Unemployment Insurance Fraud in the Debit Card Market

Umang Khetan* Jetson Leder-Luis† Jialan Wang‡ Yunrong Zhou§

November 2023

PRELIMINARY - PLEASE DO NOT CITE OR CIRCULATE

Abstract

We study fraud in the unemployment insurance (UI) system using a dataset of 35 million debit card transactions. We apply machine learning techniques to group cards into clusters corresponding to varying levels of suspicious or potentially fraudulent activity. We then conduct a triple difference-in-differences analysis based on the staggered adoption of state-level identity verification systems between 2020 and 2021 to assess the effectiveness of screening for reducing fraud. Our findings suggest that identity verification reduced payouts to suspicious cards by 40% relative to non-suspicious cards, which were largely unaffected by these technologies. Our results indicate that identity screening of new and continuing applicants may be an effective mechanism for mitigating fraud in the UI system and for ensuring the integrity of benefits programs more broadly.

Keywords: Unemployment insurance, fraud, identity theft, COVID-19, screening, debit cards, machine learning

JEL classification: J65, G51, K42, H53

*University of Iowa

†Boston University & NBER. Corresponding author. Email: jetson@bu.edu

‡University of Illinois at Urbana-Champaign & NBER

§Purdue University

We thank Cameron Ellis, Tal Gross, Mark Kutzbach, Riley League, Ryan McDevitt, Felipe Netto, Maggie Shi, conference and seminar participants at the Financial Management Association 2023 Annual Meeting and the University of Iowa for helpful comments. We are extremely grateful to Jonathan Chin and Facteus for providing data access and technical support, to Filipe Correia for making this research possible with labor-intensive work on data intake and analytics, and to the AWS Greg Gulick Honorary Award for generous support.

1. Introduction

Unemployment insurance (UI) is one of the largest financial support programs in the United States, serving millions of individuals who lose their jobs each year and boosting the aggregate economy during economic downturns. The onset of the COVID-19 pandemic and sharp economic contraction led to a large expansion in unemployment insurance, with state and federal governments disbursing over $800 billion between 2020-2021. However, the rapid expansion of UI, which prioritized the speed of delivering benefits, was also susceptible to fraud, with frequent news reports claiming diversion of funds away from the intended recipients.[1] Poor targeting of UI benefits due to fraud is not only costly to the government budget, but can also have significant economic consequences, because fraudsters are unlikely to have the same labor market and spending responses as eligible recipients.

In the case of UI fraud, and fraud in other benefits programs, an important and widely used anti-fraud tool is identity verification. Identity verification is a form of screening that requires individuals to prove their identity before receiving benefits, to avoid diversion of funds to unintended recipients. Like other policies designed to improve targeting in benefits programs, identify verification faces a trade-off between hassle costs and more precise targeting toward intended recipients (Nichols and Zeckhauser, 1982). It imposes a time and hassle cost on applicants, creates an operational burden on the system, and could potentially delay payments or even prevent eligible applicants from receiving benefits. However, a more permissive system could potentially lead to waste and fraud, allowing non-eligible or nefarious applicants to receive payments they are not entitled to. Therefore, effective anti-fraud policy in the UI system is an important area for both public administration and for understanding the broader economic effects of UI.

In this paper, we study fraud in the UI system and examine the effectiveness of anti-fraud screening measures in reducing the misallocation of unemployment benefits. In response to growing concerns about fraud, 29 states adopted identity verification technologies using third-party vendors between March 2020 and September 2021. Although the specific protocols varied by state, the vendors typically verified each applicant's identity through photo or video-based authentication. The staggered adoption of these policies provides quasi-experimental variation that we use to identify their effects on disbursement of UI to potentially fraudulent recipients.

Our results indicate that identity verification is an effective, low-cost way to reduce benefits fraud. We find that identity verification led to a 40% reduction in UI disbursed to the group of recipients we categorize as suspicious, with minimal negative impact on most other UI recipients. Our analysis suggests that adoption of increased screening can improve resource allocation in large public benefit programs.

[1] The US Government Accountability Office (GAO) released an assessment report in 2023 highlighting heightened risk of fraud (United States Government Accountability Office, 2023). It can be accessed here. Other media and policy reports highlighting fraud instances include Axios, ProPublica, CNBC, the Attorney General of Pennsylvania, and the US Department of Labor.

Measurement is a natural challenge when studying fraud, as those who commit fraud try to conceal it. We identify indicators of potentially fraudulent activity by analyzing the income and spending patterns in a data set of tens of millions of debit card transactions, where we can observe both UI deposits and consumer spending behavior. Our granular data provides us with economically meaningful dimensions to separate out potentially fraudulent recipients from among a large sample of debit card users. The debit card market is particularly relevant for the study of UI, as we estimate that about a third of unemployment insurance benefits are paid to either bank-issued or prepaid debit cards.

We deploy an unsupervised machine learning algorithm on the transactions in our database and group debit cards into clusters representing varying levels of suspicious behavior. The key features we use in the clustering algorithm include the amount of UI received by each card, cash withdrawals as a proportion of spending, the speed with which money is withdrawn as cash from ATMs, the amount and proportion of funds wired outside of the United States, and concurrent income that might indicate ineligibility for UI. Some of these features, such as the quantity and speed of cash withdrawal, have been previously documented in the fraud-detection literature as strong signals of suspicious activity (Wu, Xu, and Li, 2019). Our clustering algorithm produces five groups of cards, with two representing suspicious activity.

We validate our suspiciousness measures using details about the spending behavior of our cards, as well as comparisons to ground-truth outside data. Our machine learning approach creates clusters of cards that have low intra-cluster heterogeneity and high cross-cluster heterogeneity in key dimensions indicative of potential fraud. Cards in the suspicious group simultaneously receive abnormally large UI inflows, spend a disproportionate amount on cash withdrawals, and withdraw money from ATMs relatively quickly after the receipt of UI. Moreover, the time-series of UI disbursed to suspicious cards shows lower correlation with underlying economic trends, and these cards are concentrated in regions that report a higher share of identity theft. On the other hand, cards that form our control cluster display fewer outlier characteristics in terms of UI inflows or spending. Such cards also show higher correlation with underlying trends in unemployment levels and are geographically dispersed.

We estimate the effectiveness of anti-fraud measures by comparing the change in UI disbursements to suspicious cards before and after these measures were introduced, with the control group defined as the cluster least consistent with fraudulent activity. We create a panel of state-level identity verification measures using data from Freedom of Information Act requests we conducted on all 50 state (and the District of Columbia) unemployment agencies. We then estimate the effects of identity verification policies using a staggered triple difference-in-differences design. Our identification technique captures common time trends across states, state-specific changes in unemployment insurance payments that affect all clusters, and cluster fixed effects. Our results are robust to critiques of modern two-way-fixed-effects difference-in-difference designs.

Our findings suggest that identity verification measures primarily affect the most suspicious

clusters, with minimal effect on non-suspicious recipients. Specifically, we find that the policy intervention reduced payout to cards in the most suspicious clusters by 40% relative to the control cluster. Likewise, the number of unique suspicious cards that received UI after the implementation of these policies dropped by 36%. Moreover, we do not find a significant reduction in the UI paid to non-suspicious cards after the policy was introduced, indicating that the frictions introduced by the ID verification system had a limited negative impact on non-fraudulent applicants. Traditional difference-in-difference estimates looking at the staggered policy variation on each cluster individually further confirm the pattern that UI disbursed to suspicious cards drops, while the control group is unaffected.

Our clustering methodology also allows us to explore the extent of fraud more broadly. We extrapolate from our sample to estimate that $29.5 billion was disbursed to suspicious recipients nationwide between March 2020 and September 2021, based on a 10.17% share of suspicious cards in our sample under the two primary UI programs that collectively disbursed $289.7 billion during this period. Moreover, our treatment effects imply that, if all the states had identity verification measures in place before the start of the pandemic, then this policy could have resulted in a nationwide savings of $12.6 billion over the course of our sample period. This assessment is based on an estimated $1.8 billion reduction in the UI disbursed by treated states to suspicious recipients after the policy was implemented, and counterfactual savings of $10.8 billion if these technologies were in force at the start of the COVID-19 pandemic. Our calculations are subject to assumptions including generalizability of our treatment effects across states and the representativeness our data.

Our work also has implications for effective anti-fraud policy in other domains. Screening and regulation eliminate the need to claw-back ill-gotten gains, which can be more effective in circumstances where fraud is heavily diffuse and therefore hard to prosecute (Mookherjee and Png, 1992, Polinsky and Shavell, 2000). Moreover, when those committing fraud have limited liability, preventing the disbursement of funds is necessary because ex-post prosecution would be unlikely to recover funds (Eliason et al., 2021). These forces are currently at play, as the Department of Justice (DOJ) is involved in dozens of lawsuits nationwide to try to recover stolen UI and PPP funds.[2] In contrast, better identity verification eliminates the need for ex-post recovery. Moreover, from a political economy angle, our work speaks to the limited incentives of bureaucrats in implementing successful anti-fraud policy. Identity verification is a low-cost, high benefit program, yet it took high levels of fraud to build the political momentum to implement it.

Our work also reflects on the recent literature on the effects of unemployment insurance during the pandemic and in general, discussed below. While extensive research has shown the effects of UI on household consumption, these effects must be balanced with the caveat that a substantial share of UI spending may be lost to fraud. Moreover, future policy reforms to the UI system should consider implementing effective verification metrics, as we show they have little negative impact on legitimate users but are effective at limiting fraud.

[2]Press releases from the DOJ on enforcement actions are available here from 2021 and here from 2023.

Related Literature: Our work connects the literatures on unemployment insurance during the COVID-19 pandemic, fraud in public programs, and consumer financial markets.

A large literature has examined unemployment insurance, including recent papers examining UI during the COVID-19 pandemic. In particular, Ganong, Greig, Noel, Sullivan, and Vavra (2022) show that pandemic-era expansions of UI benefits had large impacts on spending but small impacts on job search, and Dube (2021) similarly shows that the expiration of the Federal Pandemic Unemployment Compensation (FPUC) had little impact on job finding.

Beyond the COVID-19 pandemic, household finance scholars have examined the relationship between UI and consumer finances more broadly. Ganong and Noel (2019) find that spending drops sharply when UI benefits predictably expire, consistent with behavioral models. Hsu, Matsa, and Melzer (2018) explore the interaction between UI and mortgage markets, and find that UI expansions during the Great Recession prevented more than one million foreclosures. While these papers examine the effects of UI funds on households assumed to have been qualified and seeking employment, our paper attempts to quantify the extent to which funds went to individuals and organizations on a fraudulent basis and whose spending, credit, and labor market responses may be very different from that of qualified recipients.

Fraud against the government has generated substantial interest in the literature, but little work has examined unemployment insurance fraud per se. Most closely related to our work, Griffin, Kruger, and Mahajan (2023a) estimate the size of fraud in the pandemic paycheck protection (PPP) program, which provided grants to small businesses during the pandemic, and Griffin, Kruger, and Mahajan (2023b) show that fraud across COVID-19 programs spread through social networks. Aman-Rana, Gingerich, and Sukhtankar (2022) similarly find that additional paperwork requirements for second-round PPP loans reduced blatant fraud. Fuller, Ravikumar, and Zhang (2015) discuss theoretically optimal monitoring of unemployment insurance fraud; our paper complements that work by examining real-world policies.

Other complementary research has examined waste and fraud in other US benefit programs, including federal health insurance (Leder-Luis, 2020, Howard and McCarthy, 2021) and public procurement (Liebman and Mahoney, 2017). Eliason, League, Leder-Luis, McDevitt, and Roberts (2021) show that up-front paperwork requirements are effective at eliminating fraud in unnecessary federally-funded ambulance rides, with a mechanism analogous to the screening we study in this paper. In seminal work, Nichols and Zeckhauser (1982) show that ordeals can be useful for targeting appropriate beneficiaries, and a long literature has discussed how various forms of administrative ordeals not dissimilar to identity verification have played a role in appropriately targeting beneficiaries.

A growing literature on fraud in consumer financial markets highlights its impact across domains, including investment decision-making and corporate governance. Dimmock and Gerken (2012) demonstrate the predictability of investment fraud, revealing that investors could avoid

more than 40% of total dollar losses by avoiding the riskiest 5% of firms. Extending this line of research, Dimmock, Gerken, and Graham (2018) examine the social dynamics of fraudulent behavior, showing that misconduct can be contagious as coworkers significantly influence an advisor's propensity to engage in misconduct. On the regulatory front, Gao, Pacelli, Schneemeier, and Wu (2020) delve into Suspicious Activity Reports (SARs) filed by banks, offering critical insights into institutional reporting mechanisms. More recently, Bian, Pagel, and Tang (2023) show the effect of data protections on financial fraud against consumers. Griffin and Kruger (2023) encourage more investigation in forensic finance, which attempts to detect and understand the economic consequences of these behaviors.

Extensive work in machine learning (ML) has attempted to detect different types of fraud, such as public insurance fraud, e-commerce fraud, and credit card fraud. Shekhar, Leder-Luis, and Akoglu (2023) develop novel unsupervised machine-learning tools to identify fraud against federal health insurers, that provide an 8-fold lift over random targeting. Nanduri, Jia, Oka, Beaver, and Liu (2020) show that Microsoft uses customized sequential ML models to detect both historical and emerging fraud patterns. Recent work on ML in credit card fraud mostly focuses on supervised learning (Sadineni, 2020, Melo-Acosta, Duitama-Munoz, and Arias-Londoño, 2017). Khatri, Arora, and Agrawal (2020) and Jain, Agrawal, and Kumar (2020) present comparisons of established supervised learning algorithms to differentiate between genuine and fraudulent transactions. Several studies document the use of ML to extract information on consumer financial behavior more generally. Fuster, Goldsmith-Pinkham, Ramadorai, and Walther (2022) document that ML delivers higher out-of-sample predictive accuracy for default rates but has implications on racial disparity in the distribution of these gains. Berg, Fuster, and Puri (2022) review the growth of FinTech lending and note the increasing use of ML to improve customer screening.

This paper proceeds as follows. Section 2 discusses the background and institutional details. Section 3 discusses the data, and Section 4 discusses our measurement of suspicious behavior. Section 5 presents the methodology and effects of identity verification, and Section 6 presents a broader discussion of these findings in the context of anti-fraud policy. Section 7 concludes.

2. Background and Institutional Details

Unemployment insurance (UI) is a social safety net program that provides temporary financial assistance to eligible individuals who lose their jobs through no fault of their own. The benefits are meant to provide income support for workers who are laid off or furloughed while they search for new employment opportunities. These programs are generally administered by state governments, and applicants file for claims using their state's predominantly online application portal. States verify eligibility using the demographic and economic data supplied by the applicants, which includes their identity, social security number, date of birth, and address. This system is susceptible to fraud by various means, including identity theft, wherein criminals use attributes such as the social

security number of a potentially eligible recipient to apply for and divert UI funds.

Improper payments in benefits programs including UI has been a problem historically. In 2019 for example, the Bureau of Labor Statistics reported a 9% average overpayment rate of UI, with some states as high as high as 32%.[3] UI fraud became an even more pressing concern in the wake of the COVID-19 pandemic, when unemployment insurance expanded significantly. Part of this expansion included the rise of UI for informal-sector workers through a federal program called Pandemic Unemployment Assistance (PUA). According to the Bureau of Labor Statistics (BLS), the unemployment rate in the U.S. rose from 3.5% in February 2020 to 14.8% in April 2020, the highest rate since the Great Depression. Figure 1 plots the sharp increase in UI benefits paid starting in 2020, which is mirrored in debit card data we use. The unprecedented surge in unemployment claims, relaxed eligibility criteria, and the implementation of new relief programs acted together to make it easier for fraudsters to exploit the system and obtain benefits illicitly.[4] Identity theft, where an individual claims UI benefits using stolen information such as social security and date of birth, emerged as a major mechanism of fraud at a time when it was particularly onerous for states to verify the eligibility of UI applicants (United States Government Accountability Office, 2023).

In order to detect and deter fraudulent activity, many states implemented anti-fraud measures such as identity verification. Generally contracted through third party vendors such as ID.me and LexisNexis, identity verification measures sought to eliminate fraudulent spending by ensuring that the applicant and recipient of the funds were not using a stolen identity to claim these benefits. These measures did not seek to verify eligibility based on unemployment qualification, but rather to ensure that the person claiming the benefits was in bona fide possession of the information being supplied at the time of making a UI claim. Therefore, ID verification attempts to screen out fraudulent applicants before such claims are processed, rather than ex-post identify and litigate against individuals who may have already committed fraud.

A typical identity verification process requires UI applicant to submit photos of themselves and identity documents such as a driver's license or passport via a smartphone. Protocols varied by state, but the general mechanism involved in identity verification is presented in Figure A1 for vendor ID.me. When states implemented identity verification technology, it became a requirement before funds are disbursed. This made it more challenging for identity thieves to use others' stolen identity to apply for unemployment benefits. In order to ease the operational burden on legitimate applicants, vendors allowed for both photo-based authentication and a live video-based authentication with an employee of the company. A majority of claimants were able to complete this process within 10 minutes.[5]

We conducted Freedom of Information Act (FOIA) requests on all 50 state unemployment agencies to determine the nature and timing of their unemployment insurance identity verification

[3]The Department of Labor (DOL) UI payment accuracy datasets are available here.
[4]See pages 5-7 of the US GAO assessment report here.
[5]Estimates from the US Treasury department are available here.

policies and any external vendors contracted to provide such services. We corroborated these dates with a Congressional report about UI fraud policies (Committee on Oversight and Accountability, 117th Congress, 2022) as well as news reports.[6] Table A1 provides a timeline of ID verification measures along with the vendor contracted by each state. Figure A2 presents the time variation in state-level adoption of ID verification measures as a map.

