Are Crypto Anti-Money Laundering Policies Effective?[*]

John M. Griffin[†] Kevin Mei[‡] Zirui Wang[§]

June 2, 2025

Preliminary and Incomplete. Please do not circulate or post.

Abstract

In a rapidly changing crypto landscape, it is unclear as to whether anti-money laundering policies are accomplishing their goals or futile and perfunctory activities. We measure criminal responses to recent enforcement actions in crypto. First, the 2022 OFAC sanctions against the Tornado Cash mixer resulted in an 85% decline in monthly volume, 33 basis points in incremental obfuscation costs for remaining mixer outflows, and less interactions with both Western and less regulated exchanges. Hackers and related flows switched to swaps and bridges, which are more traceable. Second, sanctions against other OFAC addresses are uncommon but have lead to over $390 million stuck on-chain. Third, over $1.3 billion in Tether in 1,798 addresses is frozen on-chain, causing criminals to move their related account activities to more costly services. Fourth, offshore exchange flows from Tornado Cash only decline significantly after Binance and OKX reach settlements with the U.S. DOJ, indicating that multiple forms of enforcement is often necessary. Overall, our results indicate that sanctions have been effective seizing funds, moving funds to more traceable and seizeable places, and raising laundering costs. Nevertheless, sanctions, freezes, and bans are relatively rare as a fraction of total criminal activity, and overseas exchanges still handle substantial non-sanctioned criminal flows. Our findings indicate specific areas for improvements that can hopefully also help guide recent crypto policy and legislative proposals.

[*]We thank seminar participants at Ohio State University and UT Austin for helpful comments, including Zahi Ben-David, Cesare Fracassi, Sam Kruger, Tom Meling, Andrey Ordin, Aaron Pancost, Alex Pettyjohn, Amin Shams, and René Stulz. We thank Juan Antonio Artero Calvo for excellent research assistance. We additionally thank Integra FEC for use of their tracing tools and for substantial crypto-research support. Griffin is an owner of Integra FEC and Integra Research Group, which engage in financial consulting, research, and recovery on a variety of issues related to the investigation of financial fraud including crypto-related activities.

[†]McCombs School of Business, University of Texas at Austin. John.Griffin@utexas.edu

[‡]McCombs School of Business, University of Texas at Austin. KevinMei@utexas.edu

[§]McCombs School of Business, University of Texas at Austin. Zirui.Wang@utexas.edu

I. Introduction

Whereas the traditional banking sector has well-established anti-money laundering (AML) procedures, policies and protocols in the crypto arena are relatively nascent and in flux. There is a general lack of understanding regarding whether crypto money laundering policies are enforced, useful, and effective in accomplishing their goals. Recent changes in the regulatory and enforcement environment provide an empirical testing ground to examine the implementation and efficacy of different forms of anti-money laundering enforcement protocols and their impact.

Money laundering is the process of moving and concealing the origins of illicit proceeds to integrate them into the legitimate financial system and enable their unrestricted use. Anti-money laundering laws are deeply ingrained in the modern financial system as controls to reduce criminal capital flows. The theoretical motivation from the Becker (1968) crime model indicates that in order to deter crime one needs to increase the probability of being caught and the fine when punished. The goal of money laundering laws and enforcement is increase the probability of criminals and money launders being caught and the punishment to the extent that it results in the forfeiture and possible fine of funds (Ferwerda, 2009). Criminals will pay high transaction costs to launder funds in such a way as to avoid detection, which ultimately reduces the profitability of the crime. Money laundering laws and enforcement can be particularly important in settings such as terrorism and crimes committed by foreign nationals where authorities may not have the ability to apprehend criminals in noncooperation nations. In these situations money laundering laws may be the main means to deter such activities. After the terrorist attacks on September 11, 2001, the U.S. extended the existing anti-money laundering framework of the Bank Secrecy Act (Cuéllar, 2002) and increased cooperation across countries through new international initiatives from the Financial Action Task Force (FATF).

Evaluating the effectiveness of this framework is difficult because of the limited availability of reliable and detailed data. Chong and Lopez-De-Silanes (2015) outline two contrasting views on whether money laundering laws matter. The conventional view that money laundering laws and enforcement reduces crime is partially outlined in the previous paragraph. In contrast, money laundering skeptics

advocate that effort is best spent on targeting the root source of crime and not on regulating money services. Firms reputational risk may also serve as a sufficient deterrent for money-laundering activities. Cuéllar (2002) argues there is tenuous relationship between anti-money laundering regulations and reducing crime, despite the large costs of implementation, reduction in privacy, and negative externalities such as reducing access to the banking system. Pol (2020) argues that there is little outcome metrics to measure AML effectiveness, criminals keep up to 99.95 percent of criminal proceeds, and that the policies actually enable all forms of serious criminal activity. Levi, Reuter, and Halliday (2017) detail how there is little reliable data to evaluate AML efforts and that regulation is largely unguided by data analysis. Within the cryptocurrencies environment, we muster a large amount of detailed data on illicit funds that is typically private in banking.

Cryptocurrency's appeal seemingly lies in its potential alternative financial system free from regulation, but to the extent that it handles money, some organizations facilitating transactions must still conform to anti-money laundering laws. To interact with the global banking system, centralized crypto exchanges must authenticate the identities of new users through anti-money laundering and know your customer (AML/KYC) processes. Most exchanges purport to monitor transactions through know your transaction (KYT) policies to avoid receiving funds from known criminals. Since investors need trust in order to deploy capital and for meaningful investments to occur, the crypto ecosystem also has an incentive to root out bad actors as it seeks to be an alternative means of raising capital for legitimate entrepreneurial activity. However, monitoring can be costly, reputational risk may be less pertinent in the crypto world, and lax monitoring may lead to more transactions and fees for the crypto industry.

Furthermore, the regulatory and enforcement environment is evolving in response to both the perceived lack of clarity surrounding regulations for this nascent technology and the growing use of cryptocurrency in transnational crime and state-sponsored cyberattacks. On August 8, 2022, the U.S. Treasury Department Office of Foreign Assets Control (OFAC) sanctioned Tornado Cash in response to its use by hackers sponsored by the Democratic People's Republic of Korea (DPRK) and thus prohibited the US financial system from accepting related funds. A lawsuit challenging the sanctions, with

the financial backing of large players in the crypto industry, argued that such precedent would put undue responsibility on developers to prevent criminals from using their services.[1] The ban was lifted on November 26, 2024 when an appellate judge ruled that Tornado Cash's immutable smart contracts are not "property" and thus OFAC did not have authority to ban their use.[2,3] On a separate track, the U.S. Department of Justice (DOJ) prosecuted and fined large international crypto exchanges such as Binance (November 21, 2023) and OKX (February 24, 2025) for their lax anti-money laundering procedures. As part of their settlements, the exchanges agreed to more rigorous monitoring procedures and DOJ oversight. We use the enforcement of these earlier regulations and other shocks to empirically evaluate whether various forms of bans and enforcement are effective in influencing the behavior of criminals and compliance by exchanges.

In this paper, we empirically evaluate the effectiveness of crypto anti-money laundering policies and procedures. We first compile a large dataset of crypto addresses used in phishing attacks, romance scams, fake projects, contract exploits, impersonation, airdrops, fake returns, SIM swaps, hacks, and ransomware and use these addresses to map the money laundering ecosystem. We use tracing analysis on the Ethereum blockchain to gain a fuller understanding of how centralized and decentralized addresses interact with addresses before and after government sanctions. The traced network is then used to analyze the Tornado Cash ban, OFAC sanctions, stablecoin seizures, and Binance and OKX settlements.

Our findings show that Tornado Cash handles considerable criminal flows from more sophisticated cyber criminal gangs, with the Lazarus group from North Korea being the largest identified user. Despite the limitations of our data and the premise of attempting to detect funds purposefully hidden, in some months more than 50% of the funds entering the mixer are from a criminal actor. Before the ban, Western centralized exchanges received around 4% of Tornado Cash outflows on average, a proportion that declined substantially afterward. Flows to overseas centralized exchange also falls

[1] As reported by Reuters and discussed in a Coinbase blog post.

[2] More information can be found in this court ruling.

[3] When the Fifth Circuit ruled, Paul Grewal, Chief Legal Officer at Coinbase, remarked that: "Privacy wins. Today the Fifth Circuit held that U.S. Treasury's sanctions against Tornado Cash smart contracts are unlawful. This is a historic win for crypto and all who cares about defending liberty. Coinbase is proud to have helped lead this important challenge." More information can be found here.

after the 2023 ban but the decline was proportionally smaller. For funds that do reach exchanges post-ban, they take on average 0.17 additional hops and incur 33 basis points in incremental transaction costs from obfuscation techniques, compared to a pre-ban baseline of approximately 50 basis points as shown in a difference-in-difference design. Additionally, the incremental costs associated with Tornado Cash outflows to exchanges are higher, and the reduction in transaction volumes more pronounced, for Western exchanges relative to overseas exchanges.[4] DeFi swaps and bridges take money directly after leaving Tornado Cash, indicating that there is likely no concern for DEXs to reject their transaction.

Additionally, we analyze the flow of hackers over the pre-ban and banning period, and find that in a difference-in-difference design that their flows to Tornado Cash decrease by 30%. Rather than using other mixers they typically move money through DeFi swaps and bridges, outlets that typically obfuscate transactions for standard tracing tools, but are nevertheless traceable. Overall, the ban on Tornado Cash seems effective in that it increases transaction costs on criminal flows and hackers switch to traceable methods where it is at least potentially possible to freeze funds.

Another important enforcement mechanism is the freezing of criminal assets. We examine additions to the OFAC sanctions list and FBI watch lists. Of the 171 Ethereum-based addresses we evaluated, we find they held approximately $840 million at the time of sanctions. We trace $335 million that were laundered through mixers after sanctions were imposed but estimate at least $390 million of sanctioned funds remain seemingly stuck on the blockchain.

Perhaps more importantly, certain law enforcement agencies and courts have been successful at requesting Tether to freeze assets. In total, we find that $1.4 billion Tether and $80 million USDC have been frozen. Of this, $650 million are frozen in addresses that appear in the traced criminal network. We also analyze other related addresses to the criminal freeze and find that their share of flows to DeFi services increase by approximately 25%, presumably to obfuscate their flows after freezes. This indicates that in addition to costing the criminals their funds, seizures can impose additional costs on criminals by incentivizing affected entities to move to costlier and less regulated services.

[4] We use Western exchanges to refer to Coinbase, Crypto.com, Gemini, and Kraken because they are some of the largest exchanges that can be accessed from North America and Europe over the 2020 to 2025 sample period.

We then consider whether DOJ settlements with exchanges deters behavior. We first revisit the Tornado Cash sanctions and observe that outflows from Tornado Cash to overseas exchanges did not immediately decline following the sanctions. However, a reduction in sanctioned flows to Binance and Huobi becomes evident only after these exchanges reached settlement agreements during our sample period. We also look at all tainted deposit addresses at exchanges and consider whether the Binance and OKX shares of tainted flows decline after their settlements. In a difference-in-differences design, we find that Binance's tainted inflow declines 18% in the year after the announcement with much of the volume substituting instead to other exchanges like OKX and bridges. However, after the OKX settlement with the DOJ, tainted deposit addresses also see a sharp negative decline in inflows. Nevertheless, when we look at a broader set of criminal activity including scamming phishing attacks, romance scams, fake projects, contract exploits, impersonation, airdrops, fake returns, SIM swaps, hacks, and ransomware, we find that exchanges still handle considerable dirty money flows after the activity. Finally, we examine the recent Bybit hack by North Korea and find that the majority of flows use the Thorchain bridge to move to Bitcoin and are storing the funds in small quantities of Bitcoin.

Overall, it appears that sanctions, seizures, and fines have been effective tools for freezing funds, keeping funds trapped on-chain, moving funds to more transparent avenues, and increasing costs for criminals.

Our findings suggest practical areas for industry improvement, enforcement, and policy to deter money laundering. First, more aggressive seizures and bans seem warranted. The use of OFAC sanctioned addresses has been low beginning in 2023, and Tether seizures, though sizable, are a seemingly small part of overall criminal activity. Second, there is a wide range of practices by crypto exchanges in handling criminal flows, thus the reputational model to deter criminal activity does not seem to be effective for many crypto exchanges. Only after settling federal charges did overseas exchanges significantly decrease their flows in OFAC sanctioned addresses. Nevertheless, overseas exchanges still handle significant tainted crypto flows, indicating that their self-policing policies may be ineffective. Third, more attention should be paid to DeFi . Criminals do not appear to be concerned that these services will reject their transactions, and criminals are much less cautious with money coming out of

DeFi services, suggesting that they believe the money is viewed as clean money. Users who use DeFi pools as "investments" by providing liquidity to these automated market makers inadvertently facilitate cheap liquidity for money laundering. Fourth, many large centralized exchanges need to do more monitoring of funds and have more rigorous KYC and KYT procedures, particularly to detect dirty money coming through mixers and swap transactions. Otherwise, criminal flows can substitute money launder destinations by using exchanges with more lax AML regulations.

Lastly, we speak to ongoing policy debates. On April 7, 2025, a DOJ order by Todd Blanche stated that the DOJ "will no longer target virtual currency exchanges, mixing and tumbling services, and offline wallets for the acts of their end users or unwitting violations of regulations," but will instead hold accountable individuals who cause harm to digital asset investors or use digital assets for various forms of organized crime. Our results show that banning mixers is effective in pushing illicit financial flow to transparent places and empowers exchanges to monitor their flows. The DOJs goal of going after organized crime including foreign actors may be substantially impeded without targeting mixing and tumbling services. Additionally, the judge issuing the recent ruling in that OFAC did not have the authority to ban code noted "OFAC's concerns with illicit foreign actors laundering funds are undeniably legitimate. Perhaps Congress will update IEEPA, enacted during the Carter Administration, to target modern technologies like crypto-mixing software." Our findings suggest that congress should consider such an update. Additionally, a bill currently debated in Congress, known as the GENIUS Act, will require stablecoin issuers to comply with law enforcement requests for asset seizure.[5] Our analysis shows that seizers can be an effective tool and this condition for stablecoin issuers is important; otherwise, criminal flows may move to stablecoin issuers who do not allow freezes.

Our paper relates to two main literatures. First, in addition to the literature outlined above, we further contribute to the literature on crime and money laundering. El Siwi (2018) notes that recognizing "money is the lifeblood of organized crime" led to the adoption of the anti-money laundering (AML) regime in Italy. Mirenda, Mocetti, and Rizzica (2022) show how organized crime utilizes cash and shell companies to obfuscate transactions entering the banking system. Moore, Clayton, and Ander-

[5] https://www.congress.gov/crs-product/IN12553

son (2009) survey the economic structure of online crime and recommend more private data sharing and police enforcement focused on online gangs.[6] Levi (2015) surveys the literature on how organized crime is financed and notes that what is known has been primarily limited to prosecuted case records. Chong and Lopez-De-Silanes (2015) find that money laundering regulations are associated with lower levels of proxies for money laundering across countries. Fracassi and Lee (2025) examine cross-country differences in anti-money laundering laws and their effectiveness. Campbell-Verduyn (2018), Al-Tawil (2022), and Wronka (2023) overview money-laundering laws and procedures, the potential challenges of application to cryptocurrencies, and variation of policies across countries. Our paper focuses on understanding the efficacy of money-laundering methods in the crypto space where criminal activities can actually be partially measured and the regulatory landscape is rapidly evolving.

