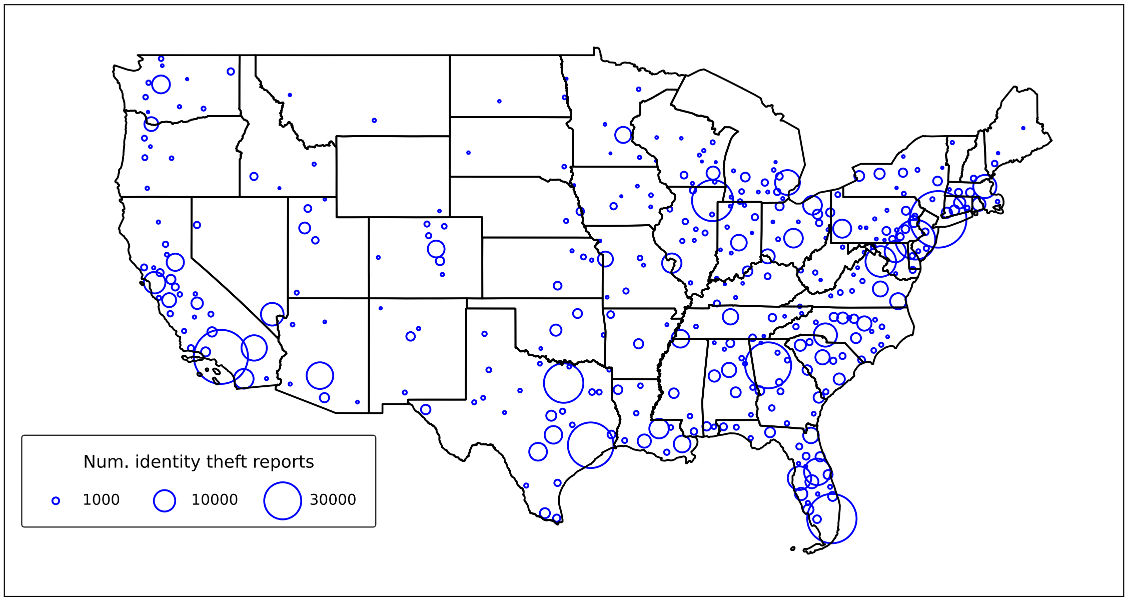

This figure plots the numbers of identity theft reports at the MSA level in the year 2022. Larger circles reflect higher numbers of identity theft reports. Data on identity theft volume are sourced from the Consumer Sentinel Network database, administered by the U.S. Federal Trade Commission.

Figure 3. Elasticity of bank branch visits

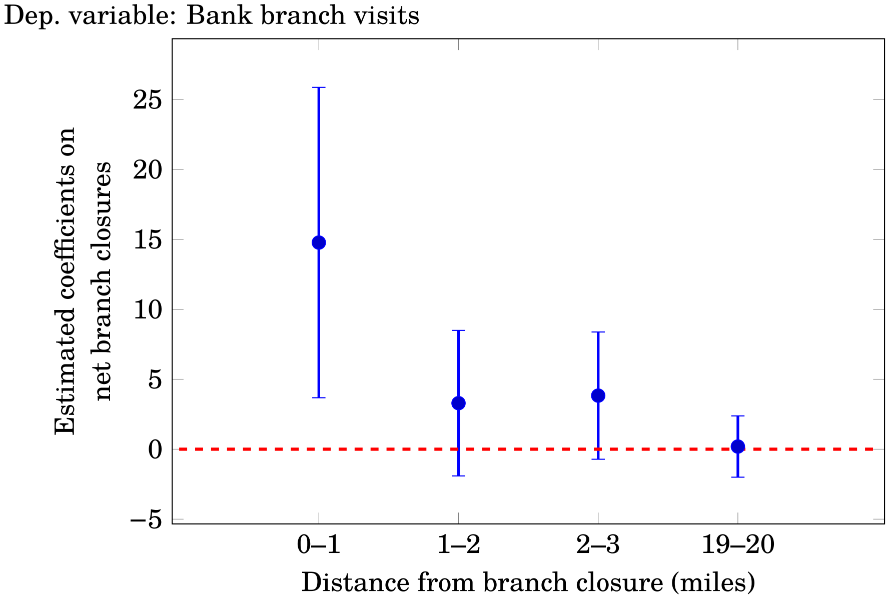

This figure presents coefficient estimates on *net branch closures* from OLS regressions of the form in equation (1). The point estimates are presented in circles and the 95% confidence intervals are represented by end points of the line. The dependent variable is *bank branch visits*, which is defined as the weekly number of visits received by a bank branch, compiled from the pass_by database. For every bank branch, we merge in the *net branch closures* of neighboring bank branches that are within a particular distance band. The variable *net branch closures* is the net decrease in the number of neighboring bank branches over the past 180 days. We calculate distances between bank branches using latitude and longitude coordinates obtained from the FDIC SOD dataset and by geocoding pass_by branch addresses with the Google Maps API.

Figure 4. Banking mobile app usage after branch closures

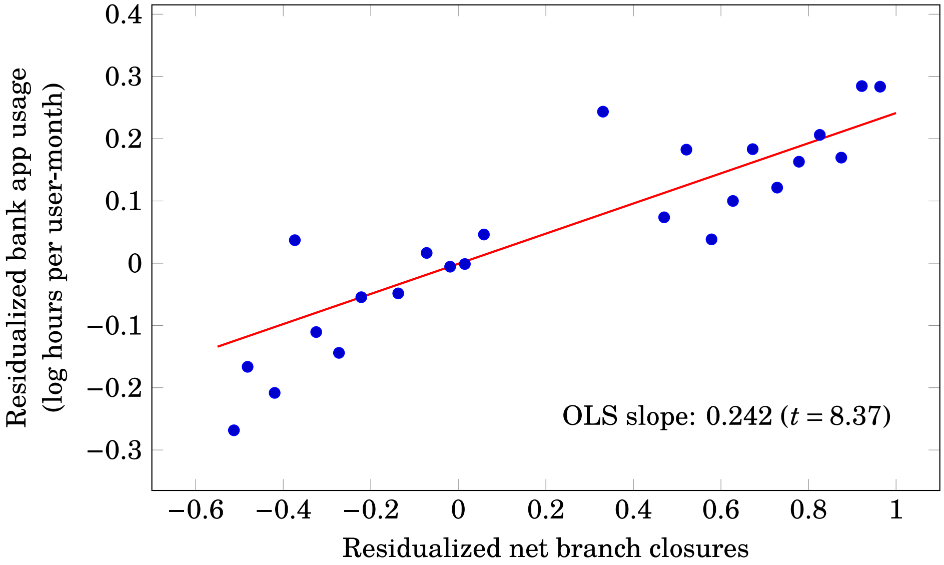

This figure presents a binned scatterplot of *bank app usage* against *net branch closures*. The variable *bank app usage* is the estimated average duration (in hours) spent by consumers in a county on the mobile app of a specific bank in a month. The variable *net branch closures* is the net decrease in the number of branches of a specific bank in a month within a county. To construct the binned scatterplot, we first regress separately *bank app usage* and *net branch closures* on county-year control variables, as well as fixed effects at the state × month and bank × month dimensions. Next, we sort the residualized *net branch closures* into bins and compute the averages of residualized *bank app usage* within these bins. Finally, we plot these residuals and the OLS best-fit line. Mobile app usage data of individual Android smartphone users is sourced from an opted-in panel from the Global Wireless Solutions Magnify database.