The implementation timing of these anti-fraud measures varied widely across states. Several states began the implementation of ID.me as early as September 2020, including Indiana and Georgia. Conversely, other states commenced this process at a later date. For instance, Massachusetts did not start using ID.me until March 2021. Several states chose to implement identity verification only for their regular (non-PUA) UI programs. Moreover, some states implemented this screening in a staggered manner (such as initially for new claimants and later to include continuing claimants). While ID.me was the dominant platform for identity verification across most states, other vendors such as LexisNexis, GIACT, and Google Analytics were also contracted.

3. Debit card transactions data

We leverage a granular data set of 85,000 unique debit card holders covering 35 million transactions between January 2019 and September 2021. These data allow us to observe inflows in the form of UI or other sources, and outflows across a wide range of spending categories, including cash, wire transfers, food and grocery purchases, and other durable and non-durable expenditures.

Our data come from Facteus, a FinTech firm that partners with banks, card issuers, and payment processors to aggregate, standardize, and anonymize transaction-level information from debit cards.[7] The data on each transaction include the amount, timing, and a brief description. Further details such as merchant category code (MCC) and card-holder ZIP code help us confirm that the data constitute a wide variety of consumer transactions and a geographically representative sample.[8] Facteus perturbs the raw data it obtains from its data partners to generate synthetic data in order to protect consumer privacy. For example, the transaction amounts we observe are perturbed by adding small mean zero noise to raw transaction amounts. Similarly, transaction time is perturbed by up to several hours around the original time. This mean zero noise is small and we ignore it in our analysis. The transaction description, MCC code, and cardholder ZIP that we rely on for key results are not perturbed by Facteus.

Debit card holders are a particularly interesting and poorly-understood population to study

[6]The Congressional report is part of a set of documents published with the press release in November 2022 and accessed by us in January 2023. A Reuters news article on the states using ID.me technology for UI verification can be found here.

[7]For other papers that use data from Facteus, see Brave, Fogarty, Aaronson, Karger, and Krane (2021), Karger and Rajan (2020) and Zhou and Correia (2022).

[8]Using card-holder ZIP codes, Zhou and Correia (2022) show that the geographic distribution of cards in Facteus sample closely resembles the population density across the United States.

in the context of public benefits. Most states allow claimants to receive benefits on state-issued prepaid debit card or pre-existing prepaid card via direct deposit.[9] Recipients who receive UI on debit cards can also engage in spending and transfers through those cards, as well as make deposits and receive other forms of income. We observe all of these types of transactions in our data. Consumers who receive government benefits via prepaid cards tend to be lower income and more likely to lack traditional bank accounts.

We estimate that about one-third of UI benefits nationwide are paid into bank-issued or prepaid debit cards. Using the annual debit card market review reports from Mercator Advisory Group, we estimate that Facteus data captures 1.23% of overall debit card spending. We scale the share of UI observed in our data, as shown in Table A2, with this coverage to calculate the overall proportion of UI payments into debit cards.[10] In addition to a large share in UI, debit card users represent an important demography of the US population. The 2021 FDIC National Survey of Unbanked and Underbanked Households reports that 32.8% of unbanked households use debit cards for key financial transactions such as receipt of income and payment of bills (Federal Deposit Insurance Corporation, 2021).[11] The survey also reports that, within the unbanked segment, 77.8% use prepaid cards for bill payments and 64.1% for income receipt, while among the banked, the usage stands at 31.2% and 34.9% respectively. This segment of the population is often overlooked in traditional data sources from banks and surveys.

Furthermore, prepaid cards may be more susceptible to fraud than bank accounts, as criminals and fraudulent recipients can quickly open one or more card accounts for the express purposes of obtaining and siphoning funds, and the banking system has much more stringent measures for fraud and account opening compared with debit cards. One limitation of our data source is that we can only connect transactions at the card level, not the individual level. Therefore, if a single individual or entity obtains UI benefits on multiple different cards, we have no way of linking them and analyzing cross-card behavior.

We use a string matching technique on transaction descriptions to identify UI inflows from state agencies. Each state agency uses a consistent transaction description for UI payments; we source these descriptions from the Washington Bankers' Association.[12] We are able to identify UI inflows for 41 of the 50 states and the District of Columbia, and Table A2 provides details on our data coverage. In four states (Arkansas, Massachusetts, Ohio, and West Virginia), we are further able to distinguish inflows from regular UI versus the Pandemic Unemployment Assistance (PUA). For

[9]The Consumer Financial Protection Bureau (CFPB) explains the methods by which UI funds can be received in a blog post that can be accessed here. A specific example of a state-issued prepaid card from the state of Michigan can be found here.

[10]For external comparison, Brave, Fogarty, Aaronson, Karger, and Krane (2021) estimate that Facteus data covers 1.25% of debit card spending based on the Monthly Retail Trade Survey (MRTS) benchmark. Further, the state of Michigan reports that about a quarter of UI claimants elect to receive benefits on debit cards. The report can be found here.

[11]FDIC survey can be accessed here.

[12]The list of transaction descriptions for each state can be found here.

these four states, we aggregate state-level UI flows into PUA and non-PUA separately.

Based on the technique above, we observe state-level UI transactions on about 127,000 debit cards beginning in January 2019 and ending in September 2021. Out of these, we focus on 84,865 cards that received at least $1,000 of UI during the sample period. For these cards, our full sample contains 2 million UI inflow transactions, 3 million non-UI inflow transactions, and 30 million outflow transactions across various spending categories. Each transaction is tagged with the merchant's MCC code, which we use to categorize into types such as income from UI or other sources, spending via cash withdrawal from ATMs or bank branches, spending on groceries or other purchases, and wire transfers.[13] Additionally, we observe transaction timestamps which allow us to measure the time gap between receipt of UI and spending through cash withdrawals.

As a validation check, we compare UI flows from Facteus to administrative data from each state's Department of Labor. Figure 1 plots the time series of UI disbursements observed in our analysis sample with the corresponding aggregates from state administrative data, and shows that the two series follow a very consistent pattern over the sample period and have a correlation of 0.95. Table A2 shows the relationship between our data and the BLS data by state.

Table 1 shows card-level and card-month level summary statistics for our sample of cards that received any UI. On average, cards received a total of $10,000 of UI benefits across 19 disbursements, at a rate of about 4 per month. Spending is broken into categories including cash withdrawals, wire transactions, and groceries. The average monthly income per card is $2,790, of which $2,126 comes from UI. Total spending has a mean of $1,807 per month, with cash spending at $449 (25%) and groceries at $341 (19%) of the total. We group spending on other categories such as liquor shops, gambling, tobacco, restaurants, and purchase of vehicles as discretionary spend and note that it represents 16% of the total.

4. Measuring Fraud with Machine Learning

We use unsupervised machine learning to categorize debit cards along varying levels of suspicious behavior. This process allows us to construct the dependent variables of UI recipiency by different clusters of cards and evaluate the impact that identity verification may have had on each of these clusters. We begin with the steps involved in this clustering procedure and then discuss attributes that validate our interpretation of these clusters in the context of unemployment insurance fraud.

4.1. Cluster Construction

We apply unsupervised machine learning techniques on our granular income and payments data to cluster cardholders into varying degrees of suspicious conduct. These algorithms require the analyst to construct feature vectors, i.e. variables, as inputs for creating clusters. Recent literature

[13] The list of MCC codes is available here.

supports the use of payments data as informative about consumers' behavior; Puri, Rocholl, and Steffen (2017) provide evidence that payment data are predictive of loan default rates, and Berg, Burg, Gombović, and Puri (2020) show that signals from individuals' digital footprints can be as informative as credit bureau scores. Within the context of fraud, Wu, Xu, and Li (2019) note that patterns in cash spending and speed of cash withdrawal can meaningfully improve the accuracy of credit-card fraud detection.

We deploy a popular unsupervised machine learning algorithm called K-Means clustering on the income and spending patterns of cards to classify them into mutually exclusive and collectively exhaustive groups. Clustering is an unsupervised form of machine learning, meaning that it does not rely on labeled training data. Unsupervised techniques have the benefit that they do not rely on successfully identified fraud, which may be non-representative (Shekhar et al., 2023). We describe below the features used in our algorithm and the K-Means clustering procedure, and Appendix A provides further details on both.

Feature construction

From our dataset of 35 million debit card transactions, we construct card-level features that can differentiate suspicious UI recipients from non-suspicious ones. In addition to the amount and speed of cash withdrawal documented in the literature to be informative, we use characteristics relevant to the context of UI fraud such as the amount of UI, length of continuous recipiency, and concurrent income. Further, Ganong and Noel (2019) document heterogeneous impact of expiration in UI benefits on various spending categories. We use some of these categories such as spending on entertainment and transport to separate individuals with behaviors distinct from general population, and sharpen the identification of suspicious cards.

- Unemployment insurance received per month (in dollars): large receipts of UI could align with the incentive of fraudsters to maximize gains from stolen information. We aggregate the UI received by a card within a month and average it over the period.
- Longest unbroken UI spell (in months): cards that show an abnormally long unbroken spell of UI could be receiving benefits using multiple applications. We measure this spell using the maximum cumulative number of months for which a card receives UI.
- Cash withdrawals as a fraction of monthly outflows: cash withdrawals are not only difficult to track, but also make it hard to claw back illegitimate gains. We identify ATM (branch) withdrawals using MCC code 6011 (6010), and express it as a fraction of total monthly spend.
- Speed of cash withdrawals after UI inflow (in hours): a consistent urgency to withdraw cash after the receipt of UI further suggests suspicious behavior. We measure this variable using the time between a UI receipt and the following withdrawal, and average it at a card level.
- International wire transfers as a fraction of monthly outflows: transfer of large sums of money outside the US could indicate theft. We locate international transfers by string-matching merchant names such as Remitly, and express it as a fraction of total monthly spend.

- Discretionary spending as fraction of monthly outflows: this variable includes spending in liquor stores, gambling/casinos, tobacco stores, purchase of vehicles, and spending in restaurants, which could represent a segment of population distinct from the average UI recipient. We express this variable as fraction of total monthly spend.
- Concurrent non-UI income (number of consecutive months): concurrent non-UI income could indicate ineligibility for UI. We measure this variable as the maximum number of consecutive months a card receives at least as much income from non-UI sources as UI benefits.

K-Means Clustering

We use K-Means clustering to partition cards into several clusters that display distinct characteristics with respect to the dimensions we identify as salient. This method confers several advantages; (i) it flags anomalous patterns in transactions without relying on *a priori* labeled training data, (ii) there is no manual specification of the number of cards that must belong to each cluster, and (iii) we do not set arbitrary cut-off points for any of the dimensions that separate one card from another.

We make a few considerations when designing the clustering algorithm. First, several features collectively indicate potential fraud, but no one feature can be used standalone as a basis for grouping cards. For example, the UI received on a card, even in excess of statutory limits, does not by itself convey sufficient information because it is possible that multiple people use the same card. However, when unusually large UI inflows are followed by substantial cash withdrawals, it could indicate suspicious motives. Second, some dimensions are informative only conditionally. For example, the speed of withdrawal is applicable to only those cards that withdraw cash after the receipt of UI. Therefore, we construct a multi-stage clustering algorithm that starts with broad clusters to separate cards along common dimensions and then creates more granular categories based on additional relevant dimensions.

Figure A3 plots a flow chart of the sequence in which clusters are created. Our main clustering algorithm focuses on around 85,000 debit cards which receive at least $1,000 in UI. We start with clustering cards into those that receive unusually large amount or abnormally long unbroken spell of UI. This step creates two clusters: "High UI" with about 24% of cards that are considered for further sub-clustering, and "Low UI" with the rest that are included in the control category.

For cards in the "High UI" cluster, the clustering algorithm simultaneously uses the shares of average monthly cash withdrawals, international wire transfers, and discretionary spending for further grouping. This step generates one cluster each for cards that withdraw an abnormally large proportion of funds in cash, those that spend a large proportion of income on categories such as alcohol and gambling, and all other cards. Within the cluster with abnormally high cash withdrawals, the algorithm further splits them into those with fast or slow withdrawals after the receipt of UI. Finally, for the cluster with remaining High UI cards, we apply the clustering algorithm to detect potentially ineligible recipients using overlap with other concurrent income. Note that the

dimension of international wires, while potentially informative, does not produce a unique cluster because very few cards in our sample both receive high UI and transmit an abnormally large proportion of their income abroad.

All cards that do not fall into any of these clusters are combined with the control cluster created earlier using "Low UI" cards. We note that about two-thirds of "High UI" cards are tagged as control because they do not display suspicious patterns pertaining to spending or other income, and therefore there is no mechanical relation between suspicious cluster and receipt of large sums of UI.

The machine learning algorithm produces five clusters that group cards along economically comparable behavior patterns. Figure A4 shows bi-variate plots with color groups representing clusters along two dimensions at a time. Table 2 shows the means and standard deviations of the feature variables associated with each cluster. Based on their characteristics, we name each of the clusters, which we use to discuss our results going forward.

1. "Suspicious (Fast Cash)": abnormally large UI, and extremely large and quick cash withdrawals. This may reflect organized criminal behavior.
2. "Suspicious (Slow Cash)": abnormally large UI and cash withdrawals, but not at the same speed as Suspicious (Fast Cash).
3. "Concurrent Income": concurrent income from other sources and high levels of UI.
4. "Discretionary Spending": high amounts of UI received, and large spending on non-necessities, such as liquor stores.
5. "Control": cards that do not fall into one of the other clusters, generically reflecting genuine, non-fraudulent behaviors.

The labels we give to each cluster reflect our interpretation of the behavior of these cards. The "Suspicious (Fast Cash)" and "Suspicious (Slow Cash)" clusters receive much higher UI than other cards, and withdraw 66% and 61% (respectively) of their money in cash. Suspicious (Fast Cash) is distinguished by having a mean time to cash withdrawal of only 14 hours, which may be consistent with organized criminal activity. Withdrawing cash removes it from the banking system and makes future transactions untraceable, as well as lowering the probability of any future clawback measures. Since the transaction timing contains some mean-zero noise, for baseline analysis, we keep both suspicious together but show robustness to separating them.

4.2. Cluster Validation

Our algorithm creates clusters that represent distinct sets of population groups within the debit card market. Below we discuss the attributes that support our interpretation that two of these clusters show suspicious activity that is most consistent with fraud. We compare these clusters based on both, observable characteristics within, and ground-truths outside of our data.

First, cards belonging to Suspicious (Fast Cash) and Suspicious (Slow Cash) clusters display behaviors most consistent with those likely to be associated with identity theft. These cards simultaneously receive abnormally large UI inflows, spend about two-thirds by way of cash withdrawal, and withdraw cash relatively quickly. For instance, one card in these clusters receives UI in excess of $20,000 in a single day, and withdraws over 80% by cash. Another card receives UI in excess of $67,000 over a year, and withdraws $57,000 in cash from various ATMs and bank branches. Notably, these cards demonstrate elevated use of cash at a time when the economy shifted towards cashless transactions in the wake of the pandemic.[14] Suspicious (Fast Cash) cluster is also characterized by same-day withdrawals on average, consistent with the idea that converting illegitimate gains into cash makes it much harder for authorities to claw back the disbursed funds.

Second, the pattern of UI disbursements to cards in the suspicious cluster diverges from underlying economic trends, unlike that of the control cluster. Figure A5 plots the correlation coefficient between unemployment levels indicated by nonfarm payrolls, and UI disbursed to the control and suspicious clusters, respectively.[15] We note that the correlation of UI amounts disbursed by all 41 states with nonfarm payroll is 0.81 for the control cluster, while it is only 0.43 for the suspicious cluster. At a state level, control cluster has a higher correlation with nonfarm payroll than suspicious cluster in all but three states. This suggests that UI disbursements to suspicious cards are decoupled from genuine unemployment trends. In five states, the correlation of UI with nonfarm payroll for suspicious cluster is either negative or near zero, in stark contrast with the control cluster, where the correlation is positive throughout, and exceeds 0.5 in all but three states.

Third, the geographical distribution of suspicious cards in our data correlates with the rate of identity theft reports as per the Federal Trade Commission. Using the card-level zip codes in our data, we map each card to a metropolitan statistical area (MSA) and calculate the proportion of potentially fraudulent cards in each MSA, see Figure A6. The MSAs with some of the highest concentration of fraudulent cards in our data include Reno-Sparks (NV), Las Vegas-Paradise (NV), Boston-Cambridge-Quincy (MA-NH), Springfield (MA), Seattle-Tacoma-Bellevue (WA) and Miami-Fort Lauderdale-Pompano Beach (FL). Many of these show up among the MSAs with above-average rates of identity theft in publicly available administrative data.[16] Formally, in Appendix A subsection A.3, we regress the share of suspicious cards in our sample on the MSA-level reports of identity theft as per the Federal Trade Commision and find a positive and significant coefficient.

Finally, in line with media reports, the share of UI disbursed to suspicious clusters increased during the pandemic. Table A3 shows the amount and share of UI that went to each cluster before and after the onset of the COVID-19 pandemic. Likewise, for the four states where we can identify

[14] Cox, Ganong, Noel, Vavra, Wong, Farrell, Greig, and Deadman (2020) document a decline in cash spending after the onset of the pandemic.

[15] Nonfarm payroll is a monthly establishment survey of the number of people employed by firms in non-agricultural sectors. Unlike UI claims or UI disbursements in the administrative data, nonfarm payroll survey is less susceptible to the same kind of identity theft-driven fraud that our suspicious cluster aims to capture.

[16] The data from Federal Trade Commission is available here.

PUA and non-PUA inflows separately, the share of UI that went to suspicious clusters was higher under the PUA program, as shown in Table A4. This is also consistent with reports that the PUA program was more prone to fraud (United States Government Accountability Office, 2023).

