

NBER WORKING PAPER SERIES

UNEMPLOYMENT INSURANCE FRAUD IN THE DEBIT CARD MARKET

Umang Khetan
Jetson Leder-Luis
Jialan Wang
Yunrong Zhou

Working Paper 32527
<http://www.nber.org/papers/w32527>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2024

We thank Cameron Ellis, Tal Gross, Mark Kutzbach, Riley League, Ryan McDevitt, Felipe Netto, Maggie Shi, David Sovich, and participants at the American Economic Association Annual Meeting, NBER Innovative Data in Household Finance Conference, Midwest Finance Association, Financial Management Association, Georgetown University, University of Iowa, and Purdue University for helpful comments. We are extremely grateful to Jonathan Chin and Facticeus 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Umang Khetan, Jetson Leder-Luis, Jialan Wang, and Yunrong Zhou. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Unemployment Insurance Fraud in the Debit Card Market
Umang Khetan, Jetson Leder-Luis, Jialan Wang, and Yunrong Zhou
NBER Working Paper No. 32527
May 2024
JEL No. G51,H53,J65,K42

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 cluster cards corresponding to varying levels of suspicious or potentially fraudulent activity. We then conduct a 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 27%, while non-suspicious cards were largely unaffected by these technologies. Our results indicate that identity screening may be an effective mechanism for mitigating fraud in the UI system and for benefits programs more broadly.

Umang Khetan
University of Iowa
S217, Tippie College of Business
E Market St
Iowa City, IA 52242
umang-khetan@uiowa.edu

Jetson Leder-Luis
Questrom School of Business
Boston University
595 Commonwealth Avenue
Boston, MA 02215
and NBER
jetson@bu.edu

Jialan Wang
Department of Finance
University of Illinois Urbana-Champaign
340 Wohlers Hall
1206 S. Sixth Street MC-706
Urbana, IL 61820
and NBER
jjalanw@gmail.com

Yunrong Zhou
Purdue University
403 W State St
West Lafayette, IN 47906
zhou1181@purdue.edu

1. INTRODUCTION

Unemployment insurance (UI) is one of the largest social safety net 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 led to a sharp economic contraction and large expansion in unemployment insurance, with state and federal governments disbursing over \$800 billion between 2020-2021. However, the rapid expansion of UI during this period, along with changes in application practices and eligibility criteria, made it susceptible to fraud. In particular, we focus in this paper on the diversion of benefits intended for eligible beneficiaries by organized criminals via identity theft.¹

Poor targeting of UI benefits due to fraud is costly to the government budget and undermines the efficacy of the UI program, as those who commit fraud are unlikely to have the same labor market and spending responses as eligible recipients. Moreover, fraud can harm the program's long-term financial stability, limiting its ability to provide support to future beneficiaries. Concern about fraud in public benefits is widespread, including such programs as the Paycheck Protection Program, cash welfare, and in-kind food transfer programs such as SNAP and WIC ([Griffin et al., 2023a](#), [U.S. Department of Justice, 2013, 2024a,b](#)). Public programs are vulnerable to fraud due to problems such as outdated IT infrastructure, insufficient cross-state coordination, the absence of a shared national database, and difficulties in verifying applications ([U.S. Department of Labor, 2023](#), [U.S. Department of Labor, Office of Inspector General, 2023](#), [Dube, 2021b](#), [Bipartisan Policy Center, 2023](#)).² These weaknesses become more pronounced during program expansions such as those during the COVID-19 pandemic, creating opportunities for fraudsters to exploit program vulnerabilities.

An important and widely-used tool to combat public benefits fraud is identity verification, a form of screening that is meant to reduce the diversion of funds paid to stolen or fake identities instead of genuine applicants. Like other policies designed to improve targeting in benefits programs, identity verification faces a trade-off between hassle costs and more precise targeting toward intended recipients ([Nichols and Zeckhauser, 1982](#)). On one hand, it imposes a time and hassle cost on applicants, creates an operational burden and financial cost on the system, and could potentially

¹[U.S. Government Accountability Office \(2023b\)](#) summarizes descriptive evidence and enforcement cases of UI fraud compiled by various government bodies, including widespread reports of identity theft, organized fraud schemes, and international criminal activity. Other media and policy reports of UI fraud include [Axios \(2021\)](#), [ProPublica \(2021\)](#), [CNBC \(2022\)](#), [Pennsylvania Office of Attorney General \(2024\)](#), and [U.S. Department of Labor, Office of Inspector General \(2021\)](#).

²As of December 2021, 32 of the 53 states and territories were still using legacy IT systems that were developed in the 1970s and 1980s to support their UI benefits system, tax system, or both. Outdated IT systems have affected states' ability to meet the needs of unemployed workers, both historically and during economic downturns ([Government Accountability Office, 2022](#)).

delay payments or even prevent eligible applicants from receiving benefits. On the other hand, a more permissive system can lead to waste and fraud by allowing individuals or criminal entities to receive payments they are not entitled to under false or stolen identities. Identity verification is part of a growing set of technological solutions to fraud and administrative problems in public benefits programs.

In this paper, we study fraud in the UI system and examine the effectiveness of identity verification as an anti-fraud screening measure to reduce the misallocation of unemployment benefits. In response to growing concerns about fraud, 30 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 UI 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 27% reduction in UI disbursements to the group of recipients we categorize as suspicious, with minimal negative impact on most other UI recipients. Our analysis suggests that adoption of identity verification screening can improve resource allocation in large public benefit programs.

Measurement is a natural challenge when studying fraud, as those who commit fraud try to conceal it, and most fraud goes undetected by law enforcement. We identify indicators of potentially fraudulent activity by analyzing the income and spending patterns in a dataset 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 a rich set of dimensions along which 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 UI benefits are paid to either bank-issued or prepaid debit cards.

We deploy k-means clustering, an unsupervised machine learning algorithm, on the transactions in our database to categorize 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 cash is withdrawn from ATMs, and concurrent payroll income that might indicate ineligibility for UI. In particular, the quantity and speed of cash withdrawal have been documented in the fraud-detection literature as strong signals of suspicious activity (Wu, Xu, and Li, 2019), while the amount of UI received by each card is a relevant feature in the context of UI fraud. Our clustering algorithm produces five groups or clusters of cards, with two representing suspicious activity most consistent with identity theft.

We validate our suspiciousness measures using comparisons to external administrative data. First, in the time-series, UI disbursed to suspicious cards sharply increased after the onset of the COVID-19 pandemic. In the cross-section, fraud was particularly pronounced under the pandemic unemployment program, consistent with trends reported by the [U.S. Government Accountability Office \(2023b\)](#). Second, we find that UI disbursed to suspicious cards shows a lower correlation with underlying economic trends compared to non-suspicious cards, which suggests that suspicious cards are not claimed by valid unemployed recipients. Third, suspicious cards are concentrated in regions of the country that report a higher incidence of identity theft in administrative data reported by the Federal Trade Commission.

We estimate the effectiveness of anti-fraud measures by measuring the response of UI disbursements to suspicious and non-suspicious card clusters after the state-level adoption of identity verification technology. We create a panel of state-level identity verification measures using data from Freedom of Information Act requests we made to the 50 state and District of Columbia unemployment agencies, corroborated with Congressional records and publicly available data. We then estimate the effects of identity verification policies using a staggered difference-in-differences design that controls for time-invariant characteristics of each state as well as national time trends.

Our findings suggest that identity verification measures primarily affect suspicious cards, with minimal effect on non-suspicious recipients. Specifically, we find that the policy interventions reduced UI payouts to cards in the suspicious clusters by 27%. The number of unique suspicious cards that received UI after the implementation of these policies dropped by 20%. This reduction is both immediate and persistent and is not driven by pre-trends, which supports a causal interpretation of our results. Moreover, we do not find a significant reduction in UI paid to non-suspicious cards after the policy was introduced, indicating that the frictions introduced by identity verification had a limited negative impact on non-fraudulent applicants.

Our results are robust to modern critiques of two-way fixed effects difference-in-differences design, as well as to alternative clustering techniques. We confirm that our baseline identification strategy is not contaminated by issues of negative weights, a concern highlighted in recent literature in the context of multiple treatments with potentially heterogeneous treatment effects. Furthermore, we find consistent results using other estimators including the semi-parametric group-time average treatment effect proposed in [Callaway and Sant’Anna \(2021\)](#), and the “stacked” difference-in-difference estimator from [Cengiz, Dube, Lindner, and Zipperer \(2019\)](#).

We also show the robustness to alternative ways of implementing the k-means clustering algorithm that identifies potentially fraudulent cards. Most importantly, we address a potential bias that may arise if identity verification alters the spending behavior of cards in our sample. Specifically, we re-deploy the unsupervised machine learning algorithm using only the sample period prior

to the treatment date, and find that our difference-in-differences estimation continues to show a sharp decline in UI disbursement to suspicious cards but not to control cards. In addition, we confirm that our results are robust to variations on the number of clusters, alternative features used to form clusters, and alternative sample restrictions that form inputs for our clustering algorithm.

Our k-means clustering methodology also allows us to explore the extent of fraud more broadly. Based on our estimate that 11.3% of UI disbursements in our sample were made to suspicious cards, and applying this estimate to the \$289.7 billion disbursed nationwide under the two primary pandemic UI programs, we extrapolate that a total of \$32.7 billion was disbursed to suspicious recipients between March 2020 and September 2021. Moreover, our treatment effects imply that identity verification technology would have resulted in nationwide savings of \$9 billion over the course of our sample period if all the states had identity verification measures in place before the start of the pandemic. This total 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 \$7.2 billion if these technologies had been implemented by all states prior to March 2020. These calculations are subject to assumptions about the consistency of treatment effects over the course of the pandemic and across states, and about the representativeness of our data for the debit card market and UI more generally.

Our work has implications for anti-fraud policies in domains outside of UI. Screening and loophole-closing regulations reduce the need to claw back ill-gotten gains after the fact, and can be more effective than enforcement measures in environments where fraud is diffuse and therefore difficult to prosecute (Mookherjee and Png, 1992, Polinsky and Shavell, 2000). Moreover, when those committing fraud have limited liability, it is more effective to prevent the disbursement of funds up front because ex-post prosecution would be unlikely to recover funds even if fraudsters could be identified (Eliason et al., 2021). While the Department of Justice (DOJ) is involved in dozens of lawsuits nationwide to try to recover stolen UI and PPP funds, our results suggest that identity verification technology could be a substitute for ex-post enforcement effort.³ Moreover, from a political economy perspective, our work informs the understanding of whether bureaucrats efficiently implement anti-fraud policy. We show that identity verification is a low-cost, high benefit program, yet it took high levels of fraud during an unprecedented pandemic 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 the labor market and household consumption, these effects must be balanced with the caveat that

³Press releases from the DOJ on enforcement actions are available in [U.S. Department of Justice \(2024c\)](#) and [U.S. Department of Justice \(2024d\)](#).

a substantial share of UI spending may be lost to fraud, and that the labor market and spending behavior of fraudulent recipients is likely to be substantially different from those of legitimate recipients. 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.

1.1. Related Literature

Our work connects the literature 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 \(2021a\)](#) 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, 2023](#), [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 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) study Suspicious Activity Reports (SARs) filed by banks, and 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. 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 predictive accuracy for default rates but has implications on racial disparity in the distribution of these gains. Berg, Fuster, and Puri (2022) note the increasing use of ML to improve customer screening in FinTech lending.

This paper proceeds as follows. Section 2 discusses the background and institutional details. Section 3 discusses the data, and Section 4 details 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. For example, in 2019 the Bureau of Labor Statistics reported a 9% average overpayment rate of UI, with some states as high as high as 32%.⁴ 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 is attributable to the introduction 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, visible in both public data and our debit card sample. 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.⁵

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 ([U.S. Government Accountability Office, 2023b](#)). [Figure A1](#) shows that government benefits programs such as unemployment insurance experienced the largest increase in reports of identity theft, eclipsing both credit card fraud and loan fraud that were dominant in the years preceding the pandemic. Notably, Government Documents or Benefits Fraud increased by 2,920% from 2019 to 2020 ([Federal Trade Commission, 2020](#)).

In order to detect and deter fraudulent activity, many states recently implemented anti-fraud measures such as identity verification. Generally contracted through third party vendors such as ID.me and LexisNexis, identity verification measures seek to reduce UI fraud by ensuring that

⁴The Department of Labor (DOL) UI payment accuracy datasets are available from [U.S. DOL \(2023\)](#).

⁵See pages 5-7 of the US Government Accountability Office assessment report from [Government Accountability Office \(2023\)](#).

the applicant is not using a stolen identity to claim these benefits. These measures do not verify eligibility (e.g., based on employment qualifications), but rather ensure that the person claiming the benefits is not using a stolen identity. Therefore, ID verification attempts to screen out fraudulent applicants before such claims are processed, rather than ex-post mechanisms to punish 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 A2](#) 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.⁶

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 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.⁷ [Table A1](#) provides a timeline of ID verification measures along with the vendor contracted by each state. [Figure A3](#) 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.

⁶Estimates from the US Treasury department are available from [U.S. Department of the Treasury \(2023\)](#).

⁷The Congressional report is part of a set of documents published with the [House Committee on Oversight and Reform \(2022\)](#) 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 from [Dave \(2021\)](#).

3. DEBIT CARD TRANSACTIONS DATA

We identify potentially fraudulent UI applicants by leveraging a granular dataset of 35 million transactions from about 84,000 unique debit card holders 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 Facticeus, a FinTech firm that partners with banks, card issuers, and payment processors to aggregate, standardize, and anonymize transaction-level information from debit cards.⁸ 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.⁹ Facticeus perturbs the raw data it obtains from its data partners 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 Facticeus.

Debit card holders are a particularly interesting population to study 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.¹⁰ 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 Facticeus 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

⁸For other papers that use data from Facticeus, see [Brave, Fogarty, Aaronson, Karger, and Krane \(2021\)](#), [Karger and Rajan \(2020\)](#) and [Zhou and Correia \(2022\)](#).

⁹Using card-holder ZIP codes, [Zhou and Correia \(2022\)](#) show that the geographic distribution of cards in Facticeus sample closely resembles the population density across the United States.

¹⁰The Consumer Financial Protection Bureau (CFPB) explains the methods by which UI funds can be received in a blog post ([Malaiyandi, 2020](#)). A specific example of a state-issued prepaid card from the state of Michigan can be found from [Michigan Department of Labor and Economic Opportunity \(2024\)](#).

of UI payments into debit cards.¹¹ 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 [Washington Bankers Association \(2021\)](#). We are able to identify UI inflows for 41 states (including the District of Columbia); [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 these four states, we aggregate state-level UI flows into the PUA and non-PUA programs separately. We use a similar technique to identify other streams of income, such as wages through payroll. [Appendix A](#) details the procedure for constructing these income streams.