Figure 5. Pre-exposure differences in MSA characteristics

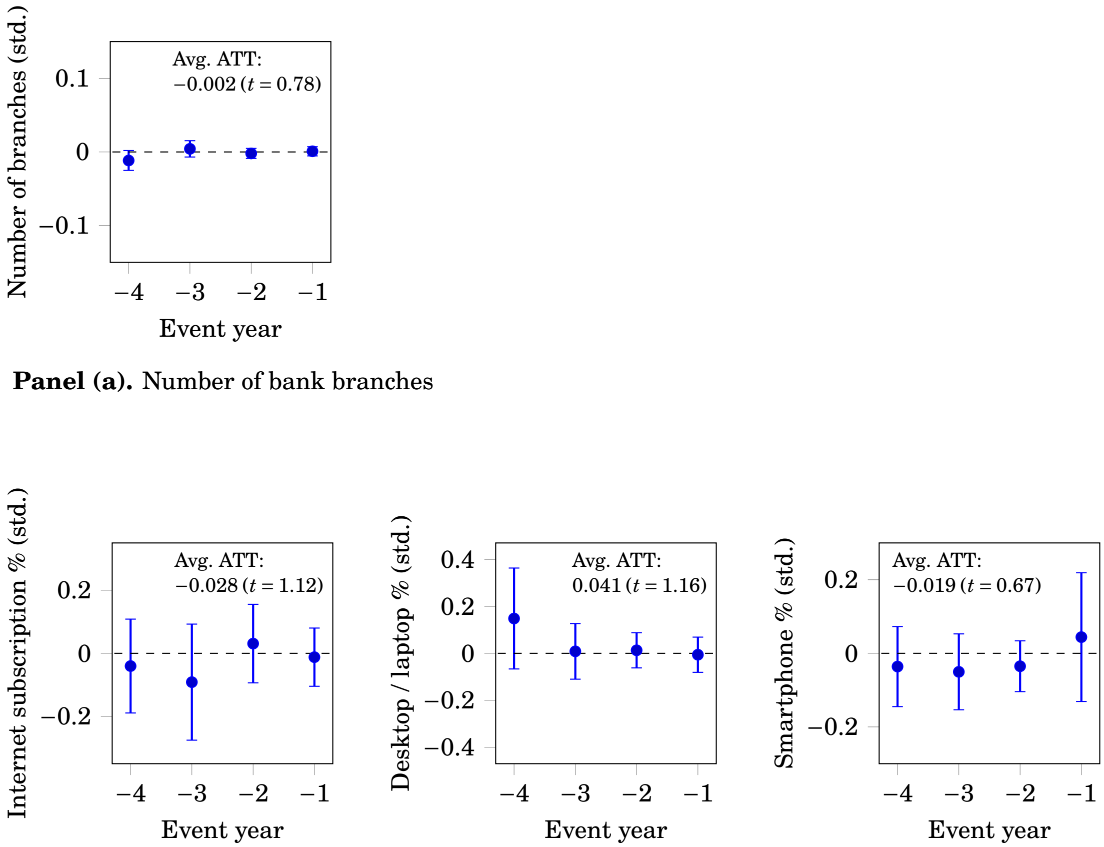

Panel (a). Number of bank branches

Panel (b). Measures of digital adoption

Figure 5. (continued)

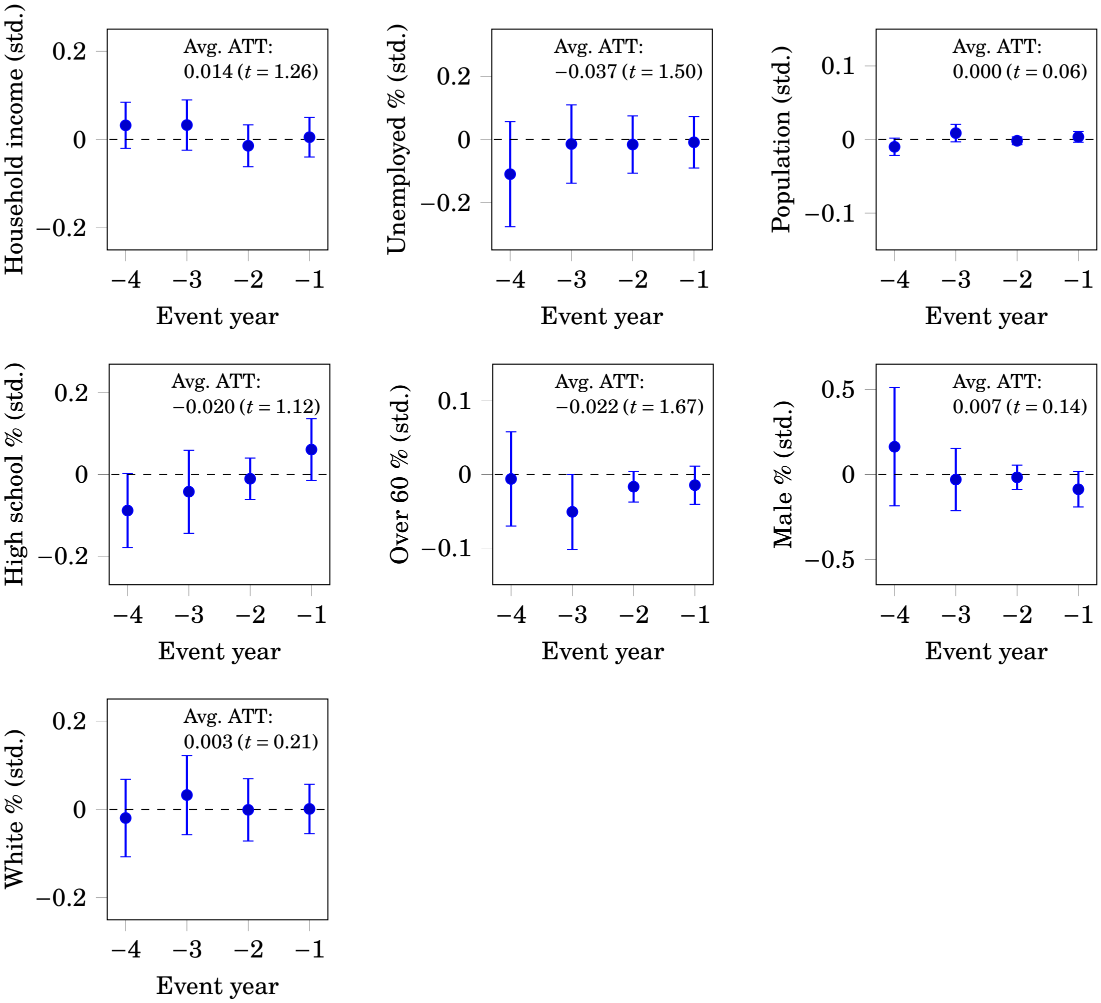

Panel (c). Demographic characteristics

This figure plots the pre-exposure differences, which are visualized by Callaway and Sant'Anna (2021) ATTs, in MSA characteristics between exposed and unexposed MSAs. An MSA is first exposed in the year when a merger occurs between large banks that both have branches in the area. The number of bank branches are obtained from the FDIC Summary of Deposits database. Demographic characteristics and measures of digital adoption are obtained from the Census Bureau. MSA characteristics are standardized to facilitate comparison. Error bars represent 95% confidence intervals. Standard errors are clustered at the MSA level.