5. The Effects of Identity Verification on UI Fraud

We construct a difference-in-differences model to estimate the effect of identity verification on state-level UI disbursement. Our main outcome variable is the UI inflows to cards based on their cluster, as measure of fraudulent or suspicious flows. We also consider the number of cards to whom UI was disbursed within a cluster as an alternate measure.

Our sample consists of the 41 states in our data where we can identify UI payments. Between January 2019 and September 2021, 22 states in our sample contracted with third-party agencies such as ID.me and LexisNexis to screen the identity of UI applicants; these states are considered treated.[17] Within these 22 treated states, the timing of adoption differs, starting from September 2020 for Georgia and Indiana to June 2021 for Delaware.[18]

We implement two specifications, each examining 6 months before and after the implementation of identity verification. First, we consider a triple-difference design, comparing the Suspicious cluster to the control (non-suspicious) cluster, before versus after the implementing of identity verification, with treated and control states. Next, we consider the classic difference-in-difference design, looking at each cluster separately. Both of these designs use the two-way-fixed-effects framework, saturating the model with state and time fixed effects. These designs allow us to examine if the more likely fraudulent cards were disproportionately affected by identity verification. Finally, we consider critiques of modern two-way fixed-effects estimators, and implement a number of robustness checks. Our estimates are robust to different specifications and are not contaminated with issues such as negative weighting.

5.1. Triple-difference specification

We estimate the model

$$y_{sct} = \sum_{cl \neq \text{“}Control\text{”}} \beta_{cl} \times \mathbb{I}_{[c=cl]} \times Post + \delta Post + \alpha_s + \alpha_c + \alpha_t + \alpha_{cs} + \alpha_{ct} + \varepsilon_{cst}, \tag{1}$$

where the dependent variable y_{sct} is the log unemployment insurance dollars paid by state s to

[17] Washington state implemented only a pilot version of identity verification on a small sample of recipients but did not extend it to the general population. We consider Washington as untreated.

[18] Some states implemented these measures for one program only (i.e. PUA or non-PUA, not both). We are able to distinguish between PUA and non-PUA inflows of UI for four states - Arkansas, Massachusetts, Ohio, and West Virginia. For those states, we construct the panel identifier as state-program rather than state.

cluster c in month t. *Post* is an indicator variable covering six months after a state s adopts identify verification technology, with six months before the implementation date being the pre-period for comparison. The treatment effect of identity verification is identified by the three β_{cl} parameters where the base cluster is "Control".

This comparison is relative to clusters with most likely genuine cards and therefore can be thought of as a triple-differences specification. The specification accounts for cluster, state, month, post, cluster-state, and cluster-month fixed effects. Standard errors are clustered by state and observations are weighted by state population. We estimate this model for Suspicious (combination of Suspicious (Fast Cash) and Suspicious (Slow Cash) clusters), Concurrent Income, and Discretionary Spending clusters. Table 3 reports the estimation result. Column (1) of Table 3 uses the log of UI dollars as the baseline dependent variable, and column (2) repeats the estimation on the log of number of cards that received UI within the cluster. (For the impact on number of cards, we count a card as a recipient of UI if it receives at least $100 because small amounts may indicate residual balances from past claims.)

Identity verification measures significantly impacted cards that displayed spending behavior most consistent with fraud. Column (1) of Table 3 shows that the treatment effect of this policy on Suspicious cluster was -0.52 log points, or a 40.5% decline, relative to the control cluster over a period of six months. This effect is observed after controlling for geographical and time-series variations in UI paid to card clusters. On the other hand, we do not observe a statistically significant impact on the other two (Concurrent Income and Discretionary Spending) clusters, indicating that these measures were targeted towards reducing specific kinds of fraud. Column (2) of Table 3 also shows a comparable decline in the number of unique cards receiving UI in the Suspicious cluster.

We next estimate a dynamic specification to validate that our results are not driven by pre-trends. The key identifying assumption is the standard parallel trends assumption – that the only differential change between suspicious and non-suspicious clusters at the time of transition that would impact UI payouts is the transition itself.[19] We estimate the model

(2)
$$y_{sct} = \sum_{\substack{\tau \in -6,6, \\ \tau \neq -1}} \sum_{cl \neq \text{"}Control\text{"}} \beta_{cl,\tau} \times \mathbb{I}_{[c=cl]} \times Reltime_{\tau} + \delta Reltime_{\tau} + \alpha_s + \alpha_c + \alpha_t + \alpha_{cs} + \alpha_{ct} + \varepsilon_{cst},$$

where the dependent variable y_{sct} is the log unemployment insurance dollars paid by state s to cluster c in month t. $Reltime_{\tau}$ refers to the number of months until treatment takes place, starting and ending six months about the month when identity verification was implemented. The treatment effect each month for each cluster is estimated using month "-1" and the "Control" cluster as base

[19] We also confirm that treated states did not introduce concurrent changes in the generosity of UI payments, such as the maximum weekly eligible amount, around the time of identity verification.

levels. Fixed effects for cluster, state, month and relative-month are included. Standard errors are clustered by state and observations are population-weighted.

Figure 2 shows the effect of identity verification on the UI received by the Suspicious cluster, relative to the control cluster. Panel (a) of Figure 2 uses log UI as the dependent variable and panel (b) uses the log of number of cards. The differences in UI payouts manifest themselves immediately upon the introduction of identity verification in (relative) month zero. ID verification reduced the UI received by the most suspicious cards by over 20% relative to the level of the control (non-suspicious) cluster immediately after its introduction.[20] This effect grows over subsequent months as more claimants are brought under the scope and payouts reduce to suspicious cards by as much as 70% in later months compared to the month before these measures were introduced. We also note the validity of parallel trends assumption before the treatment month and rule out the observed impact to any pre-existing trends.

We do not observe a similarly sharp change in the UI paid out to other clusters in Figure A8. Furthermore, we cannot rule out a downward-sloping pre-existing trend in the UI paid to the Concurrent Income cluster. These result suggest that the policy was successful in targeting claimants that sought to remove funds from the system soon after their receipt rather than individuals who may have claimed monies in excess of their eligibility or spent them on non-essential needs, but may not have conducted identity theft in the process of claiming them.

5.2. Within-cluster difference-in-differences design

Our triple difference specification compares the outcomes of the Suspicious cluster with respect to the control cluster; however, it does not tell us whether and how much the control cluster itself was impacted by this policy. We consider the possibility that both the Suspicious and control clusters may have experienced a decline in UI due to the treatment. For example, if ID verification introduced costly system-wide frictions, that would also affect genuine applicants. In contrast, if genuine applicants are unaffected, this would indicate UI identity verification was an effective form of targeting.

We conduct a classic difference-in-differences analysis to assess cluster-specific treatment effects. We estimate the following model on each of the four clusters separately:

$$y_{st} = \sum_{\substack{\tau \in -6,6, \\ \tau \neq -1}} \beta_\tau \times Reltime_\tau + \alpha_s + \alpha_t + \varepsilon_{st}. \tag{3}$$

This model is a cluster-specific version of Equation 2 and controls for state and time fixed effects. Parameter β_τ identifies the impact of identity verification on the UI paid by treated states

[20] Figure A7 shows the estimation results when the Suspicious (Fast Cash) and Suspicious (Slow Cash) clusters are kept separate, and shows similar trends.

to that cluster for each of the six months before and after implementation, based on the UI paid in month -1. Standard errors are clustered by state and observations are population-weighted. We estimate this model in turn for Suspicious, Concurrent Income, Discretionary Spending, and control clusters and show the resulting event plots in Figure 3. (Figure A9 reports the estimation using log of number of cards in each cluster that received UI instead of the log of UI dollars.)

Figure 3 show an immediate and persistent decline in UI paid to the Suspicious cluster of about 0.5 log points, or 40%, relative to the month before treatment. This magnitude is comparable to the triple-difference specification, which points to a real decline in the UI disbursed to potentially fraudulent recipients.

In contrast, Panels (b) and (c) show that ID verification does not seem to have impacted cards that received UI despite potential concurrent income or cards that spend disproportionately more on discretionary categories. This suggests that even if such cards were not necessarily eligible for UI, their primary means of claiming UI may not have been through identity theft.

Importantly, for the control cluster we do not observe a statistically significant decline in UI in any of the post-treatment months. Panel (d) of Figure 3 shows that the control cluster continued to receive UI comparable to month -1, with an upward trajectory in later months. This points to the idea that identity verification did not lower UI receipiency among non-fraudulent cardholders.

5.3. Robustness

We conduct three robustness checks to address critiques of modern two-way fixed effects designs in settings with staggered adoption of policy and potentially heterogeneous treatment effects.

5.3.1. Negative weights.

A TWFE design can give negative weights to already treated states if they are included as controls for states treated later during the sample period. This is particularly a concern when there is heterogeneity in treatment effects over time (Borusyak, Jaravel, and Spiess, 2021). In our setting, treatment effects may be heterogeneous across both time and jurisdiction. For instance, some states introduced ID verification initially for a subset of claimants (such as new applicants only) and extended it to all claimants over the following months. Similarly, the initial level of fraud in some states may have been higher than others, which could also affect the post-treatment outcomes. Roth, Sant'Anna, Bilinski, and Poe (2023) suggest that if treatment effects vary across both time and units, some observations may get negative weight in the TWFE estimand.

Following Roth et al. (2023), we check and rule out the presence of negative weights in our estimation. Under the common trends assumption, all 123 (41 states $\times$ 3 clusters) average treatment effects on the treated carry a positive weight. Consequently, the sum of the positive weights is equal to 1 and that of negative weights is equal to 0.

5.3.2. Alternative estimation of triple-differences design.

The potential inclusion of "already treated" states as control states for later treatments can lead to the violation of common trends assumptions. This is because the treated states may no longer follow the same trends as the never treated ones, due to the heterogeneity in treatment effects and a "phase-in" of treatment over time. In addition to checking for negative weights, we re-estimate both Equation 1 and Equation 2 using a "stacked" difference-in-differences approach to mitigate this concern.

To do this, we follow Cengiz, Dube, Lindner, and Zipperer (2019) in creating a stacked data set that constructs appropriate control groups for each treatment date. We follow four steps: (i) we define a window of 6 months before and after the treatment month to be included in each sub-experiment, (ii) we enumerate a total of 10 unique treatment months (or experiments) when all 22 states in our sample implemented ID verification, (iii) we define the inclusion criteria for control states in each sub-experiment to comprise of never treated and not treated during the post six month time window of that experiment, and (iv) we append the individual experiment data and re-estimate both Equation 1 and Equation 2 on this stacked data set.

The advantage of this approach is that the already treated states or the soon-to-be-treated states are not included in the evaluation of any given experiment, by design. This ensures that the estimated treatment effect is free from potential biases arising from violation of common trends. Each of these data sets is appended such that each transitioning jurisdiction appears once while each jurisdiction may appear as a control multiple times (although with different time values). In order to account for duplication of control states in the stacked dataset, we cluster standard errors by state. The resulting specification averages all of the time-varying effects into a single effect.

Table A5 shows the estimation of Equation 1 for the Suspicious, Concurrent Income, and Discretionary Spending clusters based on month -1 and the control cluster. We note that both the point estimates and the statistical significance are comparable to the baseline results of Table 3.

Figure A10 shows the dynamic event-study plots for the three clusters as in Equation 2. Again, we obtain trends consistent with the baseline TWFE version and interpret the results as causal evidence of identity verification on lower UI payout to the most likely fraudulent cards.

5.3.3. Alternative estimation of within-cluster difference-in-differences design.

We re-estimate Equation 3 using the semi-parametric group-time average treatment effect estimator proposed in Callaway and Sant'Anna (2021). This method estimates the effect of treatment separately for each group of states treated at the same time, using only those states that are never treated or not-yet treated as the control group. Using only valid comparisons for the estimation of average treatment effects, this method avoids weighting problems associated with TWFE specifications when there are multiple time-periods and variation in treatment timing.

Figure A11 shows the event study plots for all four clusters using the Callaway and Sant'Anna (2021) estimator. We continue to observe a decline in UI paid to the Suspicious cluster after the

introduction of ID verification. The pre-trends for the Suspicious cluster confirm that there was no decline observed before the implementation of policy. Further, the control cluster cards did not see a decline in the post period, consistent with the baseline results in Figure 3.

6. Estimation of Economic Magnitude

Our clustering algorithm also allows us to provide an estimate of the rough economic magnitude of UI fraud and the dollar benefit of ID verification. This is subject to assumptions such as the representativeness of our data and the generalizability of our treatment effects. We present the broad results below, and Appendix B provides detailed calculations.

We estimate that $29.5 billion was lost to fraud between March 2020 and September 2021. Table A3 shows that the share of UI disbursed to suspicious cards was 10.17% in our sample. Assuming that this share is representative, we arrive at the estimated misallocation by multiplying it with $289.7 billion, the total UI disbursed nationwide for regular and PUA programs. Furthermore, we estimate that identity verification saved $1.8 billion to treated states after implementation, and that $12.6 billion could have been counterfactually saved if all states had this policy in place at the onset of COVID-19 pandemic. These estimates are derived in three steps as detailed below.

First, we note that treated states disbursed a monthly average of $8.7 billion in the six months preceding their respective treatment month, and the (population-weighted) average number of post-treatment months was 5.1 until the end of our sample. Additionally, the share of UI paid to suspicious cards was 10.1% in the month preceding treatment. Given our treatment effect of a 40.5% reduction in UI paid to suspicious cards relative to control, identity verification saved these states $1.8 billion ($8.7 billion $\times$ 5.1 $\times$ 0.101 $\times$0.405).[21]

Second, we estimate the counterfactual savings by treated states if they had identity verification in place at the start of the COVID-19 pandemic in March 2020. As per our analysis, the pre-treatment UI disbursed by these states to suspicious cards from March 2020 through the month before treatment was $7.3 billion. Assuming that the magnitude of treatment effect remains constant at 40.5%, these states could have potentially saved another $3 billion ($7.3 billion $\times$ 40.5%) if the policy was in effect from March 2020. Finally, we extend the estimated counterfactual savings to the 19 states in our sample that did not implement identity verification, and 10 states outside our sample. These states altogether disbursed $167 billion in this period through the regular and PUA programs combined. The share of suspicious cards in this period for states that did not implement identity verification is 11.6%. Extending the same share of UI to suspicious cards among the 10 out-of-sample states, and assuming a treatment effect of 40.5% analogous to treated states, we estimate that an additional $7.8 billion could have been saved nationwide through this policy.

[21] An alternate calculation could account for the actual UI disbursed by treated states to suspicious cards after treatment, which was $3.3 billion until September 2021. If this amount was already lowered by the treatment effect of 40.5%, then the estimated savings comes to $2.2 billion ($3.3 billion $\times$ 0.405 / (1-0.405)).

We stress that these estimates rely on assumptions about our data coverage and the degree to which treatment effects observed in the states in our sample can be equally applied to all states. These are strong assumptions, particularly if there is any selection into the debit card market which we cannot observe. For example, the share of potential fraud observed in the debit card data may not equal the share of fraud through, for example, direct bank deposits and paper checks. However, in our favor, many states used debit cards to issue UI payments, regardless of the consumer's access to direct deposit, and therefore we can think of the debit card market as more representative of the population than it would be in other contexts. We estimate that about 32.5% of UI is paid into debit or prepaid cards. Thus, even within the debit card market, we estimate that $9.6 billion was lost to fraud and over $4 billion could have been saved due to identity verification.

Furthermore, we base our estimates on the two main UI programs active during the pandemic – regular state UI and the Pandemic Unemployment Assistance (PUA) for informal sector workers. The pandemic also led to an array of additional programs and enhancements to existing programs. For example, the Federal Pandemic Unemployment Compensation (FPUC) provided a federal top-up to existing benefits and was active from March through July 2020, and then again from January 2021. Since we do not observe the exact time-series of this and other programs, we do not include them in our estimation. Our estimate is therefore likely a lower bound of both potential fraud and the savings from the identity verification programs.[22]

Our estimates are presented as a point of comparison. Our conservative estimate of $12.6 billion in potential savings hints at the magnitude of economic gains that may have accrued from adoption of modern technologies to screen UI applications. The direct cost of these technologies, on the other hand, is a fraction of these savings: the Congressional documents we cite indicate that the total compensation contracted between ID.me and state agencies was under $60 million (Committee on Oversight and Accountability, 117th Congress, 2022). Overall, our analysis indicates that ID verification was a highly effective screening tool with significant positive economic benefits and low costs to the government.

7. Conclusion

Fraud is a major problem in public programs, and it is both challenging to measure as well as to eliminate. Through the COVID-19 pandemic, unemployment insurance saw a historic expansion that was plagued by fraud in the form of identity theft. This fraud threatened to divert funds from necessary recipients to potential fraudsters and to waste valuable resources in mitigating the economic effects of the recession.

We provide a first assessment of the effectiveness of identity verification in reducing fraud in UI.

[22]For comparison, the Government Accountability Office (GAO) estimates that between 11%-15% of total UI was lost to fraud during the pandemic, very close to the share of UI that we estimate went to suspicious cards. Their assessment additionally includes other pandemic era programs and therefore pegs the amount lost at $100 billion to $135 billion. The press release is accessible here.

First, we leverage machine learning tools to identify suspicious activity, clustering cards together on salient features such as how quickly cash was withdrawn, or whether the recipient had other income. Then, using data from a set of new FOIA requests, we show that the rollout of identity verification procedures was effective at reducing fraud, while sparing cards that did not engage in suspicious behavior. In all, these results suggest that identity verification measures were effective at targeting UI to needy recipients and mitigating fraud.

This study further provides new insights on how to measure fraud in benefits programs. Our clustering algorithm allows us to identify potentially fraudulent spending, which is valuable for understanding cost-benefit trade-offs in benefit expansion or the implementation of new regulations. A primary challenge in the study of fraud is that it is concealed; machine learning tools may prove broadly useful in future work to overcome these challenges.