Based on the technique above, we observe state-level UI transactions on about 105,000 debit cards beginning in January 2019 and ending in September 2021. Out of these, we focus on 84,310 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 spending transaction is tagged with a Merchant Category Code (MCC), which we use to categorize outflows into types such as spending via cash withdrawal from ATMs or bank branches, spending on groceries or food purchases, wire transfers, and spending on discretionary items such as alcohol or gambling.¹² 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 Factiveus to administrative data from each state’s Department of Labor. [Figure 1](#) plots the time series of UI disbursements observed in our 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 0.94. [Table A2](#) confirms the strong correlation for each of the 41 states individually.

[Table 1](#) shows card-level and card-month level summary statistics for our sample of 84,310 UI recipient cards. On average, cards received a total of about \$10,100 of UI benefits across 19 disbursements, at a rate of about 4 per month. UI represents 63% of the total income for a typical

¹¹For external comparison, [Brave, Fogarty, Aaronson, Karger, and Krane \(2021\)](#) estimate that Factiveus 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 ([Michigan Department of Labor and Economic Opportunity, 2021](#)).

¹²[Florida Department of Financial Services \(2021\)](#) provides the list of MCC codes.

card, and groceries represent the single largest category of spending. Panel B of [Table 1](#) shows card-month level features for the months where cards receive UI. The average income per card is \$2,790, of which \$2,125 comes from UI. Total spending has a mean of \$1,808 per month, with cash spending at \$451 (25%) and groceries at \$340 (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 find that it represents 16% of the total. Finally, we define “earnings” as income from either UI or payrolls, and note that in the months where cards receive any amount of UI, it forms 96% of their total earnings which suggests that a vast majority of cards do not receive payroll simultaneously with UI.

4. MEASURING FRAUD WITH MACHINE LEARNING

We use unsupervised machine learning to categorize debit cards into suspicious or non-suspicious categories using their income and spending behaviors. This process allows us to construct the dependent variables of UI reciprocity by different clusters of cards and evaluate the impact that identity verification may have had on each of them. Specifically, we deploy a popular 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 require labeled training data to group observations. Unsupervised techniques have the advantage that they do not rely on successfully identified fraud, which may be non-representative ([Shekhar et al., 2023](#)). Instead, this method relies only on the revealed behavior of cards and detects outliers based on the set of attributes relevant to fraud detection. Below we first describe the construction of features that form inputs to the clustering algorithm. Then, we detail the clustering procedure along with its output, and discuss the attributes that validate our interpretation of these clusters in the context of unemployment insurance fraud. [Appendix B](#) provides further details.

4.1. Feature Construction

Clustering of cards into distinct groups using unsupervised machine learning requires the analyst to construct feature vectors, i.e. variables, as inputs along which the cards will be separated. Recent literature 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. More specific to consumer fraud, [Wu, Xu,](#)

and Li (2019) note that dynamic patterns in the amount and speed of cash spending can improve the accuracy of fraudulent cash-out detection.

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 consider characteristics relevant to the context of UI fraud such as the amount of UI, length of continuous recipiency, and concurrent payroll 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. Below we detail how each of these features is constructed.

1. Unemployment insurance received (in dollars): large receipts of UI could align with the incentive of fraudsters to maximize gains from stolen information. In our main specification, we aggregate the UI received by a card throughout our sample period. As a robustness check, we use the average UI per month together with the longest unbroken spell of being on UI rolls (in months) to segregate large UI recipients from others.
2. 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 and branch withdrawals using MCC codes 6011 and 6010, and express cash withdrawal as a fraction of monthly spend.¹³
3. 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.
4. Discretionary spending as a fraction of monthly outflows: this variable includes spending in liquor stores, gambling or 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.
5. Concurrent non-UI earnings: concurrent earnings from wage or salary income could indicate potential ineligibility for UI, but should not necessarily be affected by identity verification, which does not screen on technical eligibility. We measure this variable as the monthly income from UI scaled by the total earnings from UI and payroll income.

¹³We also explore wire transfers outside of the U.S. as another potential indicator of theft. We locate international transfers by string-matching merchant names such as Western Union or Remitly. However, very few cards in our data simultaneously receive large sums of UI and transfer money abroad. Hence, we do not use this as a feature in our clustering algorithm.

4.2. *K-means Clustering*

Using these features, we deploy k-means clustering to partition cards into several clusters that display distinct characteristics. 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 cards.

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 2 plots a flow chart of the sequence in which clusters are created. Our main clustering algorithm focuses on 84,310 debit cards which receive at least \$1,000 in UI. We start with segregating cards that receive unusually large amount of UI throughout the sample period, compared to other recipients with more moderate UI inflows. This step creates two clusters: “High UI” with about 23% 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 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 (that do not display abnormal spending behaviors), we apply the clustering algorithm to detect potentially ineligible recipients using the concurrent income feature.

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 half of “High UI” cards are ultimately tagged as control because they do not display suspicious spending or other income patterns, and therefore there is no mechanical relation between suspicious cluster and receipt of large sums of UI.

Our machine learning algorithm produces five clusters that group cards along economically comparable behaviors. [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”: abnormally large UI and concurrent income from payroll.
4. “Discretionary Spending”: abnormally large UI, 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. [Table 2](#) shows that the “Suspicious (Fast Cash)” and “Suspicious (Slow Cash)” clusters receive on average twice as much UI as other clusters, are on UI rolls for more than twice as long as “Control” cards, and allocate 61% and 55% (respectively) of their total spending to cash withdrawals. 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 and also lowers the probability of any future clawback measures. For our main specification, we combine both suspicious clusters together into one group, “Suspicious” cards, but show robustness to separating them. All other clusters, including “Control”, do not reflect suspicious activity associated with identity theft.

4.3. Cluster Validation

Our algorithm creates clusters that represent distinct behaviors 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 primarily by way of cash withdrawal,

and withdraw cash relatively quickly. For instance, one card in these clusters received UI in excess of \$20,000 in a single day, and withdrew over 80% by cash. Another card received UI in excess of \$67,000 over a year, and withdrew \$57,000 in cash from various ATMs and bank branches. A one-month transaction snapshot of this card in [Figure A5](#) illustrates the pattern of UI receipts followed by same-day cash withdrawals. 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 The 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, and that the fraudsters were organized enough to quickly withdraw money from the system.

Second, the geographical distribution of suspicious cards in our data correlates with the rate of identity theft reported by the [Federal Trade Commission \(2021\)](#). Using the card-level zip codes in our data, we map each card to a metropolitan or micropolitan 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), Portland-Vancouver-Hillsboro (OR-WA), Seattle-Tacoma-Bellevue (WA) and Chicago-Naperville (IL). Many of these regions have above-average rates of identity theft in publicly available administrative data. Formally, we regress the share of suspicious cards in our sample on the MSA-level reports of identity theft as per the Federal Trade Commission, and find a positive and statistically significant relation in [Table A3](#). [Appendix B](#) provides specification details.

Third, the pattern of UI disbursements to cards in the Suspicious cluster diverges from underlying economic trends, unlike that of the Control cluster. [Figure A7](#) plots the correlation coefficient between unemployment levels indicated by monthly nonfarm payrolls and the number of cards within the Control and Suspicious clusters, respectively, to whom UI was disbursed.¹⁵ We note that the Control cluster has a higher correlation with nonfarm payroll than the Suspicious cluster in all but 3 out of the 41 states in our sample, and also at a federal level. This suggests that UI reciprocity by suspicious cards might be decoupled from genuine unemployment trends. In 9 states, the correlation of the number of UI recipients with nonfarm payroll for the Suspicious cluster is either negative or under 0.5, in stark contrast with the Control cluster, where the correlation is positive throughout, and exceeds 0.5 in all but 3 states.

Finally, we find that the share of UI to Suspicious cards follows time-series and cross-sectional

¹⁴[Cox et al. \(2020\)](#) document a decline in cash spending after the onset of the COVID-19 pandemic.

¹⁵Nonfarm 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.

patterns that are consistent with external reports of fraud in the unemployment insurance system. In line with media reports, [Table A4](#) shows that the amount and share of UI disbursed to suspicious clusters increased during the pandemic. Likewise, [Table A5](#) shows that for the four states where we can identify PUA and non-PUA inflows separately, the share of UI that went to suspicious clusters was higher under the PUA program at 20%, compared to 13.3% under the non-PUA program. This is also consistent with reports that the PUA program was more prone to fraud ([Government Accountability Office, 2023](#)).

4.4. Alternative Clustering Procedure

Our baseline unsupervised clustering algorithm draws on card features from the entire length of their presence in our sample. This includes the inflows to and spending from these cards after coming under the purview of identity verification implemented by their respective states. A potential concern with this method is that cards could change their behavior in response to identity verification, and that could bias our clustering outcomes. We address this concern by measuring fraud using pre-treatment data, and then using the full history of UI disbursed to cards to estimate the treatment effects. However, some cards start receiving UI only after the treatment date. To include such cards in the final set of clusters, we apply the same break-points that separate cards in the main clustering algorithm on these other cards that did not receive UI before the treatment date. This ensures that there is no mechanical drop in UI disbursed by states because of censoring of new applicants.

We also revise our baseline clustering to confirm that our results are robust to a range of specification and sample selection choices. First, instead of separating cards on total UI, we use UI per month and the spell (in months) for which cards receive UI as two complementary features. This method accounts for the differential length of time for which cards may be active in our sample. Second, we restrict the set of cards clustered to only those that receive UI from the 41 states finally considered for estimating treatment effects, which drops cards from states where we can identify UI but not the treatment dates. Finally, since the clustering algorithm requires us to specify the number of clusters created, we explore both manual and rule-based methods, such as binomial outcome for each feature, and find that clusters so created have comparable features.

We present results from these alternative clustering procedures as part of our robustness checks when examining the effects of identity verification.

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 to each cluster. Our main outcome variable is the UI dollars paid by states to cards based on their cluster. We also consider the number of cards to whom UI was disbursed within each cluster as an alternate measure. Our specification examines the outcome variables in each of the 6 months before and after the implementation of identity verification using a two-way fixed effects framework, saturating the model with state and time fixed effects. This design allows us to compare the effect of screening within and across clusters, and understand if the more likely fraudulent cards were disproportionately affected by this policy.

Our sample consists of 41 states where we can both identify unemployment insurance flows and the timing of identity verification. Between January 2019 and September 2021, 23 of these 41 states contracted with third-party agencies such as ID.me and LexisNexis to screen the identity of UI applicants; these states are considered treated.¹⁶ Within these 23 treated states, the timing of adoption differs, starting from September 2020 for Georgia and Indiana to June 2021 for Delaware.¹⁷ 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.

We note that the treated and control states displayed comparable socio-economic attributes over our sample period. [Table 3](#) compares treated and control states in terms of the share of suspicious cards, identity theft reports, the generosity of their state unemployment insurance policies, and other economic indicators. Further, to strengthen the exogeneity condition that any change we observe in UI disbursements is attributable solely to the policy adoption, [Table A6](#) confirms that treated states did not reduce the generosity of UI payments simultaneously with the introduction of identity verification. [Table A6](#) also shows that none of the treated states simultaneously terminated its participation in the federal PUA program.¹⁸

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 clustering methodology, which we present in [subsection 5.2](#) and [subsection 5.3](#).

¹⁶Kentucky and Washington state implemented only a pilot/partial version of identity verification on a small sample of recipients but did not extend it to the general population by the end of our sample in September 2021. We consider them as untreated.

¹⁷Some states implemented these measures for one program only (i.e. PUA or non-PUA). 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.

¹⁸[Holzer, Hubbard, and Strain \(2024\)](#) provide the list of states that chose to withdraw from the PUA program before the federally determined expiry in September 2021.

5.1. Difference-in-Differences Design

We estimate the following two-way fixed effects model to identify the effect of identity verification on each cluster separately, holding constant time-invariant characteristics of each state as well as national trends:

$$y_{st} = \beta \text{Treated}_s \times \text{Post} + \gamma \text{PUA Termination}_{st} + \alpha_s + \alpha_t + \varepsilon_{st}, \quad (1)$$

where the dependent variable y_{st} is the UI disbursed by state s in month t . *Treated* is an indicator variable that takes a value of one for the 23 states that adopted identity verification, and zero otherwise. *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. Observations outside of this window are trimmed. The treatment effect of identity verification is identified by the β parameter. *PUA Termination* is an indicator variable that takes a value of 1 when the pandemic unemployment assistance program was no longer active for state s in month t . Finally, α_s and α_t capture the state and month fixed effects, respectively. Standard errors are clustered by state and observations are population weighted.

We estimate this model for two dependent variables in turn: (log) UI dollars and (log) number of cards to whom UI is disbursed in that cluster. Furthermore, we estimate this equation for each cluster separately: First, we compare the outcomes in Suspicious (combination of Suspicious (Fast Cash) and Suspicious (Slow Cash) clusters) and Control clusters. Then, we estimate the model on Suspicious (Fast Cash) and Suspicious (Slow Cash) separately. Finally, we estimate the model on other non-suspicious clusters that receive large sums of UI but do not display behavior consistent with identity theft i.e. Concurrent Income and Discretionary Spending clusters. [Table 4](#) reports all twelve estimation results using UI dollars and number of cards as the two dependent variables.

Identity verification measures significantly impacted cards that displayed behavior most consistent with fraud. Panel A of [Table 4](#) shows that the treatment effect of this policy on the UI dollars disbursed to the Suspicious cluster was -0.312 log points, or a 27% decline, over a period of six months. Similarly, the impact of this policy on the number of Suspicious card recipients was -0.229 log points, or a 20% decline. On the other hand, we do not observe a statistically significant impact on the Control cluster in terms of UI dollars or number of recipient cards, indicating limited impact on non-suspicious recipients.

We also find that identity verification impacted the cards that are most consistent with organized criminal behavior. Panel B of [Table 4](#) shows that while both the Suspicious (Fast Cash) and Suspicious (Slow Cash) clusters witnessed a decline in UI after the policy intervention, the treatment

effects were larger for the group of cards that withdraw cash quickly after receipt of UI. The Suspicious (Fast Cash) cluster received 30% lower UI in dollar terms and 24% fewer cards, compared to 25% and 19% for Suspicious (Slow Cash). Finally, Panel C of [Table 4](#) confirms that other clusters that are distinct from the Control group but do not indicate identity theft were not impacted by identity verification. Neither the Concurrent Income nor the Discretionary Spending cluster sees a negative and statistically significant decline in UI, supporting that these measures were targeted towards reducing identity fraud, and not enforcing eligibility.