Figure 6. Dynamic effects of merger exposures on identity theft

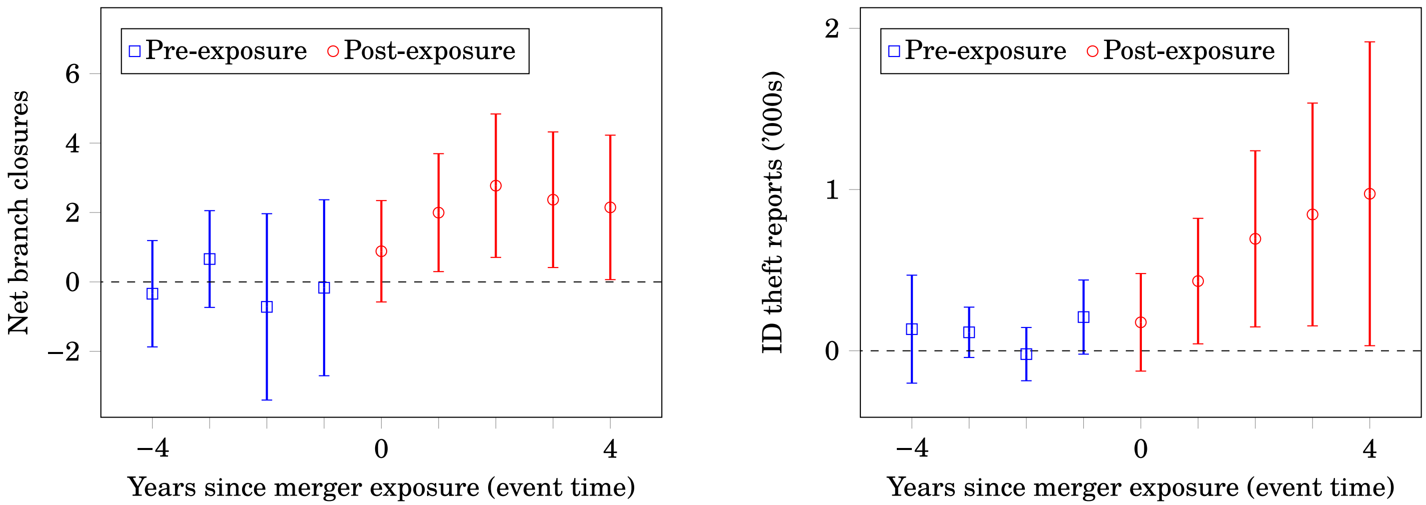

This figure plots the average treatment effects on treated (ATTs) on *net branch closures* and *ID theft reports*, and their respective 95% confidence intervals within a [−4, +4] event-time window. The variable *net branch closures* is the net year-on-year decrease in the number of bank branches within an MSA. ATTs are estimated from the Callaway and Sant'Anna (2021) doubly-robust estimator. An MSA is first exposed in the year when a merger occurs between large banks that both have branches in the area. The MSA-year number of *ID theft reports* is sourced from the Consumer Sentinel Network database, administered by the U.S. Federal Trade Commission.

Figure 7. Quantifying the effects of branch closures on identity theft

(a) Calendar ATTs and Wald estimates (MSA level)

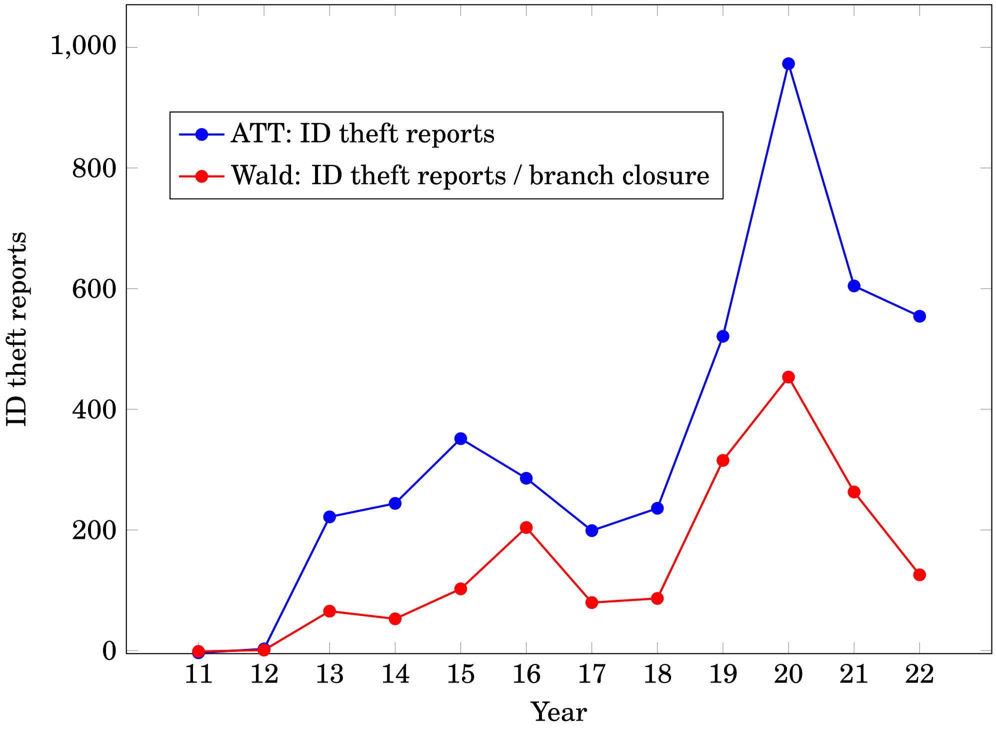

This subfigure presents the Callaway and Sant'Anna (2021) calendar ATTs for *net branch closures* (i.e., the first stage) and the number of *ID theft reports* (i.e., the second stage). The calendar ATT in a year is the average treatment effect for a MSA that is or is already exposed to large bank mergers in that year. The subfigure also presents the annual Wald estimates, which are the ratios of the second-stage estimates to the first-stage estimates. The Wald estimate represents the causal effect of one instrumented *net branch closure* on the number of *ID theft reports*.

Figure 7. (Continued)

(b) Imputed outcomes by calendar year (whole of U.S.)

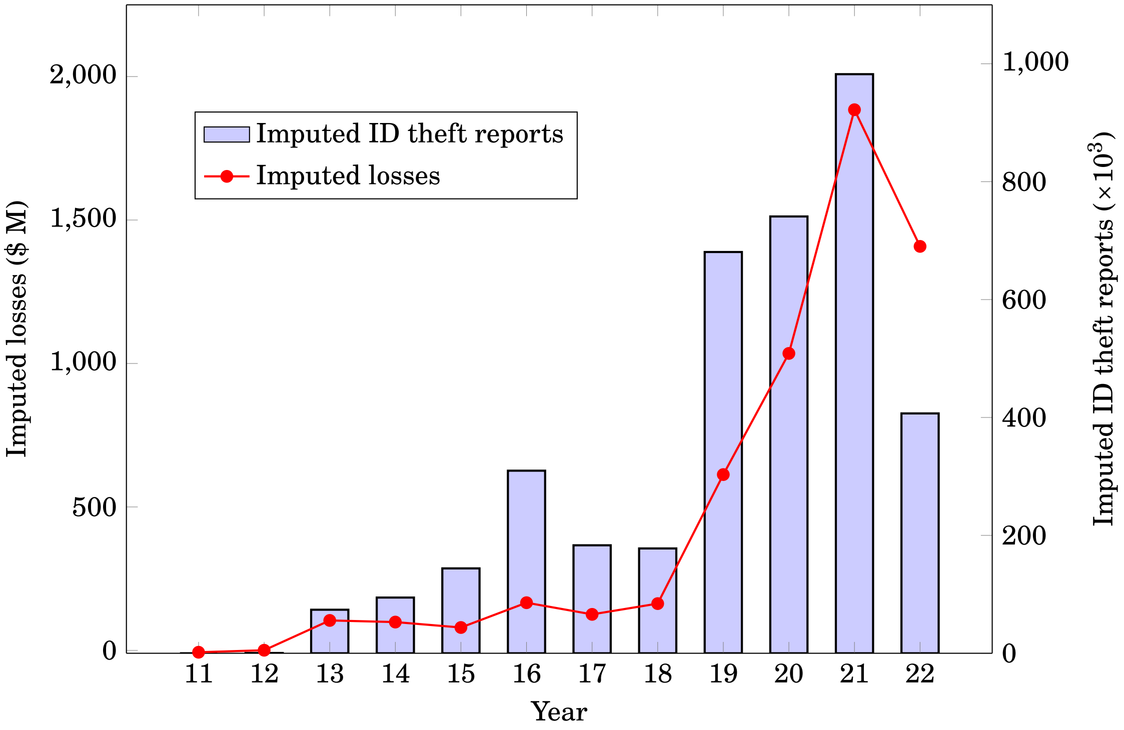

This subfigure imputes the number of *ID theft reports* by multiplying the Wald estimates in Figure 7(a) by the actual *net branch closures* in MSA-years and aggregating them to the year level. We also impute the total losses by multiplying the imputed number of ID theft reports in the MSA-year by the average loss per report in the state-year and aggregating them to the year level.