Second, there is a literature examining dark market activity in the crypto space. Meiklejohn, Pomarole, Jordan, Levchenko, McCoy, Voelker, and Savage (2013), Sokolov (2021), and Amiran, Jørgensen, and Rabetti (2022) examine the role of Bitcoin in the Silk Road (2011-2013), ransomware, and terrorism financing. Foley, Karlsen, and Putniņš (2019) find that 46% of non-exchange-related Bitcoin activity from January 3, 2009 to April 2017 is associated with darknet websites from 27 million Bitcoin users. Makarov and Schoar (2021) find only $5 billion in dark-market activities, Bitcoin mixers, and other criminal activities in 2020.[7] Griffin and Mei (2025) maps the flows of pig butchering scams by tracing the flows and showing how criminal flows are entering centralized exchanges and how these exchanges are allowing inducement payments to potential victims. Cong, Harvey, Rabetti, and Wu (2023b) show that 43 ransomware gangs carried out 2,690 attacks from May 2019 to July 2021. Cong, Grauer, Rabetti, and Updegrave (2023a) provide a useful overview as well as concrete examples of various crypto investment scams, Ponzi schemes, ransomware, money laundering, and dark markets.[8]. We

[6] Leukfeldt, Kleemans, Kruisbergen, and Roks (2019) find that technological knowledge for cybercrime in the Netherlands is often gained through a smaller set of technically skilled enablers in online marketplaces. Draca and Machin (2015) survey a growing literature on the economic incentives for crime. In terms of externalities of policies, Agca, Slutzky, and Zeume (2020) studies how anti-money laundering enforcement impacts lending by U.S. banks. Dirty money also distort macroeconomic capital allocations (Tanzi, 1996; Quirk, 1996).

[7] Chainalysis (2024) also provides a survey and examples of various types of criminal activity and calculates a total of $24.2 billion in 2023 through wallets directly identified with various identified for illegal activity though they note that their procedure undercounts because they seemingly do not count flows not other closely related addresses.

[8] This literature also fits within a larger literature of other types of nefarious trading activity in crypto, including price manipulation (Gandal, Hamrick, Moore, and Oberman, 2018; Griffin and Shams, 2020), pump-and-dump schemes (Li,

extend the literature by utilizing a comprehensive database of various types of criminal activity and examining if recent money-laundering laws are effective. This also leads to a fuller understanding of which laws might be effective and how criminals evade detection through a variety of techniques.

II. Data and Methodology

This paper develops a grouping of data sets into a unified framework for following illicit funds. We do this by collecting data on transaction flows in Bitcoin, Ethereum, and Tron. Transaction-level data are then enriched with attribution labels to assign addresses and flows to specific actors. Tracing methodologies further organize these transactions so that we can consistently follow specific funds that start at an illicit address to track their flows to subsequent destinations. This section describes the process of developing our datasets and the tracing methodology.

A. Data

We primarily use two types of data: blockchain transaction data and attribution data. Blockchain transaction data are sourced from the Bitcoin, Ethereum, and Tron blockchains. Importantly, the fields include the blockchain address of the sender and receiver, which allows us to construct paths of flows between addresses. Nevertheless, for most of our analysis we use information on the Ethereum network. This is where the Tornado Cash mixer resides and where many OFAC sanctioned addresses occur. We analyze the Bitcoin blockchain for an analysis of the Bybit hack as well as for additional robustness. When converting cryptocurrency value to dollars, we assume prices for the stablecoins Tether, USDC, and DAI are always $1. We use data from coingecko.com on end-of-day cryptocurrency prices to convert Ethereum, Wrapped Ethereum, and Bitcoin values to dollars.[9]

Beyond routine cryptocurrency transfers, transactions can also be invoked by specialized functions. We process data emitted from common functions, which allows us to follow funds that are swapped or bridged. Swapping is typically when an address uses a service such as a decentralized exchange to

Shin, and Wang, 2025; Hamrick, Rouhi, Mukherjee, Feder, Gandal, Moore, and Vasek, 2021; Phua, Sang, Wei, and Yu, 2022), insider trading (Félez-Viñas, Johnson, and Putnins, 2022), and wash trading (Pennec, Fiedler, and Ante, 2021; Cong et al., 2023b) as briefly surveyed by Griffin and Kruger (2024)

[9]This subset of currencies constitute the vast majority of cryptocurrencies we see used in our sample.

change one type of cryptocurrency to another on the same blockchain. Bridging is when an address deposits tokens into a service on one blockchain and receives that amount on another blockchain, net of any fees. We process data on swaps and bridges to further follow relevant funds. We supplement transaction data with prices based on the closing price of the transaction date.

We use a rich set of sources of attribute data. First, we use data from online sources like blockchain.com, etherscan.io, and tronscan.org to label addresses that belong to known entities. We focus on service providers such as centralized exchanges, decentralized exchanges, bridges, or mixers. Second, we used data on known illegal actors reported by various data collectors. These are mainly reported by victims or discovered by law enforcement or lawsuits. The single largest source of addresses is chainabuse.com, a leading reporting platform where victims and other users describe hacks and scams. After dropping reports of insufficient detail to categorize the type of scam, we use 227,745 reports. Table 1 presents summary statistics on these reported addresses, broken down by scam category and blockchain. We also received data on 12,554 suspicious addresses collected as part of an online publication about pig butchering scams from the United States Institute of Peace (USIP).[10]

Figure IA.1 summarizes the total number of unique addresses and total dollar inflow. We define each scam in Table IA.I. Overall, we find that addresses reported as part of a stolen funds have the largest total inflow, followed by pig butchering scams, illicit actors, and contract exploits.[11] Extortion had the most reported addresses, although only 6,000 unique active addresses. Nevertheless, it amounts to approximately $2 billion in activity. Interestingly, extortions are nearly exclusively in Bitcoin, presumably because this is the easiest crypto for non-crypto actors to access. In Table 2, we present summary statistics for each scam and averages. We count the flows into these addresses from January 2020 to November 2024. The total inflow to these addresses is a total of $51.8 billion with $22 billion from Bitcoin and $28 billion in Ethereum, and $1.2 billion in Tron. These beginning addresses should not be used to scope the total amount of activity since these incoming amounts do not

[10] We thank Jan Santiago (affiliated with PIDCO) and Raymond Hantho (Chainbrium) for sharing their data. This data was collected primarily from either interfacing with victims or probing scammer operations.

[11] In handling contract exploits, we are careful to only use the amount stolen and do not state the total amount handled. For example, if an attacker borrows a "flash loan" of $100 million to exploit a mispricing bug and steal $50 million, then we only state $50 million.

capture all the unreported scam addresses that are commonly gathered at collection points deeper in the networks. Further details about the distribution of and nature of these reports are relegated to the Internet Appendix.

B. Methodology

After identifying reported criminal transactions, we use blockchain data to construct a network of related addresses and analyze their characteristics in order to understand how the network responds to anti-money laundering regulations. In this subsection, we describe the methodology used to identify related addresses, with *tracing* serving as the primary tool and address and deposit address clustering also utilized. Later in the paper, we examine characteristics such as the transaction costs associated with using decentralized exchanges (DEXs), which will be introduced as they arise in the analyses.

B.1 Crypto Tracing

Tracing organizes transaction-level data into a framework that delineates the path of transaction flows of subsequent addresses. Tracing algorithms are an area of growing academic research (Anderson, Shumailov, and Ahmed, 2018; Möser and Narayanan, 2019; Tironsakkul, Maarek, Eross, and Just, 2022) and commonly used by law enforcement, through services providers like Chainalysis and TRM Labs, to follow capital flows. Most of our analysis is for the Ethereum blockchain which lacks the unspent transaction output clustering heuristic of Bitcoin, but nevertheless has been widely explored in tracing. We apply a suite of bulk tracing algorithms utilized and more fully described for pig butchering by Griffin and Mei (2025).[12] When tracing a given address, the first step is to collect all inflows. If outflows exist, then the tracer follows outflows to the next address. The goal is to follow tainted outflows to their end destination. Importantly, if tainted outflow funds are commingled in a downstream address that contains flows from other sources, then the forensic researcher must choose how to follow subsequent outflows.[13] We follow commingled funds on a "first-in-first-out" (FIFO) basis, a well-established and

[12] These tracing algorithms have been developed and maintained by Integra FEC. The algorithms are essentially a program with a series of step to tracing flows to stopping points such as centralized exchange hot wallets.

[13] For example, consider starting with a known criminal address at *address 0*. We follow their funds one hop to *address 1*. Suppose *address 1* also receives funds from another unknown source and it transfers all of its funds by splitting them between *address 2* and *address 3*. Both addresses are candidates for also having been tainted by the funds of *address 1*. A tracing methodology can delineate these based on the order of transactions. Therefore, tracing is more conservative than

accurate process for following specific fund transfers in cryptocurrency transaction-level data (Anderson, Shumailov, and Ahmed, 2018). By using tracing, we seek to only follow flows that are highly likely to be controlled by reported criminal addresses, instead of implicating all downstream addresses.

We trace funds to their end destinations. A trace path is terminated if it triggers any of the following criteria: (i) the flows enter an identified service, such as an exchange; (ii) the flows enter an unidentified but large address with more than 2,000 transactions; (iii) the path reaches five hops; or, (iv) traced amounts diminish to be less than 0.00001 ETH or token quantities, which are known as *dusting* transactions.[14] The most commonly triggered criteria are (i) and (ii). In the case of (i), we present results on the most common end destinations. By halting at exchanges, services, or unidentified large addresses, we avoid incorporating flows from other unrelated entities that these large addresses may have received. The tracing algorithm also handles cases where tainted flows re-enter into another reported address without double counting the ultimate destinations.

We *trace* the entire network of reported criminal flows as described in the Data subsection. This results in a network of paths from reported origins to their subsequent end destinations if they leave the blockchain. Additionally, for mixers, we trace flows that leave tainted services. For example, in the next section, we trace all Tornado Cash outflows. For services like Uniswap, Wrapped ETH, and bridge contracts, we only follow specific funds that are linked to traced transactions. Lastly, we *backtrace* or follow inflows to a tainted address back to their originating source.

B.2 Gas and Deposit Address Clustering

In addition to tracing, we use two clustering methods to extend our network to find addresses that are highly likely to be related to an address of interest: gas clustering and deposit address clustering heuristics. Gas clustering arises when two addresses may share the same "funding" address, or the first instance of receiving a small amount of Ether, the native current of Ethereum. The rationale is that all addresses need Ether because blockchain transaction costs can only be paid in Ether. Every new wallet needs Ethereum and we therefore associate the first funder as a way to link potentially

tainting all downstream addresses.

[14] Dusting transactions may obfuscate tracing by creating many more paths for a researcher to follow.

related addresses. This idea is incorporated in services like etherscan, the most popular service to view Ethereum transactions, where they display the first gas funder for every wallet. Deposit address clustering arises when two exchange deposit addresses receive funds from the same sender. The rationale is that deposit addresses are sensitive information like a bank account number and therefore if one sender transfers funds to two different deposit addresses, then likely that the two deposit addresses are linked, as discussed in Victor (2020). [15] Common gas funders and deposit address senders can then be used to link one address of interest to another in their nearby network. We drop any gas funders or senders that have more than 2,000 transactions and drop any contracts to avoid linking clusters through services. We use gas funding to define related addresses to serve as a "treated" group for stablecoin seizures and deposit address clustering to measure inflows to potentially tainted addresses after exchanges settle with the U.S. Department of Justice.

III. Tornado Cash Sanctions

Tornado Cash is the most well-known mixer on Ethereum and the most widely used money laundering service by illicit actors, particularly in hacks and smart contract exploits.[16] We begin by examining how other crypto actors interact with this service, focusing on flows into and out of Tornado Cash to better understand how both centralized and decentralized exchanges respond to transactions involving this well-known service. Further, on August 8, 2022, the U.S. Treasury Department sanctioned Tornado Cash, prohibiting the U.S. financial system from accepting related funds. This ban was lifted on November 26, 2024, when a U.S. judge ruled that smart contracts or computer code could not be sanctioned, as they do not constitute the property of a foreign national or entity. We use this regulatory timeline to study how crypto activity interfaces with Tornado Cash across two main periods: before the ban and during the ban. We also show some time-series analysis after the ban, though there is limited history and more limited activity. We begin with an overview of aggregate fund flows, then address three key questions: (1) Do user flows to Tornado Cash decline following the sanctions? (2)

[15] On most centralized exchanges, each deposit address is linked to a specific customer account, and users can often generate new deposit addresses at no cost. These addresses are sensitive because if tainted funds are traced to one, law enforcement can subpoena the exchange, which may then be legally obligated to disclose the customer's identity.

[16] Figure IA.9 shows destinations cross-tabulated by crime type.

Has criminal inflow to Tornado Cash decreased? (3) Are centralized exchanges effective in enforcing sanctions?

A. Overview of Flows

We first provide an overview of Tornado Cash flows. Panel A of Figure 1 shows the type of identifiable addresses that use Tornado Cash. To keep the exercise manageable, we sample the 2,500 largest nodes within the Tornado Cash network. To the left of the center Tornado Cash node, we notice that many of the addresses that remit funds to Tornado Cash are recognizable hacks and illicit organizations including the Lazarus Group. Among paths entering Tornado Cash, many are intertwined in more complicated networks, while outgoing paths are more discernibly clear. Users appear to funnel more money from decentralized exchanges (i.e., Uniswap, 1inch) than centralized exchanges. Further, Tornado Cash withdrawals are commonly funneled back to DeFi within one hop. Funds also move to centralized exchanges. If exchanges enforced sanctions, then we should expect this flow to cease after sanctions.

B. Do User Flows to Tornado Cash Decline Following the Sanctions?

Figure 2 shows the time series of all flows to Tornado Cash, along with totals traced from reported criminal flows. The August 2022 ban seems to have been effective in reducing volume to the mixer. Before the ban, Tornado Cash was handling more than $400 million per month, but after the ban by October 2022, volume is less than $50 million per month, or a more than 85% decline. Interestingly, the mixer becomes more attractive for criminals with the large volume as it is easier to plausibly deny that the outflows one receives are different from the input transactions. The volume in the mixer stays low until the ban is lifted. Nevertheless, the volume post-ban does not rise to pre-ban levels.

We also plot the time series of the tainted criminal address activity within the total activity. Despite the limited nature of our data and that we are missing criminal activity, the totals vary widely by month but at times show a sizable fraction of the flows due to criminal activity. For example, in April 2022, approximately 28% of the flows are due to traced criminal flows to Tornado Cash. In July 2022 just prior to the ban approximately 45% of the funds are due to criminal flows. These totals are high

in March 2024. The decline in late 2024 is likely more of a by-product of our data reporting since reported criminal addresses are gathered with a considerable lag as previously discussed. Nevertheless, the numbers are likely understated since criminal flows are coming from sources which either cannot or which we have not previously traced, including Tornado Cash itself, bridges, and the Wrapped Ether contract. When one examines the total flows into Tornado Cash as a fraction of the total flows for which there could be attribution, the percentage of the flows that is from criminal activity jumps considerably.