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 there was no pre-existing decline in the UI paid to the Suspicious cluster before the policy intervention which would confound the impact of the intervention itself. We estimate the model

$$y_{st} = \sum_{\substack{\tau \in \{-6, 6, \\ \tau \neq -1\}}} \beta_{\tau} \times Reltime_{\tau} + \gamma PUA Termination_{st} + \alpha_s + \alpha_t + \varepsilon_{st}, \quad (2)$$

where the variable *Reltime* is the number of months since the policy was introduced in the respective treated state. For never-treated states, *Reltime* = -1 throughout. We center the implementation month at *Reltime* = 0 and consider six months before and after the implementation to evaluate the validity of parallel trends assumption. Parameter β_{τ} identifies the impact of identity verification on the UI paid by treated states to the respective cluster for each of the six months before and after implementation, based on the UI paid in *Reltime* = -1. The specification controls for state and time fixed effects, as well as an indicator for months *t* when a state *s* did not participate in the federal PUA program. Standard errors are clustered by state and observations are population weighted. Similar to [Equation 1](#), we estimate this model in turn for Suspicious, Concurrent Income, Discretionary Spending, and Control clusters.

[Figure 3](#) shows the event plot of β_{τ} estimates for the Suspicious cluster. The decline in UI paid to Suspicious cluster began immediately upon the implementation of identity verification, with persistent negative effects for the rest of the post-period. This immediate and persistent decline is visible in both UI dollars (panel a) and number of cards (panel b). The effect grew over the months following implementation, as more claimants were brought under the scope, and payouts were reduced to suspicious cards by as much as 40% 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. [Figure A8](#) splits the event study for the Suspicious (Fast Cash) and Suspicious (Slow Cash) clusters and confirms that both the sub-categories within the suspicious cards displayed a sharp decline in UI post treatment,

with the parallel trends assumption continuing to hold.

Importantly, we do not observe a statistically significant decline in UI in any of the post-treatment months for the Control cluster. [Figure 4](#) reports the estimated β_τ in panel (a) for UI dollars and panel (b) for number of cards in the Control cluster. We note that these cards continued to receive UI comparable to the pre-period, with a slight upward trajectory in later months. This points to the idea that identity verification did not lower UI reciprocity among non-fraudulent cardholders. Similarly, in [Figure A9](#) we do not observe any sharp change in the UI paid out to the Concurrent Income and Discretionary Spending clusters. This indicates that those clusters, which display distinct spending and eligibility patterns from the control cluster but are not suspicious in terms of identity theft, are also not affected, in line with our expectations. Overall, our difference-in-difference estimation suggests that identity verification reduced the UI disbursed to cards that displayed behaviors most consistent with identity theft, while leaving other types of recipients generally unaffected.

5.2. Robustness to Two-Way Fixed Effects Specification

We conduct a series of robustness checks to address critiques of modern two-way fixed effects designs in settings with staggered adoption of policy and potentially heterogeneous treatment effects. We show that our results are robust to these alternative methodologies, and that our main finding – that identity verification affected suspicious cards only – persists. In [Appendix C](#), we present additional variations of our baseline specification.

5.2.1. Negative weights.

A recent literature has highlighted the potential problems associated with negative weights in two-way fixed effects designs ([De Chaisemartin and d’Haultfoeuille, 2020](#), [Borusyak, Jaravel, and Spiess, 2021](#), [Roth, Sant’Anna, Bilinski, and Poe, 2023](#)). This is particularly a concern when there is heterogeneity in treatment effects over time, which may be the case in our setting.

Following [De Chaisemartin and d’Haultfoeuille \(2020\)](#) and [Roth et al. \(2023\)](#), we check and rule out the presence of negative weights in our estimation. We find that, under the common trends assumption for both the Suspicious and the Control clusters, all average treatment effects on the treated carry a positive weight. Consequently, the sum of the positive weights is throughout equal to 1 and that of negative weights is equal to 0.

5.2.2. Alternative Two-Way-Fixed-Effects Estimation.

We re-estimate [Equation 1](#) using the method 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, where each group represents the time period when units are first treated. 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. We implement this method using the not-yet treated states as controls.

Panel A of [Table A7](#) reports [Callaway and Sant’Anna \(2021\)](#) estimator for the β coefficient in [Equation 1](#). Consistent with the baseline TWFE specification, all observations are population-weighted and standard errors are clustered by state. For the Suspicious cluster, we continue to observe a sharp decline in both UI dollars and number of cards to whom UI is paid after the introduction of identity verification. The size of coefficients is comparable to baseline estimates reported in [Table 4](#), while the statistical significance is larger. On the other hand, there is no significant decline in the UI disbursed to the Control cluster.

[Figure A10](#) shows the event study plots (in green) of the β_τ coefficients in [Equation 2](#) for the Suspicious and the Control clusters using [Callaway and Sant’Anna \(2021\)](#) estimator. We continue to observe a decline in UI paid to the Suspicious cluster after the introduction of identity verification. The pre-trends for the Suspicious cluster confirm that there was no decline observed before the implementation of policy. Further, the Control cluster did not see a decline in the post-period. Both the results are consistent with the baseline results reported in [Figure 3](#) for the Suspicious and [Figure 4](#) for the Control cluster.

5.2.3. *Alternative estimation using “stacked” difference-in-differences.*

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 the treatment effects using a “stacked” difference-in-differences approach to mitigate this concern. Following [Cengiz, Dube, Lindner, and Zipperer \(2019\)](#), we create a stacked dataset of treatment months that each includes only the group of states that were treated in that month and a set of control states that were either (i) never treated, or (ii) not treated up to six months after the respective treatment month. We re-estimate [Equation 1](#) and [Equation 2](#) on this stacked dataset with rest of the specifications analogous to our baseline two-way fixed effects model.

Panel B of [Table A7](#) shows the estimation of the stacked difference-in-difference for the Suspicious and the Control cluster. We note that for both the dependent variables, the UI dollars and number of cards, there is a significant decline in UI disbursed to suspicious cards of magnitude comparable to the baseline results of [Table 4](#). However, the Control cluster does not see any

such reduction, validating our interpretation that this policy targeted the most likely fraudulent recipients.

Figure A10 shows the dynamic event study plots (in blue) of the β_τ coefficients in Equation 2 for the suspicious and control clusters using the “stacked” difference-in-differences approach. 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. Robustness to Alternative Cluster Construction

Next, we show our treatment effects are robust to a number of alternative ways of clustering cards in our sample. Most importantly, we restrict the income and spending features of cards to months before treatment for their UI-paying state, and re-create all clusters following the same steps as in the baseline method. We then estimate Equation 1 and Equation 2 on these clusters. Panel A of Table A8 reports β estimates under Equation 1, and Figure A11 plots the event study versions of β_τ specified in Equation 2. Both exhibits report treatment effects for the Suspicious and Control clusters, with UI dollars and number of cards as dependent variables in turn.

We find consistent evidence that identity verification reduced payouts to the Suspicious cluster but not to the Control cluster. The magnitude and statistical significance of treatment effects on these revised clusters confirms that identity verification sharply reduced payouts to suspicious cards, while leaving control cards largely unaffected.

Table A8 also shows robustness of our results to a number of other choices made during the clustering process. Panel B of Table A8 confirms that separating High UI and Low UI cards using UI per month and spell (in months) of being on UI rolls instead of total UI does not alter the treatment effects. Panel C shows that specifying the number of clusters using a binomial outcomes rule produces similar outcomes as specifying the number of clusters using the economic properties of the features used.¹⁹ Finally, panel D shows that restricting the clustering sample to only those cards that receive UI from the 41 states considered for the final analysis has a treatment effect comparable to the baseline where we clustered using all UI-receiving cards in our sample.

¹⁹Under a binomial outcomes rule, the total number of clusters created at each step equals 2^N , where N is the number of features used in the k-means algorithm. In our baseline clustering, we specify the number of clusters as three when separating “High UI” cards using two features: cash withdrawals as a fraction of total outflows, and discretionary spending as a fraction of total outflows. Under the binomial rule, we specify the number of clusters as four.

6. ESTIMATION OF ECONOMIC MAGNITUDE

Our methods also allow us to provide an estimate of the rough economic magnitude of UI fraud and the dollar benefit of identity 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 D](#) provides detailed calculations.

We estimate that \$32.7 billion was lost to fraud between March 2020 and September 2021. [Table A4](#) shows that the share of UI disbursed to suspicious cards was 11.3% in our sample. Assuming that this share is representative, we arrive at the estimated misallocation by multiplying it by \$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 a total of \$9 billion could have been counterfactually saved if all states had this policy in place at the onset of the COVID-19 pandemic. These estimates are derived in three steps detailed below.

First, we note that treated states disbursed a monthly average of \$8.9 billion in the six months preceding their respective treatment month, and the (population-weighted) average number of post-treatment months was 5.7 until the end of our sample. Additionally, the share of UI paid to suspicious cards was 13.2% in the month preceding treatment. Given our treatment effect of a 27% reduction in UI paid to suspicious cards, identity verification saved these states an estimated \$1.8 billion ($\$8.9 \text{ billion} \times 5.7 \times 0.134 \times 0.27$).²⁰

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 \$9.2 billion. Assuming that the magnitude of treatment effect on suspicious cards remains constant at 27%, these states could have potentially saved another \$2.4 billion ($\$9.2 \text{ billion} \times 27\%$) if the policy was in effect from March 2020. Finally, we extend the estimated counterfactual savings to the 18 states in our sample that did not implement identity verification, and 10 states outside of our sample. These states altogether disbursed \$159.7 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 11.2%. Extending the same share of UI to suspicious cards among the 10 out-of-sample states, and assuming a treatment effect of 27% analogous to treated states, we estimate that an additional \$4.8 billion could have been saved nationwide through this policy.²¹

²⁰An alternate calculation could account for the actual UI disbursed by treated states to suspicious cards after treatment, which was \$5.6 billion until September 2021. If this amount was already lowered by the treatment effect of 27%, then the estimated savings comes to \$2.1 billion ($\$5.6 \text{ billion} \times 0.27 / (1-0.27)$).

²¹It is possible that the 10 states outside of our sample may have had different levels of fraud. If we exclude

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 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, solely within the debit card market, we estimate that \$10.6 billion was lost to fraud and that \$2.9 billion could have been saved if identity verification had been implemented more broadly and more quickly.

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

In contrast to the billions of dollars in savings that identity verification provides, the economic costs are quite modest. Congressional documents indicate that the total compensation contracted between ID.me and state agencies was under \$60 million ([Committee on Oversight and Accountability, 117th Congress, 2022](#)). The hassle costs are also quite low: at a conservative 15 minutes per applicant at a wage of \$25 per hour, identity verification on 30 million UI applications would only have an opportunity cost of \$187.5 million. This is reflected, in part, by the lack of impact on the Control cluster we observe, validating that hassle costs were not a major factor for that group.²³ Overall, our analysis indicates that identity verification was not only powerful at eliminating public benefits fraud, but also highly cost-effective.

those states in our estimation, we arrive at the counterfactual savings of \$1.9 billion for non-treated states.

²²For comparison, the US Government Accountability Office 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 ([U.S. Government Accountability Office, 2023a](#)).

²³[Figure A12](#) confirms that there was no decline in state-wide UI disbursements.

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. 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). ⁶
- Axios (2021, June). Pandemic unemployment fraud: Benefits stolen. <https://www.axios.com/2021/06/10/pandemic-unemployment-fraud-benefits-stolen>. ²
- Baker, S. R. and C. Yannelis (2017). Income changes and consumption: Evidence from the 2013 federal government shutdown. *Review of Economic Dynamics* 23, 99–124. ⁴¹
- 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. ¹²
- Berg, T., A. Fuster, and M. Puri (2022). Fintech lending. *Annual Review of Financial Economics* 14, 187–207. ⁷
- Bian, B., M. Pagel, and H. Tang (2023). Consumer surveillance and financial fraud. Technical report, National Bureau of Economic Research. ⁷
- Bipartisan Policy Center (2023). Integrity data hub: A multi-state solution to unemployment insurance fraud. <https://bipartisanpolicy.org/blog/integrity-data-hub-multi-state-solution-unemployment-insurance-fraud/>. ²
- Borusyak, K., X. Jaravel, and J. Spiess (2021). Revisiting event study designs: Robust and efficient estimation. *arXiv preprint arXiv:2108.12419*. ²¹
- 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. ^{10, 11}
- Callaway, B. and P. H. Sant’Anna (2021). Difference-in-differences with multiple time periods. *Journal of econometrics* 225(2), 200–230. ^{4, 21, 22, 58, 68}
- 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. ^{4, 22, 58, 68}
- CNBC (2022). Fraud had significant role in \$163 billion leak from pandemic-era unemployment system. <https://www.cnn.com/2022/05/27/fraud-had-significant-role-in-163-billion-leak-from-pandemic-era-unemployment-system.html>. ²
- Committee on Oversight and Accountability, 117th Congress (2022, November 17). Congressional Press Release. ^{9, 25}

- 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. ¹⁶
- Dave, P. (2021). States using id.me, rival identity check tools for jobless claims. <https://www.reuters.com/business/states-using-idme-rival-identity-check-tools-jobless-claims-2021-07-22/>. ⁹
- De Chaisemartin, C. and X. d’Haultfoeuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review* 110(9), 2964–2996. ²¹
- Dimmock, S. G. and W. C. Gerken (2012). Predicting fraud by investment managers. *Journal of Financial Economics* 105(1), 153–173. ⁷
- 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. ⁷
- Dube, A. (2021a). 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. ⁶
- Dube, A. (2021b). A plan to reform the unemployment insurance system in the united states. *Hamilton Project Policy Proposal* 3. ²
- 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. ^{5, 7}
- Federal Trade Commission (2020). Consumer sentinel network data book 2020. https://www.ftc.gov/system/files/documents/reports/consumer-sentinel-network-data-book-2020/csn_annual_data_book_2020.pdf. ⁸
- Federal Trade Commission (2021). Id theft by metro area. <https://public.tableau.com/app/profile/federal.trade.commission/viz/FraudandIDTheftMaps/IDTheftbyMetroArea>. ¹⁶
- Florida Department of Financial Services (2021). Merchant category codes (mccs). <https://fs.fldfs.com/iwpapps/pcard/docs/MCCs.pdf>. ¹¹
- Fuller, D. L., B. Ravikumar, and Y. Zhang (2015). Unemployment insurance fraud and optimal monitoring. *American Economic Journal: Macroeconomics* 7(2), 249–290. ⁶
- 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. ⁷