Table 1. Descriptive statistics

Panel A: Summary statistics ($N = 4{,}266$)

	Mean	S.D.	P10	P50	P90
ID theft reports ($\times 10^3$)	1.42	5.32	0.08	0.28	2.47
Net branch closures	2.9	14.2	–1	1	8
Δ Branch (−)	3.9	12.0	0	1	8
Δ Branch (+)	1.0	7.0	0	0	1
Over 60	18.4	7.1	10.5	17.2	26.8
Male	50.1	1.3	48.6	50.1	51.6
Unemployed	6.6	2.9	3.6	6.0	10.5
White	76.6	13.5	60.0	80.0	90.0
High school	88.2	5.4	82.0	89.3	93.5
Household income ($\times 10^3$)	56.1	13.0	41.8	53.8	72.8
Population ($\times 10^5$)	7.3	15.8	1.2	2.7	15.6

Panel B: Pairwise correlations (% pts.)

		A	B	C	D	E	F	G	H
A	ID theft reports								
B	Net branch closures	57							
C	Over 60	4	10						
D	Male	(9)	(11)	(40)					
E	Unemployed	1	0	(17)	1				
F	White	(28)	(16)	5	7	(3)			
G	High school	0	4	28	(11)	(2)	22		
H	Household income	28	22	29	(21)	0	(23)	36	
I	Population	76	57	(6)	(2)	1	(27)	(1)	31

Panel A reports the summary statistics of the variables used in our main analysis. The MSA-year number of *ID theft reports* is compiled by the Consumer Sentinel Network of the U.S. Federal Trade Commission. Branch closures are computed from the Summary of Deposits files of the Federal Deposit Insurance Corporation. MSA-year demographic data are collected from the U.S. Census Bureau. The statistics for the variables *over 60*, *male*, *unemployed*, *white*, and *high school* are reported in percentage points. Panel B reports the Pearson pairwise correlation coefficients between the variables used in the main test. Correlation coefficients are rounded to their nearest integers and expressed in percentage points. Negative values are contained in parentheses.

Table 2. Bank branch closures and identity theft

Dependent variable: ID theft reports

	(1)	(2)	(3)	(4)
Net branch closures	213.42***	88.04***	145.22***	145.02***
	(5.91)	(3.21)	(5.48)	(5.46)
Cybersecurity				–903.51*
				(1.89)
Over 60		–104.14**	–106.65	–104.02
		(1.97)	(1.61)	(1.57)
Male		–162.83***	–167.60***	–166.67***
		(3.74)	(3.48)	(3.46)
Unemployed		0.00***	0.00*	0.00*
		(2.64)	(1.91)	(1.88)
White		–63.49**	–67.87**	–67.84**
		(2.39)	(2.30)	(2.30)
High school		–77.15***	–91.05***	–91.72***
		(2.92)	(2.62)	(2.63)
log(Household income)		759.87	275.03	232.82
		(0.62)	(0.17)	(0.15)
log(Population)		674.88	1,102.31	1,014.13
		(0.49)	(0.49)	(0.45)
Sample	All	All	2014–22	2014–22
# Obs.	4,266	4,266	3,126	3,126
R^2	0.324	0.690	0.718	0.719
MSA FE		✓	✓	✓
Year FE		✓	✓	✓

This table presents estimates from OLS regressions. The dependent variable is the number of *ID theft reports* at the MSA-year level, collected from the Consumer Sentinel Network database of the U.S. Federal Trade Commission. The key independent variables are *net branch closures* and *cybersecurity*. The variable *net branch closures* is the net year-on-year decrease in the number of bank branches within an MSA. The variable *cybersecurity* is the average number of bot and fraud detection technologies deployed by banks that operate branches within an MSA. We collect data on banks' technological deployments from the BuiltWith database. Standard errors are clustered at the MSA level. Absolute values of t-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 3. Merger exposures and identity theft

Average treatment effects on treated (ATTs)

	(1)	(2)
Outcome:	Net branch closures	ID theft reports
Pre-exposure	−0.21	30.27
	(0.21)	(1.44)
Post-exposure	3.51***	726.89***
	(3.15)	(2.65)
Overall	2.79***	455.56***
	(2.90)	(2.72)
Wald estimate	163.28***	
	(4.75)	

This table presents ATTs on *net branch closures* and *ID theft reports* from the Callaway and Sant'Anna (2021) difference-in-differences estimator. The variable *net branch closures* is the net year-on-year decrease in the number of bank branches within an MSA. The MSA-year number of *ID theft reports* is collected from the Consumer Sentinel Network database of the U.S. Federal Trade Commission. The *pre* (*post*)*-exposure averages* are the aggregated ATTs before (after) merger exposures. An MSA is first exposed in the year when a merger occurs between large banks that both have branches in the area. Standard errors are clustered at the MSA level. Standard errors of the Wald estimates are computed using the delta method. Absolute values of t-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 4. (continued)

Panel A. Comparison of MSAs with(out) branch closures upon merger exposure

	Conditional on exposure:			
Characteristics	Branch closures > 0	Branch closures $= 0$	Diff.	*t*-stat
Digital adoption (having:)				
Internet subscription	75.8	76.1	–0.2	(0.17)
Desktop / laptop	77.6	77.5	0.2	(0.11)
Smartphone	74.4	74.2	0.2	(0.18)
Demographics				
Over 60	11.7	11.0	0.7	(0.87)
Male	50.8	50.7	0.1	(0.54)
Unemployed	10.6	9.7	0.9	(1.56)
White	78.8	78.0	0.8	(0.31)
High school	86.3	87.0	–0.7	(0.69)
Household income ($\times 10^3$)	47.0	46.9	0.1	(0.08)
Population ($\times 10^5$)	5.84	6.05	–0.21	(0.14)
Num. branches	165.5	174.0	–8.5	(0.21)

This table compares the characteristic means between MSAs that, conditional on exposure to large bank mergers, (i) have at least one branch closure, and (ii) have no branch closures. We report the differences in characteristic means and their corresponding *t*-statistics. Demographic characteristics are taken from the year 2010. Measures of digital adoption are taken from the year 2015 (earliest year availabl

Table 4. (continued)

Panel B. Complier characteristics analysis

Characteristics	Full sample	Compliers	Never-takers	Always-takers
Digital adoption (having:)				
Internet subscription	74.7	76.2	76.0	73.6
Desktop / laptop	76.6	77.0	77.4	76.1
Smartphone	72.8	72.6	73.8	72.3
Demographics				
Over 60	11.7	11.2	11.0	12.1
Male	50.8	50.5	50.7	50.9
Unemployed	10.2	11.6	9.7	10.1
White	80.2	83.7	78.0	80.3
High school	86.1	84.0	87.0	86.2
Household income ($\times 10^3$)	45.9	46.6	46.9	45.2
Population ($\times 10^5$)	4.15	4.94	6.05	3.04
Num. branches	120.9	119.8	174.0	95.8
Proportion of sample (% pts.)	100	19	26	55

This table compares the characteristic means of complier MSAs with those of the full sample and non-compliers. Compliers are MSAs that (do not) encounter at least one branch closure due to the (absence) presence of merger exposures. Always-takers (Never-takers) are MSAs that would (not) have closed branches regardless of merger exposures. Characteristic means of compliers and non-compliers are computed using the methodology outlined in Marbach and Hangartner (2020). Demographic characteristics are taken from the year 2010. Measures of digital adoption are taken from the year 2015 (earliest year available). We report in percentage points the proportions of the sample made up of compliers, never-takers, and always-takers.