C. Has Criminal Inflow to Tornado Cash Declined?

We also find that most of the criminal inflows into Tornado Cash are due to technically sophisticated groups contract exploits (such as those who exploit features of smart contracts or inject code to gain access to wallets), stolen funds, or illicit actors. and not due to various sorts of scam activity.[17] We broadly label these these criminals as "hackers" and, in this subsection, examine if hackers change their behavior around the Tornado Cash sanctions.[18]

We use data on 5,867 unique hacker reports, identified from addresses labeled as contract exploits, stolen funds, or illicit actors. For each report in our dataset, we trace monthly flows both to Tornado Cash as well as to all other destinations on Ethereum. We construct a panel dataset where the unit of observation are the flows from each labeled hack, to each destinations, per month.[19] We plot the average flows to different destinations in Figure IA.3. In this figure, we observe that Tornado Cash and other services have similar inflows from hackers before June 2022. While one might expect any large drop to occur only after the August ban, Tornado Cash's inflows actually begin to fall below those of other services starting in June 2022, as seen in Figure 2. In May 2022, the U.S. Treasury sanctioned a mixer for the first time, Blender.io.[20] Shortly after, in June 24, 2022, a U.S.-based service named Har-

[17] As seen in Figure IA.9

[18] We choose to focus on all hackers because this group has shown more technological sophistication and have used Tornado Cash before the sanctions. Importantly, we cannot simply take a sample of each user that have a transaction history with Tornado Cash because Tornado Cash users have a tendency to rotate wallets and avoid re-using the same wallet in the future. Therefore, for many users, usage will mechanically decline after the first transaction.

[19] This is a balanced panel in that if the reported hacker address i does not send any flows to destination d in month t, then we encode this as zero. The alternative would be to have an unbalanced panel with missing data.

[20] https://home.treasury.gov/news/press-releases/jy0768

mony lost nearly $100 million after criminals injected malicious software into the company's services. The criminals then used Tornado Cash to hide their proceeds. Harmony claimed to be working with US law enforcement to recover their funds, and this theft of was stated as one of the reasons for the ban.[21] With these events unfolding, the ban on Tornado Cash may have been anticipated. Criminals may have feared that if they continued using the service, their funds could be trapped should sanctions be imposed.

To formally test the effect of the anticipated ban, we estimate a difference-in-differences (DiD) analysis on hacker flows. We define the traced flows going to Tornado Cash as the treatment group, while those going to all other services as the control group. Specifically, we estimate a regression of the form as follows:

$$\log(1 + TaintedFlows)_{i,d,t} = \sum_{t \neq June2022} \beta_t \times \mathbb{1}(Month = t) \times Tornado_d + \delta \times Tornado_d + \mu_i + \gamma_t + \varepsilon_{i,d,t}$$

where $\log(1 + TaintedFlows)_{i,d,t}$ is the natural logarithm of 1+ the dollar amount of flows traced from hacker report i to destination d in month t, $\mathbb{1}(\text{Month} = t)$ is an indicator for calendar month t, $Tornado_d$ is an indicator equal to one if the destination is Tornado Cash, and zero otherwise, and μ_i and γ_t are hacker report and month fixed effects, respectively. Figure 3 plots the estimated DiD coefficients (β_t), which capture how the difference in inflows between Tornado Cash and other destinations evolves in each month compared to the baseline month. We find no statistically significant difference in months leading up to June 2022. In July, right before the official sanction, usage of Tornado Cash drops sharply relative to other services. This is consistent with our prior that hackers reacted to the possibility of a ban and shifted funds elsewhere. After the ban in August, Tornado Cash usage falls further and remains approximately 40% lower for about a year.

[21] See details of the heist at this link and details of the ban at this link.

D. Are Centralized Exchanges Effective in Enforcing Sanctions?

After OFAC sanctioned Tornado Cash, money services businesses, including exchanges, were prohibited from processing transactions associated with it. To assess whether exchange users responded to these restrictions, we trace the flows exiting Tornado Cash to determine whether users avoided transferring those funds to centralized exchanges.

We identify and follow transaction paths from Tornado Cash to their eventual destinations. Figure 4 summarizes these paths using Sankey diagrams that visualize outflows from Tornado Cash before and after the August 8, 2022 sanction. Flows originate from Tornado Cash, pass through intermediate hops, and ultimately reach centralized exchanges in blue, decentralized exchanges in red, bridges in purple, or wrapped ETH in gray. When flows reach a decentralized exchange, we continue to follow the funds in the new cryptocurrency to identify their final destinations. Before the ban, we observe relatively large direct flows to centralized exchanges. After the ban, we observe a sizable shift, with fewer flows reaching centralized exchanges and more funds routed through DEXs and bridges.

Figure 5 compares these paths in time series. The bars in each sub-panel represent a monthly distribution, showing the share of Tornado Cash outflows that terminate in Western centralized exchanges (Panel A), overseas centralized exchanges (Panel B), and cross-chain bridges (Panel C). Throughout the paper, we include Coinbase, Crypto.com, Gemini, and Kraken as Western exchanges, primarily because these are the main exchanges that can be accessed from the US. All other exchanges are included as overseas exchanges. Blue bars represent flows that move directly from Tornado Cash to the destination, while yellow bars indicate flows that are first routed through a decentralized exchange before reaching the final destination. The red line shows the average number of hops in each path, and the blue line shows the average duration in days between the exit from Tornado Cash and arrival at the destination.

We find that flows to Western centralized exchanges decreased significantly after the ban. The average monthly share of outflows to Western CEXs declined from 3.99% in the twelve months before the ban to 2.14% in the twelve months after, a 46.20% reduction. This difference is statistically significant at the 1% level (t-statistics = 3.57, p-value = 0.0017). The decline is driven specifically by the

reduction in direct flows to Western centralized exchanges. The share of direct flows fell from 2.60% to 0.78% over the same period, a 69.80% reduction, which is statistically significant at the 0.1% level (t-statistics = 5.86, p-value < 0.001). The dollar-weighted share of paths that routed through a swap or other DEX also increased but is not statistically significant. Overseas centralized exchange flow also falls by the end of 2023, and bridges become the dominant destination for Tornado Cash flows. These paths also have more costly characteristics in that they require more hops and are more likely to use a DEX after the ban. Figure IA.5 presents transaction costs before and after the sanctions were imposed, split by Western and overseas exchanges. The transaction costs include transaction gas fees paid and costs from swaps.[22] While paths previously required an average of 50 basis points before the ban, the paths that entered Western exchanges averaged 1.71% or 171 basis points in transactions during the ban. However, costs have only increased from 38 to 66 basis points for overseas exchanges.

Table 3 formally tests whether transfer paths to Western centralized exchanges became more complex and costly after the Tornado Cash ban, using a difference-in-differences regression that compares Western exchanges (treatment group) and overseas exchanges (control group). Western CEXs are treated because the OFAC sanction is a U.S. action, and compliance is expected to be more strictly enforced by exchanges with U.S. regulatory exposure. The regressions are estimated at the path-exit level, where each path represents a sequence of transfers originating from a single withdrawal from Tornado Cash that passes through one or more intermediate hops. Each path can have multiple exits, where each exit is a distinct cash-out event to a CEX and is counted separately. We find that the dollar-weighted average path that terminated in domestic exchanges used 0.29 more hops, required about 120 more days, and incurred between 77-87 basis points in additional transaction cost. Users, therefore, are increasingly unwilling to take funds from Tornado Cash to Western centralized exchanges. Those that do incur almost twice as much cost compared to transaction paths before the ban.

More broadly, we also test whether paths from Tornado Cash incurred more cost to reach centralized exchanges, regardless of jurisdiction. We compare characteristics of paths leaving Tornado Cash and entering centralized exchanges to all tainted paths from the traced network to centralized

[22] For Tornado Cash paths, we follow swapped funds through Uniswap, 1inch, 0x, fixedfloat, paraswap, and curve.fi, which accounts for 85% of the Tornado Cash outflows to DEXs.

exchanges. The results presented in Table 4 shows that Tornado Cash outflows incurred 0.12 more average hops, required 130 more days, and 33 basis points in additional cost compared to other tainted paths after the sanctions were imposed.

In summary, the study of Tornado Cash shows how various crypto players interact with a service known to handle dirty money. We briefly note that Tornado Cash handles considerable criminal flows from more sophisticated cyber-criminal gangs, and we explore this in more detail in Section 7. The main takeaway is that the Tornado Cash inflow has declined after the sanctions. A large share of the remaining flows to Tornado Cash are associated with reported crimes. However, when considering flows of reported hackers, we find that criminal flows to Tornado Cash have declined. Tornado Cash outflows increasingly do not enter centralized exchanges in a straightforward manner and those that do incur greater cost, likely in an effort to obfuscate the source of capital. Overall, the ban appears to have reduced the effectiveness of the mixer.

IV. OFAC Sanctions

The U.S. Treasury Department's Office of Foreign Assets Control (OFAC) began listing digital currency addresses on its sanctions list in 2018. Initially, these sanctions targeted individuals linked to state-sponsored cyberattacks, but over time, the list has expanded to include cryptocurrency exchanges and services associated with illicit activity. Once an address is sanctioned, U.S.-based entities are legally prohibited from transacting with it. To study the effects of these sanctions, we compile a dataset of cryptocurrency addresses designated by OFAC, supplemented with addresses published by the Federal Bureau of Investigation (FBI), resulting in a total of 546 addresses.[23] We focus on Ethereum-based addresses, excluding 59 addresses related to Tornado Cash, yielding 171 addresses for analysis. In this section, we evaluate whether these sanctions are effective at inhibiting the flows of these designated addresses.

[23] An example of FBI-designated addresses can be found in this press release.

A. Are OFAC Sanctioned Addresses Able to Convert to Fiat?

In Figure 6, we plot the balances of sanctioned addresses and the destinations of their trace outflows. Panel A shows the time-series of the number of sanctioned addresses and balance of sanctioned addresses before sanctions are imposed in gray, and after sanctions are imposed in red. We exclude Tornado Cash because we analyzed this in more detail in the prior section. There are only 28 days were sanctions were introduced totaling to 171 addresses, a relatively low number of addresses compared to the over 200,00 addresses reported for scams, phishing, ransomware, and other cyber crime. Most of the addresses are clustered in a few key days, which correspond to events such as when the FBI releases a network of addresses linked to specific hacks. In gray, we plot the sum of balances over time, where we observe that the highest balance of all 171 addresses peaked in March 2022 with $538 million. In red, we plot the sum of balances of addresses after they have been sanctions. We can see that these entities easily move funds out of their addresses, with only $37 million still held in sanctioned addresses. Thus, sanctions are frequently imposed only after addresses attain their maximum observed balance.

Across all addresses, summing a cross-sectional snapshot of the day prior to sanctions, these addresses in total held $840 million in current balances. However, the median account at the time of sanction only held $77,000. Altogether, many of these addresses carry relatively small balances and are never used against after sanctions are imposed. However, for some large addresses, more than $300 million of cryptocurrency has flowed out of sanctioned addresses to other addresses in their deeper network.

Panel B plots outflows from sanctioned addresses to destinations indexed by time to sanction events. We find that prior to sanctions, these addresses have established accounts as a few centralized exchanges in blue. After being sanctioned, transactions to centralized exchanges total only $6 million. We additionally trace $53 million to decentralized exchanges that seem to provide services to sanctioned flows. We follow $110 million to large unidentified addresses where it is unclear if these funds found an ultimate end destination to convert to fiat.[24] The largest single destination are other

[24] Large unidentified wallets are large entities where we lack data to attribute the address to a known exchange, but we do not continue tracing out of conservatism.

mixers, such as Tornado Cash, where $334 million has been laundered after the funds were already sanctioned, or 40% of the $840 million in sanctioned funds. We calculate the remaining funds on chain as between $390-500 million (46-60% of $800 million), as determined by the funds that did not exit to a centralized exchange nor a mixer.[25]

In summary, sanctions appear effective. Of those that are sanctioned, between 46-60% of funds remains on-chain and thus effectively frozen if no centralized exchange willingly accepts those funds. However, sanctions on cryptocurrency addresses are relatively rare, with only 28 instances between 2020-2025. Mixers are the most popular destination to launder funds after an address has been sanctioned.

V. Freezing Stablecoins

The risk of asset seizure is a critical consideration in money laundering schemes. Stablecoin issuers like Tether (USDT) and Circle (USDC) have the technical ability to freeze these assets and prohibit specific users from future transactions with that stablecoin. However, stablecoin issuers have historically not been held to the standards of traditional banks and "the absence of a regulated financial institution, subject to AML/CFT obligations can limit authorities' collection of and access to information. It can also reduce the effectiveness of preventive measures" (US Department of the Treasury, 2024). In this section, we consider how illicit flows respond when the expected probability of asset seizure may increases due to related freezes.

Figure 7 presents data on 1,798 instances where Ethereum addresses were prohibited from future Tether interactions. In total, almost $1.4 billion in Tether is held in frozen addresses.[26] Of the $1.4 billion frozen, we find that approximately $600 million also appears in the traced network of reported criminal flows. In Table 5, we present statistics on the 396 frozen addresses that also appear in this network. As also seen in Panel A, these addresses are most often found in the network of pig butchering,

[25] $840 million subtracted by $334 million to mixers and $6 million to centralized exchanges yields $500 million. An additional $110 million may have exited if these funds were sent to a shadow exchange, but data on these addresses does not exist.

[26] We drop instances where an address was added to the list of frozen assets and then later removed. To be consistent with the tracing framework that conservatively excludes large addresses, we only show addresses that are frozen with less than 2,000 transactions. This ensures that we do not trace paths associated with shadow exchanges.

scams, impersonation, and phishing addresses. This may be indicative that stablecoin issuers are more likely to respond to freeze requests from law enforcement in investigations tight to scams with larger average losses. In the last row, we denote the aggregate overlap and consider the implications. The 396 frozen addresses are downstream from 3,179 reported criminal origins. In total these origins have transferred $4.9 billion in total outflow. Of the original capital leaving these addresses, $555 million are received by the addresses that have been frozen, compared to the $603 million frozen.[27] Therefore, one back-of-the-envelope calculation is that, conditional on a network being correctly targeted for seizure, only 12% of capital ($600 million out of $4.9 billion) in these addresses are currently seized. In total, we find that $1.4 billion Tether and $80 million USDC have been frozen. Of this, $650 million are frozen in addresses that appear in the traced network.

In Panel B of Figure 7, we plot the activity of the addresses leading up to the asset seizure. These seizure have a wide variety in number of days active and balance at the time of being frozen. We see that many were active for years prior to being frozen. However, most of them cease activity upon seizure, with a few exceptions using USDC and Ether. In the next subsections, we consider how related addresses react when shocked by a plausibly random seizure, and if higher risk leads to higher cost.

A. How Do Related Addresses React to Asset Freezes?

Figure 8 plots flows in the hours immediately before and after the freeze. We find that affected addresses quickly swapped their remaining stablecoins into other assets within the next 24 hours. Frozen addresses transferred out their USDC and Ether, in part because they were unable to move any frozen Tether. We then examine addresses related to the frozen ones through gas clustering, as described in the methodology section. Panel B shows that these related addresses also rapidly moved USDC and DAI shortly after the initial asset seizure, including after swapping unfrozen Tether into these currencies.[28] We formally test this behavior using a difference-in-differences framework. The treated group consists of the frozen addresses and their related counterparts, while for every treated address, the

[27] Interestingly, $109 million of blacklisted funds are destroyed out of the total frozen. The term "blacklist" is used because that is the name of asset seizure function in the Tether contract.

[28] While Tether and USDC can be frozen by their respective stablecoin issuers, the DAI smart contract does not have the functionality to restrict future usage. Notably, the issuer of DAI has recently shifted its focus to a new stablecoin with built-in seizure capabilities.

control group is a random sample of 20 addresses that also received Tether in the preceding three days. The estimated coefficients are plotted in Figure IA.7, where the outcome variables are transaction value and transaction count. Table IA.II tabulates results and finds that the related addresses transferred an average 450% more dollar flow in 1.6 greater number of transactions in the 24 hours after a freeze than the control group, indicating that the market participants are concerned that additional funds may be frozen.