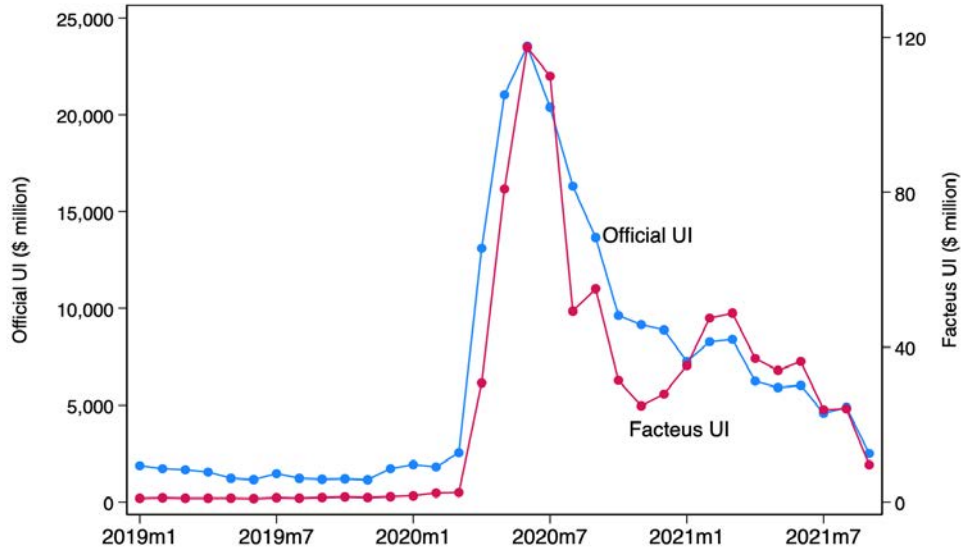
- Ganong, P., F. E. Greig, P. J. Noel, D. M. Sullivan, and J. S. Vavra (2022). Spending and job-finding impacts of expanded unemployment benefits: Evidence from administrative micro data. Technical report, National Bureau of Economic Research. ⁶
- Ganong, P. and P. Noel (2019). Consumer spending during unemployment: Positive and normative implications. *American Economic Review* 109(7), 2383–2424. ^{6, 13}
- Gao, J., J. Pacelli, J. Schneemeier, and Y. Wu (2020). Dirty money: How banks influence financial crime. Available at SSRN 3722342. ⁷
- Government Accountability Office (2022). Transformation needed to address program design, infrastructure, and integrity risks. <https://www.gao.gov/assets/gao-22-105162.pdf>. ²
- Government Accountability Office (2023). DOL needs to address substantial pandemic UI fraud and reduce persistent risks. <https://www.gao.gov/assets/gao-23-106586.pdf>. ^{8, 17}
- Griffin, J. M. and S. Kruger (2023). What is forensic finance? Available at SSRN 4490028. ⁷
- Griffin, J. M., S. Kruger, and P. Mahajan (2023a). Did fintech lenders facilitate ppp fraud? *The Journal of Finance*. ^{2, 6}
- 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)*. ^{6, 44}
- 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. ⁴³
- Holzer, H. J., G. Hubbard, and M. R. Strain (2024). Did pandemic unemployment benefits increase unemployment? evidence from early state-level expirations. *Economic Inquiry* 62(1), 24–38. ^{18, 67}
- House Committee on Oversight and Reform (2022). Chairs maloney, clyburn release evidence facial recognition company id.me downplayed excessive wait times for americans seeking unemployment relief funds. <https://oversightdemocrats.house.gov/news/press-releases/chairs-maloney-clyburn-release-evidence-facial-recognition-company-idme>. ⁹
- Howard, D. H. and I. McCarthy (2021). Deterrence effects of antifraud and abuse enforcement in health care. *Journal of Health Economics* 75, 102405. ⁶
- 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. ⁶
- 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. ⁷

- Karger, E. and A. Rajan (2020). Heterogeneity in the marginal propensity to consume: evidence from covid-19 stimulus payments. ¹⁰
- 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. ⁷
- Leder-Luis, J. (2023). Can whistleblowers root out public expenditure fraud? evidence from medicare. *Review of Economics and Statistics*, 1–49. ⁶
- 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. ⁷
- Magnac, T. (2017). Atm foreign fees and cash withdrawals. *Journal of Banking & Finance* 78, 117–129. ⁴²
- Malaiyandi, S. (2020). You have options for how to receive your unemployment benefits. <https://www.consumerfinance.gov/about-us/blog/receive-your-unemployment-benefits-options/>. ¹⁰
- 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. ⁷
- Michigan Department of Labor and Economic Opportunity (2021). Unemployment insurance claimants issued new debit cards. <https://www.michigan.gov/leo/news/2021/08/31/unemployment-insurance-claimants-issued-new-debit-cards>. ¹¹
- Michigan Department of Labor and Economic Opportunity (2024). Uia debit card information. <https://www.michigan.gov/leo/bureaus-agencies/uia/tools/publications/uia-debit-card>. ¹⁰
- Mookherjee, D. and I. P. Png (1992). Monitoring vis-a-vis investigation in enforcement of law. *The American Economic Review*, 556–565. ⁵
- 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. ⁷
- Nichols, A. L. and R. J. Zeckhauser (1982). Targeting transfers through restrictions on recipients. *The American Economic Review* 72(2), 372–377. ^{2, 7}
- Pahlka, J. (2023). *Recoding America*. Metropolitan Books. ²⁶

- Pennsylvania Office of Attorney General (2024). Caregiver pleads guilty to stealing intellectually disabled client’s personal information to get nearly \$90k in unemployment benefits. <https://www.attorneygeneral.gov/taking-action/case-update-caregiver-pleads-guilty-to-stealing-intellectually-disabled-clients-personal-information-to-get-nearly-90k-in-unemployment-benefits>. ²
- Polinsky, A. M. and S. Shavell (2000). The economic theory of public enforcement of law. *Journal of economic literature* 38(1), 45–76. ⁵
- ProPublica (2021). How unemployment insurance fraud exploded during the pandemic. <https://www.propublica.org/article/how-unemployment-insurance-fraud-exploded-during-the-pandemic>. ²
- 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. ¹²
- 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. ⁴³
- 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*. ²¹
- 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. ⁷
- Shekhar, S., J. Leder-Luis, and L. Akoglu (2023). Unsupervised machine learning for explainable health care fraud detection. ^{7, 12}
- U.S. Department of Justice (2013). Wic fraud complaints. <https://www.justice.gov/archive/usao/nys/pressreleases/November13/WICFraudArrests/WIC%20Fraud%20Complaints.pdf>. ²
- U.S. Department of Justice (2024a). 32 wic participants convicted for selling their wic vouchers for cash. <https://www.justice.gov/usao-sdga/pr/32-wic-participants-convicted-selling-their-wic-vouchers-cash>. ²
- U.S. Department of Justice (2024b). Hyde park man pleads guilty to covid relief and federal assistance benefit fraud. <https://www.justice.gov/usao-ma/pr/hyde-park-man-pleads-guilty-covid-relief-and-federal-assistance-benefit-fraud>. ²

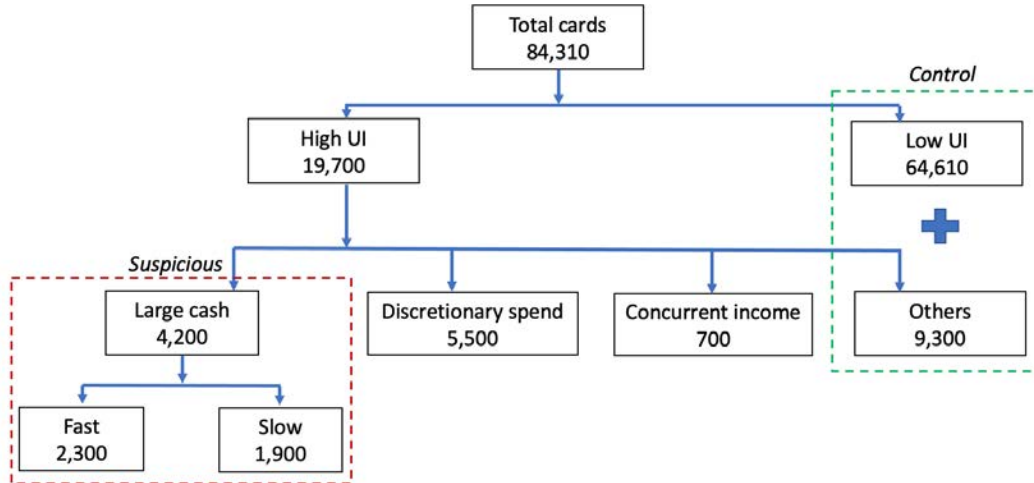
- U.S. Department of Justice (2024c). Justice department takes action against covid-19 fraud. <https://www.justice.gov/opa/pr/justice-department-takes-action-against-covid-19-fraud>. ⁵
- U.S. Department of Justice (2024d). Leader of \$20m covid-19 relief fraud ring sentenced to 15 years. <https://www.justice.gov/opa/pr/leader-20m-covid-19-relief-fraud-ring-sentenced-15-years>. ⁵
- U.S. Department of Labor (2023). Ui modernization 2023 strategy. <https://www.dol.gov/agencies/eta/ui-modernization/2023-strategy>. ²
- U.S. Department of Labor, ETA (2024). Unemployment insurance data dashboard. <https://oui.doleta.gov/unemploy/DataDashboard.asp>. ⁶¹
- U.S. Department of Labor, Office of Inspector General (2021). The ETA needs to ensure state workforce agencies implement effective ui program fraud controls for high risk areas. <https://www.oig.dol.gov/public/reports/oa/2021/19-21-002-03-315.pdf>. ²
- U.S. Department of Labor, Office of Inspector General (2023). Oversight work on unemployment insurance. <https://www.oig.dol.gov/doloiguioversightwork.htm>. ²
- U.S. Department of the Treasury (2023). Id.me instructions for CERTS portal. <https://home.treasury.gov/system/files/136/ID.me-Instructions-for-CERTS-Portal.pdf>. ⁹
- U.S. DOL (2023). Unemployment insurance payment accuracy data. <https://www.dol.gov/agencies/eta/unemployment-insurance-payment-accuracy/data>. ⁸
- U.S. Government Accountability Office (2023a). More fraud has been found in federal covid funding—how much was lost under unemployment insurance programs. <https://www.gao.gov/blog/more-fraud-has-been-found-federal-covid-funding-how-much-was-lost-under-unemployment-insurance-programs>. ²⁵
- U.S. Government Accountability Office (2023b). DOL needs to address substantial pandemic UI fraud and reduce persistent risks. <https://www.gao.gov/assets/gao-23-106586.pdf>. ^{2, 4, 8}
- Washington Bankers Association (2021). State ui ach id list. https://wabankers.com/userfiles/uploads/State_UI_ACH_ID_List.xlsx. ^{11, 41}
- Wu, Y., Y. Xu, and J. Li (2019). Feature construction for fraudulent credit card cash-out detection. *Decision Support Systems* 127, 113155. ^{3, 12}
- Zhou, Y. and F. Correia (2022). Are we friends? cross-predictability of stock returns. ¹⁰

Figure 1: Unemployment Insurance in Administrative and Facticeus Data




Notes: This figure plots monthly unemployment insurance (UI) disbursements from administrative (left axis) and Facticeus debit card data (right axis) between January 2019 and September 2021 across 41 states. Administrative UI data are sourced from the [US Department of Labor](#) and include both the regular and pandemic unemployment assistance (PUA) programs. Facticeus UI data are extracted by string-matching UI payment descriptions from debit card transaction records and aggregated across the 41 states where we can observe these flows. The list of states is available in [Table A1](#).

Figure 2: Summary of K-means Clustering Algorithm




Notes: This figure depicts the order in which cards are grouped into clusters representing varying levels of suspicious or fraudulent behavior. The top branch splits the universe of 84,310 cards into those with high and low UI based on total UI received by each card over the sample. Within cards receiving “High UI”, the algorithm uses spending on cash withdrawals and discretionary spending on categories such as alcohol and gambling to create two clusters: “Suspicious” and “Discretionary spend.” For “High UI” cards that do not fall into the two categories above, the algorithm evaluates concurrent non-UI income to create the “Concurrent Income” cluster. All other un-categorized “High UI” cards are combined with “Low UI” cards to create a “Control” cluster. Cards with large cash spending are further split based on the speed of withdrawal after the receipt of UI. The labels show the approximate number of cards in each split.

Figure 3: Impact of Identity Verification on Suspicious Cluster




(a) UI Dollars




(b) Number of Cards

Notes: This figure plots estimates of β_τ from a model of the form in equation 2 for the suspicious cards cluster. These event studies plot coefficients for UI disbursements up to 6 months before and after the implementation of identity verification. Panel (a) shows the results for log UI dollars as the dependent variable and panel (b) shows the results for log number of cards receiving UI as the dependent variable. Error bars provide the 95% confidence intervals for each estimate.

Figure 4: Impact of Identity Verification on Control Cluster



(a) UI Dollars



(b) Number of Cards

Notes: This figure plots estimates of β_τ from a model of the form in Equation 2 for the control cards cluster. These event studies plot monthly UI disbursements up to 6 months before and after the implementation of identity verification. Panel (a) considers (log) UI dollars as the dependent variable and panel (b) considers (log) number of cards to whom UI is disbursed as the dependent variable. Error bars provide the 95% confidence interval for each estimate.

Table 1: Descriptive Statistics of Debit Card Sample

Panel A: Card-level	Mean	SD	p25	p50	p75	N
Total UI (\$)	10,113	8,970	3,251	7,515	14,112	84,310
# UI (count)	19	21	4	12	26	84,310
UI per trans (\$)	1,011	1,524	383	577	952	84,310
UI as % of income	0.63	0.32	0.34	0.68	0.97	84,310
Total income (\$)	22,923	23,814	6,660	15,383	31,129	84,310
Total spend (\$)	14,953	15,900	3,598	10,095	21,035	84,310
Cash spend (\$)	2,885	5,340	0	356	3,622	84,310
Wire spend (\$)	372	1,797	0	0	0	84,310
International wire (\$)	40	733	0	0	0	84,310
Groceries (\$)	3,167	5,450	250	1,265	3,729	84,310
Discretionary spend (\$)	2,565	3,585	285	1,271	3,461	84,310
Panel B: Card-month level	Mean	SD	p25	p50	p75	N
Total UI (\$)	2,125	2,106	824	1,701	2,697	401,249
# UI (count)	4	3	2	4	5	401,249
UI per trans (\$)	698	1,007	261	452	770	401,249
UI as % of income	0.82	0.26	0.68	0.98	1.00	401,249
UI as % of earnings	0.96	0.14	1.00	1.00	1.00	401,249
Total income (\$)	2,790	2,592	1,204	2,193	3,474	401,249
Total spend (\$)	1,808	1,852	563	1,375	2,470	401,249
Cash spend (\$)	451	947	0	0	529	401,249
Wire spend (\$)	52	311	0	0	0	401,249
International wire (\$)	4	92	0	0	0	401,249
Groceries (\$)	340	641	11	133	415	401,249
Discretionary spend (\$)	288	413	39	165	391	401,249

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. We include cards that receive at least \$1,000 in unemployment insurance during the sample period from January 2019 to September 2021. Panel B further conditions on those months where a card receives any amount of UI. We identify UI inflows using string matching of transaction descriptions and categorize spending using merchant transaction codes. We group spending on liquor shops, gambling, tobacco, restaurants, and purchase of vehicles into discretionary spending.