Our work also opens a number of questions about the political economy of benefits programs, which may prove fruitful for future work. The technology required for UI verification was available before the pandemic, but it took potentially billions of dollars of fraud to occur before it was implemented. This highlights agency issues in the administration of UI fraud and other benefits programs, where technological adoption has largely been inefficiently slow (Pahlka, 2023). Given that identity verification was so valuable, our study highlights the fact that it could have been implemented before the expansion of UI, eliminating even more fraud.

References

Aman-Rana, S., D. Gingerich, and S. Sukhtankar (2022). Screen now, save later? the trade-off between administrative ordeals and fraud. *The Trade-Off between Administrative Ordeals and Fraud (August 18, 2022)*. [5]

Berg, T., V. Burg, A. Gombović, and M. Puri (2020). On the rise of fintechs: Credit scoring using digital footprints. *The Review of Financial Studies 33*(7), 2845–2897. [11]

Berg, T., A. Fuster, and M. Puri (2022). Fintech lending. *Annual Review of Financial Economics 14*, 187–207. [6]

Bian, B., M. Pagel, and H. Tang (2023). Consumer surveillance and financial fraud. Technical report, mimeo. [6]

Borusyak, K., X. Jaravel, and J. Spiess (2021). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*. [18]

Brave, S. A., M. Fogarty, D. Aaronson, E. Karger, and S. D. Krane (2021). Tracking us consumers in real time with a new weekly index of retail trade. [8, 9]

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of econometrics 225*(2), 200–230. [19, 49]

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics 134*(3), 1405–1454. [19, 48, 49, 54]

Committee on Oversight and Accountability, 117th Congress (2022, November 17). Congressional Press Release. [8, 21]

Cox, N., P. Ganong, P. Noel, J. Vavra, A. Wong, D. Farrell, F. Greig, and E. Deadman (2020). Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data. *Brookings Papers on Economic Activity 2020*(2), 35–82. [14]

Dimmock, S. G. and W. C. Gerken (2012). Predicting fraud by investment managers. *Journal of Financial Economics 105*(1), 153–173. [5]

Dimmock, S. G., W. C. Gerken, and N. P. Graham (2018). Is fraud contagious? coworker influence on misconduct by financial advisors. *The Journal of Finance 73*(3), 1417–1450. [6]

Dube, A. (2021). Aggregate employment effects of unemployment benefits during deep downturns: Evidence from the expiration of the federal pandemic unemployment compensation. Technical report, National Bureau of Economic Research. [5]

Eliason, P. J., R. J. League, J. Leder-Luis, R. C. McDevitt, and J. W. Roberts (2021). Ambulance taxis: the impact of regulation and litigation on health care fraud. Technical report, National Bureau of Economic Research. [4, 5]

Federal Deposit Insurance Corporation (2021). 2021 FDIC national survey of unbanked and underbanked households. [9]

Fuller, D. L., B. Ravikumar, and Y. Zhang (2015). Unemployment insurance fraud and optimal monitoring. *American Economic Journal: Macroeconomics 7*(2), 249–290. [5]

Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther (2022). Predictably unequal? the effects of machine learning on credit markets. *The Journal of Finance 77*(1), 5–47. [6]

Ganong, P., F. E. Greig, P. J. Noel, D. M. Sullivan, and J. S. Vavra (2022). Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data. Technical report, National Bureau of Economic Research. [5]

Ganong, P. and P. Noel (2019). Consumer spending during unemployment: Positive and normative implications. *American economic review 109*(7), 2383–2424. [5, 11]

Gao, J., J. Pacelli, J. Schneemeier, and Y. Wu (2020). Dirty money: How banks influence financial crime. *Available at SSRN 3722342*. [6]

Griffin, J. M. and S. Kruger (2023). What is forensic finance? *Available at SSRN 4490028*. [6]

Griffin, J. M., S. Kruger, and P. Mahajan (2023a). Did fintech lenders facilitate ppp fraud? *The Journal of Finance*. [5]

Griffin, J. M., S. Kruger, and P. Mahajan (2023b). Is fraud contagious? social connections and the looting of covid relief programs. *Social Connections and the Looting of COVID Relief Programs (October 12, 2023)*. [5, 36]

Hartigan, J. A. and M. A. Wong (1979). Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics) 28*(1), 100–108. [35]

Howard, D. H. and I. McCarthy (2021). Deterrence effects of antifraud and abuse enforcement in health care. *Journal of Health Economics 75*, 102405. [5]

Hsu, J. W., D. A. Matsa, and B. T. Melzer (2018). Unemployment insurance as a housing market stabilizer. *American Economic Review 108*(1), 49–81. [5]

Jain, V., M. Agrawal, and A. Kumar (2020). Performance analysis of machine learning algorithms in credit card fraud detection. In *2020 8th International Conference on Reliability, Infocom Technologies and Optimization (Trends and Future Directions)(ICRITO)*, pp. 86–88. IEEE. [6]

Karger, E. and A. Rajan (2020). Heterogeneity in the marginal propensity to consume: evidence from covid-19 stimulus payments. [8]

Khatri, S., A. Arora, and A. P. Agrawal (2020). Supervised machine learning algorithms for credit card fraud detection: a comparison. In *2020 10th international conference on cloud computing, data science & engineering (confluence)*, pp. 680–683. IEEE. [6]

Leder-Luis, J. (2020). Can whistleblowers root out public expenditure fraud? evidence from medicare. [5]

Liebman, J. B. and N. Mahoney (2017). Do expiring budgets lead to wasteful year-end spending? evidence from federal procurement. *American Economic Review 107*(11), 3510–3549. [5]

Melo-Acosta, G. E., F. Duitama-Munoz, and J. D. Arias-Londoño (2017). Fraud detection in big data using supervised and semi-supervised learning techniques. In *2017 IEEE Colombian conference on communications and computing (COLCOM)*, pp. 1–6. IEEE. [6]

Mookherjee, D. and I. P. Png (1992). Monitoring vis-a-vis investigation in enforcement of law. *The American Economic Review*, 556–565. [4]

Nanduri, J., Y. Jia, A. Oka, J. Beaver, and Y.-W. Liu (2020). Microsoft uses machine learning and optimization to reduce e-commerce fraud. *INFORMS Journal on Applied Analytics 50*(1), 64–79. [6]

Nichols, A. L. and R. J. Zeckhauser (1982). Targeting transfers through restrictions on recipients. *The American Economic Review 72*(2), 372–377. [2, 5]

Pahlka, J. (2023). *Recoding America*. Metropolitan Books. [22]

Polinsky, A. M. and S. Shavell (2000). The economic theory of public enforcement of law. *Journal of economic literature 38*(1), 45–76. [4]

Puri, M., J. Rocholl, and S. Steffen (2017). What do a million observations have to say about loan defaults? opening the black box of relationships. *Journal of Financial Intermediation 31*, 1–15. [11]

Rao, A. R., S. Garai, D. Clarke, and S. Dey (2018). A system for exploring big data: an iterative k-means searchlight for outlier detection on open health data. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. [35]

Roth, J., P. H. Sant'Anna, A. Bilinski, and J. Poe (2023). What's trending in difference-in-differences? a synthesis of the recent econometrics literature. *Journal of Econometrics*. [18]

Sadineni, P. K. (2020). Detection of fraudulent transactions in credit card using machine learning algorithms. In *2020 Fourth International Conference on I-SMAC (IoT in Social, Mobile, Analytics and Cloud)(I-SMAC)*, pp. 659–660. IEEE. [6]

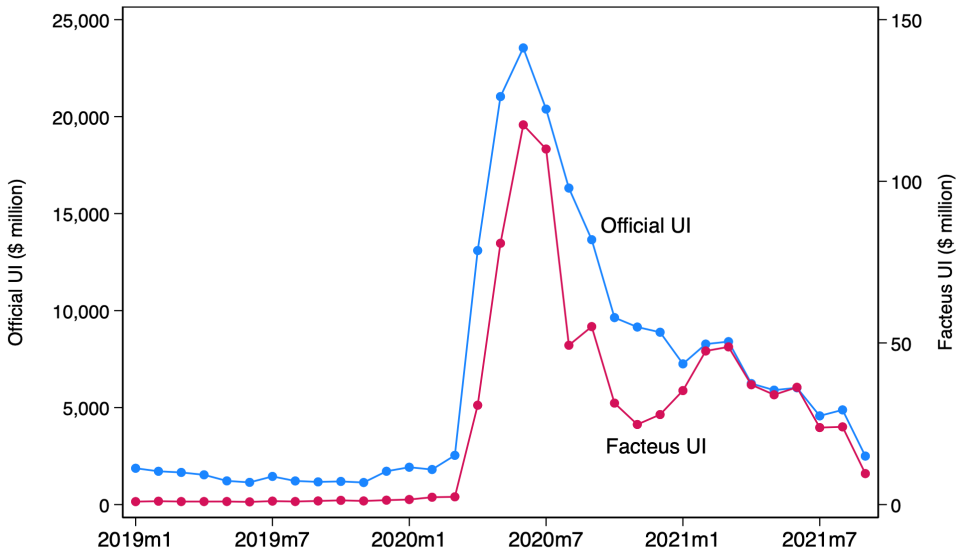
Shekhar, S., J. Leder-Luis, and L. Akoglu (2023). Unsupervised machine learning for explainable health care fraud detection. [6, 11]

United States Government Accountability Office (2023). Unemployment insurance: DOL needs to address substantial pandemic UI fraud and reduce persistent risks. Technical Report GAO-23-106586. [2, 7, 15]

Wu, Y., Y. Xu, and J. Li (2019). Feature construction for fraudulent credit card cash-out detection. *Decision Support Systems 127*, 113155. [3, 11]

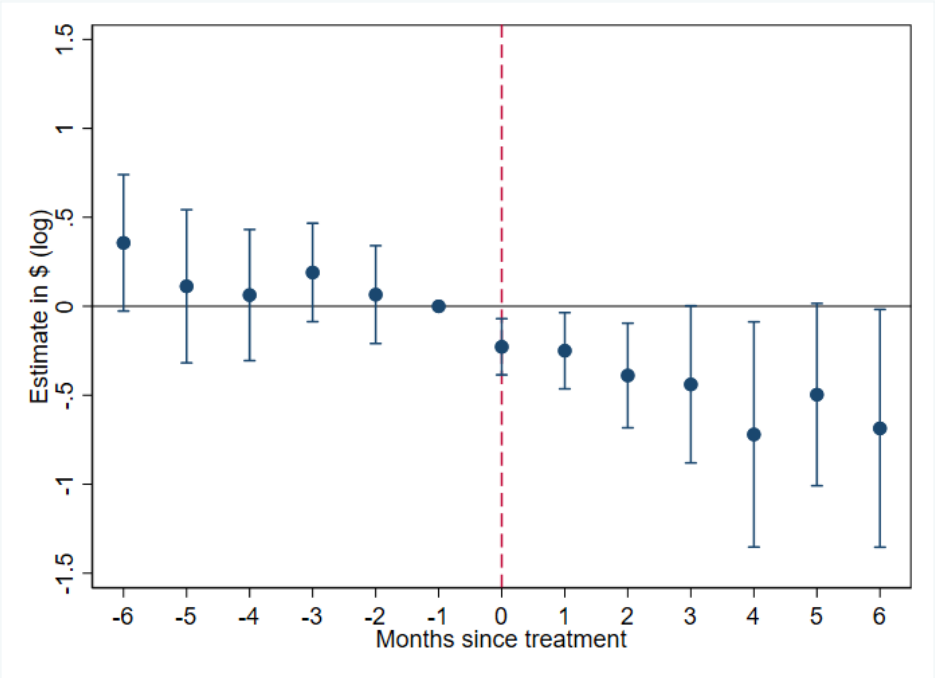
Zhou, Y. and F. Correia (2022). Are we friends? cross-predictability of stock returns. [8]

Figure 1: Monthly UI Disbursements in Administrative and Facteus Data

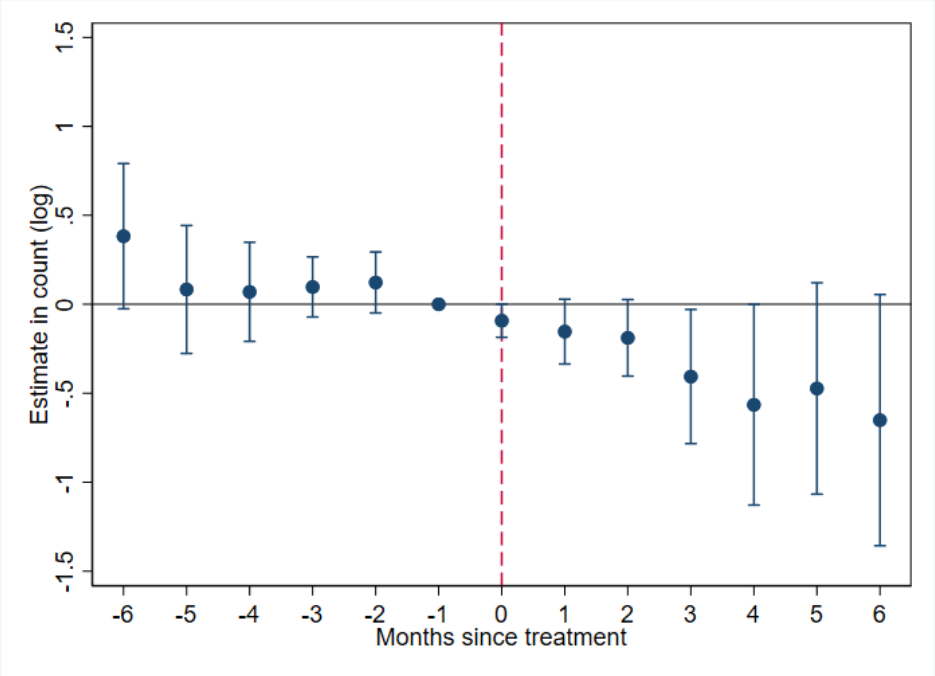


Notes: This figure plots total monthly UI disbursements from administrative data (left axis) and from Facteus debit card data (right axis). The sample period runs from January 2019 through September 2021, and comprises 40 states and the District of Columbia. Administrative data are sourced from the US Department of Labor and can be accessed here.

Figure 2: Impact of Identity Verification on Suspicious Cluster



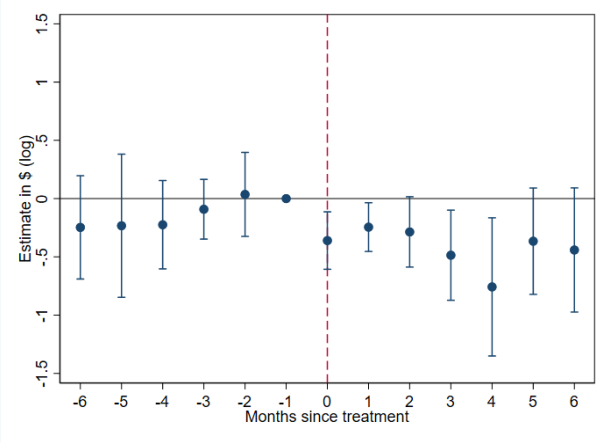
(a) UI dollars



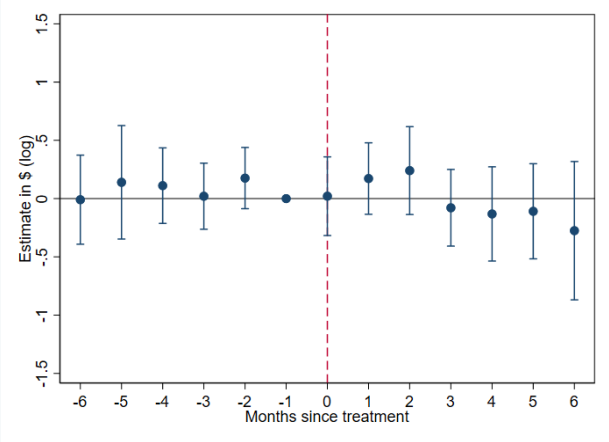
(b) Number of cards

Notes: This figure plots an event-study version of the triple difference Equation 2 for the suspicious cluster. Treatment refers to the implementation of identity verification measure at a state level in month 0. The treatment effect is compared to month -1 and to the control cluster. The dependent variable in plot (a) is log(UI) and in plot (b) is the log of number of cards in each cluster.

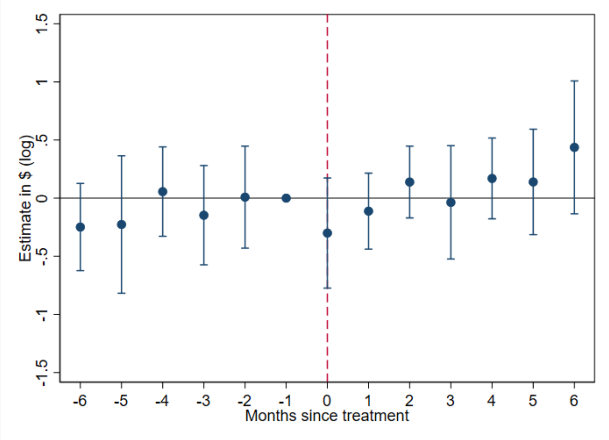
Figure 3: Within-cluster Impact of Identity Verification on UI dollars



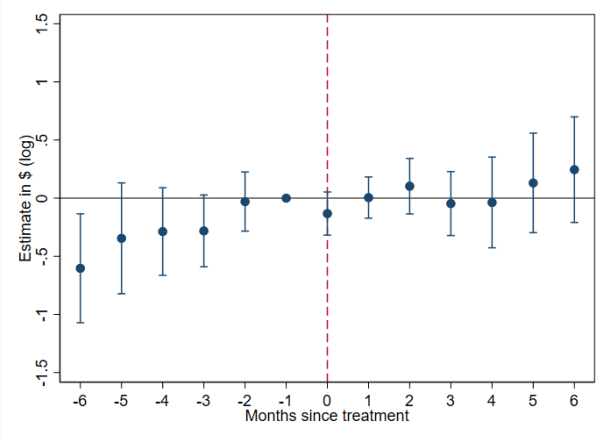
(a) Suspicious



(b) Concurrent Income



(c) Discretionary Spending



(d) Control

Notes: This figure plots the event study version of Equation 3, the difference-in-difference specification, for all four clusters individually. While our clusters reflecting concurrent income, discretionary spending, or control (non-suspicious) spending are unaffected, the suspicious cluster shows an immediate and persistent drop in UI disbursement following identity verification.