Next, we further investigate whether treated addresses shift their activities toward DeFi services after their Tether balances are frozen. To formally test it, we employ a difference-in-differences design at the group-cohort-month level. Each cohort represents one seizure event and consist of a treatment group and a control group. The treated group consists of the frozen address and its related addresses identified via gas clustering. For every freeze event, the control group is a random sample of 20 addresses that received inflows of at least $100 within the seven days prior to the freeze. For each treatment and control group within every cohort, we define the dependent variable *DeFi Share* as follows:

$$\textit{DeFi Share}_{g,c,t} = \frac{\sum_{i \in g} DeFi\ flows_{i,g,c,t}}{\sum_{i \in g} All\ flows_{i,g,c,t}},$$

where $DeFi\ flows_{i,g,c,t}$ denotes the dollar amount transferred to DeFi services by address i in treatment or control group g in cohort c in month t. Then, we run a regression of the form:

$$\textit{DeFi Share}_{g,c,t} = \sum_{t \neq t_{freeze}} \beta_t \times \mathbb{1}(Month = t) \times Treat_g + \mu_c + \gamma_t + \varepsilon_{g,c,t}$$

where μ_c are cohort fixed effects, γ_t are month fixed effects. Figure 9 plots the coefficients by event time and Table 6 tabulates the coefficients. Overall, we find the share of flows to DeFi services by frozen addresses and those linked to them increases by approximately 25% following Tether's freezing of their assets. This is evidence in support of the idea that asset seizures as part of anti-money laundering enforcement can incentivize affected entities to turn to costlier and less regulated services, such as DeFi protocols, potentially to obfuscate their flows.

B. What Money Laundering Patterns are Associated with Asset Seizure?

After observing that asset seizure leads related addresses to use costlier obfuscation services, a natural economic question is whether greater obfuscation efforts are associated with a lower probability of seizure. In Figure 10, we plot the dollar-weighted probability that Tether freezes an address in our traced network. We sort paths into quintiles based on two metrics: total transaction cost and the duration of the path. Total transaction cost includes the transaction cost fees paid at each hop and sums the total fees paid for each hop on the path. It also calculates the spread lost from swaps where one cryptocurrency is converted into another cryptocurrency.[29] The spread is calculated based on the dollar value of crypto of input compared to the amount received as output received. Duration of the path is calculated as the time leaving an origin compared to the time when funds were received at a centralized exchange.

We find that the paths most likely to contain a frozen address fall into the highest quintiles of both transaction cost and path duration, with a freeze probability of 7.8%. In contrast, paths in the lowest cost and shortest duration quintiles have only a 0.4% chance of containing a frozen address. One must be careful in evaluating this figure because it reflects only a correlation and not a causal relationship. A key endogeneity concern is that law enforcement may focus more heavily on sophisticated actors who have obtained larger illicit proceeds. These entities may delay their actions because they are aware that depositing large sums too quickly could attract attention. As a result, they may be unable to pool with the faster, low-cost transaction paths used by smaller actors, and instead must store funds on-chain for longer periods, increasing their exposure to detection.

VI. U.S. Settlements with Exchanges

The most critical money laundering defense are services that allow on-chain funds to be converted into fiat currency and reintegrated into the traditional financial system. For most users, exchanges offer the deepest liquidity and greatest breadth of features for offboarding on-chain funds. Exchanges are also

[29] For reported address flows, we follow swapped funds through Uniswap, 1inch, Tokenlon and curve.fi, which accounts for 92% of the DEX activity.

the best positioned player with rails to the traditional financial system to correctly perform know-your-transaction monitoring. However, some exchanges have historically maintained weak compliance processes, which create opportunities for money launderers to convert crypto into fiat undetected.

On November 21, 2023, Binance pleaded guilty to anti-money laundering and sanctions violations as part of a settlement with the U.S. Department of Justice.[30] As part of the agreement, the company's founder and CEO resigned, and Binance paid over $4 billion in penalties. At the time, Attorney General Merrick Garland stated, "Binance became the world's largest cryptocurrency exchange in part because of the crimes it committed. Now it is paying one of the largest corporate penalties in U.S. history." On May 17, 2024, the DOJ appointed two independent compliance monitors to oversee Binance for a three-year term. Similarly, on February 24, 2025, OKX pleaded guilty to violating U.S. anti-money laundering laws.[31] We use these announcements to evaluate whether customer flows to these exchanges changed following the pleas.[32] We focus on changes around the settlement of two tainted flows: Tornado Cash and victim-reported illicit flows.

A. Does Tornado Flows decrease after Exchange Settlements?

A sharp test of Binance's sanctions compliance is whether it continued to receive flows from Tornado Cash after the settlement. Figure 11 presents data in the same format as Figure 5, but focuses specifically on flows to Binance compared to all other overseas exchanges around November 2023. The blue and yellow bars plot the monthly share of Tornado Cash outflows reaching Binance and all other overseas exchanges, with yellow denoting paths that swapped through a DEX. In 2023 and before the DOJ settlement with Binance, Binance received an average of 3.98% of all Tornado Cash outflows to overseas centralized exchanges. In the months after the DOJ's settlement with Binance and before the Tornado Cash ban was lifted in November 2024, this share declined sharply to 0.75%, representing an 81.29%

[30] https://www.justice.gov/archives/opa/pr/binance-and-ceo-plead-guilty-federal-charges-4b-resolution

[31] https://www.justice.gov/usao-sdny/pr/okx-pleads-guilty-violating-us-anti-money-laundering-laws-and-agrees-pay-penalties

[32] Customers may choose not to remit funds to these exchanges either because of the announcement effect, or due to direct tighter post-plea compliance measures that deter future transactions.

reduction, which is statistically significant at 0.1% level (t-statistic = 5.61). The decline is even larger when measured after the start of the compliance monitorship in June 2024, falling further to 0.54%. In contrast, flows to other overseas exchanges also declined but less dramatically, from 11.80% before the settlement to 7.64% post-settlement. This is primarily due to Tornado Cash users redirecting their funds toward bridges at the same period as shown in Figure 5.

Additionally, the red line in Panel A shows that addresses that continued to reach Binance did so with more intermediate hops after the monitorship begins. We formally test this pattern using a dynamic difference-in-differences framework, with estimated monthly coefficients plotted in Panel B. The outcome variable is the number of intermediate hops for transfer paths from Tornado Cash to centralized exchange destinations. The treatment group consists of flow paths to Binance, and the control group consists of flow paths to other overseas centralized exchanges. The estimates show no significant change in the number of hops immediately after the settlement, but a clear increase following the start of the compliance monitorship. This pattern suggests that users who continued sending Tornado Cash flows to Binance began adopting more complex routing paths to obscure the origin of funds, and that users became more hesitant to send funds to Binance. Notably, the number of hops declined after the Tornado Cash ban was lifted, consistent with reduced incentives for concealment once the sanction was removed.

Overall, the fact that Tornado Cash to overseas exchanges did not decline sharply after Tornado Cash sanctions but only after the exchanges settled with the DOJ indicates that the crypto industry does not always self-police and that regulatory enforcement can be effective at ensuring compliance.

B. Do Other Criminal Flows decrease after Exchange Settlements?

As discussed, Tornado Cash and OFAC sanctioned flows are a relatively small set of activity. We next consider whether the share of tainted flows from various forms of criminal activity appears to shift from Binance or OKX to other exchanges after their respective announcements. For this section, we use deposit address clustering to find related addresses that remain active over different time periods

than those directly appear in our traced network, as discussed in the methodology section.[33]

To investigate this shift, Panel A of Figure 12 compares monthly inflows to tainted deposit addresses at Binance (solid line) and at other exchanges (dashed line), with vertical red dashed lines indicating the timing of the DOJ settlement and the start of the compliance monitorship. It shows that while inflows to tainted addresses are rising for both Binance and other exchanges over time, the increase is notably sharper for other exchanges after the settlement, indicating a relative decline in Binance's share of tainted inflows.

To formally test for a differential shift, we estimate a difference-in-differences (DID) regression at the deposit address-month level. Specifically, the regression is of the form:

$$\log\left(1 + Total\ Inflow_{i,t}\right) = \sum_{t \neq t_{Nov2023}} \beta_t \times \mathbb{1}(Month = t) \times Treat_i + \mu_i + \gamma_e + \eta_t + \varepsilon_{i,t}$$

where $\log\left(1 + Total\ Inflow_{i,t}\right)$ denotes the log of one plus the total tainted inflow received by deposit address i in exchange e in month t. We include μ_i for deposit address fixed effects, γ_e for exchange fixed effects, and η_t for month fixed effects. The treatment group consists of tainted Binance deposit addresses, while the control group is all tainted deposit addresses at other exchanges. Panel B plots differences-in-difference coefficients. While there is no immediate impact in the months just after the settlement, we find that the inflow to tainted Binance deposit addresses begins to fall relative to controls starting in early 2024 and continues to decline after the monitorship is announced. Table 7 reports regression results for both the Binance and OKX samples using the same difference-in-differences specification. Column (1) presents estimates for the Binance sample, where the *post* period is defined as months after November 2023 (following the DOJ settlement). The results show that inflows to tainted Binance deposit addresses declined by 18.4% relative to other exchanges after the settlement, consistent with the visual pattern in Figure 12.

We also examine the effects of the OKX settlement, comparing tainted flows to OKX deposit addresses versus other exchanges around its settlement date, using the same difference-in-differences

[33] Another important challenge is that the network of tainted flows suffers from a sample bias such that tainted addresses are reported with a lag. However, we can overcome this bias by comparing the magnitude of tainted paths to Binance and tainted paths to other exchanges in the same month, assuming that the reporting lag is independent of their downstream exchange destination.

design. Column (2) of Table 7 reports the regression results, and Figure IA.8 visualizes the corresponding time series patterns. The results show a decline in tainted flows to OKX as well. However, given the recency of the announcement, the aggregate time series bears some sample bias as discussed earlier. The difference-in-difference also suggests a decline flows relative to all tainted flows, but more data will be needed to see the full effect. Overall, this suggests that tainted flows to both Binance and OKX have modestly declined some after their respective settlements. Nevertheless, these exchanges still handle large fractions and amounts of criminal flows indicating that the bulk of their efforts is focused on the official OFAC flows.

VII. Aggregate Trends

The prior sections explore an arsenal of anti-money laundering regulations used to deter illicit financial flows between 2020 and 2025. These actions have increased the cost of money laundering and may have heightened expectations around the probability of asset seizure. This section explores: what substitutes are available to users as the costs of money laundering increase?

Having established that hackers reduced Tornado Cash usage, we next investigate where funds switched. Figure 13 compares two cohorts: (i) hackers active before the ban, and (ii) those who started moving the funds after the ban. We plot the traced dollar flows of each group to various destinations, highlighting sizable shifts in blue and entirely new top destinations in red. The pre-ban group relied heavily on Tornado Cash. By contrast, the post-ban group uses Tornado Cash at lower rates and shifts to other major services. We see that Wrapped Ether experiences the largest jump. Hackers appear to convert their stolen Ether into Wrapped Ether, possibly as an intermediary step before onward transfers. Among the newly popular destinations, Thorchain stands out the most. Thorchain is a bridging protocol that enables cross-chain transfers without going through centralized intermediaries. Hackers likely hope that moving funds across different blockchains via bridges will obscure the trail and make their activities more difficult to trace. Our results indicate that while the ban on Tornado Cash reduced its usage, hackers adapt quickly and adopt other mixing or bridging services once any single laundering route becomes riskier.

North Korean hackers are an emblematic example that encapsulates how the environment has evolved. Not only were North Korea prolific users of Tornado Cash before the software was sanctioned, but a high-profile North Korea hack was also the precipitating event that led to the 2022 sanctions.[34] In February 2025, North Korea is alleged to have stolen $1.4 billion from the exchange Bybit, marking the largest crypto heist in history. Figure 14 plots these flows, with stolen funds initially exiting Bybit on Ethereum and being forwarded by hackers through Ethereum-based wallets and services, represented in the center. Within Ethereum, we see $58 million sent to OKX, but the majority of funds, or $967 million, moved swiftly through Thorchain to addresses on the Bitcoin blockchain, likely to minimize seizure risk. We calculate the transaction costs of this operation: transfers on Ethereum cost $12 thousand, using Thorchain to bridge to Bitcoin incurred $3.79 milion (42 basis points), and subsequent Bitcoin transfers added another $67 thousand in miner fees, for a combined $3.87 million. Notably, North Korea seems to have substituted from Tornado Cash, a service with relatively low cost (30 basis points), to bridges, which are more costly. The excessive splitting into over 34,000 transactions was likely an attempt to cheaply obfuscate their flows. Additionally, speed appeared to be a priority, given that, after thousands of transactions, the Ethereum proceeds were bridged to Bitcoin within 24 hours following the hack. Once on the Bitcoin blockchain, the funds continue to circulate through a wide network. Approximately $28 million have been deposited in Freebitco.in, a Bitcoin-based wallet and gambling service. However, the vast majority of funds otherwise remain dormant on the Bitcoin blockchain.

The Bybit hack illustrates that storing assets in Bitcoin, despite the structural disadvantages of blockchain transparency and the price volatility of floating cryptocurrencies, can be preferable to the relatively high risk of asset seizure on Ethereum. While it was technically feasible to route the stolen funds through a mixer, the limited liquidity in such protocols is unlikely to support the laundering of over $1 billion within a short time horizon without attracting enforcement action. Instead, the hackers bridged the funds into Bitcoin, which, owing to its decentralized design and the absence of custodial intermediaries, offers greater protection against seizure. The North Korean Lazarus Group

[34] North Korea hacked Harmony Bridge and used Tornado Cash to launder the funds in June 2022 and Tornado Cash was sanctioned August 2022.

may be assessing opportunities to exchange Bitcoin for fiat currency or real-world goods. Ultimately, the episode highlights that, without substantial liquidity in mixers, even high profile and extremely sophisticated attacks like this are potentially traceable.

VIII. Conclusion

We provide the first empirical investigation of money laundering policies in the crypto arena. We observe decreased illicit volume following the sanctioning of Tornado Cash, other OFAC designations, and major exchange settlements. We also find evidence that users respond to these restrictions by switching to higher-cost methods. Tornado Cash users, for instance, incur higher transaction costs when attempting to access centralized exchanges, and addresses linked to stablecoin freezes show increased reliance on decentralized exchanges. Sanctions appear effective in making it more difficult for sanctioned entities to find destinations to off-board funds. Overall, our findings indicate that crypto asset freezes and anti-money laundering enforcement have been effective with the ban on Tornado Cash and Tether asset freezes the most effective.

Nevertheless, our analysis indicates many areas for improvement. Even though OFAC sanctions kept funds on-chain, relatively few addresses and dollar amounts are sanctioned. Additionally, while overseas exchanges see reduced sanction-related flows after enforcement actions, they continue to receive some illicit volume and have not experienced a substantial reduction in tainted flows from other forms of criminal activity such as scamming. The sanctions against Tornado Cash were also insufficient to deter some overseas exchanges until additional fines were assessed against exchanges, indicating that multiple enforcement actions may be necessary to plug enforcement weakspots. Therefore, reputational risk and the goodwill of exchanges to stop crime does not appear to be an effective deterrent in the crypto space. Reducing flows to the Tornado Cash mixers led to much more blockchain transparency; thus, if mixer popularity increases as sanctions are lifted, then efforts to freeze criminal proceeds may be substantially more difficult. Additional research should also consider the costs and benefits of monitoring. Our research demonstrates how the blockchain makes tracing and monitoring more straightforward than in other contexts such that large-scale monitoring could presumably be

automated for reasonable costs. We hope additional research will further analyze crypto crime and enforcement to help enact optimal policies that deter crime.