Table 2: Summary Statistics by Cluster

	Suspicious (Fast Cash)	Suspicious (Slow Cash)	Concurrent Income	Discretionary Spending	Control
Unemployment Insurance (\$/month)	2,464 (2,566)	2,378 (1,798)	1,128 (647)	1,437 (934)	1,297 (1,646)
Longest UI spell (months)	8.3 (4.0)	8.2 (4.0)	8.7 (3.6)	8.9 (4.1)	3.7 (3.1)
Cash withdrawal (share of spend)	0.61 (0.17)	0.55 (0.18)	0.08 (0.09)	0.03 (0.07)	0.18 (0.25)
Discretionary spend (share of spend)	0.10 (0.08)	0.11 (0.08)	0.16 (0.07)	0.38 (0.10)	0.20 (0.16)
Time to cash (hours)	14 (5)	51 (75)	197 (504)	193 (664)	241 (912)
UI as % of earnings	1.00 (0.06)	1.00 (0.07)	0.70 (0.12)	1.00 (0.07)	1.00 (0.11)
Stimulus checks (\$/month)	51 (112)	52 (112)	90 (139)	48 (98)	44 (133)
Tax refunds (\$/month)	36 (102)	41 (112)	90 (157)	49 (114)	50 (163)

Notes: This table shows the mean and standard deviations of features displayed by cards within each cluster created by the k-means algorithm. Standard deviations are reported in parentheses below the means. Suspicious cards are characterized by larger sums of UI per month, longer unbroken spell on UI rolls, and larger proportion of spending via cash withdrawals compared to Control cluster. Further, cards in the Suspicious (Fast Cash) cluster withdraw cash within a day of receipt of UI. Cards in the Concurrent Income cluster simultaneously receive, on average, 30% of their monthly income from non-UI payrolls. Discretionary spend includes spending on liquor shops, gambling, tobacco, restaurants, and purchase of vehicles.

Table 3: Comparison between Treated and Control States

	Treated States			Control States			Diff.	Data Source
	Mean	SD	p50	Mean	SD	p50	p-val	
Share of Suspicious UI (\$)	12.2%	8.4%	11.8%	10.2%	7.6%	9.5%	0.444	Facteus
Share of Suspicious cards (#)	10.4%	7.2%	9.1%	9%	6.8%	7.4%	0.529	Facteus
ID theft reports (2019)	5.4	1.4	5.6	5.2	1.8	5.1	0.682	FTC
ID theft reports (2020)	143.3	270.2	42.4	169	271.7	46	0.765	FTC
Max. eligible UI (\$ per week)	496	223	463	550	167	529	0.379	BLS
Max. eligible UI (# of weeks)	26	2	26	25	4	26	0.287	BLS
Unemployment rate	6.4%	3.3%	5.6%	6.6%	3.3%	5.9%	0.329	BLS
Labor force participation rate	62.7%	3.6%	62.6%	62.5%	4.8%	63.3%	0.506	BLS
Household income (\$, 2019)	67,760	9,194	70,030	71,854	13,814	72,275	0.288	USCB
Education (\geq Bachelor) (2019)	32.1%	5.4%	31.3%	32.9%	8.6%	31.9%	0.726	USCB

Notes: This table compares the 23 treated and 18 control states across unemployment and macro-economic variables. (Table A1 provides the list of treated and control states.) We calculate the share of suspicious UI and cards based on our unsupervised machine learning algorithm that groups debit cards into various clusters. We source the state-level number of identity theft reports from the 2019 and 2020 Consumer Sentinel Network Data Books of the [Federal Trade Commission \(FTC\)](#). ID theft reports are per 100,000 people. We source monthly statistics on the state-level UI eligibility, unemployment rates, and labor force participation rate from the website of the [Bureau of Labor Statistics \(BLS\)](#). Annual statistics on state-level median household income and education levels are sourced from the [United States Census Bureau \(USCB\)](#). For monthly series, we use the data as of one month prior to the adoption of identity verification for treated states and a time-series average from March 2020 through September 2021 for control states. We report the p-values associated with a two-sided t-test on the difference in means of treated and control states in the second to last column on the right.

Table 4: Treatment Effects of Identity Verification

Panel A: Suspicious v/s Control cluster	Suspicious		Control	
	UI Dollars	Number of Cards	UI Dollars	Number of Cards
Treated \times Post	-0.312** (0.123)	-0.229** (0.107)	0.045 (0.185)	0.067 (0.192)
N	643	643	920	920
Panel B: Within Suspicious cluster	Suspicious (Fast Cash)		Suspicious (Slow Cash)	
	UI Dollars	Number of Cards	UI Dollars	Number of Cards
Treated \times Post	-0.358** (0.136)	-0.272** (0.115)	-0.287** (0.112)	-0.214** (0.102)
N	570	570	597	597
Panel C: Other non-suspicious clusters	Concurrent Income		Discretionary Spending	
	UI Dollars	Number of Cards	UI Dollars	Number of Cards
Treated \times Post	-0.113 (0.199)	-0.133 (0.140)	0.017 (0.198)	-0.025 (0.156)
N	565	565	734	735
Controls	Y	Y	Y	Y
State, Month FE	Y	Y	Y	Y

Notes: This table shows that identity verification impacted UI disbursement to suspicious cards. The table reports estimates of β from a two-way fixed effects model of the form in [Equation 1](#). The model is estimated separately for each cluster, and uses (log) UI Dollars and (log) Number of Cards as dependent variables in turn. The post period starts from the month in which identity verification was adopted by treated states and includes up to six months after that, until the end of our sample period in September 2021. The pre-period includes six months before the implementation month. Standard errors are clustered by state and reported in parentheses. Observations are weighted by state population. [Figure 3](#) shows this table’s event study version for the Suspicious cluster, [Figure 4](#) for the Control cluster, [Figure A8](#) for the sub-categories of Suspicious cluster, and [Figure A9](#) for other non-suspicious 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

May 2024

A. INCOME CLASSIFICATION

We construct income streams from card transactions that can be used as dimensions for clustering cards. The primary income indicators used for clustering are unemployment insurance (UI) and income from wages or salaries (Payroll). We also look for other income streams such as tax refunds, stimulus checks, and child tax credit to generally characterize cards in the sample, though these one-off credits are not used for clustering cards.

Unlike merchant category codes (MCC) for spending, income categories are not directly identifiable using pre-coded categories. Hence, we deduce the type of income using transaction descriptions. To do this, we read the transaction descriptions and match them to strings that identify UI or payroll income. The exact procedure entails the following steps.

1. Retain all transactions that are not coded as “Spend” under the transaction type column.
2. Locate UI using the ACH transaction descriptions file ([Washington Bankers Association, 2021](#)).
3. Tag income as Payroll where the merchant name contains string “PAYROLL”, “SALARY”, or “PR PAYMENT”.²⁴ Following [Baker and Yannelis \(2017\)](#), we also apply a minimum \$100 and maximum \$50,000 cutoff to transactions while tagging Payroll. Finally, to reduce the likelihood that payroll contains peer-to-peer fund transfers, we exclude strings containing “FUNDS TRANSFER”.
4. Drop cards that receive over \$1,000,000 in total income over the sample period because these could be business-owned cards.²⁵
5. Calculate the share of UI out of all earnings as $UI / (UI + Payroll)$, and term it “UI as a % of Earning.” This ratio is defined only for months where cards receive any amount of UI.
6. Identify other income sources using “IRS TREAS 310”: combined with “TAX REF” for Tax Refund, “CHILDCTC” for Child Tax Credit, and “TAXEIP” for Economic Impact Payments or stimulus.

²⁴[Baker and Yannelis \(2017\)](#) provide a list of strings that identify income in their transaction data; “PAYROLL” and “SALARY” are the most common of them. “PR PAYMENT” refers to wages paid by discount chains such Dollar General.

²⁵The Census Bureau defines small business as those with at least \$1,000,000 in annual income. See [here](#).

In total, we create five income streams: UI, Payroll, Tax refunds, Stimulus, and Child tax credit. In our data, the major vendors associated with Payroll are Walmart, Dollar General, Texas Roadhouse, Pilot Oil, Target, and Home Depot. We use UI as a % of earnings as a feature in our k-means clustering algorithm to locate cards that concurrently receive large sums of UI and non-UI wages, and separate them from cards with only one stream of income. The former group of cards may be more consistent with ineligibility to receive UI rather than fraud through identity theft. [Table 1](#) provides card and card-month level summary statistics, and shows that in the month with non-zero UI (Panel B), UI forms 82% of the total income of an average card, and 96% of total earnings.

B. CLUSTERING

B.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 into a card-month total. The second step is to calculate the averages per month for spending variables and scale them by total spending in that month. Other variables such as speed of withdrawal and the longest unbroken spell of UI are calculated at a card level. Specific details for each feature are described below.

- Unemployment insurance received (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. In the baseline version, we aggregate the UI received by a card throughout the sample period. In a robustness exercise, we use the average UI 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 and use it along with UI \$ per month in the robustness exercise.
- 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, multiple ATM withdrawals within a month and from different banks could attract withdrawal charges ([Magnac, 2017](#)), which should disincentivize legitimate recipients from spending via cash. We identify ATM withdrawals using MCC code 6011, bank branch withdrawals using MCC code 6010, and other types of cash withdrawal using codes 6050 and 6051, add them together at a card-month level, and convert it to a fraction of total monthly spend 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 a 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.
- Discretionary spending as a 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 a percentage of total spending at the card-month level.
- UI as a % of earnings: by design, UI compensates individuals for loss of wage income. Concurrent non-UI income could indicate ineligibility to receive UI. We identify concurrent income by finding inflows to the cards that contain words including “salary” or “payroll”, although this potentially misses inflows that are unlabeled or not labeled as such. By inspection, Payroll income sources include major employers of low-income beneficiaries such as Walmart and discount chains. We measure this variable as $UI/(UI+Payroll)$ in the months when UI is received.

B.2. *K-means Clustering*

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 in informative features is large.

The first split of cards using the k-means algorithm is based on the total UI received by cards. This feature separates cards based on reciprocity, which is relevant to the economic context of fraud. Using total UI as a feature, we specify that cards be split into two groups along this dimension. Out of 84,310 cards, 19,700 fall in the high UI bucket. These cards are further split into three groups using two spending categories – average fraction of spending on cash withdrawals, and on discretionary items. We specify that cards be split into three groups where

one group shows abnormally large cash withdrawals, the other spends abnormally large amount on discretionary items, and a third exhibits no outlier behavior along either of these two dimensions. For the third group of cards that do not show outlier spending patterns but nevertheless receive large sums of UI, we apply the clustering algorithm to detect potentially ineligible recipients using the variable UI as a % of earnings. We specify that these cards be split into two groups - those that receive concurrent non-UI Payroll income, and others that do not. The latter group of cards is combined with the Control cluster created earlier using “Low UI” cards. We note that about half of “High UI” cards are ultimately tagged as Control because they do not display suspicious spending or other income patterns, and therefore large UI receipts do not automatically designate cards as suspicious. Finally, within the Suspicious group of cards, we use the speed of cash withdrawal after UI receipt to split them into two sub-categories: Suspicious (Fast Cash) and Suspicious (Slow Cash).

The algorithm creates five clusters as shown in [Figure 2](#). The first two are suspicious in that they simultaneously receive large sums of UI and spend, on average, 55% to 61% in cash. The next two clusters represent either spending on discretionary items, or concurrent income from other sources, indicating some likelihood of ineligibility to receive UI. Concurrent income indicates potential ineligibility, but does not point towards identity theft. Discretionary spending could point towards a distinct set of people who are also less likely to commit identity theft, but have different responses to UI and have been studied in other literature on UI. All other cards with the least fraudulent behavior form part of the Control cluster – which form the bulk of our sample and do not display any combination of suspicious activity.

B.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

$$\text{Share of Suspicious Cards}_i = \beta \text{Identity Theft Reports}_i + \varepsilon_i, \quad (3)$$

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. The regressor is the number of identity theft reports (per capita) reported by the Federal Trade Commission in that MSA in the year 2020. 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 observations. Standard errors are clustered by state and observations are weighted by MSA population. [Table A3](#) reports the estimation results, with the number of cards added as an additional control in column (2) to account for a potential correlation between MSA size and incidence of theft.

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.4. Alternative Clustering Procedure

Our baseline clustering algorithm classifies all 84,310 cards in our sample that receive at least \$1,000 in unemployment insurance using the full history of their income and spending between January 2019 and September 2021. However, a potential concern with this method is that cards could change their behavior *in response to identity verification*, and that could bias our clustering outcomes. For example, it is conceivable that identity verification changed the spending habits of cards that were adversely affected by this policy. Similarly, if identity verification indeed reduced the UI disbursed by states to suspicious cards only, then such cards might be less likely to be classified as suspicious ex-ante, thereby downward biasing our estimates of the effectiveness of this policy. We address this concern by measuring fraud using pre-treatment data and then using the full history of UI disbursed to cards for estimating the treatment effects.

We re-construct card clusters by drawing features only from the months before each card is subject to treatment. To do this, we first tag each card to a “payer state” from which it received UI. For cards that receive UI from multiple states, we retain the one where they receive most UI from. (About 1% of cards in our sample receive UI from multiple states throughout the sample period.) Then, we set a cut-off date for each card which is the month of treatment for its respective state, and retain all the relevant card-month variables before that month only. Cards belonging to control or untreated states require no adjustment; we use their full available history. We then re-calculate all the card and card-month level features using the pre-treatment month observations. The rest of the clustering algorithm proceeds as per the baseline methodology. Note that there are 8,800 cards that start receiving UI only after the treatment date. We do not include these cards in the main clustering algorithm, but apply the same break-points that separate and categorize other cards into their respective clusters.

The next set of robustness involves changing some of the specifications used in our baseline clustering algorithm. First, we use UI per month and the spell of being on UI rolls as two features instead of only total UI received by cards, to group them into High UI or Low UI cards. This alternative method accounts for the differential length of time that cards are present in our sample. Second, we specify the number of clusters at each step using a binomial outcomes rule, where for each feature inputted, there should be two clusters created. This changes the number of clusters in the second step, where we use cash spending and discretionary spending as two features, from three to four. Finally, we drop those cards from clustering that receive UI from states such as Iowa, where we can identify UI inflows but not the treatment date. This ensures that the sample of cards clustered is exactly the same as that considered for treatment effects evaluation.

C. ROBUSTNESS TO TREATMENT EFFECTS

The paper describes robustness to modern critiques of staggered two-way fixed effects designs, using alternative estimation techniques. In this appendix, we detail two additional checks. First, we re-estimate [Equation 1](#) and [Equation 2](#) using equally weighted observations (as opposed to population weighted), which treats all states equally. Results in panel C of [Table A7](#) continue to show a negative impact on UI disbursed to Suspicious cluster.

Next, we re-estimate [Equation 1](#) in a way that accounts for internal zeroes i.e. months where a state does not pay any UI to a cluster. In the baseline, the use of log treats those observations as missing. In this robustness exercise, we use $\log(1+UI)$ and $\log(1+\text{Number of Cards})$ as alternate dependent variables to include months with zero UI by a state to a cluster. Panel D of [Table A7](#) reports the estimation results.