Table 5. Bank branch closures and consumer complaints

Dependent variable ($\times 10^2$): 1(Complaint)

	(1)	(2)	(3)
Type of complaint	Any	ID theft	ID theft
Net branch closures	3.26***	0.34***	0.17**
	(9.78)	(3.71)	(2.00)
Over 60	0.19***	0.00	
	(5.99)	(0.30)	
Male	–0.06**	–0.03***	
	(2.54)	(4.13)	
Unemployed	0.26***	0.00	
	(4.60)	(0.48)	
White	–0.16***	–0.01***	
	(9.75)	(8.09)	
High school	0.07***	0.01***	
	(2.80)	(3.83)	
log(Household income)	7.36***	–0.18**	
	(8.12)	(2.33)	
log(Population)	4.90***	0.12***	
	(28.07)	(6.47)	
# Obs.	1,374,000	1,374,000	1,374,000
R^2	0.115	0.011	0.251
County FE	✓	✓	
Month FE	✓	✓	
Bank-County FE			✓
County-Month FE			✓
Bank-State-Month FE			✓

This table presents estimates from OLS regressions. The dependent variable is 1(*complaint*), an indicator that switches on if a bank receives a CFPB complaint in a county-month. In columns 2 and 3, we use the bart-large-mnli model (Yin, Hay, Roth, 2019) to identify CFPB complaints that are related to identity theft. The key independent variable is *net branch closures*, which is the net decrease in the number of bank branches over the past 180 days in the county. All standard errors are clustered at the county level. Standard errors in column 3 are additionally clustered at the bank-month level. Absolute values of t-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 6. Heterogeneous effects of bank branch closures on identity theft

Average treatment effects on treated (ATTs)				
	(1)	(2)	(3)	(4)
	Net branch closures	ID theft reports	Net branch closures	ID theft reports
Panel A. Bank mergers by branch reliance:				
	High		Low	
Pre-exposure	0.47	64.28**	0.42**	53.93*
	(1.14)	(2.25)	(2.26)	(1.81)
Post-exposure	19.52***	4,444.48***	1.70**	656.88***
	(2.76)	(2.58)	(2.05)	(3.03)
Overall	11.67***	2,362.36***	1.03	440.27***
	(2.70)	(2.92)	(1.48)	(3.09)
Wald estimate	202.43***		428.33	
	(8.76)		(1.70)	
Panel B. Bank mergers by digital focus:				
	Low		High	
Pre-exposure	0.15	12.15	0.27	5.35
	(0.49)	(1.38)	(1.29)	(0.72)
Post-exposure	3.27***	555.73***	2.66	119.29
	(2.79)	(2.78)	(1.59)	(1.01)
Overall	1.93**	281.70**	2.66*	229.06*
	(2.20)	(2.30)	(1.72)	(1.94)
Wald estimate	145.96***		86.11*	
	(2.52)		(1.80)	

(Continued next page)

Table 6. (continued)

Average treatment effects on treated (ATTs)

	(1)	(2)	(3)	(4)
	Net branch closures	ID theft reports	Net branch closures	ID theft reports
Panel C. Bank mergers by consumer tech-savviness:				
	Low		High	
Pre-exposure	0.15	–4.34	0.43	45.67
	(0.37)	(0.97)	(1.55)	(1.50)
Post-exposure	2.47**	534.21**	5.21***	1,120.46***
	(2.20)	(2.17)	(3.19)	(2.74)
Overall	1.91**	398.31**	3.79***	668.13***
	(2.02)	(2.31)	(2.79)	(2.62)
Wald estimate	208.54***		176.29***	
	(10.56)		(3.03)	

This table presents ATTs on *net branch closures* and *ID theft reports* from the Callaway and Sant'Anna (2021) difference-in-differences estimator. The variable *net branch closures* is the net year-on-year decrease in the number of bank branches within an MSA. The MSA-year number of *ID theft reports* is collected from the Consumer Sentinel Network database of the U.S. Federal Trade Commission. The *pre* (*post*)*-exposure averages* are the aggregated ATTs before (after) merger exposures. An MSA is first exposed in the year when a merger occurs between large banks that both have branches in the area. We classify bank mergers based on the (i) *branch reliance* of the acquirer (Panel A), (ii) *digital focus* of the acquirer (Panel B), and (iii) *consumer tech-savviness* in the exposed MSA (Panel C). The *branch reliance* of a bank is the ratio of its number of branches to its total deposits. The *digital focus* of a bank is the ratio of the download volume of its mobile app on the Google Play Store to its number of branches. The *consumer tech-savviness* of an MSA is the proportion of consumers who have internet subscriptions. Standard errors are clustered at the MSA level. Standard errors of the Wald estimates are computed using the delta method. Absolute values of *t*-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 7. Merger exposures and consumer-level mobile app usage

Average treatment effects on treated (ATTs)

	(1)	(2)
Outcome:	Net branch closures	Mobile app usage
Pre-exposure	0.002	0.311
	(0.15)	(0.41)
Post-exposure	0.125***	2.285***
	(9.82)	(2.76)
Overall	0.098***	2.446**
	(8.73)	(2.11)
Wald estimate	24.91**	
	(2.04)	

This table presents ATTs on *net branch closures* and *mobile app usage* from the Callaway and Sant'Anna (2021) difference-in-differences estimator. The variable *mobile app usage* is the number of hours spent on all mobile applications, except mobile banking apps, in a month by each Android smartphone user in an opted-in panel from the Global Wireless Solutions Magnify database. The variable *net branch closures* is the net decrease in the number of bank branches in the month, matched to the consumer's residence at the county level. The *pre* (*post*)*-exposure averages* are the aggregated ATTs before (after) merger exposures. An MSA is first exposed in the monrh when a merger occurs between large banks with branches in the area. Standard errors are clustered at the consumer level. Standard errors of the Wald estimates are computed using the delta method. Absolute values of t-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 8. Merger exposures and consumption patterns

Average treatment effects on treated (ATTs)

	(1)	(2)	(3)
Outcome:	Net branch closures	Online spending gap	Online transaction gap
Pre-exposure	0.213**	0.000	–0.002
	(1.98)	(0.08)	(1.11)
Post-exposure	0.415***	0.016***	0.015**
	(4.42)	(2.72)	(2.30)
Overall	0.311***	0.011**	0.009**
	(3.79)	(2.39)	(2.01)
Wald estimate	—	0.036***	0.030***
	—	(2.90)	(2.58)