References

Agca, Senay, Pablo Slutzky, and Stefan Zeume, 2020, The Weight of Compliance: Anti-Money Laundering Enforcement, Bank Composition, and Lending, *Working Paper* .

Al-Tawil, Tareq Na'el, 2022, Anti-money laundering regulation of cryptocurrency: Uae and global approaches, *Journal of Money Laundering Control* 26, 1150–1164.

Amiran, Dan, Bjørn N. Jørgensen, and Daniel Rabetti, 2022, Coins for bombs: The predictive ability of on-chain transfers for terrorist attacks, *Journal of Accounting Research* 60, 427–466.

Anderson, Ross, Ilia Shumailov, and Mansoor Ahmed, 2018, Making Bitcoin Legal, *Security Protocols XXVI* 11286, 243–253.

Becker, Gary S., 1968, Crime and punishment: An economic approach, *Journal of Political Economy* 76, 169–217.

Campbell-Verduyn, Malcolm, 2018, Bitcoin, crypto-coins, and global anti-money laundering governance, *Crime, Law and Social Change* 69, 283–305.

Chainalysis, 2024, The 2024 crypto crime report.

Chong, Alberto, and Florencio Lopez-De-Silanes, 2015, Money laundering and its regulation, *Economics amp; Politics* 27, 78–123.

Cong, Lin, Kimberly Grauer, Daniel Rabetti, and Henry Updegrave, 2023a, Blockchain forensics and crypto-related cybercrimes.

Cong, Lin William, Campbell R. Harvey, Daniel Rabetti, and Zong-Yu Wu, 2023b, An anatomy of crypto-enabled cybercrimes.

Cuéllar, Mariano-Florentino, 2002, The tenuous relationship between the fight against money laundering and the disruption of criminal finance, *J. Crim. L. & Criminology* 93, 311.

Draca, Mirko, and Stephen Machin, 2015, Crime and Economic Incentives, *Annual Review of Economics* 7, 389–408.

El Siwi, Yara, 2018, Mafia, money-laundering and the battle against criminal capital: the italian case, *Journal of Money Laundering Control* 21, 124–133.

Ferwerda, Joras, 2009, The economics of crime and money laundering: Does anti-money laundering policy reduce crime?, *Review of Law & Economics* 5, 903–929.

Foley, Sean, Jonathan R Karlsen, and Tālis J Putniņš, 2019, Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?, *The Review of Financial Studies* 32, 1798–1853.

Fracassi, Cezare, and Eric Lee, 2025, The (in-)effectiveness of anti-money laundering, *Working Paper* .

Félez-Viñas, Ester, Luke Johnson, and Talis J. Putnins, 2022, Insider trading in cryptocurrency markets, *SSRN Electronic Journal* .

Gandal, Neil, JT Hamrick, Tyler Moore, and Tali Oberman, 2018, Price manipulation in the bitcoin ecosystem, *Journal of Monetary Economics* 95, 86–96.

Griffin, John M, and Samuel Kruger, 2024, What is Forensic Finance?, *Foundations and Trends® in Finance* 14, 137–243.

Griffin, John M., and Kevin Mei, 2025, How Do Crypto Flows Finance Slavery? The Economics of Pig Butchering, *Working Paper* .

Griffin, John M., and Amin Shams, 2020, Is bitcoin really untethered?, *The Journal of Finance* 75, 1913–1964.

Hamrick, J.T., Farhang Rouhi, Arghya Mukherjee, Amir Feder, Neil Gandal, Tyler Moore, and Marie Vasek, 2021, An examination of the cryptocurrency pump-and-dump ecosystem, *Information Processing and Management* 58, 102506.

Leukfeldt, E. Rutger, Edward R. Kleemans, Edwin W. Kruisbergen, and Robert A. Roks, 2019, Criminal networks in a digitised world: on the nexus of borderless opportunities and local embeddedness, *Trends in Organized Crime* 22, 324–345.

Levi, Michael, 2015, Money for crime and money from crime: Financing crime and laundering crime proceeds, *European Journal on Criminal Policy and Research* 21, 275–297.

Levi, Michael, Peter Reuter, and Terence Halliday, 2017, Can the aml system be evaluated without better data?, *Crime, Law and Social Change* 69, 307–328.

Li, Tao, Donghwa Shin, and Baolian Wang, 2025, Cryptocurrency Pump-and-Dump Schemes, *Journal of Financial and Quantitative Analysis* Forthcoming.

Makarov, Igor, and Antoinette Schoar, 2021, Blockchain Analysis of the Bitcoin Market, Working Paper.

Meiklejohn, Sarah, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage, 2013, A fistful of bitcoins: characterizing payments among men with no names, *Proceedings of the 2013 conference on Internet measurement conference* 127–140.

Mirenda, Litterio, Sauro Mocetti, and Lucia Rizzica, 2022, The economic effects of mafia: Firm level evidence, *American Economic Review* 112, 2748–2773.

Moore, Tyler, Richard Clayton, and Ross Anderson, 2009, The Economics of Online Crime, *Journal of Economic Perspectives* 23, 3–20.

Möser, Malte, and Arvind Narayanan, 2019, Effective cryptocurrency regulation through blacklisting, *Preprint* .

Pennec, Guénolé Le, Ingo Fiedler, and Lennart Ante, 2021, Wash trading at cryptocurrency exchanges, *Finance Research Letters* 43, 101982.

Phua, Kenny, Bo Sang, Chishen Wei, and Gloria Yang Yu, 2022, Don't trust, verify: The economics of scams in initial coin offerings.

Pol, Ronald F., 2020, Anti-money laundering: The world's least effective policy experiment? together, we can fix it, *Policy Design and Practice* 3, 73–94.

Quirk, Peter J., 1996, Macroeconomic implications of money laundering, *IMF Working Papers* 96, 1.

Sokolov, Konstantin, 2021, Ransomware activity and blockchain congestion, *Journal of Financial Economics* 141, 771–782.

Tanzi, Vito, 1996, Money laundering and the international financial system, *IMF Working Papers* 1996, 1.

Tironsakkul, Tin, Manuel Maarek, Andrea Eross, and Mike Just, 2022, Context matters: Methods for bitcoin tracking, *Forensic Science International: Digital Investigation* 42–43, 301475.

US Department of the Treasury, 2024, 2024 National Money Laundering Risk Assessment.

Victor, Friedhelm, 2020, Address clustering heuristics for ethereum.

Wronka, Christoph, 2023, Financial crime in the decentralized finance ecosystem: new challenges for compliance, *Journal of Financial Crime* 30, 97–113.

Figure 1: Tornado Cash Transaction Networks

This figure visualizes flows involving Tornado Cash. It illustrates the network of flows involving Tornado Cash, with senders positioned on the left and receivers on the right. Edges concave down represent flows moving from left to right (e.g., the curve moves as if going from 9 o'clock to 3 o'clock), while edges concave up indicate flows moving from right to left (e.g., from 3 o'clock to 9 o'clock). This is a sample constructed by selecting the largest 2,500 nodes within 5 hops from Tornado Cash and keeping connected paths. This was selected to reflect the largest ∼10,000 edges related to Tornado Cash.

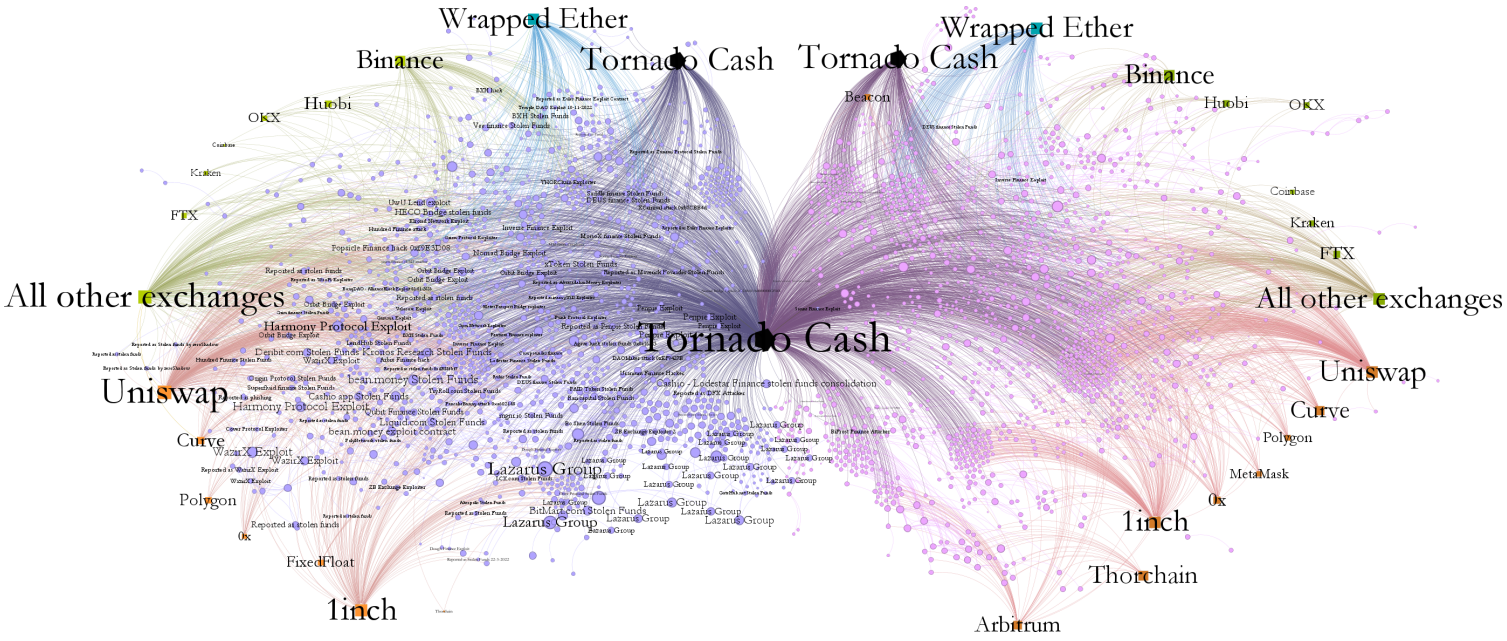

Figure 2: Inflows to Tornado Cash

This figure shows the total inflows and traced criminal inflows to Tornado Cash from 2021 to 2024. The blue line indicates the total monthly inflows to Tornado Cash. The red line represents traced monthly criminal inflows from all scam types. The dashed vertical lines indicate the period during which Tornado Cash was sanctioned by the U.S. Treasury, beginning on August 8, 2022, and lifted on November 26, 2024. The legend provides aggregate statistics for the entire period.

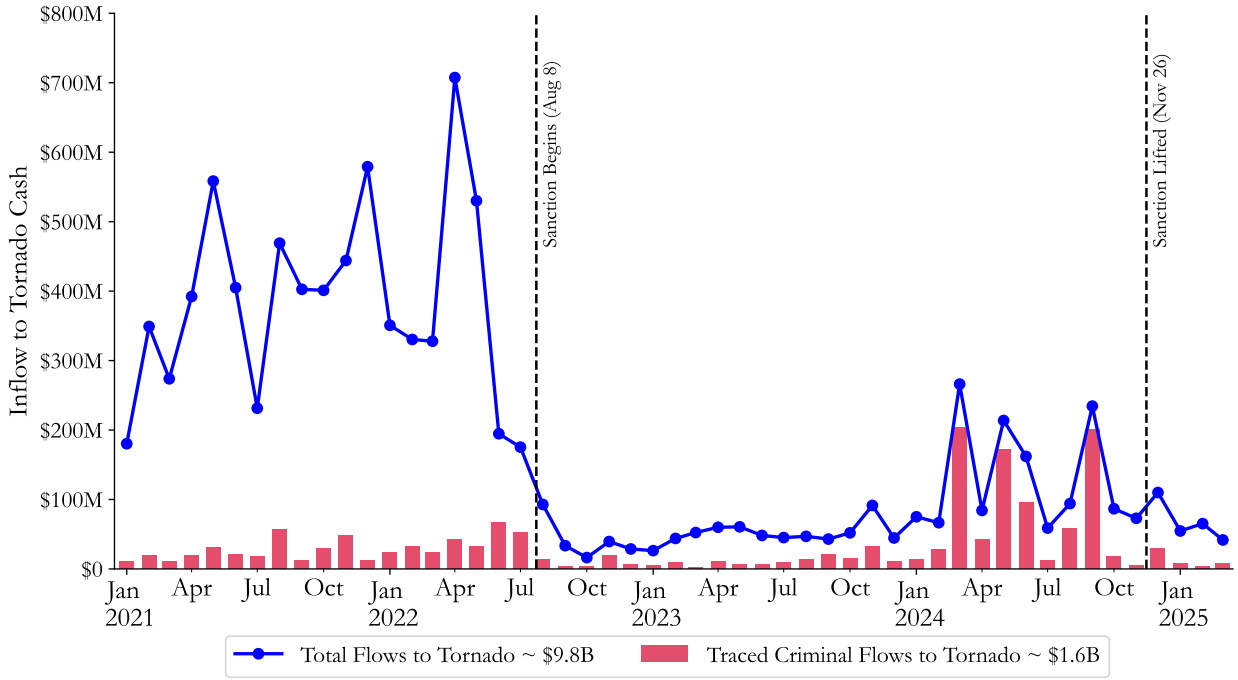

Figure 3: Effects of U.S. Treasury Sanctions Against Tornado Cash on Hacker Flows

This figure examines the change in hacker flows around the U.S. Treasury sanctions on Tornado Cash. It plots monthly difference-in-differences (DiD) coefficients from the regression

$$\log(1 + TaintedFlows)_{i,d,t} = \sum_{t \neq June2022} \beta_t \times \mathbb{1}(Month = t) \times Tornado_d + \delta \times Tornado_d + \mu_i + \gamma_t + \varepsilon_{i,d,t}$$

where $\log(1+TaintedFlows)_{i,d,t}$ is the natural logarithm of $1+$ the dollar flows traced from reported hacker address i to destination d in month t. Hacker report and month fixed effects are included. Standard errors are double-clustered by hacker report and month. Vertical dashed lines mark the May 6, 2022 sanction of Blender.io and the August 8, 2022 sanction of Tornado Cash.