We find that including all the zeroes no longer shows any impact of the policy on Suspicious cluster. However, we also find that a number of small states (e.g., Alabama, Delaware, Idaho) have a majority of their time-series observations as zero, suggesting that our cluster-level sample is too small to capture meaningful trends in those states. When we exclude those states in Panel E of [Table A7](#), we find the results closer to our baseline specification.

D. 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 11.3%, as per [Table A4](#). 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 \$21.8 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 \$32.7 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 4](#) shows that identity verification declined the monthly UI paid out to suspicious clusters by 27%. 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 \$8.9 billion per month, in the six months preceding the treatment. Further, the (population-weighted) average number of post-treatment months was 5.7 until September 2021 when our sample ends. Finally, the share of UI paid to suspicious cards was 13.2% in the month preceding treatment. This gives us the counterfactual UI of \$6.7 billion ($\$8.9 \text{ billion} \times 5.7 \times 13.2\%$), which would have been paid to suspicious cluster in the absence of identity verification. Given the treatment effect of a 27% reduction in UI paid to suspicious cards, the savings from identity verification comes to \$1.8 billion ($\$6.7 \text{ billion} \times 0.27$) 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 \$5.6 billion. If this amount was

already lowered by the treatment effect of 27%, 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 \$2.1 billion ($\$5.6 \text{ billion} \times 0.27 / (1-0.27)$).

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 \$88 billion (\$59 billion under non-PUA and \$29 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 10.4%.²⁶ Therefore, the UI amount allocated to suspicious cards was \$9.2 billion ($\$88 \text{ billion} \times 10.4\%$). Assuming that the magnitude of treatment effect remains constant at 27%, these states could have potentially saved \$2.4 billion ($\$9.2 \text{ billion} \times 27\%$) 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 18 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 identify UI stream in Factiveus data, and 1 state (Iowa) for which the timeline of identity verification is unclear. These states altogether disbursed \$159.7 billion in this period (\$84.7 billion through non-PUA and \$75 billion under the PUA program). The share of suspicious cards in this period for states that did not implement identity verification is 11.2%. Assuming that (i) the same share applies to the 10 states outside our sample, and (ii) identity verification would have reduced this share by 27% analogous to treated states, we estimate that an additional \$4.8 billion could have been saved by these 28 states ($\$159.7 \text{ billion} \times 11.2\% \times 0.27$). The total counterfactual savings from ID verification for the full period and all states put together comes to a little over \$9 billion ($\$1.8 \text{ billion} + \$2.5 \text{ billion} + \4.8 billion).

E. ANALYSIS USING PUBLIC DATA

As a complementary analysis, we evaluate the effect of ID verification using public data obtained from US Department of Labor (DOL) for the period between January 2019 and December 2021. The DOL regularly releases UI-related summary report, which includes state-level information on claims activities, as well as the number and amount of payments under State unemployment insurance laws. In this analysis, we focus on initial claims and continued claims. Continued claims are reported on a weekly basis, and we select the last week’s observations of each month to convert them into the monthly level. Initial claims are at a monthly level originally.

To assess the treatment effects of ID verification on claims activities, we conduct a classic difference-in-differences analysis, using timing closely aligning with our main specification.²⁷ We estimate the model of [Equation 2](#) for each type of claim activity separately. This model controls for both state and time fixed effects. The parameter β_τ identifies the impact of identity verification on the UI paid by treated states, including a six-month

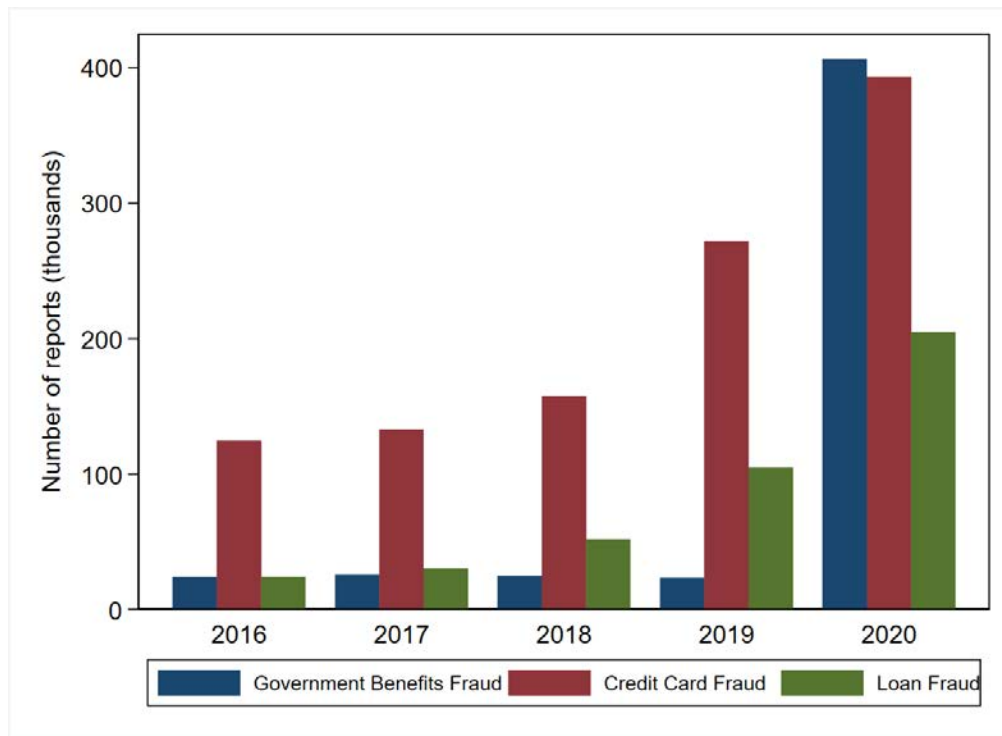
²⁶Consistent with the reports of a rise in fraud after the pandemic, the average share of suspicious cards is lower on average leading up to the treatment than it is in the month immediately preceding the treatment.

²⁷There are exceptions in three states where ID verification timings for regular UI and PUA diverge. In these instances, we apply specific timing for regular UI. The affected states are Arizona (2021m3), Nevada (2021m5), and Pennsylvania (2021m7).

period both before and after the 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 and show the resulting event plots in [Figure A13](#).

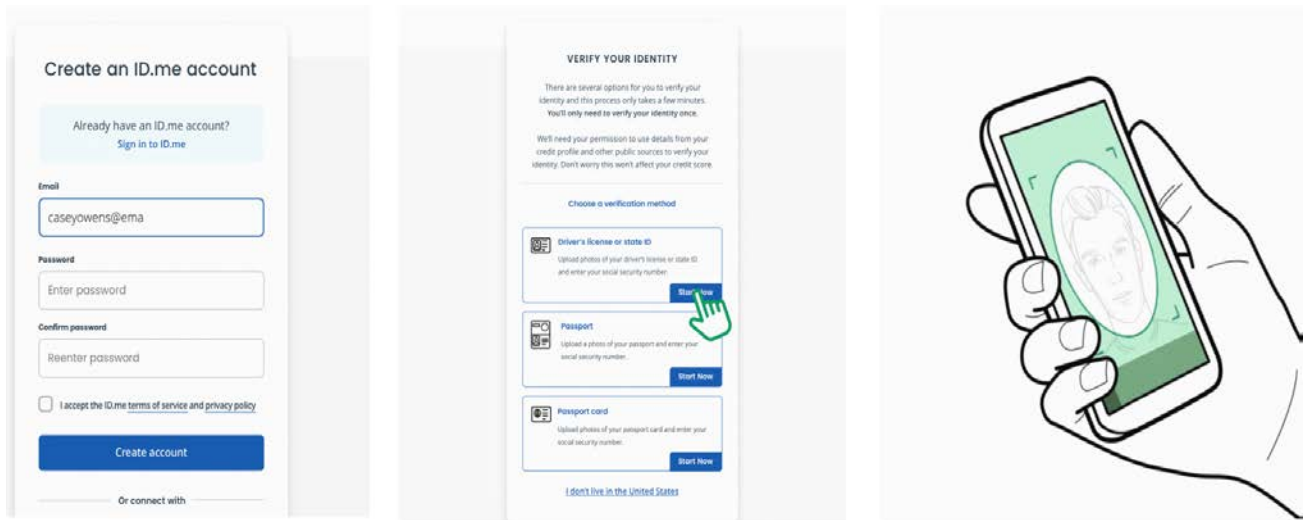
Panel (a) of [Figure A13](#) shows a persistent decline in the log initial claims. Before the implementation of the treatment, the coefficients are not statistically significant, though they exhibit a modest upward trend. However, treatment reverses that trend, and the post-treatment estimates are statistically significantly negative, indicating decline in the volume of initial claims due to the treatment. Similarly, Panel (b) of [Figure A13](#) presents a decline in log continued claims. Similar to the pattern observed in Panel (a), a slight upward trend is observed before the treatment, but negative point estimates appear for the entire post-period implying a real decline in the number of continued claims.

Figure A1: Identity Theft Reports



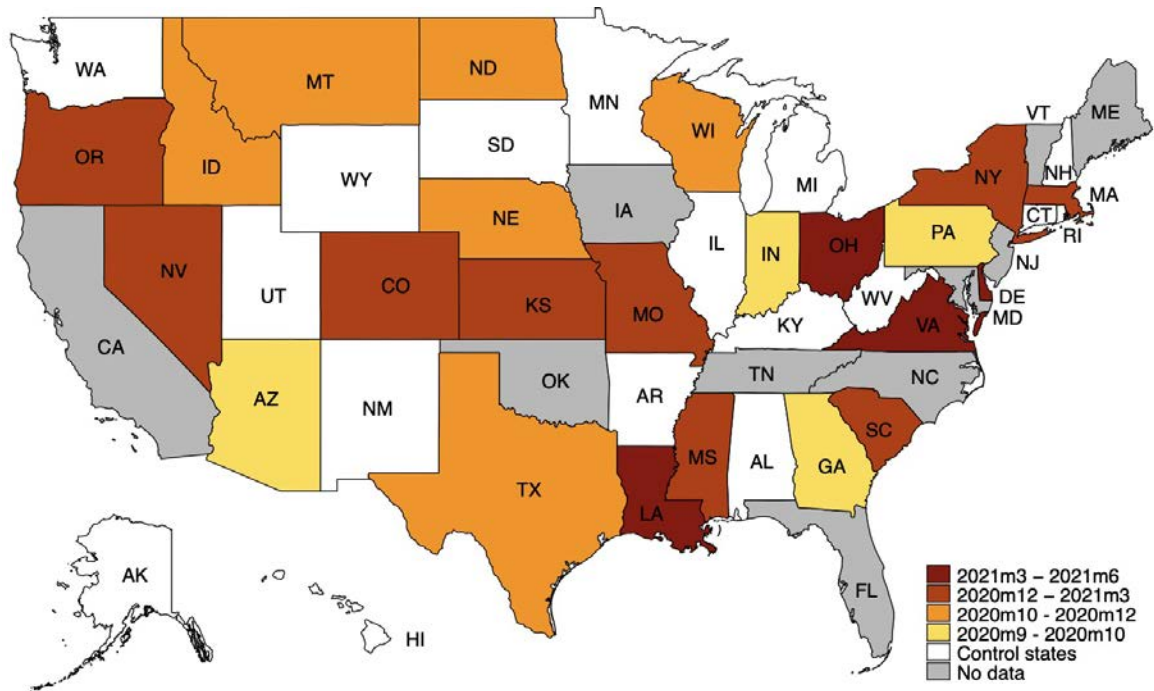
Notes: This figure plots the annual number of identity theft reports between 2016 and 2020 across all states by the category of theft. Identity theft data are sourced from the 2020 Consumer Sentinel Network Data Book of the [Federal Trade Commission](#). Identity theft reports increased across all categories in 2020 but the largest increase was recorded in the case of government benefits fraud.

Figure A2: Identity Verification Process



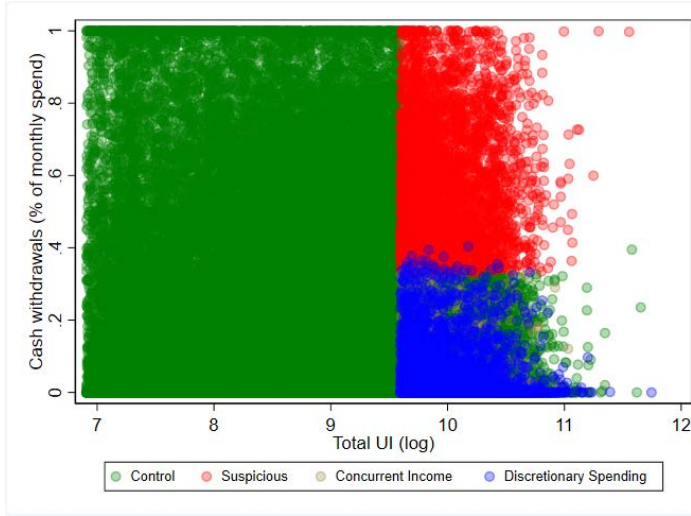
Notes: This figure shows the general steps involved in the “self-service identity verification” process that applicants undergo when claiming unemployment insurance benefits. The applicant is required to select a picture identity document, such as a state-issued driver’s license or passport, upload its photo, and take a photograph of themselves to confirm that they are in bona fide possession of the document. For applicants who cannot complete the “self-service”, the vendor provides a video-call based authentication with their agent. This figure has been sourced and adapted from ID.me’s [website](#).

Figure A3: Identity Verification Adoption Timeline




Notes: This map shows treated states, control states, and the states for which we cannot identify unemployment insurance or the timing of identity verification adoption. States that implemented the policy are shaded based on the timing of adoption between January 2019 and September 2021. We obtain data on adoption timeline from the Freedom of Information Act (FOIA) responses, Congressional records, or publicly-available information.

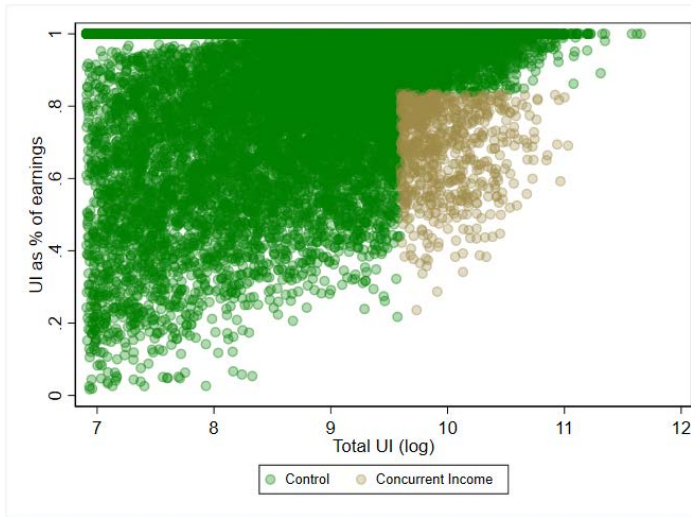
Figure A4: Bivariate Plots of Cluster Features




(a) Cash withdrawals (%) & Total UI (log)



(b) Discretionary spending (%) & Total UI (log)



(c) UI as % of earnings & Total UI (log)



(d) Cash withdrawals (%) & Time to cash (hours)


Notes: This figure scatters cards along two dimensions at a time, with dot colors representing the cluster to which each card belongs. Darker regions show greater density of cards at that point. Panel (a) shows that suspicious cards (in red) simultaneously receive large amounts of UI and spend primarily via cash withdrawals. Panel (b) similarly characterizes the Discretionary Spending cluster using spending on categories such as alcohol and tobacco. Panel (c) shows that cards with concurrent income are different from cards in the Control cluster because they receive large sums of UI simultaneously with payroll income. Within suspicious cards, panel (d) shows that cards that withdraw cash within a day of UI receipt are classified as Suspicious (Fast Cash) and others as Suspicious (Slow Cash). Time to cash in panel (d) is censored at 100 hours for ease of representation.