This table presents ATTs on *net branch closures*, *online spending gap*, and *online transaction gap* from the Callaway and Sant'Anna (2021) difference-in-differences estimator. The variable *online spend gap* (*online transaction gap*) is the monthly dollar value (number) of online transactions less that of offline transactions in an MSA, scaled by the monthly dollar value (number) of all transactions in an MSA. Transactions data are collected from the SafeGraph database. The variable *net branch closures* is the net decrease in the number of bank branches in the month within an MSA. The *pre* (*post*)-*exposure averages* are the aggregated ATTs before (after) merger exposures. An MSA is first exposed in the month when a merger occurs between large banks that both have branches in the area. Standard errors are clustered at the MSA level. Standard errors of the Wald estimates are computed using the delta method. Absolute values of t-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 9. Merger exposures and unwanted calls

Average treatment effects on treated (ATTs)

	(1)	(2)	(3)	(4)
Outcome:	Net branch closures	Unwanted calls	Marketing calls	Unwanted calls (net)
Pre-exposure	0.052	–0.838	–0.232	–0.719
	(0.31)	(0.78)	(0.44)	(0.64)
Post-exposure	1.134***	2.109**	–1.308	2.692**
	(4.81)	(2.03)	(1.37)	(2.33)
Overall	1.135***	2.105**	–1.271	2.690**
	(4.81)	(2.02)	(1.35)	(2.33)
Wald estimate	—	1.855*	–1.120	2.370**
	—	(1.86)	(1.30)	(2.10)

This table presents ATTs on *net branch closures*, *unwanted calls*, *marketing calls*, and *unwanted calls (net)* from the Callaway and Sant'Anna (2021) difference-in-differences estimator. The monthly numbers of *unwanted calls* and *marketing calls* are reported by consumers in a county to the the U.S. Federal Trade Commission. The variable *unwanted calls (net)* is defined as the number of *unwanted calls* less that of *marketing calls*. The variable *net branch closures* is the net decrease in the number of bank branches in a county over the past 180 days. The *pre* (*post*)*-exposure averages* are the aggregated ATTs before (after) merger exposures. A county is first exposed in the month when a merger occurs between large banks that both have branches in the area. Standard errors are clustered at the county level. Standard errors of the Wald estimates are computed using the delta method. Absolute values of t-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Table 10. Bank branch closures and phishing sites

	(1)	(2)	(3)	(4)
	In-sample		Out-of-sample	
Dependent variable:	Phishing exposure	ID theft reports ($\times 10^3$)	Phishing exposure	ID theft reports ($\times 10^3$)
Net branch closures	0.024***		0.058***	
	(6.09)		(9.08)	
MSA phishing		14.59*		4.05***
		(1.95)		(2.91)
Over 60		0.25**		0.06
		(2.33)		(0.71)
Male		0.07		–0.12
		(0.72)		(1.31)
Unemployed		0.00		0.00
		(1.21)		(0.14)
White		–0.05**		0.02
		(2.28)		(1.00)
High school		–0.16***		0.11
		(2.69)		(1.24)
log(Household income)		–1.44		0.63
		(0.99)		(0.46)
log(Population)		–0.88		14.36*
		(0.41)		(1.74)
Unit of obs.	Bank-MSA-Month	MSA-Year	Bank-MSA-Month	MSA-Year
Sample period	2018–22	2018–22	2018–19	2020–22
# Obs.	49,888	1,725	12,723	1,063
R^2	0.001	0.834	0.006	0.946
MSA FE		✓		✓
Year FE		✓		✓

This table presents estimates from OLS regressions. The dependent variable *phishing exposure* in column 1 is the month-over-month change in the number of websites that use the favicon of a legitimate bank website but lack secure-sockets-layer (SSL) certification, weighted by the population of an MSA where the bank operates. The independent variable *net branch closures* is the net month-on-month decrease in the number of bank branches within an MSA. The dependent variable in column 2 is the number of *ID theft reports*, collected from the Consumer Sentinel Network database of the U.S. Federal Trade Commission. The key independent variable *MSA phishing* is the weighted sum of predicted *phishing exposure*, aggregated to the MSA-year level. The sum is weighted by the number of branches operated by banks in a given MSA-year. Columns 3 and 4 are the out-of-sample analogs of columns 1 and 2. Standard errors are clustered at the MSA level. Absolute values of t-statistics are reported in parentheses. *, **, *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

SMU
SINGAPORE MANAGEMENT UNIVERSITY
Sim Kee Boon Institute for Financial Economics

Internet Appendix to:

Grand Theft Identity: The Privacy Costs of Digitalization

Kenny Phua Chishen Wei Gloria Yang Yu

Abstract

The Internet Appendix contains supplementary information and additional tests for the paper "Grand Theft Identity: The Privacy Costs of Digitalization". The contents of the Internet Appendix are organized as follows. Section IA.1 reports detailed statistics of identity theft reports. Section IA.2 tabulates the list of entities that contribute data to the FTC Consumer Sentinel Network database. Section IA.3 provides additional details on the consumer response to branch closures. Section IA.4 presents additional details on the complier characteristics analysis. Section IA.5 details annual imputed financial losses from identity theft due to branch closures. Section IA.6 provides causal evidence on the effect of branch closures on identity theft from synthetic difference-in-differences estimations.

Internet Appendix to:

Grand Theft Identity: The Privacy Costs of Digitalization

Abstract

The Internet Appendix contains supplementary information and additional tests for the paper "Grand Theft Identity: The Privacy Costs of Digitalization". The contents of the Internet Appendix are organized as follows. Section IA.1 reports detailed statistics of identity theft reports. Section IA.2 tabulates the list of entities that contribute data to the FTC Consumer Sentinel Network database. Section IA.3 provides additional details on the consumer response to branch closures. Section IA.4 presents additional details on the complier characteristics analysis. Section IA.5 details annual imputed financial losses from identity theft due to branch closures. Section IA.6 provides causal evidence on the effect of branch closures on identity theft from synthetic difference-in-differences estimations.

IA.1 Detailed statistics on identity theft reports

We present detailed statistics on identity theft reports from the Federal Trade Commission (FTC) Consumer Sentinel Network database in 2023. Panel A provides a breakdown of these reports by the types and subtypes of identity theft. Panel B provides a breakdown of these reports by identity theft types and age groups.

- Table IA.1 here -

IA.2 Data contributors of the FTC Consumer Sentinel Network database

We provide the list of organizations that provide data to the FTC Consumer Sentinel Network database.

- Table IA.2 here -

IA.3 Details on consumer response to branch closures.

Table IA.3 contains the full regressions results that underpin Figure 3 in the main text.

- Table IA.3 here -

IA.4 Details on the complier characteristics analysis

In this section, we provide details on our complier characteristics analysis. To fix ideas, every MSA has two unobserved potential treatment indicators $D(0)$ and $D(1)$ that manifest an observed treatment $D \in \{0,1\}$ representing the presence of branch closures. The indicator Z switches on if an MSA is exposed to mergers. The matrix below maps the treatment responses to Z of the subpopulations.

	$D(Z=0)$	$D(Z=1)$
Compliers	0	1
Always-Takers	1	1
Never-Takers	0	0
Defiers	1	0

Because we only observe the realized treatment D but not $D(0)$ and $D(1)$, we cannot classify individual MSAs into subpopulations. Specifically, compliers and always-takers assigned to the exposure group ($Z = 1$) are observably identical. The same applies to compliers and never-takers assigned to the control group ($Z = 0$). Marbach and Hangartner (2020) proposes a framework to estimate the mean of a characteristic X within subpopulations by imposing four assumptions.

Assumption 1 (Monotonicity)**.** $D(1) \geq D(0)$.

Assumption 2 (Independence of instrument)**.** $D(0), D(1), X \perp\!\!\!\perp Z$.

Assumption 3 (Relevance condition)**.** $E[D \mid Z=1] \neq E[D \mid Z=0]$.

Assumption 4 (Probability bounds on assignment)**.** $0 < \Pr(Z=1) < 1$.