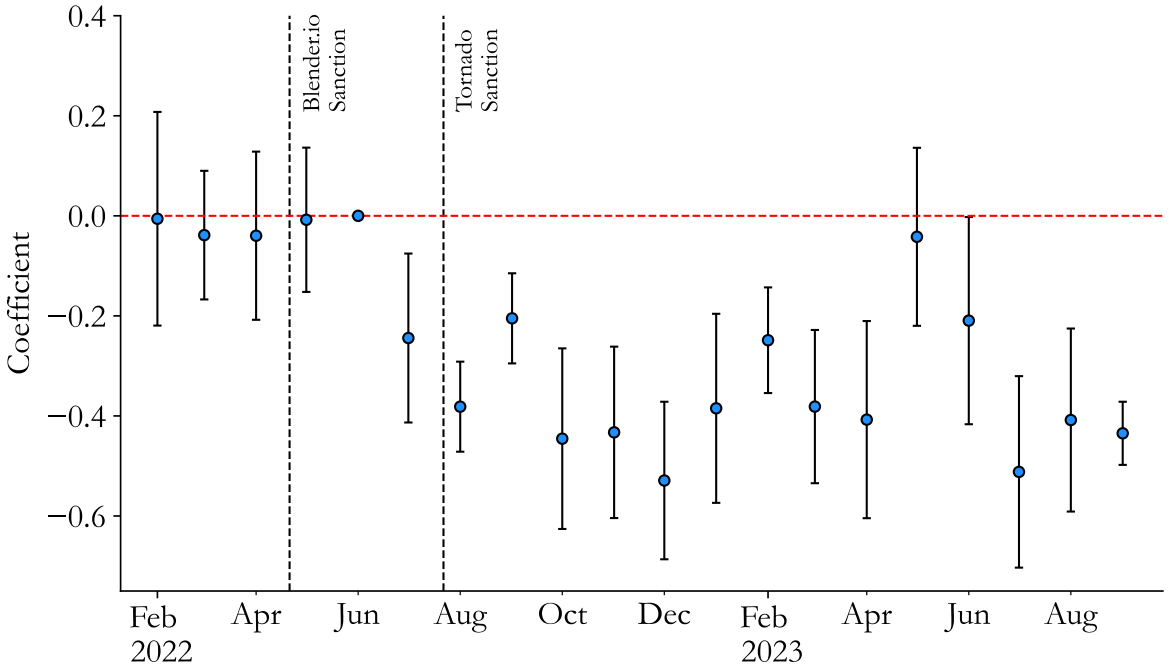

Figure 4: Tornado Cash Outflows to Destinations by Intermediate Hop

This figure visualizes Tornado Cash outflows and their destinations before and after the Tornado Cash ban on August 8, 2022, using Sankey diagrams. Panel A shows outflows before the ban, and Panel B shows outflows after the ban. In each diagram, flows originate from Tornado Cash on the left and proceed through intermediate hops (e.g., Hop 1, Hop 2, shown in green) before reaching final destinations on the right. Flows into centralized exchanges (CEXs) are shown in blue, decentralized exchanges (DEXs) in red, blockchain bridges in purple, and wrapped ETH (WETH) in gray. When flows arrive at a DEX, funds are traced through additional hops to identify their ultimate destinations after swaps. The width of each flow indicates the relative volume of funds moving along that path. Western exchanges are Coinbase, Crypto.com, Gemini, and Kraken, and all other exchanges are included as overseas exchanges.

Panel A: Before Tornado Cash Ban

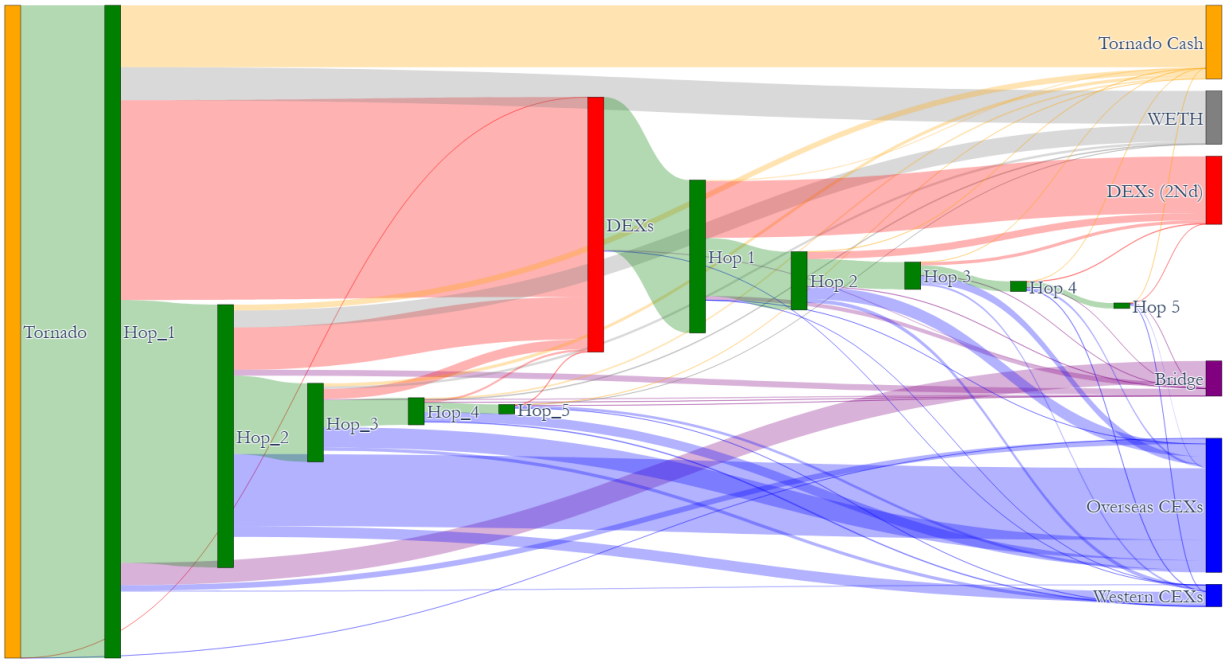

Panel B: After Tornado Cash Ban

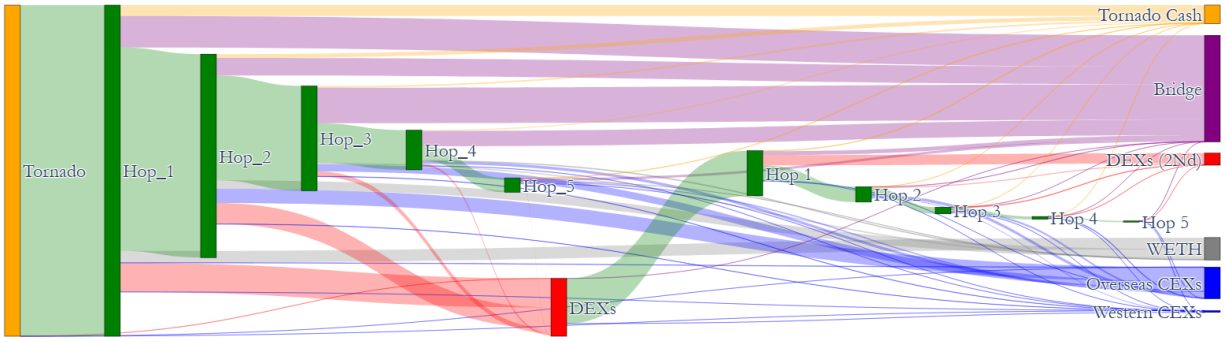

Figure 5: Effects of Tornado Cash Ban on Its Outflows to Different Destinations

This figure examines the effects of the Tornado Cash ban on its outflows to different destinations. It shows Tornado Cash outflows to Western CEXs in Panel A, overseas CEXs in Panel B, and bridges in Panel C. In each panel, bars represent the percent of monthly outflows to the destination category: blue bars represent flows that moved from Tornado Cash to the destination without passing through any decentralized exchanges (DEXs), while yellow bars represent flows that were first sent from Tornado Cash to DEXs, and then arrived at CEXs after being traced through DEXs. The red line plots the average number of hops before funds arrive, and the dark blue line plots the average number of days for transfers. Western exchanges are Coinbase, Crypto.com, Gemini, and Kraken, and all other exchanges are included as overseas exchanges. Two black vertical dashed lines mark the sanction of Tornado Cash on August 8, 2022, and its lifting on November 26, 2024.

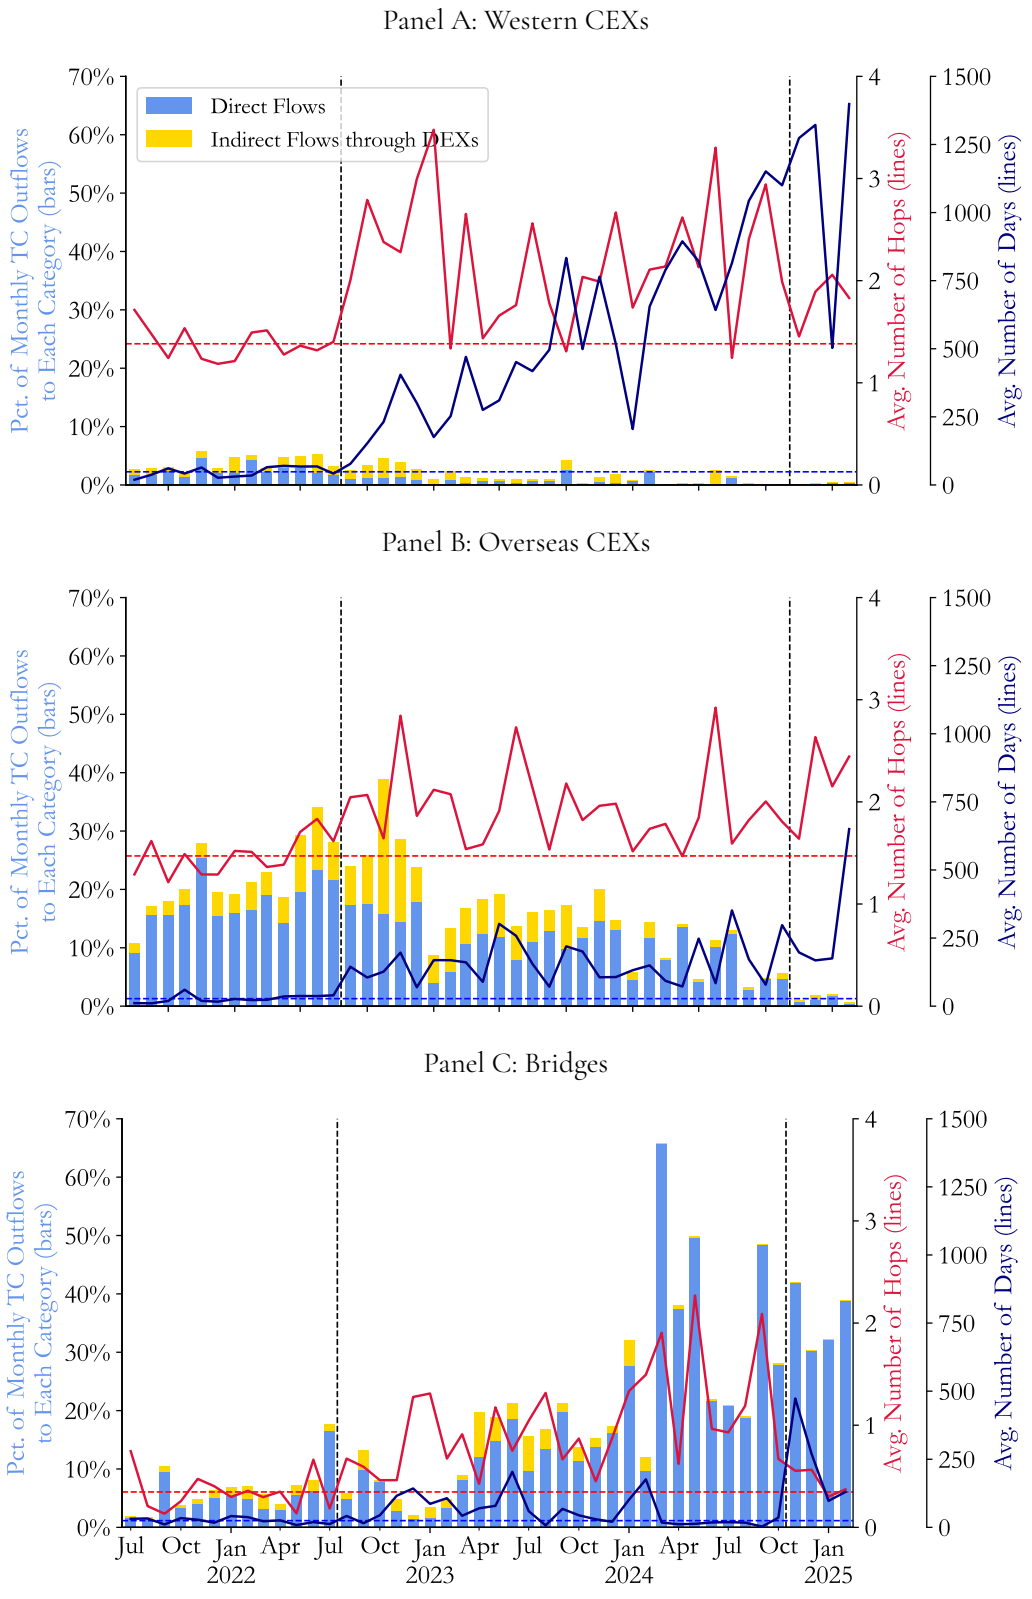

Figure 6: Effect of OFAC Sanctions

This figure plots OFAC sanction activity. Panel A shows, on the left-hand y-axis, the balances of all addresses over time, with gray if the address had not yet been subject to sanctions, and red after sanctions. The right-hand y-axis plots the number of addresses that have been sanctioned over time in dashed lines and circles that denote the date of new sactions. Panel B plots destinations of flows leaving addresses sanctioned by OFAC, indexed by months since sanctions were imposed.