Figure A5: Debit Card Transactions of a Suspicious Card

318.15	Load	SCESC-UIBENEFITS UI BENEFITICYT				.	11/9/20 4:57
323.11	Load	SCESC-UIBENEFITS UI BENEFIT121U3				.	11/9/20 4:54
6345.1	Load	SCESC-UIBENEFITS UI BENEFITV2ZJ2				.	11/9/20 5:07
503.01	Spend	BANK OF AMERICA HQBB5	COLUMBIA	SC		6011	11/9/20 20:42
202.99	Spend	BANK OF AMERICA JQM5B	COLUMBIA	SC		6011	11/9/20 21:04
1500	Spend	BANK OF AMERICA	COLUMBIA	SC		6010	11/9/20 20:57
502.99	Spend	WELLS FARGO BAN5ZGSK	COLUMBIA	SC		6011	11/10/20 5:38
1500.01	Spend	WELLS FARGO C/A	COLUMBIA	SC		6010	11/10/20 21:53
203.01	Spend	WELLS FARGO BANBRK3W	COLUMBIA	SC		6011	11/10/20 5:35
999.99	Spend	WELLS FARGO C/A	COLUMBIA	SC		6010	11/12/20 20:27
326.66	Load	SCESC-UIBENEFITS UI BENEFIT111GG				.	11/16/20 5:18
318.83	Load	SCESC-UIBENEFITS UI BENEFITD4M1D				.	11/16/20 5:04
503	Spend	WELLS FARGO BANVF112	COLUMBIA	SC		6011	11/16/20 5:42
327.23	Load	SCESC-UIBENEFITS UI BENEFITBMKLB				.	11/23/20 4:58
318.59	Load	SCESC-UIBENEFITS UI BENEFIT1A3CB				.	11/23/20 5:09
142.99	Spend	1ST CITIZENS FIRIP	AIKEN	SC		6011	11/23/20 17:50
503	Spend	1ST CITIZENS YVP4J	AIKEN	SC		6011	11/23/20 17:44
325.8	Load	BASE GREEN DOT-RETAIL LOAD0S4EK				.	11/29/20 22:57
324.56	Load	SCESC-UIBENEFITS UI BENEFITH1YYK				.	11/30/20 5:10
330.32	Load	SCESC-UIBENEFITS UI BENEFITPPLE2				.	11/30/20 4:55
139.99	Spend	PALMETTO BHSKD	South Congare	SC		6011	11/30/20 21:17
499.99	Spend	PALMETTO JBM1Z	South Congare	SC		6011	11/30/20 21:32
99.47	Load	BASE GREEN DOT-RETAIL LOADHEDGP				.	12/1/20 22:38
326.22	Load	SCESC-UIBENEFITS UI BENEFIT5LPZK				.	12/7/20 5:11
317.88	Load	SCESC-UIBENEFITS UI BENEFITJK0BM				.	12/7/20 5:19
503	Spend	BANK OF AMERICA A5RJR	COLUMBIA	SC		6011	12/7/20 20:15
122.15	Load	BASE GREEN DOT-RETAIL LOAD2L4IU				.	12/8/20 22:30
142.99	Spend	BANK OF AMERICA H3KE5	COLUMBIA	SC		6011	12/7/20 20:27
405.24	Load	BASE GREEN DOT-RETAIL LOADGDHAT				.	12/11/20 20:24
326.91	Load	SCESC-UIBENEFITS UI BENEFITB13FF				.	12/14/20 4:59
330.88	Load	SCESC-UIBENEFITS UI BENEFITH0TKR				.	12/14/20 5:17
500.01	Spend	PALMETTO OVICH	South Congare	SC		6011	12/14/20 21:43
140	Spend	PALMETTO KMSAH	South Congare	SC		6011	12/14/20 21:42

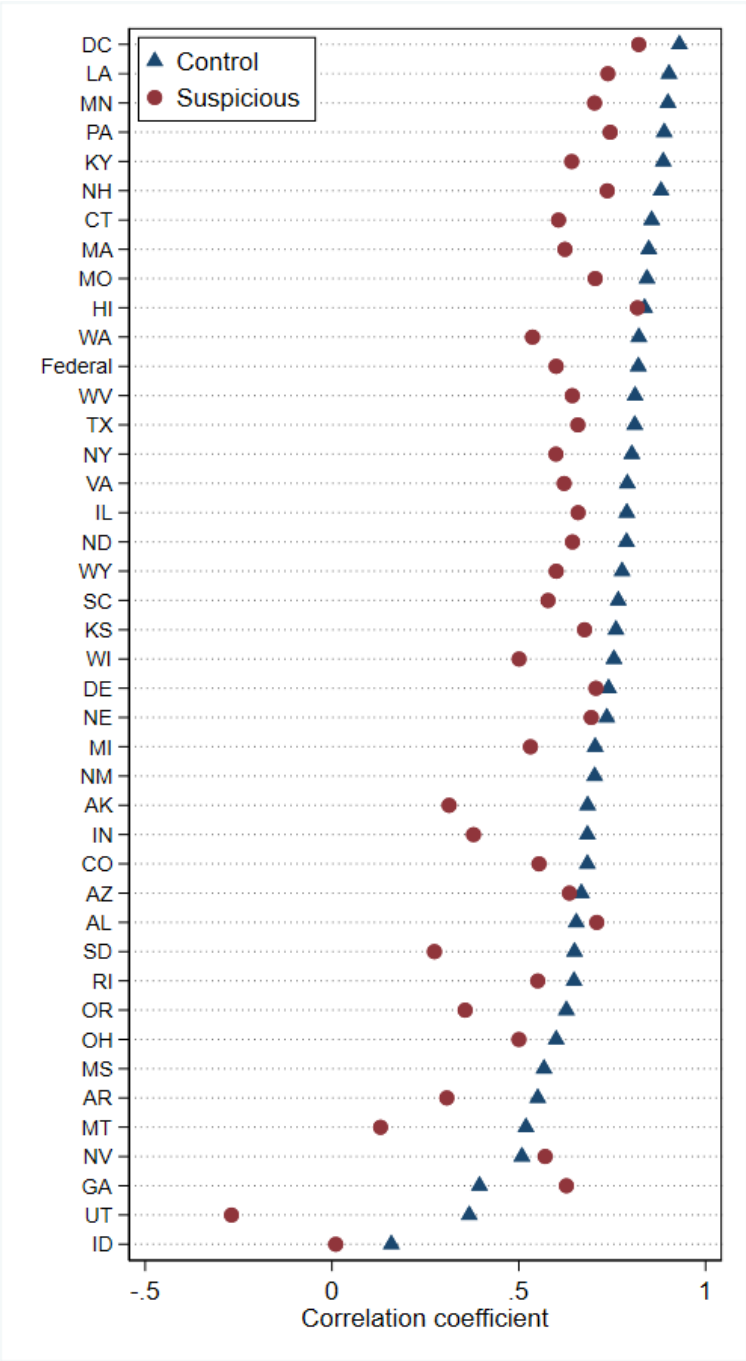
Notes: This figure illustrates suspicious transaction patterns using a one-month snapshot of a card that our algorithm categorizes as Suspicious. First column lists transaction amounts in dollars, second column indicates whether the transaction is a load or a spend, third column describes the transaction, fourth and fifth columns indicate the location and state where the transaction took place, second to the last column contains merchant category codes (MCCs), and the last column is a transaction timestamp. The card receives unemployment insurance benefits from the state of South Carolina and withdraws money from ATMs (MCC 6011) and bank branches (MCC 6010) soon after. This card receives a total of \$67,000 in unemployment insurance over a 12 month period and spends \$57,000 in cash withdrawals. Card ID and account numbers are omitted for confidentiality.

Figure A6: Geographic Concentration of Suspicious Cards



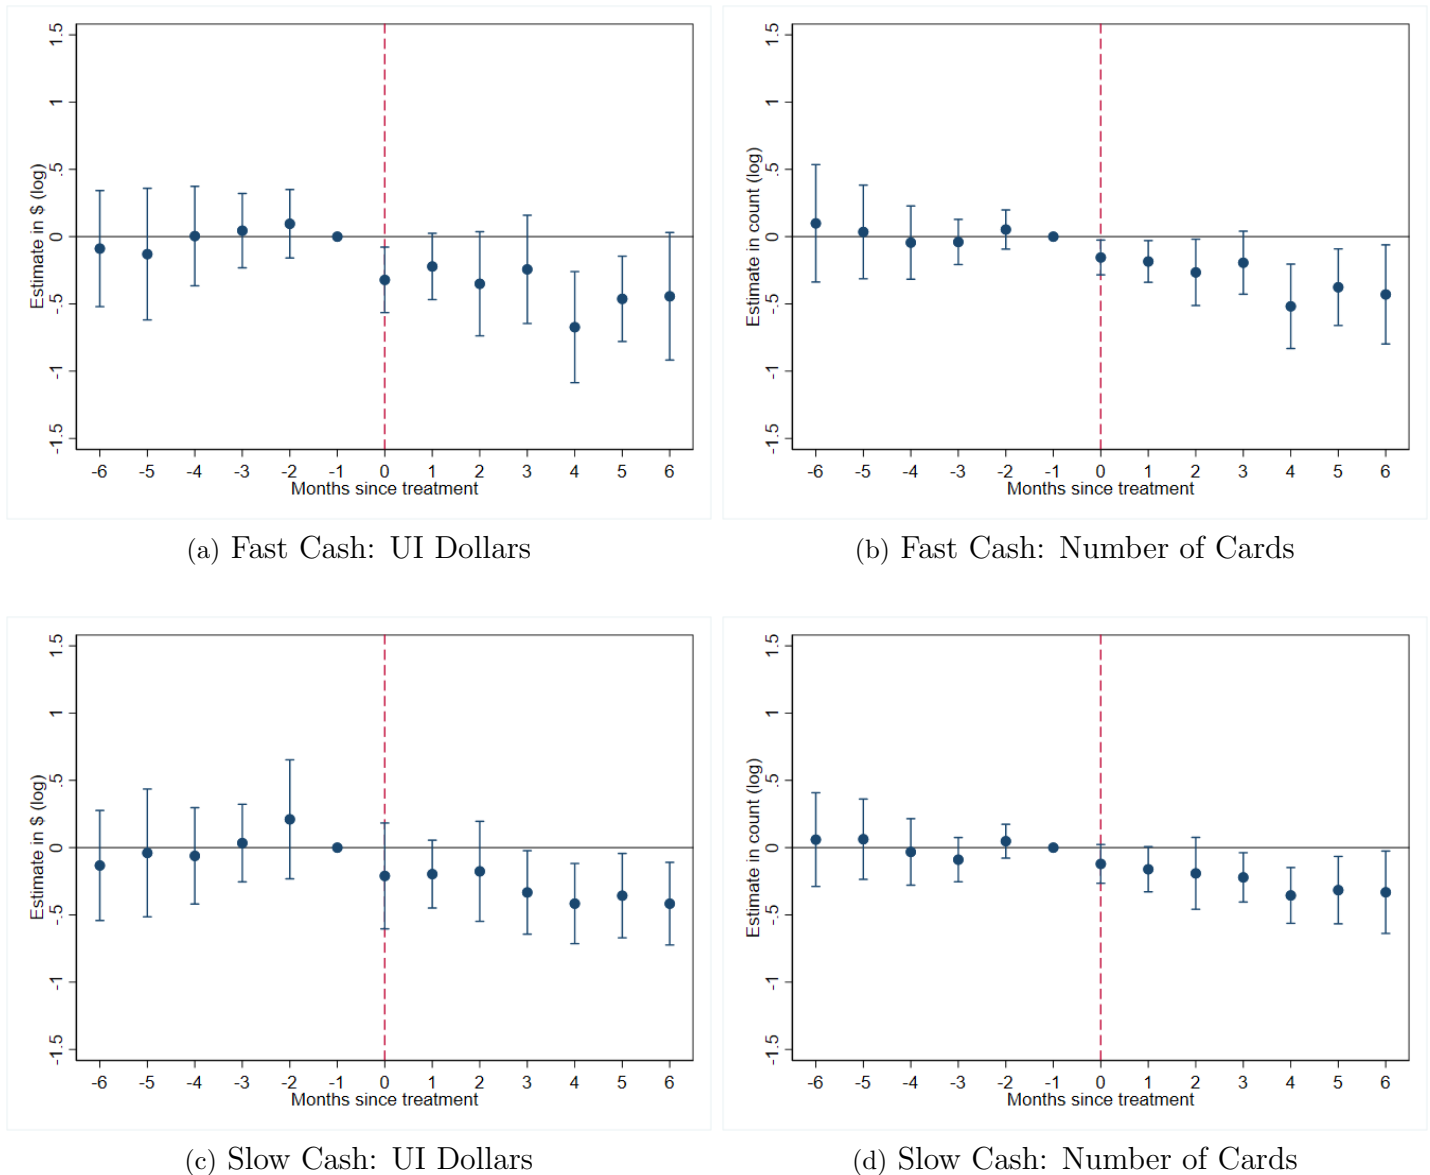
Notes: This figure shows the geographic distribution of cards belonging to the Suspicious cluster as a share of all the cards at a Metropolitan or Micropolitan Statistical Area (MSA) level. The location of each card is deduced from their ZIP codes. We find commonalities in the MSAs with higher concentration of suspicious cards under our method with areas that reported higher per capita identity thefts as per the [Federal Trade Commission](#). [Table A3](#) 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: Correlation between Unemployment Insurance and Nonfarm Payroll




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

Figure A8: Impact of Identity Verification on Sub-Categories of Suspicious Cluster




Notes: The figure plots estimates of β_τ from a model of the form in Equation 2 for the Suspicious (Fast Cash) cluster in panels (a) and (b), and Suspicious (Slow Cash) cluster in panels (c) and (d). These event studies plot monthly UI disbursements up to 6 months before and after the implementation of identity verification. Panels (a) and (c) consider (log) UI dollars as the dependent variable, and panels (b) and (d) consider (log) number of cards to whom UI is disbursed as the dependent variable. Error bars provide the 95% confidence interval for each estimate. We observe an immediate and persistent decline in UI disbursement to both the sub-categories of suspicious cluster following the introduction of identity verification.