Table 1: Descriptive Statistics of Debit Card Sample

Panel A: Card-level	Mean	SD	p25	p50	p75	N
Total UI ($)	10,104	8,957	3,254	7,512	14,093	84,865
# UI (count)	19	21	4	12	26	84,865
UI per trans ($)	1,008	1,519	382	575	949	84,865
UI as % of income	0.63	0.32	0.34	0.68	0.97	84,865
Total income ($)	22,942	23,810	6,675	15,406	31,161	84,865
Total spend ($)	14,956	15,915	3,602	10,090	21,032	84,865
Cash spend ($)	2,875	5,347	0	343	3,594	84,865
Wire spend ($)	371	1,794	0	0	0	84,865
International wire ($)	40	731	0	0	0	84,865
Groceries ($)	3,176	5,464	252	1,271	3,738	84,865
Discretionary spend ($)	2,570	3,589	287	1,275	3,469	84,865
Panel B: Card-month level	Mean	SD	p25	p50	p75	N
Total UI ($)	2,126	2,104	825	1,703	2,700	403,349
# UI (count)	4	3	2	4	5	403,349
UI per trans ($)	697	1,005	261	451	769	403,349
UI as % of income	0.82	0.26	0.68	0.98	1.00	403,349
Total income ($)	2,790	2,589	1,205	2,196	3,476	403,349
Total spend ($)	1,807	1,849	563	1,374	2,468	403,349
Cash spend ($)	449	945	0	0	524	403,349
Wire spend ($)	52	310	0	0	0	403,349
International wire ($)	4	92	0	0	0	403,349
Groceries ($)	341	641	11	133	416	403,349
Discretionary spend ($)	288	413	40	166	391	403,349

Notes: This table shows card-level (panel A) and card-month level (panel B) income and spending statistics from our sample of transaction-level debit card data conditional on receiving at least $1,000 of total unemployment insurance (UI) during the sample period from January 2019 to September 2021.

Table 2: Summary Statistics by Cluster

	Suspicious (Fast Cash)	Suspicious (Slow Cash)	Concurrent Income	Discretionary Spending	Control
Unemployment Insurance ($/month)	3,122	3,015	2,132	2,618	1,200
	(2,412)	(1,823)	(458)	(1,473)	(1,527)
Longest UI spell (months)	4.7	4.5	10.9	5.0	4.2
	(3.7)	(3.7)	(4.1)	(4.1)	(3.5)
Cash withdrawal (share of spend)	0.66	0.61	0.09	0.02	0.13
	(0.18)	(0.19)	(0.11)	(0.05)	(0.21)
International wire (share of spend)	0.002	0.001	0.003	0.0004	0.002
	(0.02)	(0.02)	(0.03)	(0.01)	(0.02)
Non-essentials (share of spend)	0.07	0.07	0.19	0.52	0.22
	(0.06)	(0.06)	(0.09)	(0.13)	(0.15)
Time to cash (hours)	14	78	412	1,371	322
	(6)	(173)	(754)	(921)	(738)
Concurrent non-UI income (months)	0.2	0.3	4.7	0.5	0.6
	(0.8)	(1)	(2.1)	(1.2)	(1.2)
Stimulus checks ($, thousands)	0.2	0.2	2.1	0.2	0.6
	(0.9)	(0.9)	(2.6)	(0.8)	(1.5)
Tax refunds ($, thousands)	0.2	0.2	1.4	0.2	0.8
	(0.7)	(1)	(2.7)	(0.8)	(2.3)

Notes: This table shows the mean and standard deviations of card features at a cluster level. Clusters were created using an unsupervised K-Means algorithm. Standard deviations are reported in parentheses below the means.

Table 3: Treatment Effect of Identity Verification

	log(UI)	log(number of cards)
	(1)	(2)
Suspicious × Post	-0.520**	-0.441**
	(0.240)	(0.217)
Concurrent Income × Post	-0.309	-0.331
	(0.205)	(0.211)
Discretionary Spending × Post	-0.045	-0.137
	(0.137)	(0.176)
Observations	2,271	2,268
Adj. R^2	0.91	0.95
Cluster, State, Month, Post, Cluster-State, Cluster-Month FE	Y	Y

Notes: This table reports the results from a triple difference estimation from Equation 1 of the impact of identity verification measures on the UI received by various card clusters. In column (1), the dependent variable is log(UI), the log of unemployment benefits paid by each state to a cluster in a month, and in column (2) it is the log of number of cards in each cluster to whom these benefits were paid. The control group is the cluster with non-suspicious cards. The post period includes the six months after this policy was introduced for the respective state. Standard errors are clustered by state, and reported in parentheses. Identity verification reduces UI disbursed to the Suspicious cluster. Figure 2 shows this table's event-study version for the Suspicious cluster and Figure A8 for other clusters. $^{*}p < 0.1;^{**}p < 0.05;^{***}p < 0.01$.

Appendix for "Unemployment Insurance Fraud in the Debit Card Market"

Umang Khetan Jetson Leder-Luis Jialan Wang Yunrong Zhou

November 2023

A. Clustering

A.1. Feature construction

From our dataset of 35 million debit card transactions, we construct card-level features that can differentiate suspicious UI recipients from non-suspicious ones. We start with transaction-level data for all the debit cards in our sample and aggregate key features to card or card-month level to be used as dimensions for unsupervised clustering. The first step in this process is to aggregate all income and spending to a card-month total. The second step is to calculate the average per month in dollar terms for income and spending variables. We use per-month averages rather than per-card average to account for the differential length of time for which cards could be present in our sample. For spending variables such as cash withdrawal, we scale them by total spending in that month. Other variables such as speed of withdrawal, longest unbroken spell of UI and the number of months of concurrent non-UI income are calculated at a card level. Specific details for each feature are described below.

- Unemployment insurance received per month (in dollars): large amounts of UI flowing into a single card could align with the incentive of fraudsters to maximize gains from stolen information. Cards with abnormally high UI receipts could also indicate centralization of credit from multiple stolen identities. We sum the UI receipts for each card at a month level and convert it into a monthly average for all the months that a card is present in our sample.
- Longest unbroken UI spell (in months): most states limit the longest unbroken spell of UI to encourage individuals to find jobs. Cards that show an abnormally long unbroken spell of UI could be receiving benefits using multiple applications. We construct a time-series of UI receipts at a card-month level and calculate the number of consecutive months for which it received UI. We then retain the longest such spell for the card.
- Cash withdrawals as a fraction of monthly outflows: spending by means of cash withdrawal is not only difficult to track, but also makes it hard to claw back illegitimate receipts. Furthermore, most ATM withdrawals incur a fee, which should dis-incentivize legitimate recipients from spending via cash. We identify ATM withdrawals using MCC code 6011 and bank branch withdrawals using MCC code 6010, and add them together at a card-month level. Then, we convert it to a percentage of total spend by that card in that month so that it is comparable across cards and times.
- Speed of cash withdrawals after UI inflow (in hours): a consistent urgency to withdraw cash after the receipt of UI further suggests suspicious behavior. We measure this variable using the timing difference between UI receipt and the next immediate withdrawal, and then average this at a card-level. The timestamps are expressed up to the second but contain mean-zero noise to protect user privacy. For timing difference, we calculate the time elapsed from the last UI hit to the next ATM or bank branch withdrawal. We take care to include only those cases where cash was withdrawn after the receipt of UI and not before.
- International wire transfers as a fraction of monthly outflows: UI funds are intended to cushion against economic shocks. Transfer of large sums of money outside the country could be indicative of ineligibility and theft. We locate international wire transfers using string-matching of merchant names such as Remitly,

and express this variable as percentage of total spend at card-month level.[23]

- Discretionary spending as fraction of monthly outflows: this variable includes spending in liquor stores, gambling/casinos, tobacco stores, purchase of vehicles, and spending in restaurants. While spending on these categories may not indicate fraud, these could point towards distinct groups of people that could be separated from and lead to sharper identification of fraudsters using identity theft. We express this variable as percentage of total spend at card-month level.
- Concurrent non-UI income (number of consecutive months): by design, UI compensates individuals for loss of wage income. Concurrent non-UI income could indicate ineligibility to receive UI. We measure this variable as the number of consecutive months a card receives at least as much income from non-UI sources as UI benefits and average it at a card-level. We identify non-UI income as the residual from total income after subtracting UI, and one-off payments such as stimulus checks or tax credits.

A.2. Classification procedure

Clustering algorithms are machine learning techniques that group observations by their related values. Rather than extrapolating from known fraudulent patterns, unsupervised learning groups observations by similar patterns among variables chosen by the researcher. A K-means algorithm chooses centroids for each dimension and creates clusters based on the proximity of each card to the centroid along that dimension (Hartigan and Wong, 1979). The K-Means algorithm also normalizes all features across cards so they are weighted equally in each step.

We follow a multi-step clustering procedure that isolates cards with outlier characteristics across multiple dimensions. A multi-step or iterative K-Means algorithm allows us to apply a common set of informative dimensions in each step and narrow down to the set of truly outlier cards. For example, Rao, Garai, Clarke, and Dey (2018) apply iterative K-means in two steps: first, over the entire dataset, and second, over cluster subsets to further elaborate any dimensions that contain outliers. This process allows the uncovering of outliers within what might appear to be broadly homogeneous clusters in the initial step. The algorithm updates the centroids at each iteration to more accurately identify homogeneous clusters. As a result, the within cluster heterogeneity is small and the cross-cluster difference on informative features is large.

The first K-means algorithm is based on the average UI received per month by cards and the longest unbroken spell of UI. These features separate cards based on recipiency, which is relevant to the economic context of fraud. Out of about 85,000 cards, 20,000 fall in the high UI bucket. These cards are further split based on spending categories such as average fraction of spending on cash withdrawals, international transfers, and on discretionary items. For cards that do not show outlier spending patterns but nevertheless receive a large sum of UI, we apply the clustering algorithm to detect potentially ineligible recipients using overlap with other concurrent income. All cards that do not fall into any of these clusters are combined with the control cluster created earlier using "Low UI" cards. We note that about two-thirds of "High UI" cards are tagged as control because they do not display suspicious patterns pertaining to spending or other income, and therefore there is no mechanical relation between suspicious cluster and receipt of large sums of UI.

[23] A recent case of UI wire fraud by foreign nationals carried out using debit cards is available here.

The algorithm creates five clusters. The first two are suspicious in that they simultaneously receive large sums of UI and spend, on average, two-thirds of it in cash. The next two clusters represent either spending on non-essentials or concurrent income from other sources, indicating some likelihood of ineligibility to receive UI. Concurrent income indicates ineligibility but does not point towards identity theft as strongly as the suspicious clusters. Discretionary spending could point towards a distinct set of people who are also less likely to commit identity theft because these expenses are traceable to the cardholder. All other cards with the least fraudulent behavior form part of the control cluster – who form the bulk of our sample and do not display any combination of suspicious activity.

A.3. Cluster validation

We perform geographical validation of the clusters generated by our machine learning algorithm using data on official identity theft reports provided by the Federal Trade Commission. These data are available at an MSA (Metropolitan Statistical Area) level. Similar to Griffin, Kruger, and Mahajan (2023b) who regress excess UI claims in a county on their indicators of PPP (Paycheck Protection Program) fraud, we estimate the model

$$\text{Share of Suspicious Cards}_i = \beta \text{Identity Theft Reports}_i + \gamma \text{Population}_i + \varepsilon_i, \tag{4}$$

where the dependent variable is the fraction of cards in our data at an MSA that we classify as suspicious. We map the ZIP codes of each card to arrive at the MSA they belong to. To reduce noise in this estimation, we retain MSAs where we have at least 100 cards. Because there is no time variation in our ZIP code level measure, we use 2020 as the year when most states had not yet implemented ID verification and gets us the largest set of observation. Standard errors are clustered by state. Table A6 reports the estimation results, with population added as an additional control in column (2).

Our measure of suspicious cards strongly correlates with official reports of identity theft. In both the columns, the coefficient attached to Identity Theft Reports is positive and significant, indicating that the cards we categorize as suspicious belong to MSAs that reported a higher population-adjusted share of identity theft reports.

B. Estimation of economic magnitude

Our calculations of the economic magnitude of fraud and savings from screening technologies are based on administrative data on unemployment insurance (UI) payments (sourced from the US Department of Labor), the share of UI to suspicious cards in our data, and treatment effects implied by our two-way fixed-effects model.

We start by noting that the share of total UI received by suspicious cards between March 2020 and September 2021 was 10.17%, as per Table A3. Administrative data indicate that the total UI paid under the pandemic unemployment assistance (PUA) and regular state UI (non-PUA) programs by these 41 states was $192.8 billion during the same period. Under the assumption that the share of fraud observed in our data is representative, the UI allocated to potentially fraudulent cards amounts to $19.6 billion among these 41 states. Furthermore, we can

extrapolate this number to include the remaining 10 states that do not constitute our sample. The total UI paid by all 50 states (and District of Columbia) between March 2020 and September 2021 was $289.7 billion ($161.2 billion under non-PUA and $128.5 under PUA program). This brings us to an estimated $29.5 billion paid to suspicious cards over 19 months starting March 2020.

Next, we estimate the dollar savings from ID verification programs to treated states. Table 3 shows that identity verification declined the monthly UI paid out to suspicious clusters by 40.5%. There are two ways of estimating the dollar savings to treated states. The first method extrapolates both the amount of UI and the share of suspicious cluster from the pre-treatment period to the post-treatment period. Under this method, we note that the average UI disbursed by treated states was $8.7 billion per month, in the six months preceding the treatment. Further, the (population-weighted) average number of post-treatment months was 5.1 until September 2021 when our sample ends. Finally, the share of UI paid to suspicious cards was 10.1% in the month preceding treatment. This gives us the counterfactual UI of $4.5 billion ($8.7 billion $\times$ 5.1 $\times$ 10.1%), which would have been paid to suspicious cluster in the absence of identity verification. Given the treatment effect of a 40.5% reduction in UI paid to suspicious cards, the savings from identity verification comes to $1.8 billion ($4.5 billion $\times$ 0.405) for the states that implemented this policy.

The second method does not rely on the pre-treatment amount of UI and share of fraud clusters, but instead accounts for the actual UI paid by treated states and scales it by the treatment effect. The actual UI disbursed to suspicious cards by treated states after treatment until September 2021 was $3.3 billion. If this amount was already lowered by the treatment effect of 40.5%, then we can derive the estimated savings by inflating this number by the treatment effect and calculating the difference between the two. This method provides us with an estimated savings of $2.2 billion ($3.3 billion $\times$ 0.405 / (1-0.405)).

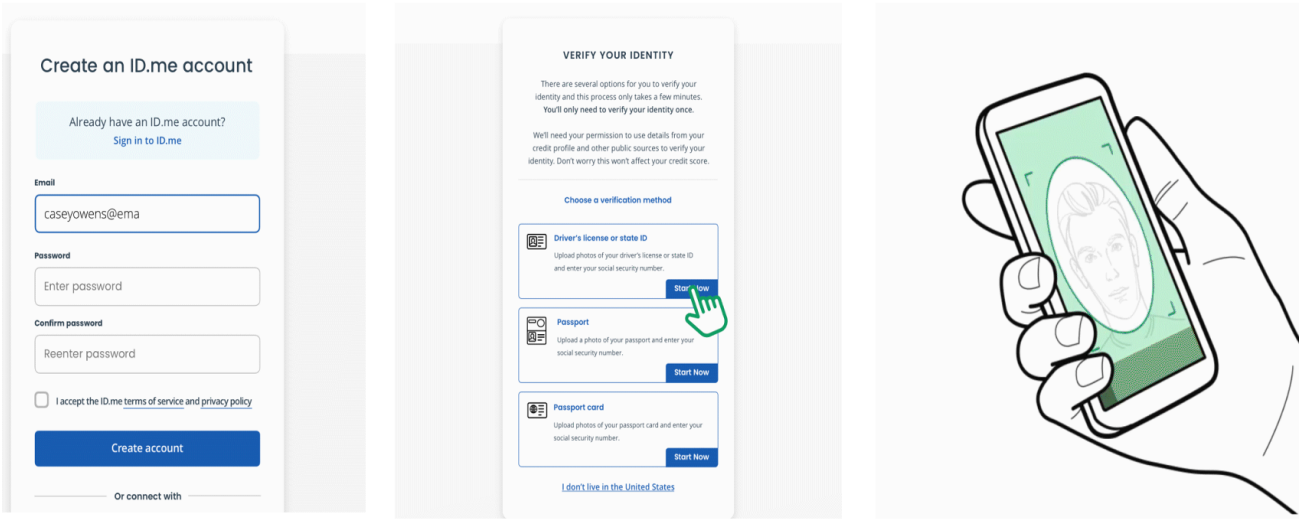
We now calculate the counterfactual savings to treated states that would have accrued if ID verification was in place at the start of the COVID-19 pandemic in March 2020. The pre-treatment UI disbursed by these states in total was $90.4 billion ($60 billion under non-PUA and $30.4 billion under PUA program). As per our analysis, the average share of UI to suspicious cards from March 2020 through the month before treatment was 8.1%.[24] Therefore, the UI amount allocated to suspicious cards was $7.3 billion ($90.4 billion $\times$ 8.1%). Assuming that the magnitude of treatment effect remains constant at 40.5%, these states could have potentially saved $3 billion ($7.3 billion $\times$ 40.5%) if identity verification was in place from the beginning.