Panel A: Balance of addresses before and after sanctions

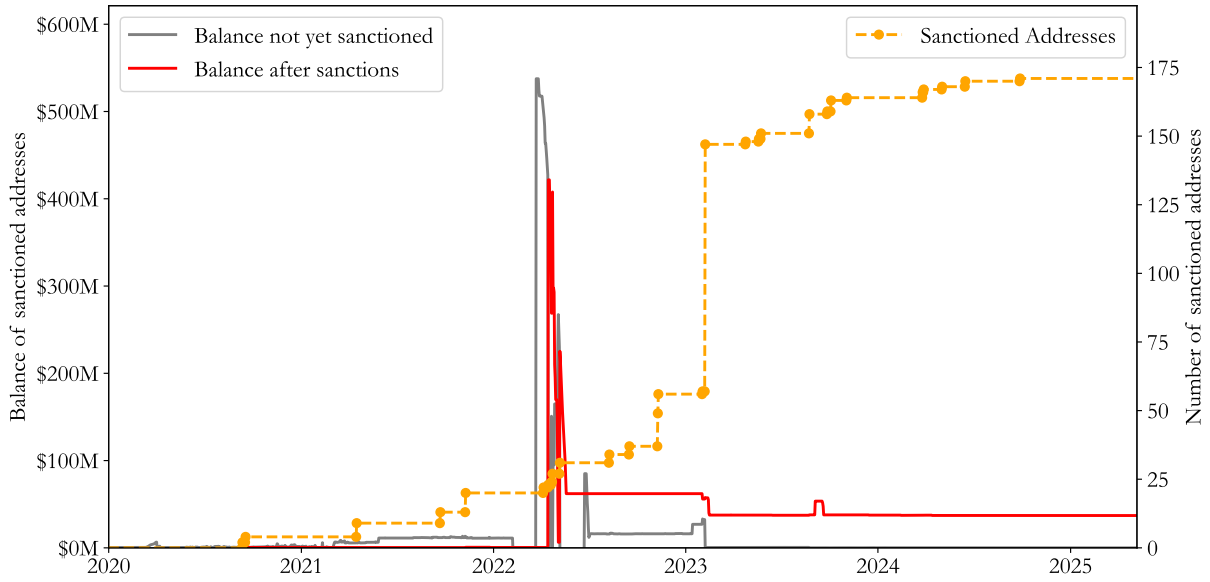

Panel B: Outflows traced from sanctioned addresses to end destinations

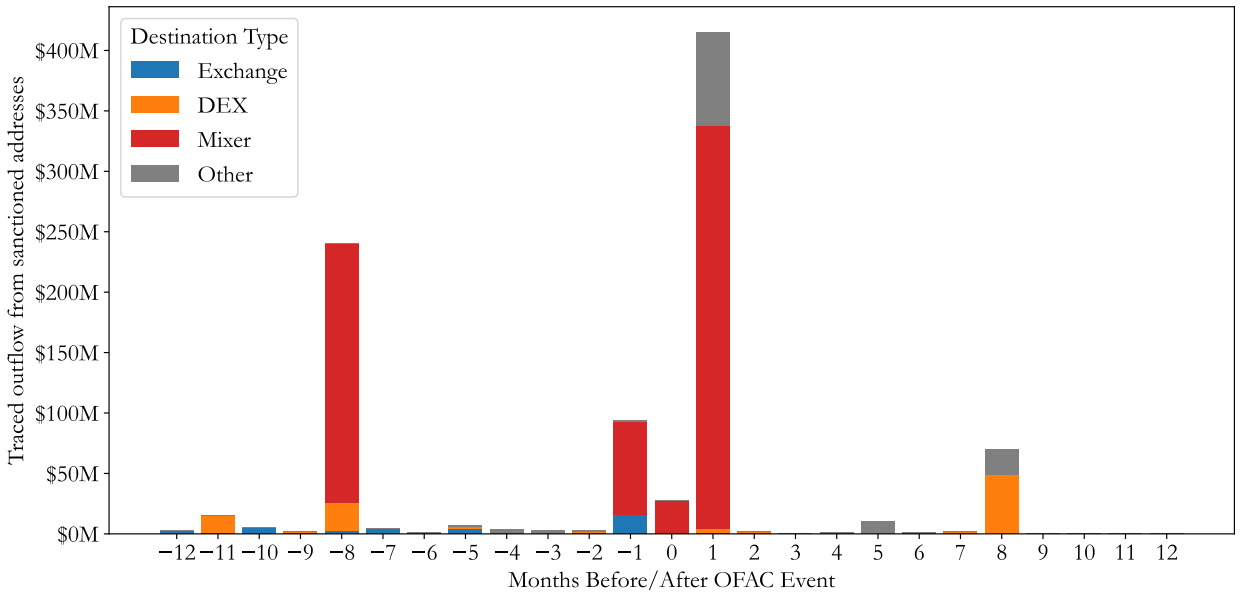

Figure 7: Tether Freezes

This figure presents the scope of Tether freezes and the pre-seizure activity of frozen addresses. Panel A summarizes enforcement scale over time: the red line (left axis) plots the cumulative number of blacklisted addresses, while the stacked bar (right axis) shows the cumulative dollar value of Tether frozen. Colored segments of each bar indicate frozen addresses appearing in the traced network of reported criminal flows and show the corresponding scam type, while white segments indicate addresses outside that network. Panel B illustrates inflow trajectories for frozen addresses prior to seizure. This panel plots a select sample of frozen addresses that have received more than $10K in total inflows. The horizontal axis reports the number of days each address was active prior to being frozen, and the vertical axis (log scale) shows the total dollar inflow to that address. Each dot represents a transfer event, with dot size reflecting the cumulative inflow received by the address at that point in time. Multiple dots forming a line represent inflow trajectories for a single address. Red crosses mark the date each address was frozen.

Panel A: Cumulative Freezes Over Time

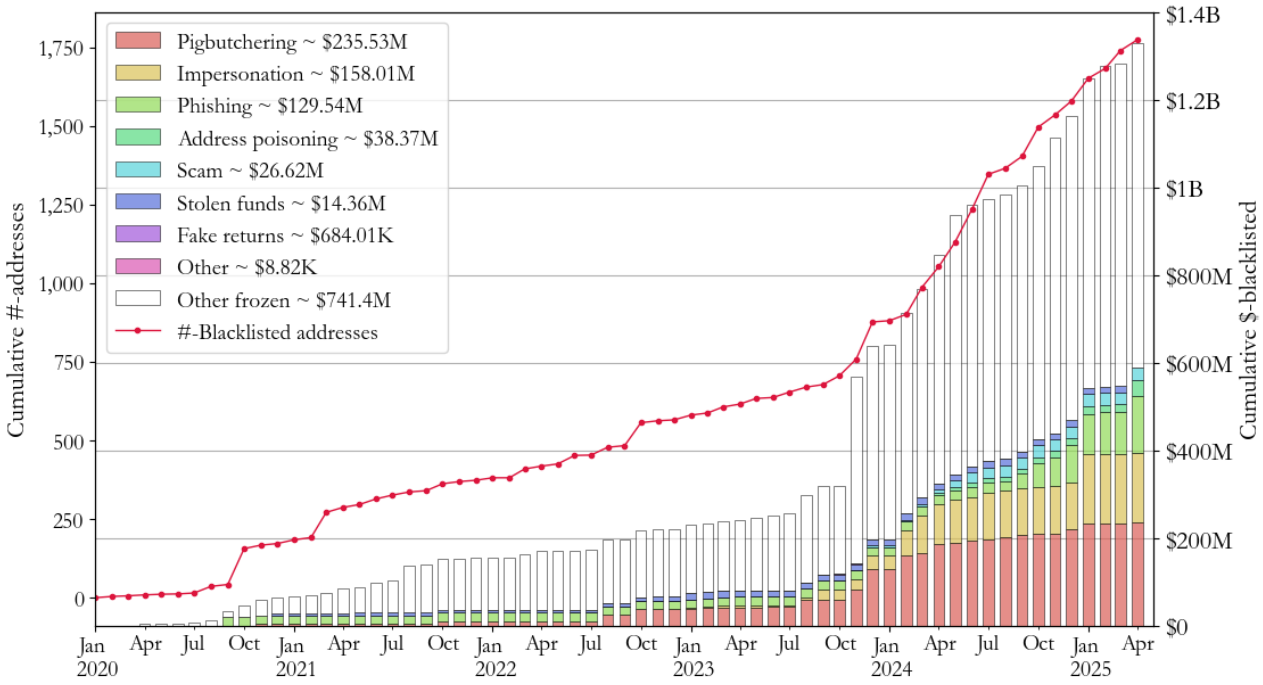

Panel B: Activity of Frozen Addresses Prior to Seizure

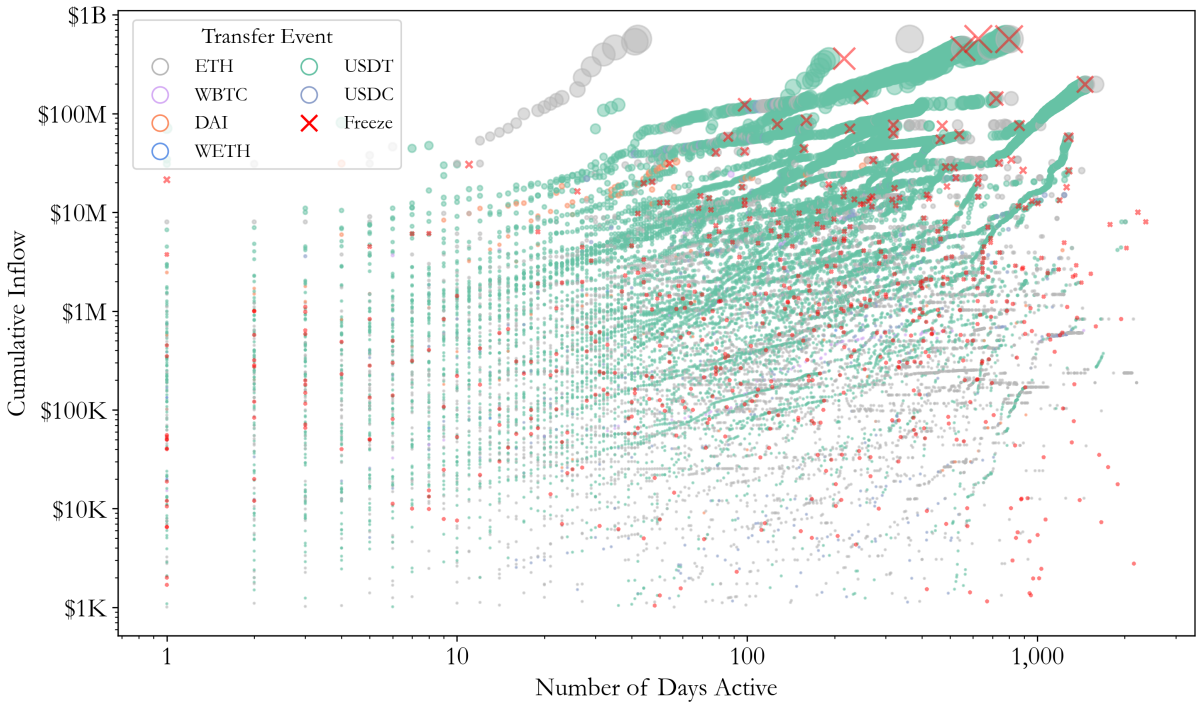

Figure 8: Tether Seizure: Immediate Effects

This figure presents outflows indexed by the hours prior to and after the Tether freeze. The type of cryptocurrency is denoted by color. Panel A shows the outflows of addresses that had assets frozen. Panel B shows the outflows of addresses related to addresses that were frozen.

Panel A: Frozen addresses

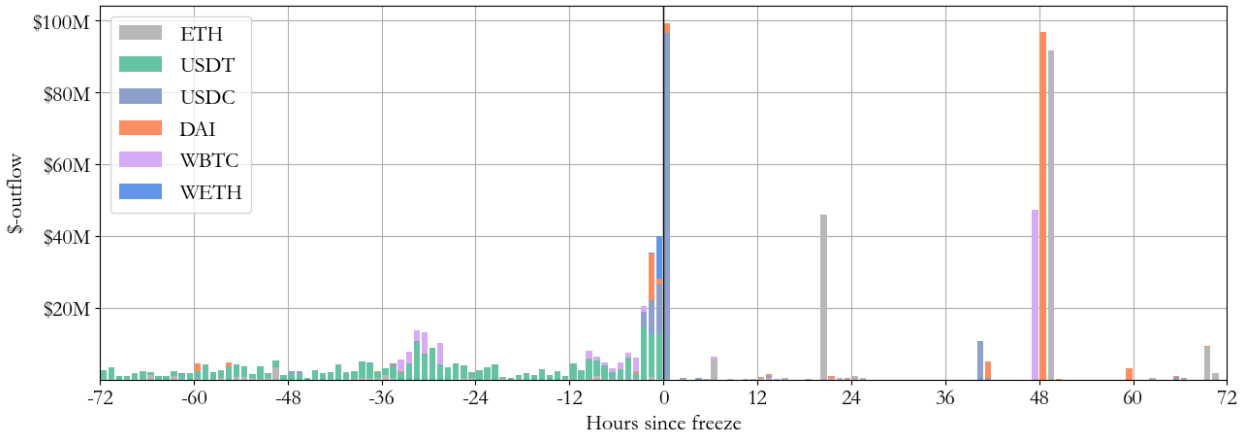

Panel B: Related addresses

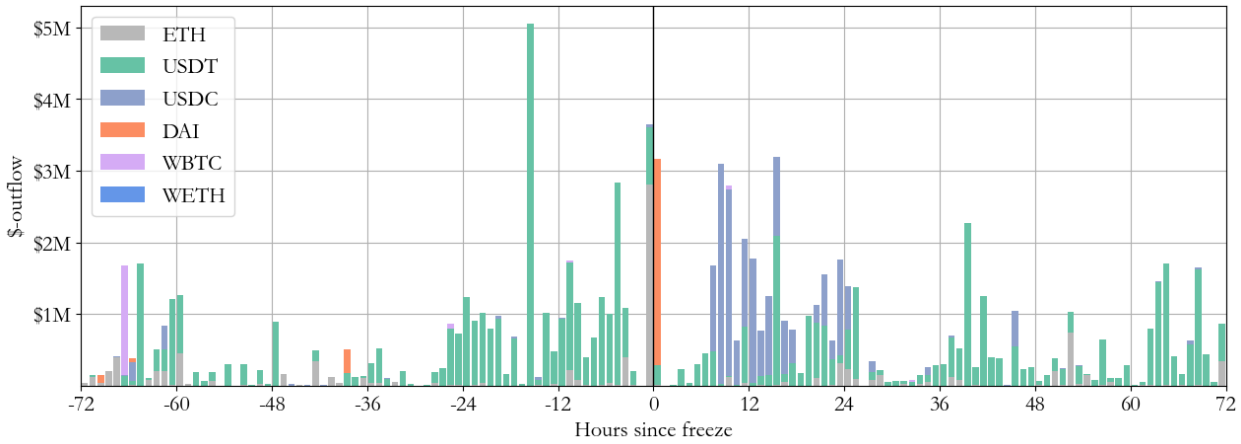

Figure 9: Tether Seizure: Long Term Effects

This figure examines whether blacklisted addresses and their related addresses increase their use of DeFi services following an asset freeze. It plots coefficients from a difference-in-differences regression of the DeFi share of outflows at the group-cohort-month level. Each cohort represents one seizure event and consist of a treatment group and a control group. The treated group consists of the frozen address and its related addresses. For every freeze event, the control group is a random sample of 20 addresses that received inflows of at least $100 within the seven days prior to the freeze. Specifically, the following regression is estimated.

$$\textit{DeFi Share}_{g,c,t} = \sum_{t \neq t_{freeze}} \beta_t \times \mathbb{1}(Month = t) \times Treat_g + \mu_c + \gamma_t + \varepsilon_{g,c,t}$$

where $\textit{DeFi Share}_{g,c,t} = \frac{\sum_{i \in g} DeFi\ flows_{i,g,c,t}}{\sum_{i \in g} All\ flows_{i,g,c,t}}$ and it denotes the share of flows sent to DeFi services for the treatment or control group g in cohort c in month t. The regression includes the cohort fixed effects and event time fixed effects. The sample period is 12 months before and after the freeze date of each cohort. Standard errors are clustered by asset seizure cohort and event time.