Figure A9: Impact of Identity Verification on Other Non-Suspicious Clusters




(a) Concurrent Income: UI Dollars



(b) Concurrent Income: Number of Cards




(c) Discretionary Spending: UI Dollars




(d) Discretionary Spending: Number of Cards

Notes: This figure plots estimates of β_τ from a model of the form in Equation 2 for the Concurrent Income cluster in panels (a) and (b), and Discretionary Spending cluster in panels (c) and (d). Cards in these two clusters exhibit distinct income and spending patterns from the Control cluster, but they are not consistent with unemployment insurance fraud through identity theft. These event studies plot monthly UI disbursements up to 6 months before and after the implementation of identity verification. Panels (a) and (c) consider (log) UI dollars as the dependent variable, and panels (b) and (d) consider (log) number of cards to whom UI is disbursed as the dependent variable. Error bars provide the 95% confidence interval for each estimate.


Figure A10: Impact of Identity Verification (Alternative Estimation)




(a) Suspicious: UI Dollars



(b) Suspicious: Number of Cards




(c) Control: UI Dollars




(d) Control: Number of Cards

Notes: This figure plots estimates of β_τ from a model of the form in Equation 2, re-estimated for the Suspicious cluster in panels (a) and (b) and control cluster in panels (c) and (d) using two alternative methods. These event studies plot monthly UI disbursements up to 6 months before and after the implementation of identity verification based on the technique proposed in Callaway and Sant'Anna (2021) (in green) or “stacked” difference-in-differences proposed in Cengiz et al. (2019) (in blue). Panels (a) and (c) consider (log) UI dollars as the dependent variable, and panels (b) and (d) consider (log) number of cards to whom UI is disbursed as the dependent variable. Error bars provide the 95% confidence interval for each estimate.


Figure A11: Impact of Identity Verification (Pre-Treatment Clustering)




(a) Suspicious: UI Dollars



(b) Suspicious: Number of Cards




(c) Control: UI Dollars




(d) Control: Number of Cards

Notes: This figure shows robustness of our baseline results to considering only pre-treatment months in the clustering algorithm. The figure plots estimates of β_τ from a model of the form in Equation 2 for the Suspicious cluster in panels (a) and (b), and Control cluster in panels (c) and (d). Clusters are created using card features drawn from the pre-treatment part of our sample, with the same break-points applied to cards that start receiving UI only after the treatment month. These event studies plot monthly UI disbursements up to 6 months before and after the implementation of identity verification. Panels (a) and (c) consider (log) UI dollars as the dependent variable, and panels (b) and (d) consider (log) number of cards to whom UI is disbursed as the dependent variable. Error bars provide the 95% confidence interval for each estimate.

Figure A12: Impact of Identity Verification on State-wide Disbursement




(a) UI Dollars




(b) Number of Cards

Notes: This figure plots estimates of β_τ from a model of the form in Equation 2 for total state-wide UI disbursements i.e, to all clusters put together. These event studies plot monthly UI disbursements up to 6 months before and after the implementation of identity verification. Panel (a) considers (log) UI dollars as the dependent variable, and panel (b) considers (log) number of cards to whom UI is disbursed as the dependent variable. Error bars provide the 95% confidence interval for each estimate.

Figure A13: Impact of Identity Verification in Administrative Data



(a) Initial Claims



(b) Continued Claims

Notes: This figure plots estimates of β_τ from a model of the form in Equation 2 using administrative data sourced from U.S. Department of Labor, ETA (2024). Initial claims refer to the number of unemployment insurance claims made by applicants at the start of a new spell, while continuing claims refer to follow-on claims after a new spell is initiated. These event studies plot monthly UI claims up to 6 months before and after the implementation of identity verification. The dependent variables in plot (a) and (b) are log initial claims and log continued claims, respectively. Error bars provide the 95% confidence interval for each estimate.

Table A1: 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
Virginia (VA)	May-21	ID.me	Congress/Public
Wisconsin (WI)	Nov-20	Google Analytics	FOIA/Public
Untreated states			
Alabama (AL)	Hawaii (HI)	New Hampshire (NH)	Washington (WA)
Alaska (AK)	Illinois (IL)	New Mexico (NM)	West Virginia (WV)
Arkansas (AR)	Kentucky (KY)	Rhode Island (RI)	Wyoming (WY)
Connecticut (CT)	Michigan (MI)	South Dakota (SD)	
District of Columbia (DC)	Minnesota (MN)	Utah (UT)	

Notes: This table provides the treatment timeline for states where we can identify both unemployment insurance disbursements in the Facteus data and obtain the adoption date from FOIA responses, Congressional records, or publicly-available information. We use the earliest implementation date for states that adopted identity verification over a period of time for multiple programs or types of claimants. We consider states that adopted only pilot programs (e.g., KY and WA) as untreated.

Table A2: Disaggregation of Unemployment Insurance Data by State

	Correlation	Share in %		Correlation	Share in %
Federal (41 states)	0.94	0.40	Mississippi (MS)	0.72	0.01
Alaska (AK)	0.83	0.36	Montana (MT)	0.89	1.12
Alabama (AL)	0.79	1.11	North Dakota (ND)	0.86	0.45
Arkansas (AR)	0.83	1.02	Nebraska (NE)	0.81	0.60
Arizona (AZ)	0.74	0.08	New Hampshire (NH)	0.93	0.25
Colorado (CO)	0.91	0.30	New Mexico (NM)	0.83	0.14
Connecticut (CT)	0.89	0.45	Nevada (NV)	0.72	0.17
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	1.07
Georgia (GA)	0.72	0.05	Oregon (OR)	0.57	0.14
Hawaii (HI)	0.80	0.22	Pennsylvania (PA)	0.91	0.16
Idaho (ID)	0.79	0.43	Rhode Island (RI)	0.86	0.58
Illinois (IL)	0.84	0.44	South Carolina (SC)	0.88	1.19
Indiana (IN)	0.88	1.33	South Dakota (SD)	0.92	0.45
Kansas (KS)	0.87	0.81	Texas (TX)	0.92	0.31
Kentucky (KY)	0.87	0.92	Utah (UT)	0.89	0.36
Louisiana (LA)	0.89	0.83	Virginia (VA)	0.89	0.62
Massachusetts (MA)	0.93	0.23	Washington (WA)	0.89	0.27
Michigan (MI)	0.88	0.71	Wisconsin (WI)	0.80	0.82
Minnesota (MN)	0.86	0.64	West Virginia (WV)	0.93	0.50
Missouri (MO)	0.91	0.59	Wyoming (WY)	0.93	0.35

Notes: This table reports the representativeness of unemployment insurance observed in our data. The table shows the time-series correlation between the state-level unemployment insurance disbursements we observe in Factiveus data and the administrative UI data obtained from the US Department of Labor, for the period between January 2019 and September 2021. We also report the estimated share of UI captured in our debit card data with official disbursements made across all payment methods.

Table A3: Geographical Validation of Suspicious Cards with Identity Theft Reports

	Share of Suspicious Cards	
	(1)	(2)
Identity Theft Reports (per capita)	5.819** (2.494)	5.370** (2.546)
Number of Cards (thousands)		0.008* (0.005)
N	275	275
Adj. R^2	0.09	0.10
Unit of observation	MSA	MSA

Notes: This table shows the relationship between suspicious cards in our data and geographic areas that report a higher incidence of identity thefts. The table reports results from an ordinary least squares estimation for a model of the form in [Equation 3](#). The dependent variable is the share of suspicious cards out of all cards in our data at the level of Metropolitan or Micropolitan Statistical Area (MSA). We use card-level ZIP code to map each card to its MSA. The regressor of interest is the number of identity theft reports (per capita) in 2020 as per Federal Trade Commission data. Column (2) additionally controls for the number of cards in our data. Observations are weighted by the population of the MSA, which we obtain from the Census Bureau. Standard errors clustered by state are reported in parentheses. [Figure A6](#) shows the share of suspicious cards in each MSA in the form of a map. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A4: UI \$ and % Share to Clusters Before and After the Onset of COVID-19 Pandemic

Category	Pre-pandemic (Jan 2019 - Feb 2020)		Pandemic (Mar 2020 - Sep 2021)	
	\$ million	% of total	\$ million	% of total
Suspicious (Fast Cash)	0.05	0.3	51.17	6.2
Suspicious (Slow Cash)	0.07	0.4	42.31	5.1
Concurrent Income	0.29	1.7	15.78	1.9
Discretionary Spending	1.31	7.9	123.94	15.0
Control	14.83	89.7	592.4	71.8

Notes: This table shows how UI paid to suspicious cards changed after the onset of the COVID-19 pandemic. The table reports cluster-level UI (\$ million) and share of total UI paid before and starting March 2020, which separates the two periods. These data are aggregated across all 41 states in our sample.

Table A5: UI \$ and % Share to Clusters Under PUA and Non-PUA Programs

Category	PUA		Non-PUA	
	\$ million	% of total	\$ million	% of total
Suspicious (Fast Cash)	3.39	12.6%	8.48	8.3%
Suspicious (Slow Cash)	1.97	7.4%	5.07	5.0%
Concurrent Income	0.27	1.0%	1.68	1.6%
Discretionary Spending	2.09	7.8%	13.08	12.8%
Control	19.10	71.2%	73.89	72.3%

Notes: The table reports the cluster-level UI (\$ million) and share of total UI 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 A6: Unemployment Insurance Generosity by Treated States

	Maximum UI per week		Maximum number of weeks		PUA termination
	Pre-6 months	Post-6 months	Pre-6 months	Post-6 months	Months>treatment
Arizona (AZ)	240	240	26	26	11
Colorado (CO)	649	693	26	26	8
Delaware (DE)	400	400	26	26	3
Georgia (GA)	365	365	18	26	9
Idaho (ID)	458	463	26	21	6
Indiana (IN)	390	390	26	26	12
Kansas (KS)	503	540	26	26	7
Louisiana (LA)	247	247	26	26	2
Massachusetts (MA)	1,282	1,282	26	28	6
Mississippi (MS)	235	235	26	26	3
Missouri (MO)	320	320	20	20	3
Montana (MT)	572	598	28	28	7
Nebraska (NE)	451	456	26	26	6
Nevada (NV)	491	533	26	26	6
New York (NY)	504	504	26	26	7
North Dakota (ND)	640	652	26	26	6
Ohio (OH)	672	672	26	26	5
Oregon (OR)	683	733	26	26	6
Pennsylvania (PA)	584	591	26	26	11
South Carolina (SC)	326	326	20	20	3
Texas (TX)	528	535	26	26	7
Virginia (VA)	378	378	26	26	4
Wisconsin (WI)	370	370	26	26	10

Notes: This table confirms that states did not simultaneously reduce the generosity of their UI programs when introducing identity verification measures. The first four columns report the eligibility criteria for treated states under the regular state UI programs: the maximum weekly amount that recipients could claim at a weekly level, and the maximum number of consecutive weeks that recipients could remain on UI rolls. We present the means of both measures over the 6 months before and after the adoption of identity verification. We consider only regular state UI because the eligibility criteria for the pandemic unemployment assistance (PUA) were determined by the federal government. With the exception of Idaho’s reduction of maximum number of weeks, treated states did not reduce the generosity of UI payments simultaneously with the introduction of screening measures. Data are obtained from the US Department of Labor and can be accessed [here](#). The last column reports the number of months after treatment when each state ended participation in the PUA program (which had a federal ending date of September 2021). We obtain these dates from [Holzer et al. \(2024\)](#).

Table A7: Robustness to Two-Way Fixed Effects Specification

	Suspicious		Control	
	UI Dollars	Number of Cards	UI Dollars	Number of Cards
Panel A: Estimation using Callaway and Sant'Anna (2021)				
Treated \times Post	-0.345*** (0.110)	-0.227*** (0.081)	-0.107 (0.144)	-0.028 (0.156)
N	599	599	675	675
Panel B: Estimation using “stacked” difference-in-differences				
Treated \times Post	-0.301** (0.147)	-0.239* (0.124)	0.095 (0.227)	0.109 (0.240)
N	2,648	2,648	2,953	2,953
Panel C: Baseline TWFE on equally-weighted observations				
Treated \times Post	-0.267* (0.142)	-0.237* (0.119)	0.132 (0.165)	0.155 (0.175)
N	643	643	920	920
Panel D: Accounting for internal zeroes for all states				
Treated \times Post	0.314 (0.358)	-0.085 (0.150)	0.352 (0.386)	0.160 (0.241)
N	893	893	939	939
Panel E: Accounting for internal zeroes excluding states with majority zeroes				
Treated \times Post	-0.233* (0.136)	-0.227** (0.102)	0.394 (0.384)	0.160 (0.240)
N	551	551	926	926
Controls	Y	Y	Y	Y
State, Month FE	Y	Y	Y	Y

Notes: This table shows robustness estimates from our baseline two-way fixed effects model with alternative estimators. Panel A re-estimates [Equation 1](#) using the procedure in [Callaway and Sant'Anna \(2021\)](#). Panel B uses a “stacked” difference-in-differences procedure similar to [Cengiz et al. \(2019\)](#). Panel C reports β estimates using equally-weighted observations as opposed to population-weighted in the baseline. Panel D uses $\log(1+\text{UI Dollars})$ and $\log(1+\text{Number of Cards})$ as dependent variables to account for zeroes in the state-month panel, and Panel E does the same after excluding states where a majority of observations are zero for the respective cluster. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Table A8: Robustness to Clustering Choices

	Suspicious		Control	
	UI Dollars	Number of Cards	UI Dollars	Number of Cards
Panel A: Clustering using pre-treatment card features				
Treated \times Post	-0.613*** (0.159)	-0.420*** (0.140)	0.153 (0.181)	0.165 (0.198)
N	663	663	918	918
Panel B: Clustering using UI per month and longest spell of being on UI rolls				
Treated \times Post	-0.372** (0.154)	-0.268* (0.154)	0.032 (0.185)	0.085 (0.199)
N	562	562	912	912
Panel C: Specifying the number of clusters as the square of number of features				
Treated \times Post	-0.312** (0.122)	-0.213* (0.111)	0.048 (0.179)	0.074 (0.188)
N	572	572	920	920
Panel D: Considering cards that receive UI only from the states used in final analysis				
Treated \times Post	-0.310** (0.123)	-0.230** (0.107)	0.045 (0.185)	0.068 (0.192)
N	643	643	920	920
Controls	Y	Y	Y	Y
State, Month FE	Y	Y	Y	Y

Notes: This table presents our baseline two-way fixed effects model with alternative clustering choices. Panel A re-estimates Equation 1 using card-month features drawn only from pre-treatment sample period. Panel B uses UI per month and longest spell as alternative features to total card-level UI. Panel C shows results when we always specify the number of clusters as binomial outcomes. Panel D shows the estimates when we restrict clustering of cards to only those states that are considered in final analysis. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.