The last step is to extend the estimated counterfactual savings to states that do not form our treated states sample. This includes the 19 states in our sample that did not implement identity verification between March 2020 and September 2021, 9 states outside our sample for which we are unable to identity UI stream in Facteus data, and 1 state (Iowa) for which the timeline of identity verification is unclear. These states altogether disbursed $167 billion in this period ($87 billion through non-PUA and $80 billion under the PUA program). The share of suspicious cards in this period for states that did not implement identity verification is 11.6%. Assuming that (i) the same share applies to the 10 states outside our sample, and (ii) identity verification would have reduced this share by 40.5% analogous to treated states, we estimate that an additional $7.8 billion could have been saved by

[24]Consistent with the reports of a rise in fraud after the pandemic, the average share of suspicious cards is lower on average leading up to the treatment than it is in the month immediately preceding the treatment.

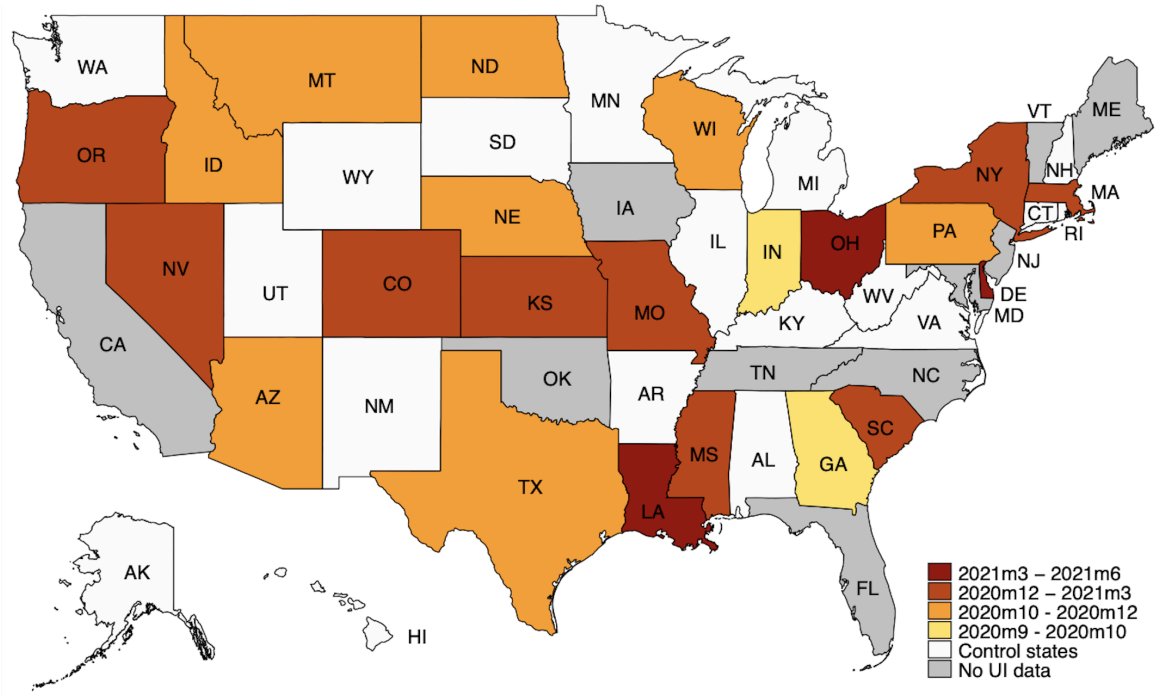
these 29 states ($167 billion $\times$ 11.6% $\times$ 0.405). The total counterfactual savings from ID verification for the full period and all states put together comes to $12.6 billion ($1.8 billion + $3 billion + $7.8 billion).

Figure A1: Identity Verification Process



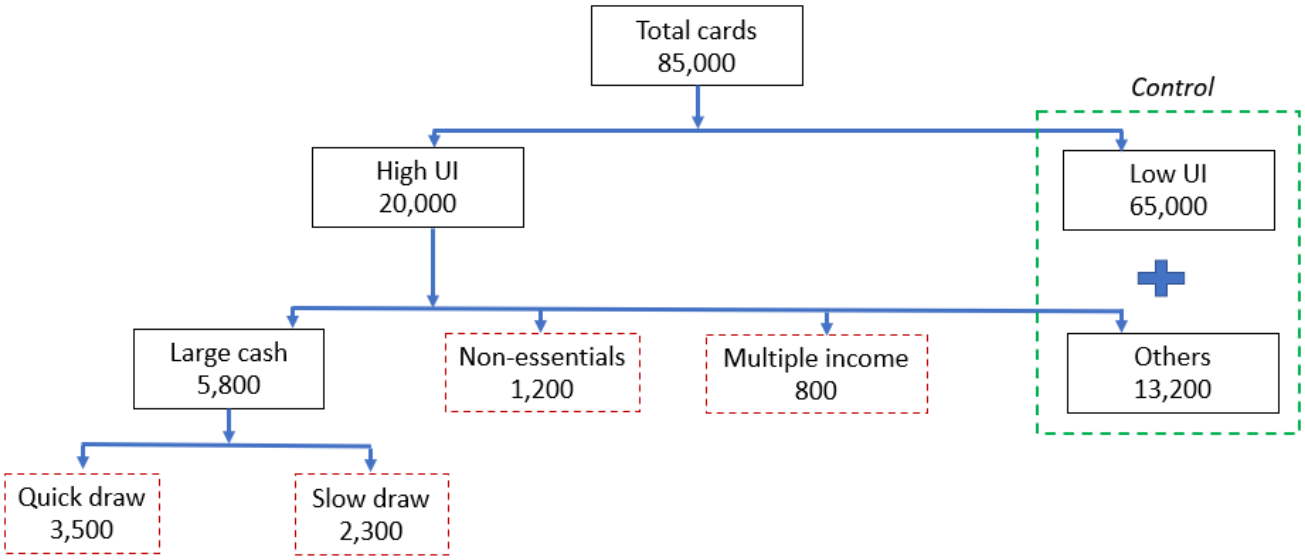
Notes: This figure shows three general steps involved in verifying an applicant's identity when claiming UI benefits. The figure has been adapted from ID.me's website.

Figure A2: Treatment Month for Each State



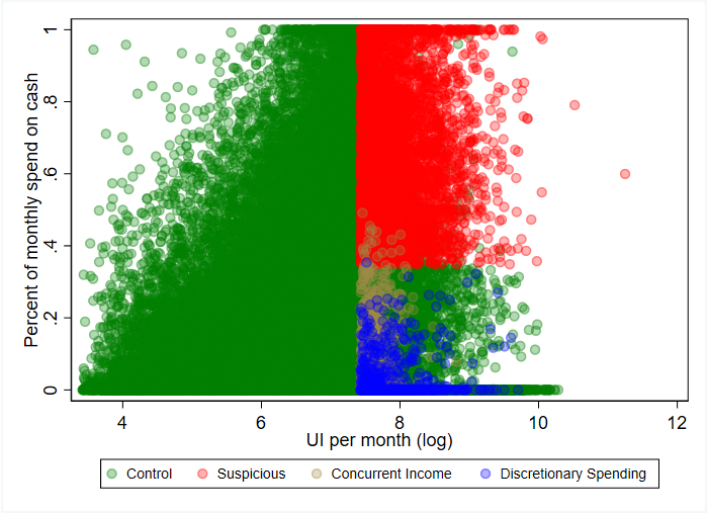
Notes: This figure shows treated states, control states, and the states for which we cannot identify UI or the timing of ID verification. This figure also shows the timing of the adoption of identity verification by each treated state.

Figure A3: Summary of K-Means Clustering

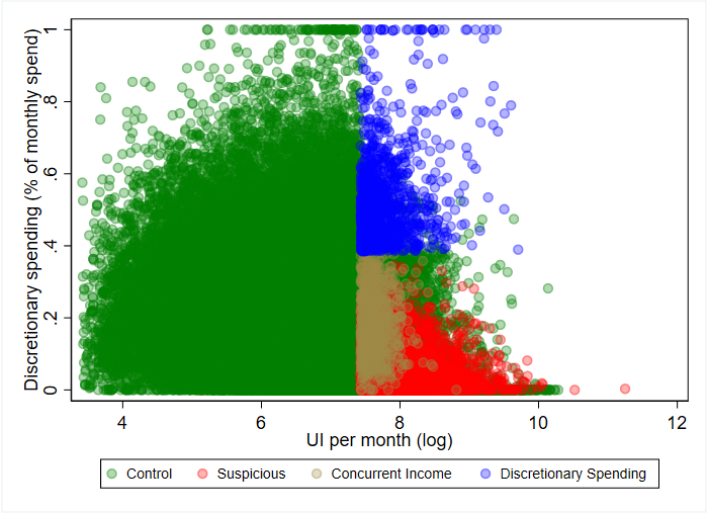


Notes: This figure depicts the order in which cards are grouped into clusters representing varying levels of fraudulent behavior. The top branch splits about 85,000 cards into those with high and low UI using monthly UI received by each card and the longest spell of being on UI rolls. Within cards receiving high levels of UI, the algorithm uses spending on cash withdrawals, non-essential categories such as alcohol and gambling, and concurrent non-UI income to create three clusters. Cards with large cash spending are further split based on the speed of withdrawal after the receipt of UI. Cards with low UI or high UI but without outsized spending in the above-mentioned categories are collected under the control cluster. The labels show the approximate number of cards in each split.

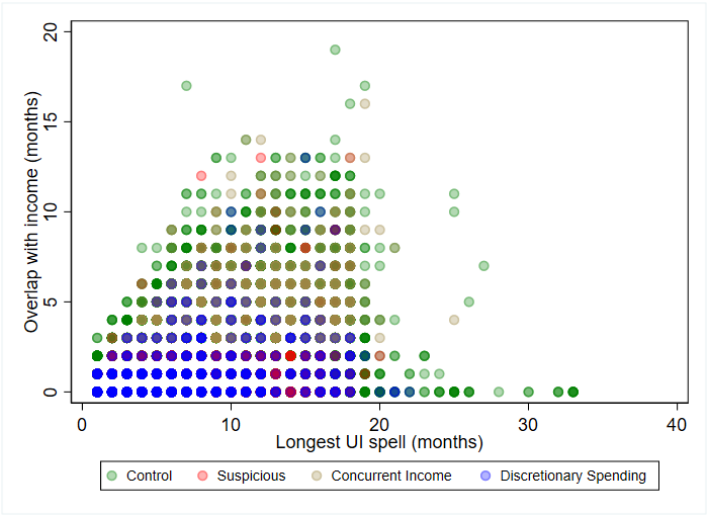
Figure A4: Bivariate Plots of K-Means Clusters



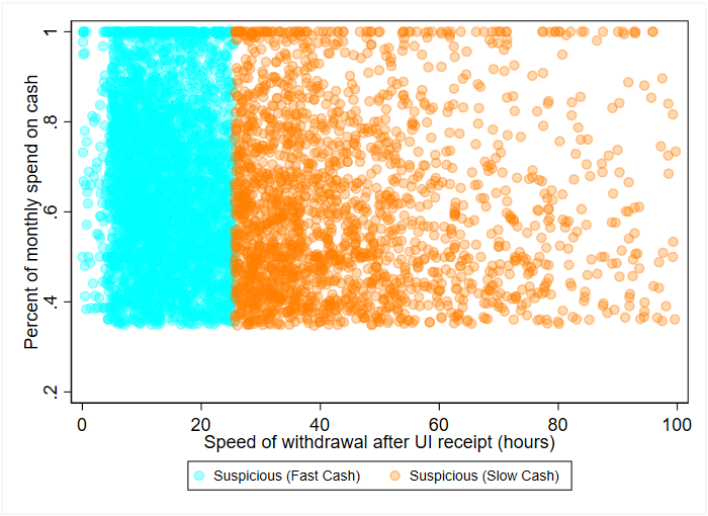
(a) Cash withdrawals (%) v/s UI per month (log)



(b) Discretionary spending (%) v/s UI per month (log)



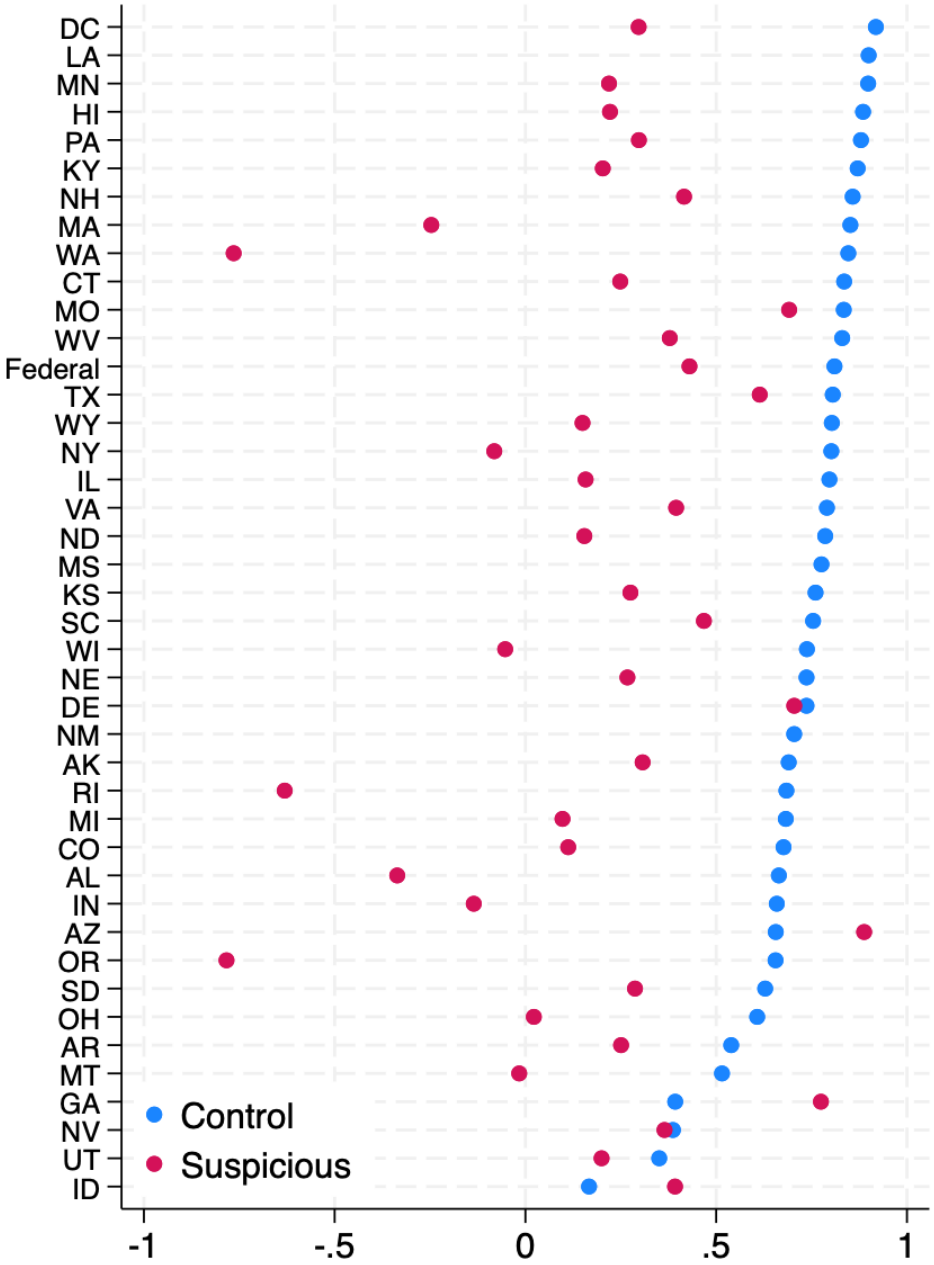
(c) Other income overlap v/s longest UI spell



(d) Cash withdrawals (%) v/s Time to cash (hours)

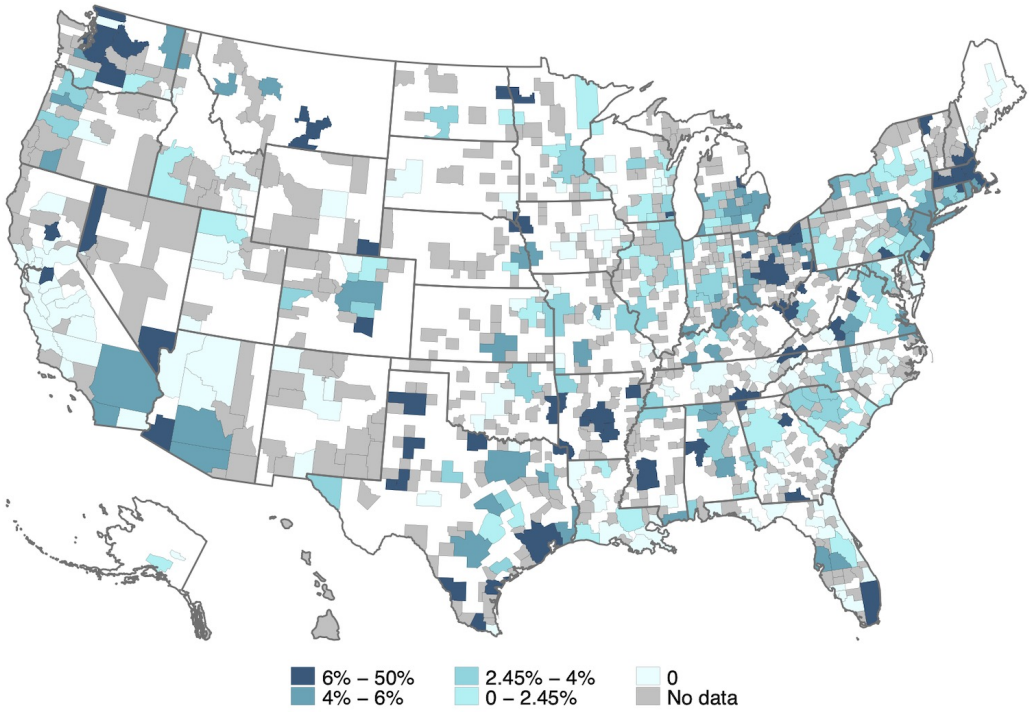
Notes: This figure shows a scatter plot of each card and the cluster it belongs to using two dimensions at a time. Darker regions show greater density of cards at that point.