Assumption 1 is standard in IV analysis and posits that the instrument cannot have an opposite effect on any subpopulation, thereby ruling out defiers (Angrist, Imbens, Rubin, 1996). Assumption 2 implies the independence of Z with both X and $D(Z)$, which holds if merger exposures are randomly assigned across MSAs. Assumption 3 is the relevance condition, which guarantees that the fraction of compliers is nonzero. The assumption 4 strictly bounds the probability of assignment between 0 and 1 to ensure there is variation in Z across MSAs.

Under assumptions 1 and 2, observable and unobservable always-takers draw from the same distribution of X. So, we can profile the characteristic mean for always-takers by focusing on the observable subset of nonencouraged ($Z = 0$) MSAs that experience branch closures:

$$E[X \mid D(0) = D(1) = 1] = E[X \mid D = 1, Z = 0] \tag{IA.1}$$

By the same logic, we can profile the characteristic mean for never-takers by focusing on encouraged ($Z = 1$) MSAs that do not experience branch closures:

$$E[X \mid D(0) = D(1) = 0] = E[X \mid D = 0, Z = 1] \tag{IA.2}$$

We cannot immediately estimate the characteristic means for compliers because they are observably identical to always-takers and never-takers when $Z = 1$ and $Z = 0$, respectively. Instead, we first decompose the population mean into a linear combination of subpopulation means by the Law of Iterated Expectations:

$$\begin{aligned} E[X] = \quad & E[X \mid D(1) > D(0)] \cdot \Pr[D(1) > D(0)] \\ + \quad & E[X \mid D(1) = D(0) = 1] \cdot \Pr[D(1) = D(0) = 1] \\ + \quad & E[X \mid D(1) = D(0) = 0] \cdot \Pr[D(1) = D(0) = 0] \end{aligned} \tag{IA.3}$$

Substituting in equations (IA.1) and (IA.2) and expanding all conditionals, we can express the characteristic mean for compliers as a function of observables.

$$\begin{aligned} E[X \mid D(1) > D(0)] = & \left(E[X] - \frac{E[X1_{\{D=0,Z=1\}}]}{\Pr[Z=1]} - \frac{E[X1_{\{D=1,Z=0\}}]}{1 - \Pr[Z=1]} \right) \\ & \left(1 - \frac{\Pr[D=0, Z=1]}{\Pr[Z=1]} - \frac{\Pr[D=1, Z=0]}{1 - \Pr[Z=1]} \right)^{-1} \end{aligned} \tag{IA.4}$$

IA.5 Details on quantifying the effects of branch closures on identity theft

Table IA.4 tabulates the annual statistics used to impute financial losses from identity theft due to branch closures.

- Table IA.4 here -

IA.6 Synthetic difference-in-differences estimation

To create a more credible counterfactual trajectory, we estimate synthetic difference-in-differences models (Arkhangelsky, Athey, Hirshberg, Imbens, et al., 2021) that optimally weight unexposed MSAs to create granular counterfactuals. Specifically, we create synthetic control units by optimally weighting never-exposed MSAs to match exposed MSAs on pre-treatment residualized outcomes, obtained after partialling out the covariates listed in Table 2. This procedure is equivalent to matching directly on pre-exposure outcomes and covariates, as described in Section 4.2.2 of the main text.

- Table IA.5 here -

The results in Table IA.5 are consistent with our findings in Section 4.2 of the main text. We find no statistically significant preexposure differences in branch closures and identity theft between exposed and unexposed MSAs, suggesting that merger exposures are as good as randomly assigned with respect to these outcomes. Column 1 shows that MSAs exposed to large bank mergers encounter 1.13 ($t = 2.13$) more branch closures, on average. The reduced-form model in Column 2 shows that such MSAs also have 144.02 ($t = 2.82$) more cases of identity theft. Taking the overall ATTs in both columns together, a branch closure leads to an increase of 127.45 ($t = 2.22$) identity theft reports, a magnitude similar to that estimated in Table 3.

Table IA.1. Detailed statistics on identity theft reports

Panel A. Breakdown of identity theft reports by types

Identity theft type	Subtype	Num. reports
Credit card	New accounts	381,122
	Existing accounts	44,855
Loan or lease	Apartment or house rented	13,201
	Auto loan/lease	52,070
	Business/personal loan	81,342
	Federal student loan	6,815
	Non-federal student loan	10,921
	Real estate loan	7,551
Bank account	Debit cards, electronic funds transfer, or ACH	42,148
	Existing accounts	18,723
	New accounts	84,335
Govt. documents or benefits	Driver’s license issued/forged	8,977
	Govt. benefits applied for/received	82,419
	Other govt. documents issued/forged	9,096
	Passport issued/forged	1,623
Employment or tax-related	Employment or wage-related	31,207
	Tax	60,970
Phone or utilities	Landline telephone – existing accounts	1,125
	Landline telephone (new accounts)	4,578
	Mobile telephone (existing accounts)	7,853
	Mobile telephone (new accounts)	43,225
	Utilities (existing accounts)	1,896
	Utilities (new accounts)	28,725
Other identity theft	Email or social media	19,534
	Evading the law	5,526
	Insurance	11,402
	Medical services	13,683
	Online shopping or payment account	18,058
	Other	205,505
	Securities accounts	5,513

(Continued next page)

Table IA.1. (continued)

Panel B. Breakdown of identity theft reports by age groups

Age group	< 19	20-29	30-39	40-49	50-59	60-69	70-79	> 80
Identity theft type								
Bank account	1,833	20,558	36,859	28,804	20,589	14,339	7,076	2,132
Credit card	2,501	71,900	122,246	84,604	53,438	27,974	10,899	2,852
Employment or tax-related	13,774	16,826	17,827	13,765	10,869	7,877	3,899	1,261
Govt. documents or benefits	1,989	11,373	21,791	19,095	16,061	9,692	3,274	987
Loan or lease	874	26,152	44,611	30,730	20,437	11,264	4,117	1,069
Other identity theft	2,415	42,673	66,702	45,545	27,768	13,665	5,174	1,417
Phone or utilities	622	13,653	23,066	16,770	11,801	6,974	2,582	650

This table presents detailed statistics on identity theft reports from the Federal Trade Commission (FTC) Consumer Sentinel Network database in 2023. Panel A provides a breakdown of these reports by the types and subtypes of identity theft. Panel B provides a breakdown of these reports by identity theft types and age groups. A single consumer report can span more than one type/subtype of identity theft.

Table IA.2. Data contributors to the FTC Consumer Sentinel Network database

AARP Fraud Watch Network
Alaska Attorney General
Apple Inc.
Arvest Bank
AT&T Corporation
Australian Competition and Consumer Commission
Best Buy Co. Inc.
Canada Competition Bureau
Capital One Bank
Colorado Attorney General
Comcast Corporation
Connecticut Department of Consumer Protection
Consumer Financial Protection Bureau
Corporation for National and Community Service
Costco Wholesale Corporation
Craigslist
Cybercrime Support Network
Discover Bank
Dominion Energy
eBay
FedEx
First National Bank of Omaha
Florida Attorney General, Office of Citizen Services
Florida Department of Agriculture and Consumer Services
Grants.gov
Handshake
Hawaii Office of Consumer Protection
Hewlett-Packard
Idaho Attorney General
Indeed
Indiana Attorney General
International Association of Better Business Bureaus
Internet Crime Complaint Center
Iowa Attorney General
Iowa, Clinton County Sheriff's Office
JPMorgan Chase & Co.
LinkedIn
Los Angeles County Department of Consumer Affairs
Louisiana Attorney General
Maine Attorney General
Massachusetts Attorney General
MasterCard International
Michigan Attorney General
Microsoft Corporation Cyber Crime Center
Mississippi Attorney General
MoneyGram International