Figure A5: Correlation of Monthly UI with Unemployment implied by Nonfarm Payroll



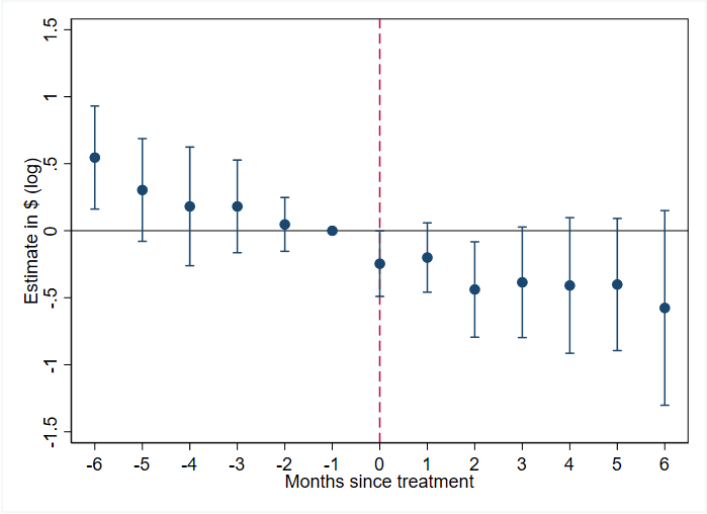
Notes: This figure plots the correlation between (inverse of) nonfarm payroll and the number of cards to whom UI was disbursed to suspicious cluster (in red) and control cluster (in blue). Number of suspicious cards to whom UI was disbursed has a lower correlation with underlying economic trends. Non-farm payrolls are derived from establishment surveys and likely less susceptible to distortions from identity-theft, such as actual UI claims or disbursements. Federal includes 41 states in our sample. The time-series runs from January 2019 through September 2021.

Figure A6: Share of Suspicious Cards at the MSA level

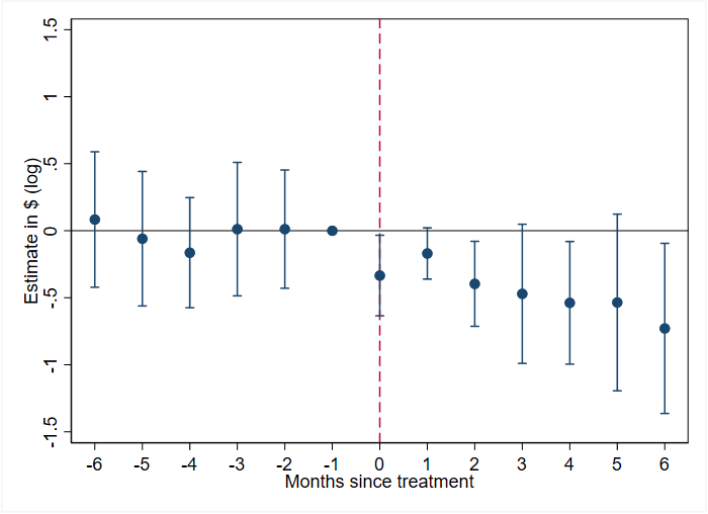


Notes: This figure shows the geographic distribution of cards belonging to the Suspicious cluster as a share of all the cards at a Metropolitan Statistical Area (MSA) level. The location of each card is deduced from their ZIP codes. We find commonalities in the MSAs with higher concentration of fraud with areas that reported higher per capita identity thefts as per the Federal Trade Commission. Table A6 reports a formal estimate of the correlation between MSA-level suspicious cards in our sample with the official data on identity theft reports from the Federal Trade Commission.

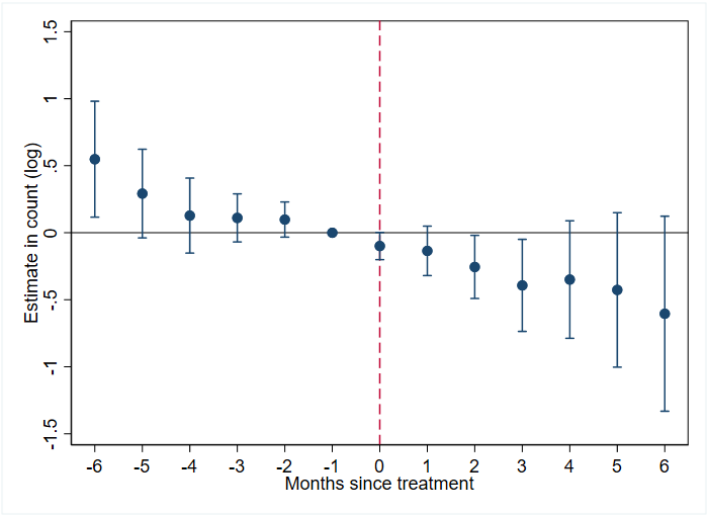
Figure A7: Event-study for Suspicious (Fast Cash) and Suspicious (Slow Cash) Clusters



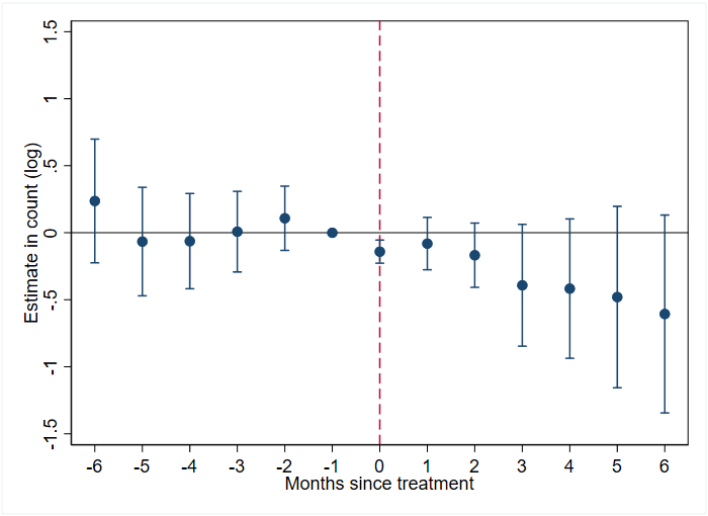
(a) Suspicious (Fast Cash): UI dollars



(b) Suspicious (Slow Cash): UI dollars



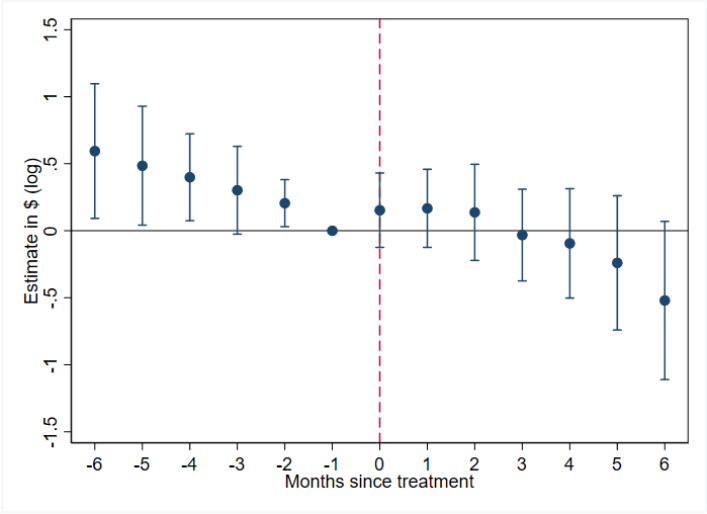
(c) Suspicious (Fast Cash): number of cards



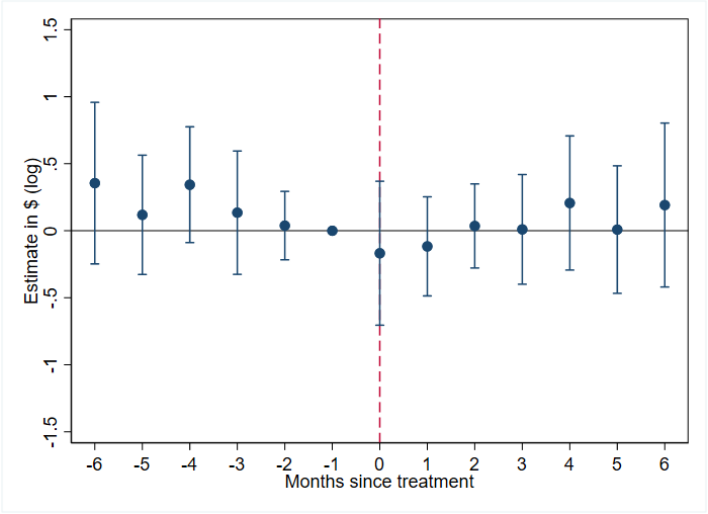
(d) Suspicious (Slow Cash): number of cards

Notes: This figure reports the results of Equation 2 for Suspicious (Fast Cash) and Suspicious (Slow Cash) clusters separately. The dependent variable in plots (a) and (b) is log(UI), the log of unemployment benefits paid by each state to a cluster in a month. The dependent variable in plots (c) and (d) is the log of number of cards in each cluster to whom these benefits were paid by each state in a month.

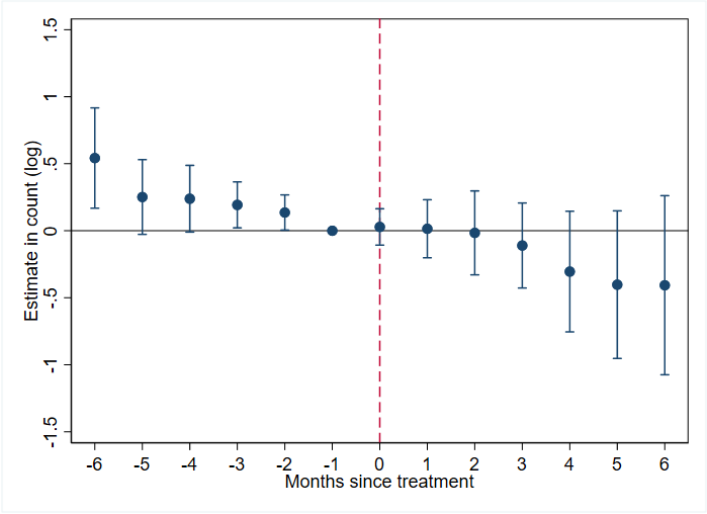
Figure A8: Event-study for Concurrent Income and Discretionary Spending Clusters



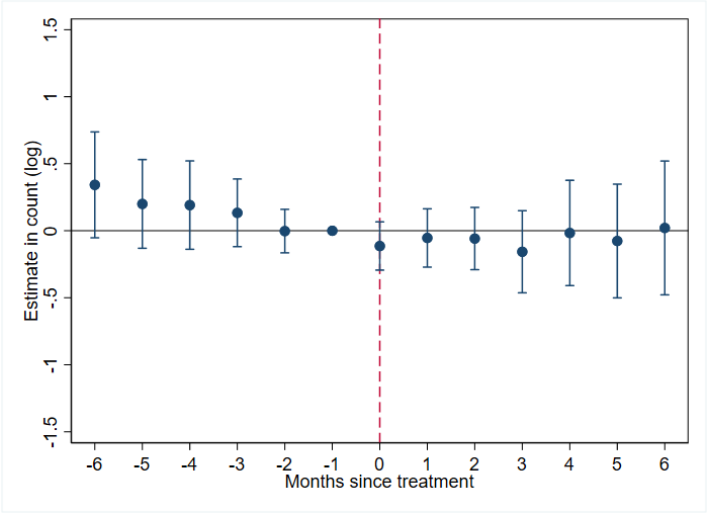
(a) Concurrent Income: UI dollars



(b) Discretionary Spending: UI dollars



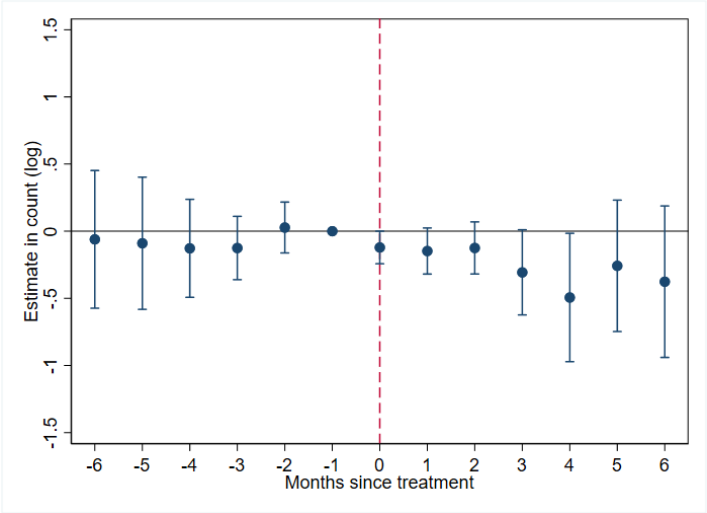
(c) Concurrent Income: number of cards



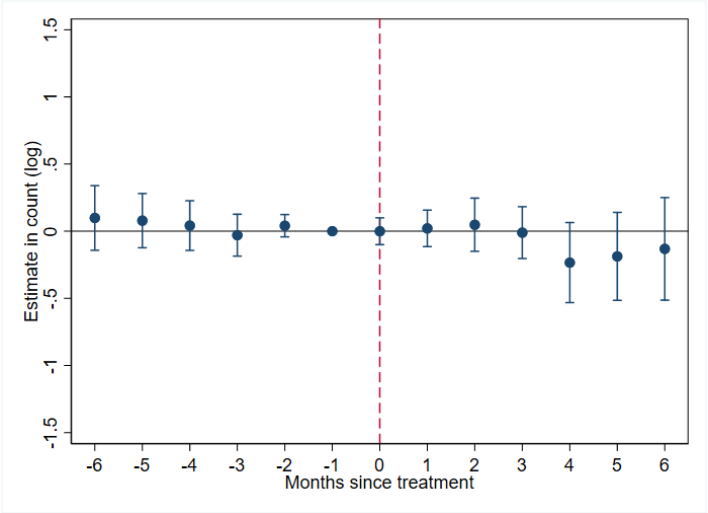
(d) Discretionary Spending: number of cards

Notes: This figure reports the results of Equation 2 for Concurrent Income and Discretionary Spending clusters. The dependent variable in plots (a) and (b) is log(UI), the log of unemployment benefits paid by each state to a cluster in a month. The dependent variable in plots (c) and (d) is the log of number of cards in each cluster to whom these benefits were paid by each state in a month.

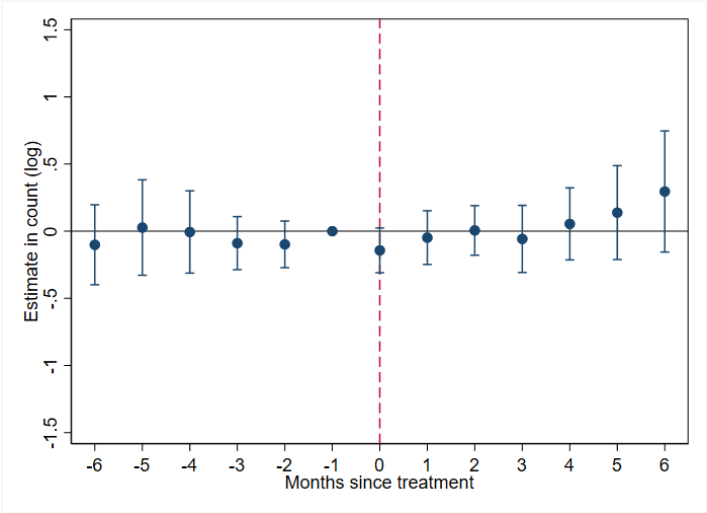
Figure A9: Within-cluster Impact of Identity Verification on Number of Cards



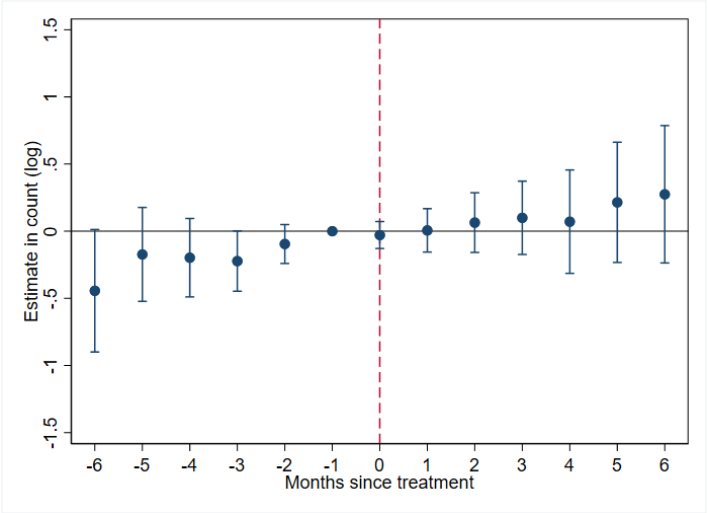
(a) Suspicious



(b) Concurrent Income



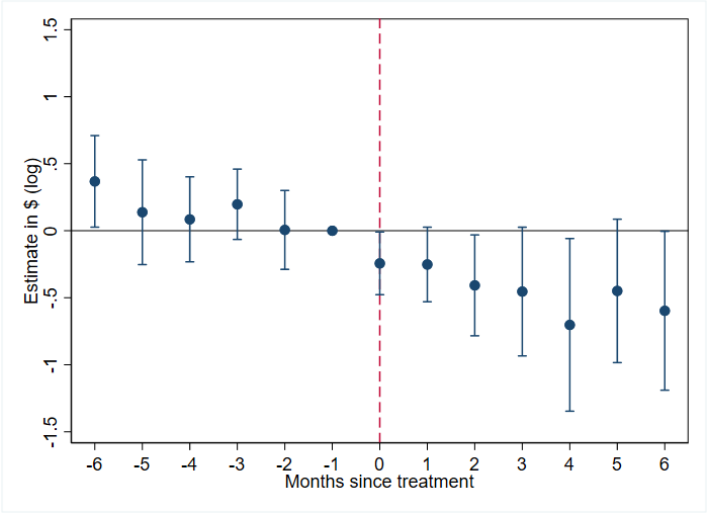
(c) Discretionary Spending



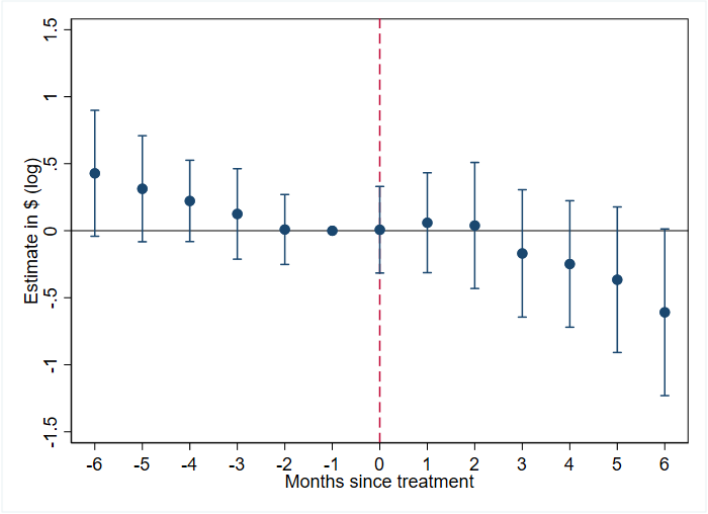
(d) Control

Notes: This figure plots the event-study version of Equation 3, the difference-in-difference specification, for all four clusters individually. The dependent variable in all plots is the log of number of cards in each cluster to whom these benefits were paid by each state in a month. While our clusters reflecting concurrent income, discretionary spending, or control (non-suspicious) spending are unaffected, the suspicious cluster shows an immediate and persistent drop in the number of cards to whom UI benefits were disbursed following identity verification.

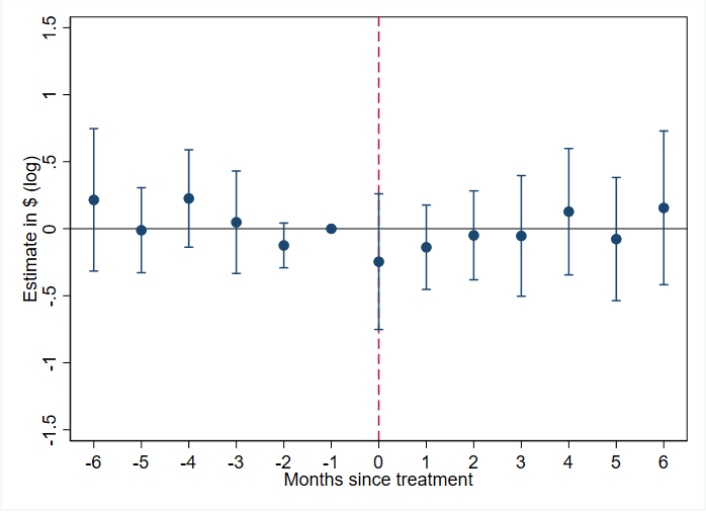
Figure A10: Event-study using Stacked Difference-in-differences



(a) Suspicious



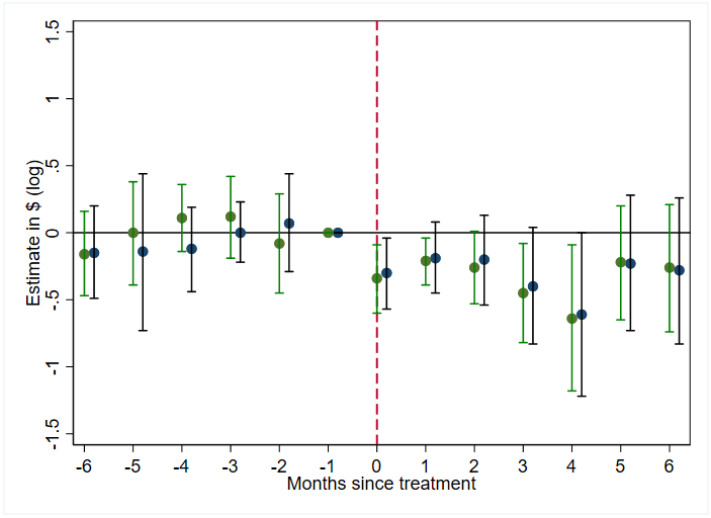
(b) Concurrent Income



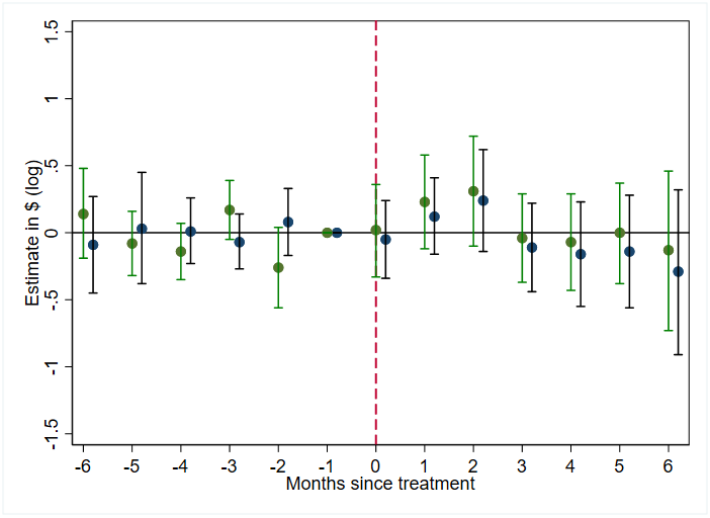
(c) Discretionary Spending

Notes: This figure plots the event-study of the form in Equation 2 estimated using the stacked difference-in-differences method as in Cengiz et al. (2019). The dependent variable in each plot is log(UI), the log of unemployment benefits paid by each state to a cluster in a month.

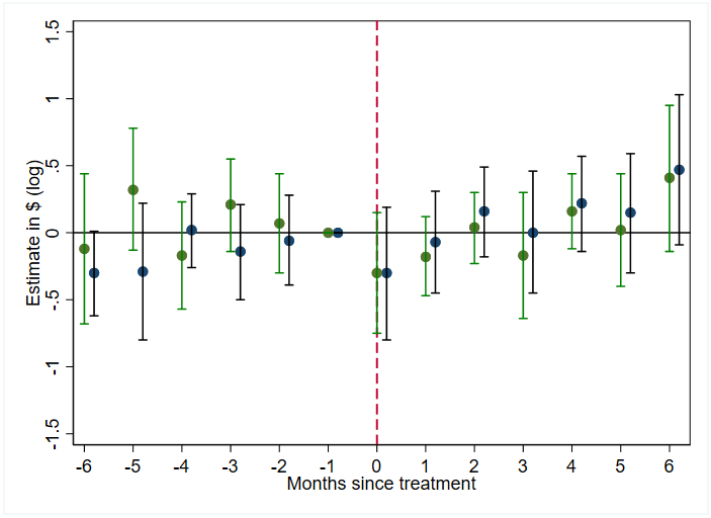
Figure A11: Within-cluster Impact of Identity Verification - Alternative Estimation



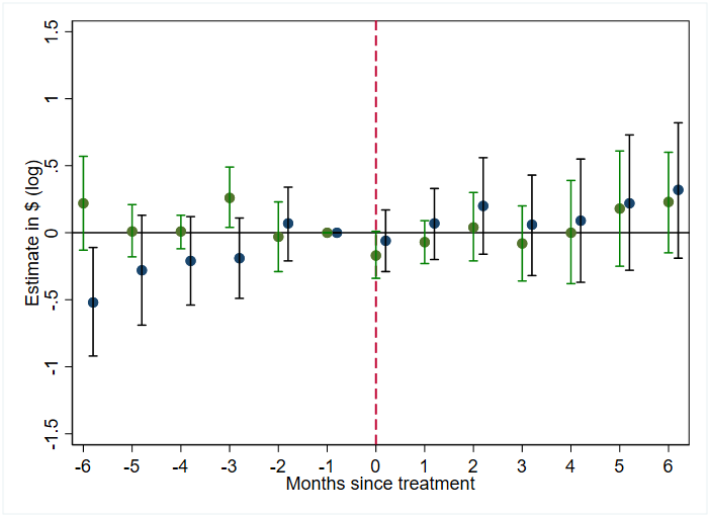
(a) Suspicious



(b) Concurrent Income



(c) Discretionary Spending



(d) Control

Notes: This figure plots the estimation of event studies of the form in Equation 3 based on the estimation technique in Callaway and Sant'Anna (2021) (in green) or Cengiz et al. (2019) (in blue). The dependent variable in each plot is log(UI), the log of unemployment benefits paid by each state to a cluster in a month.

Table A1: Timing of Identity Verification Adoption by State

Treated state	Treatment	Agency/vendor	Data source
Arizona (AZ)	Oct-20	ID.me	FOIA
Colorado (CO)	Jan-21	ID.me	FOIA
Delaware (DE)	Jun-21	ID.me	FOIA
Georgia (GA)	Sep-20	ID.me	Congress/Public
Idaho (ID)	Dec-20	ID.me	FOIA
Indiana (IN)	Sep-20	ID.me	FOIA/Congress
Kansas (KS)	Feb-21	LexisNexis	FOIA
Louisiana (LA)	May-21	ID.me	Congress/Public
Massachusetts (MA)	Mar-21	ID.me	FOIA
Mississippi (MS)	Mar-21	ID.me	Congress/Public
Missouri (MO)	Mar-21	ID.me	FOIA
Montana (MT)	Nov-20	ID.me	Congress
Nebraska (NE)	Dec-20	GIACT	FOIA
Nevada (NV)	Mar-21	ID.me	FOIA
New York (NY)	Feb-21	ID.me	Congress
North Dakota (ND)	Dec-20	ID.me	Congress
Ohio (OH)	Apr-21	LexisNexis	FOIA/Public
Oregon (OR)	Mar-21	ID.me	Congress/Public
Pennsylvania (PA)	Oct-20	ID.me	FOIA
South Carolina (SC)	Mar-21	ID.me	Congress/Public
Texas (TX)	Nov-20	ID.me	FOIA
Wisconsin (WI)	Nov-20	Google Analytics	FOIA/Public
Untreated states			
Alabama (AL)	Hawaii (HI)	New Hampshire (NH)	Virginia (VA)
Alaska (AK)	Illinois (IL)	New Mexico (NM)	Washington (WA)
Arkansas (AR)	Kentucky (KY)	Rhode Island (RI)	West Virginia (WV)
Connecticut (CT)	Michigan (MI)	South Dakota (SD)	Wyoming (WY)
District of Columbia (DC)	Minnesota (MN)	Utah (UT)	

Notes: This table lists states for which we can identify both UI disbursements in the Facteus data and the timing of ID verification adoption based on FOIA, publicly-available information, and congressional records. For states that implemented multiple programs, we use the earliest date of implementation.

Table A2: Disaggregation of UI Data by State

	Correlation	Share in %		Correlation	Share in %
Federal (41 states)	0.95	0.40	Mississippi (MS)	0.75	0.02
Alaska (AK)	0.83	0.37	Montana (MT)	0.89	1.17
Alabama (AL)	0.78	1.15	North Dakota (ND)	0.86	0.45
Arkansas (AR)	0.83	1.04	Nebraska (NE)	0.82	0.60
Arizona (AZ)	0.74	0.08	New Hampshire (NH)	0.93	0.26
Colorado (CO)	0.91	0.30	New Mexico (NM)	0.83	0.14
Connecticut (CT)	0.89	0.45	Nevada (NV)	0.72	0.20
District of Columbia (DC)	0.55	0.18	New York (NY)	0.87	0.22
Delaware (DE)	0.91	0.33	Ohio (OH)	0.83	1.10
Georgia (GA)	0.73	0.05	Oregon (OR)	0.56	0.14
Hawaii (HI)	0.80	0.22	Pennsylvania (PA)	0.91	0.16
Idaho (ID)	0.81	0.57	Rhode Island (RI)	0.86	0.59
Illinois (IL)	0.84	0.44	South Carolina (SC)	0.88	1.21
Indiana (IN)	0.88	1.34	South Dakota (SD)	0.92	0.47
Kansas (KS)	0.87	0.85	Texas (TX)	0.92	0.31
Kentucky (KY)	0.87	0.92	Utah (UT)	0.89	0.36
Louisiana (LA)	0.89	0.84	Virginia (VA)	0.95	0.40
Massachusetts (MA)	0.93	0.23	Washington (WA)	0.89	0.27
Michigan (MI)	0.88	0.71	Wisconsin (WI)	0.80	0.84
Minnesota (MN)	0.86	0.66	West Virginia (WV)	0.93	0.51
Missouri (MO)	0.91	0.60	Wyoming (WY)	0.93	0.37

Notes: This table characterizes our UI data by state, showing the correlation of our time series data with that state's official UI data, as well as the estimated share of UI data we can observe from that state.

Table A3: UI $ and % Share to Clusters Before and After the Onset of COVID-19 Pandemic

	UI received ($, million)		Share of UI received (%)	
Category	Jan 2019-Feb 2020	Mar 2020-Sep 2021	Jan 2019-Feb 2020	Mar 2020-Sep 2021
Suspicious (Fast Cash)	0.01	50.62	0.08	6.13
Suspicious (Slow Cash)	0.01	33.39	0.08	4.04
Concurrent Income	0.20	22.26	1.19	2.69
Discretionary Spending	0.11	19.38	0.64	2.35
Control	16.2	700.55	98.00	84.79

Notes: This table describes UI inflows by cluster, including the total dollars and share of total dollars paid to each cluster before and after the onset of COVID-19 pandemic.

Table A4: UI $ and % Share to Clusters Under PUA and Non-PUA Programs

	PUA		Non-PUA	
Category	$ million	% of total	$ million	% of total
Suspicious (Fast Cash)	11.78	13.4%	1.68	4.1%
Suspicious (Slow Cash)	6.27	7.2%	0.58	1.4%
Concurrent Income	1.02	1.2%	1.02	2.5%
Discretionary Spending	1.77	2.0%	0.57	1.4%
Control	66.79	76.2%	37.54	90.7%

Notes: This table reports the cluster-level dollars and share of total dollars paid under the PUA and non-PUA programs in our data for four states where we can separate the two flows: Arkansas, Massachusetts, Ohio, and West Virginia. The sample period for this comparison runs from April 2020 through May 2021, when both programs were fully active.

Table A5: Treatment Effects of Identity Verification - Stacked Difference-in-differences

	log(UI)	log(number of cards)
	(1)	(2)
Suspicious × Post	-0.527**	-0.460*
	(0.244)	(0.238)
Concurrent Income × Post	-0.323	-0.350
	(0.226)	(0.236)
Discretionary Spending × Post	-0.020	-0.119
	(0.146)	(0.187)
Observations	9,739	9,727
Adj. R^2	0.92	0.94
Cluster, State, Month, Post, Cluster-State, Cluster-Month FE	Y	Y

Notes: This table reports the results from a triple difference-in-differences estimation of the form in Equation 1 using the estimation method in Cengiz et al. (2019), where stacked datasets are created for each treatment date. In column (1), the dependent variable is log(UI), the log of unemployment benefits paid by each state to a cluster in a month, and in column (2) it is the log of number of cards in each cluster to whom these benefits were paid. The control group is the cluster with non-suspicious cards. Standard errors clustered by state are reported in parentheses. Figure A10 shows this table's event-study version. $^{*}p < 0.1;^{**}p < 0.05;^{***}p < 0.01$.

Table A6: Geographical Validation of Suspicious Cards with Identity Theft Reports

	Share of Suspicious Cards	
	(1)	(2)
Identity Theft Reports (per capita)	2.569**	2.493**
	(1.166)	(1.195)
Population (2020, tens of million)		0.002
		(0.007)
Observations	94	94
Adj. R^2	0.05	0.04
Unit of observation	MSA	MSA

Notes: This table reports results from an ordinary least squares estimation for a model of the form in Equation 4. The dependent variable is the share of suspicious cards out of all cards in our data at the level of Metropolitan Statistical Area (MSA). We map the ZIP code associated with cards with the MSA they belong to. The regressor of interest is the number of identity theft reports in 2020 as per Federal Trade Commission data, scaled by the MSA population. Column (2) additionally controls for the 2020 population of the MSA. To reduce noise, we include MSAs with at least 100 cards in our sample. Standard errors clustered by state are reported in parentheses. Figure A6 shows the ratio in at MSA level in the form of a map. $^{*}p < 0.1;^{**}p < 0.05;^{***}p < 0.01$.