(Continued next page)

Table IA.2. (Continued)

National Consumers League
National Council on Aging
National Grid
Nebraska Attorney General
Nevada Attorney General
Nevada Department of Business and Industry
New Mexico, Albuquerque
New York State Attorney General
North Carolina Department of Justice
Ohio Attorney General
Ohio, Cuyahoga County Department of Consumer Affairs
Oregon Department of Justice
Pennsylvania Attorney General
PeopleClaim
PepsiCo, Inc.
Petscams.com
PrivacyStar
Prosperity Bank
Publishers Clearing House
Rent Group, Inc.
Sages Theater, Inc.
Scam Advisor
Scam Detector
Society of Citizens Against Relationship Scams
South Carolina Department of Consumer Affairs
Tennessee Division of Consumer Affairs
U.S. Bureau of Prisons
U.S. Citizenship and Immigration Services
U.S. Customs and Border Protection
U.S. Department of Defense
U.S. Department of Education
U.S. Department of Health and Human Services, Office of Inspector General
U.S. Department of Justice, Consumer Protection Branch
U.S. Department of Justice, Disaster Fraud Task Force
U.S. Department of Justice, Elder Fraud Hotline
U.S. Department of Justice, Executive Office for Immigration Review
U.S. Department of Justice, Task Force on Market Integrity and Consumer Fraud
U.S. Department of the Treasury, Internal Revenue Service
U.S. Department of Veterans Affairs
U.S. Drug Enforcement Administration
U.S. Federal Bureau of Investigation
U.S. Federal Communications Commission
U.S. Marshals Service
U.S. Parole Commission
U.S. Patent and Trademark Office
U.S. Postal Inspection Service
U.S. Social Security Administration

(Continued next page

Table IA.2. (Continued)

United Parcel Service
USA.gov
Utilities United Against Scams
Valve Corporation
Verizon Wireless
Walmart Corporation
Washington State Attorney General
Western Union Company
Wisconsin Department of Agriculture, Trade and Consumer Protection
Xerox Corporation
Zelle
Zillow Group

This table contains the list of data contributors to the Federal Trade Commission (FTC) Consumer Sentinel Network database.

Table IA.3. Elasticity of bank branch visits

Dependent variable: Bank branch visits

	(1)	(2)	(3)	(4)
Distance band	0–1 miles	1–2 miles	2–3 miles	19–20 miles
Net branch closures	14.77***	3.29	3.83	0.19
	(2.61)	(1.24)	(1.65)	(0.17)
Over 60	−2.37***	−2.40***	−2.39***	−2.42***
	(4.66)	(4.63)	(4.63)	(4.71)
Male	0.13	0.19	0.20	0.20
	(0.17)	(0.24)	(0.26)	(0.26)
Unemployed	0.00***	0.00***	0.00***	0.00***
	(6.35)	(6.33)	(6.32)	(6.33)
White	−0.01***	−0.01***	−0.01***	−0.01***
	(5.71)	(5.73)	(5.73)	(5.71)
High school	−1.26**	−1.20**	−1.21**	−1.16**
	(1.97)	(1.93)	(1.93)	(1.82)
log(Household income)	−97.66***	−97.99***	−97.91***	−98.73***
	(6.41)	(6.44)	(6.48)	(6.45)
log(Population)	49.47***	48.66***	48.60***	48.64***
	(8.49)	(8.05)	(8.06)	(8.06)
# Obs.	11,676,567	11,676,567	11,676,567	11,676,567
R^2	0.595	0.595	0.595	0.595
County-Week FE	✓	✓	✓	✓
Bank-Week FE	✓	✓	✓	✓
County cluster	✓	✓	✓	✓
Bank cluster	✓	✓	✓	✓
Week cluster	✓	✓	✓	✓
Implied elasticity	18.1%	5.1%	7.8%	0.7%

This table presents estimates from OLS regressions. The dependent variable is *bank branch visits*—the weekly number of visits received by a bank branch, compiled by the pass_by database. For every bank branch, we merge in the *net branch closures* of neighboring bank branches that are within the distance band stated in each column. The variable *net branch closures* is the net decrease in the number of bank branches over the past 180 days. We calculate the distances separating bank branches using their latitude and longitude coordinates. These coordinates are sourced from the FDIC SOD dataset and through geocoding the pass_by branch addresses via the Google Maps geocoding API. t-statistics are reported in parentheses.

Table IA.4. Quantifying the effects of branch closures on identity theft

Imputed effects from calendar ATTs

Year	(1) ATT (Net branch closures)	(2) ATT (ID theft reports)	(3) Wald estimate = (2)/(1)	(4) Total num. reports	(5) Total losses (U.S.$ million)
2011	2.60	−3.60	−1.38	4,973	6.40
2012	3.40	2.79	0.82	476	0.60
2013	3.39	221.65	65.37	73,868	104.46
2014	4.63	244.01	52.66	94,525	98.76
2015	3.43	351.36	102.35	144,005	79.83
2016	1.40	285.66	204.04	309,740	166.27
2017	2.50	198.85	79.61	183,186	125.60
2018	2.73	235.93	86.46	177,838	163.15
2019	1.65	521.01	315.28	680,686	612.84
2020	2.15	972.88	453.48	740,994	1,035.07
2021	2.30	604.49	262.94	982,351	1,884.40
2022	4.41	554.20	125.62	406,888	1,408.09

This table presents effects of branch closures on identity theft imputed from the Callaway and Sant'Anna (2021) calendar ATTs. The calendar ATT in a year is the average treatment effect for a MSA that is or is already exposed to large bank mergers in that year. Columns 1 and 2 present the calendar ATTs for *net branch closures* (i.e., the first stage) and the number of *ID theft reports* (i.e., the second stage), respectively. Column 3 presents the Wald estimates, which are the ratios of the second-stage estimates to the first-stage estimates. The Wald estimate represents the causal effect of one instrumented *net branch closure* on the number of *ID theft reports*. Column 4 imputes the number of *ID theft reports* by multiplying the Column 3 estimates by the actual *net branch closures* in MSA-years and aggregating them to the year level. Column 5 imputes the total losses by multiplying the imputed number of ID theft reports in the MSA-year by the average loss per report in the state-year and aggregating them to the year level.

Table IA.5. Synthetic difference-in-differences: Merger exposures and identity theft

Average treatment effects on treated (ATTs)

	(1)	(2)
Outcome:	Net branch closures	ID theft reports
Pre-exposure	0.03	–5.72
	(0.20)	(0.93)
Post-exposure	1.72***	177.42***
	(5.44)	(4.36)
Overall	1.13**	144.02***
	(2.13)	(2.82)
Wald estimate	127.45**	
	(2.22)	

This table presents ATTs on *net branch closures* and *ID theft reports* from the Arkhangelsky, Athey, Hirshberg, Imbens, et al. (2021) synthetic difference-in-differences estimator. Synthetic control units are constructed by optimally weighting never-exposed MSAs to match exposed MSAs on pre-treatment residualized outcomes, obtained after partialling out the covariates listed in Table 2. This procedure is equivalent to matching directly on pre-exposure outcomes and covariates, as described in Section 4.2.2 of the main text. The variable *net branch closures* is the net year-on-year decrease in the number of bank branches within an MSA. The MSA-year number of *ID theft reports* is collected from the Consumer Sentinel Network database of the U.S. Federal Trade Commission. The *pre* (*post*)*-exposure averages* are the aggregated ATTs before (after) merger exposures. An MSA is first exposed in the year when a merger occurs between large banks that both have branches in the area. Standard errors are clustered at the MSA level. Standard errors of the Wald estimates are computed using the delta method. Absolute values of t-statistics are reported in parentheses.