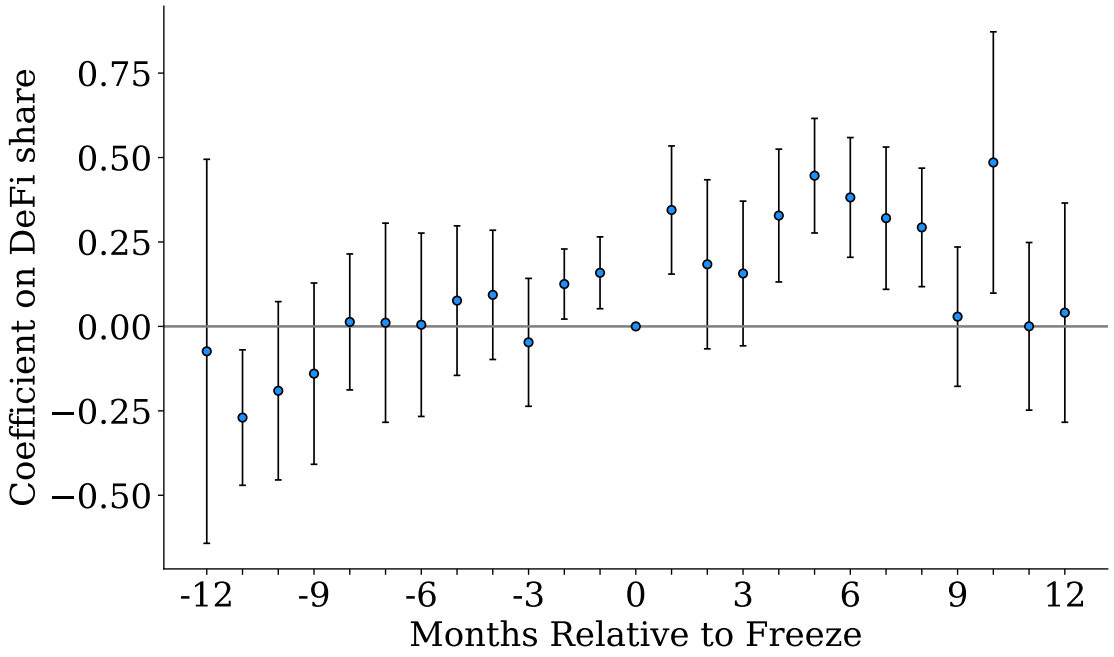

Figure 10: Risk of Seizure versus Transaction Cost and Speed

This figure examines how transaction costs and exit speed correlate with the risk of account freezing for funds originating from reported criminal addresses. The sample includes all identified transfer paths from reported criminal addresses to CEXs. Each path is categorized along two dimensions: quintiles of transaction costs (x-axis) and quintiles of the number of days required to exit to a CEX (y-axis). For each cell, the height of the bar indicates the probability that at least one address along the path was subsequently frozen by Tether (z-axis). The probabilities are calculated as weighted averages across paths belonging to the category, weighted by the dollar amount of funds entering CEXs. Transaction cost is expressed as a percentage by summing transaction cost fees paid at each hop and the spread lost from swaps, then dividing by the funds ultimately deposited into CEXs.

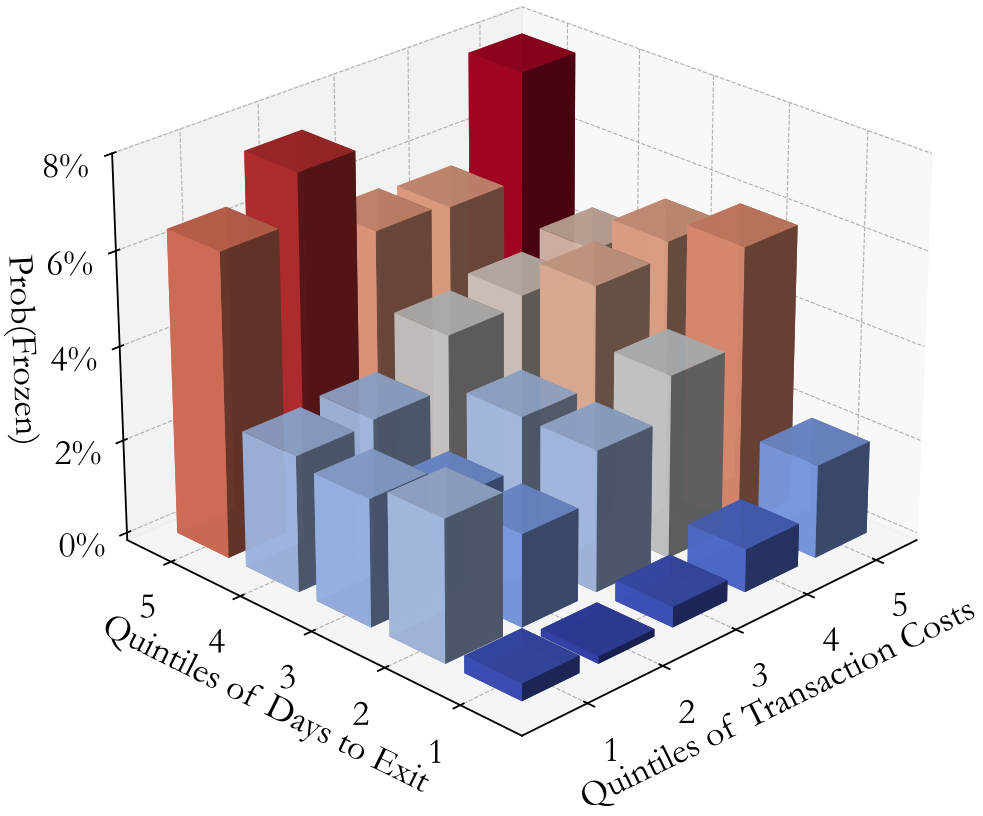

Figure 11: Tornado Cash Outflows to Binance around Binance Settlement

This figure examines the effects of the Binance DOJ settlement and subsequent compliance monitorship on Tornado Cash outflows. Panel A presents Tornado Cash outflows to Binance (left subpanel) and to other overseas CEXs (right subpanel). In each subpanel, bars represent the percent of monthly outflows to the destination category: blue bars represent flows that moved from Tornado Cash to the destination without passing through any decentralized exchanges (DEXs), while yellow bars represent flows that were first sent from Tornado Cash to DEXs, and then arrived at CEXs after being traced through DEXs. The red line shows the average number of hops before reaching the destination. Panel B compares the number of hops for transfers from Tornado Cash to Binance versus other overseas CEXs using a difference-in-differences (DID) regression, with the estimated coefficients plotted over time. Western exchanges are Coinbase, Crypto.com, Gemini, and Kraken, and all other exchanges are included as overseas exchanges. Two red vertical dashed lines indicate key regulatory events at Binance: the DOJ settlement on November 21, 2023, and the announcement of the compliance monitorship on May 17, 2024. Two black vertical dashed lines mark the Tornado Cash sanction on August 8, 2022, and the lifting of the sanction on November 26, 2024.

Panel A: Tornado Cash Outflows

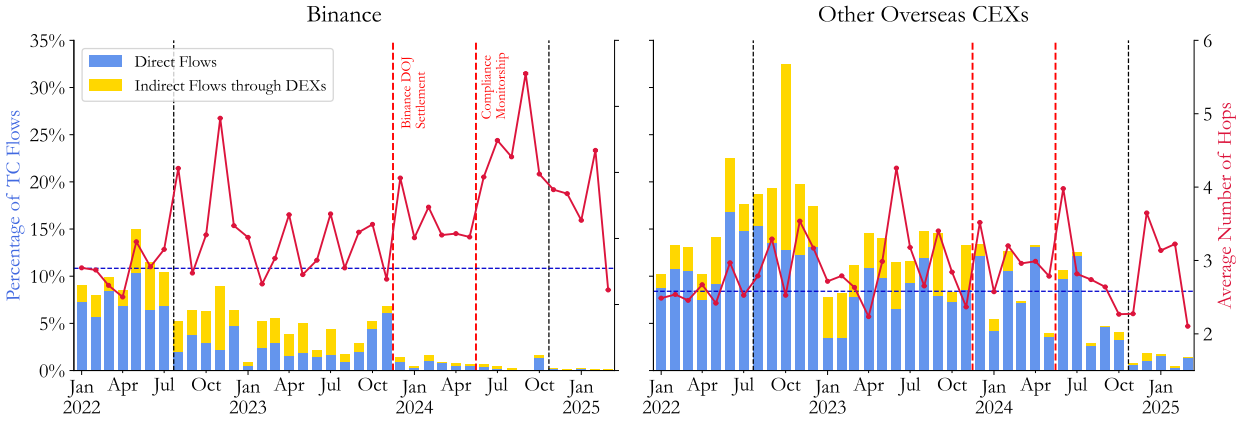

Panel B: DID Analysis on Number of Hops

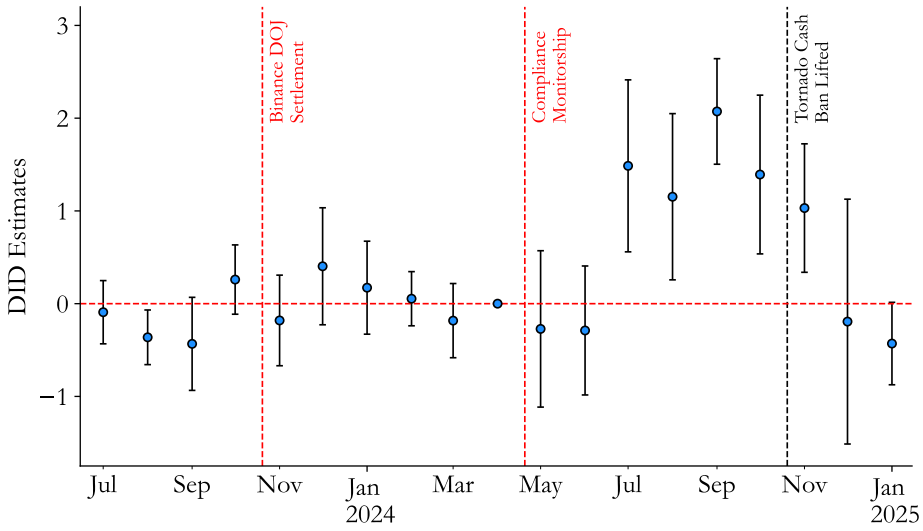

Figure 12: Tainted Flows to Binance around Binance Settlement

This figure examines whether tainted flows to Binance declined following the Binance DOJ settlement and subsequent compliance monitorship. Panel A shows monthly inflows to Binance tainted deposit addresses (solid line) compared to inflows to all other tainted deposit addresses (dashed line). Panel B presents the estimated coefficients from a difference-in-differences regression at the deposit address-month level. The treatment group is composed of tainted Binance deposit addresses, and the control group consists of tainted deposit addresses at all other exchanges. The regression is of the form:

$$\log\left(1 + Total\ Inflow_{i,t}\right) = \sum_{t \neq t_{Nov2023}} \beta_t \times \mathbb{1}(Month = t) \times Treat_i + \mu_i + \gamma_e + \eta_t + \varepsilon_{i,t}$$

where $\log\left(1 + Total\ Inflow_{i,t}\right)$ denotes the log of one plus the total tainted inflow received by deposit address i in exchange e in month t. The regression includes deposit address fixed effects, exchange fixed effects, and month fixed effects. The first red dashed line corresponds to the DOJ settlement date, and the second red dashed line marks the announcement of the compliance monitorship. Standard errors are clustered by deposit address and month.

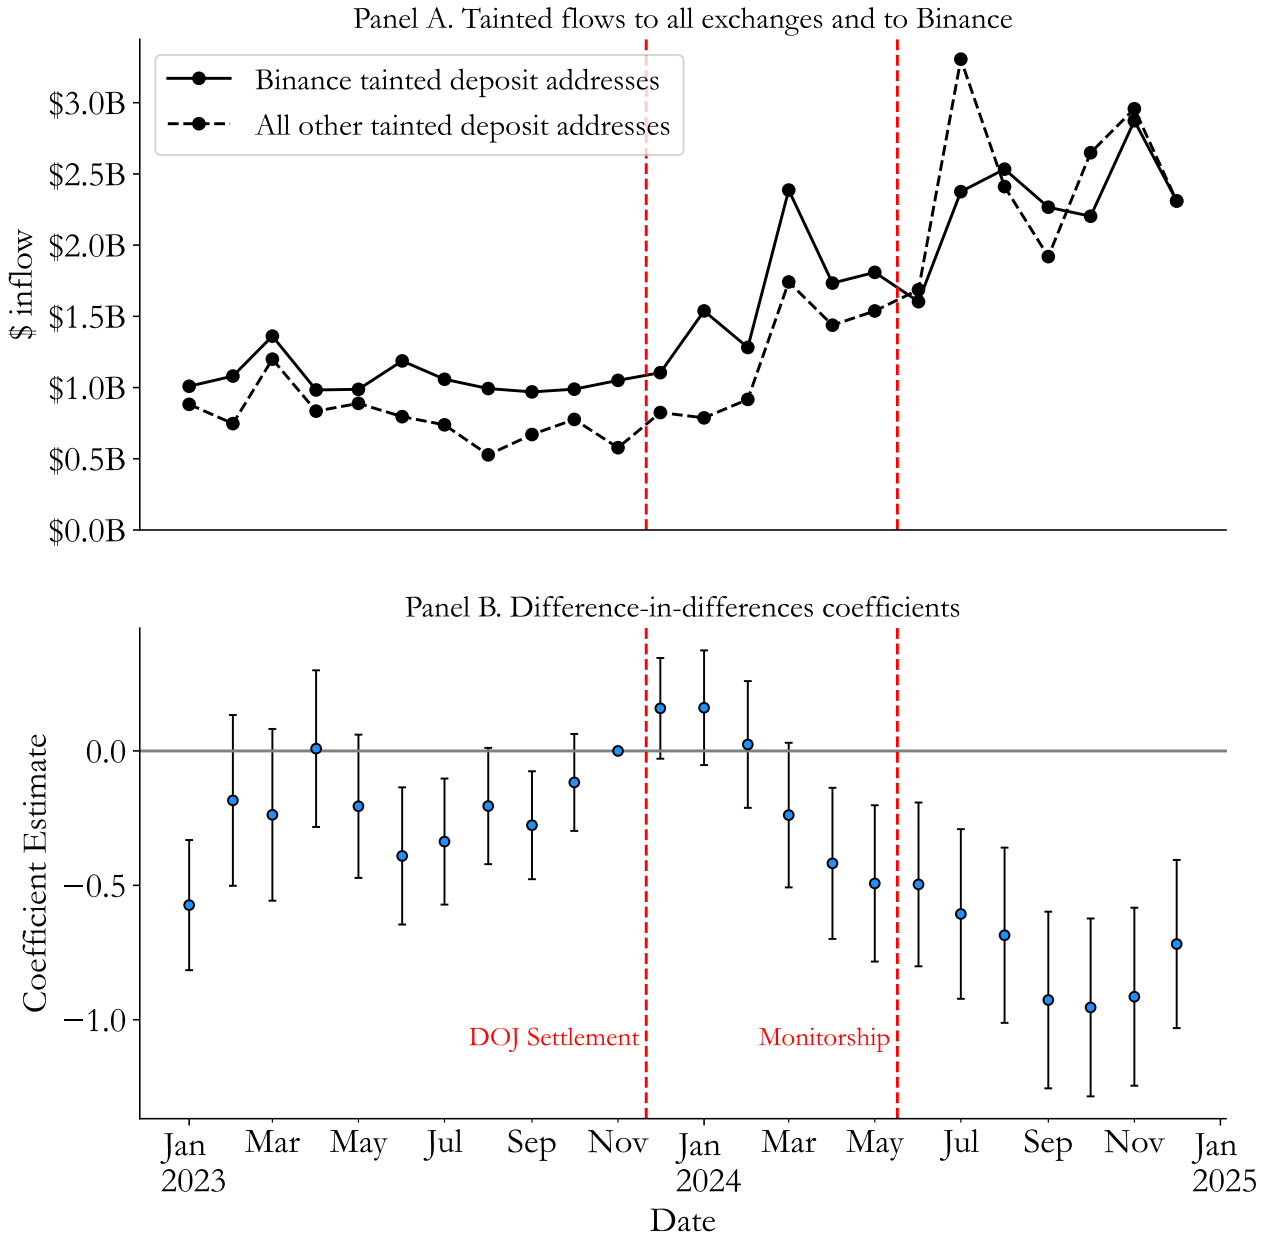

Figure 13: Effects of Tornado Cash Ban on Hacker Flows Destinations

This figure examines the change in hacker flows around the Tornado Cash ban. It displays flows of reported hackers to various destinations for hackers who started moving criminal funds before (left) and after (right) the ban. It highlights sizable shifts in blue and entirely new top destinations in red.

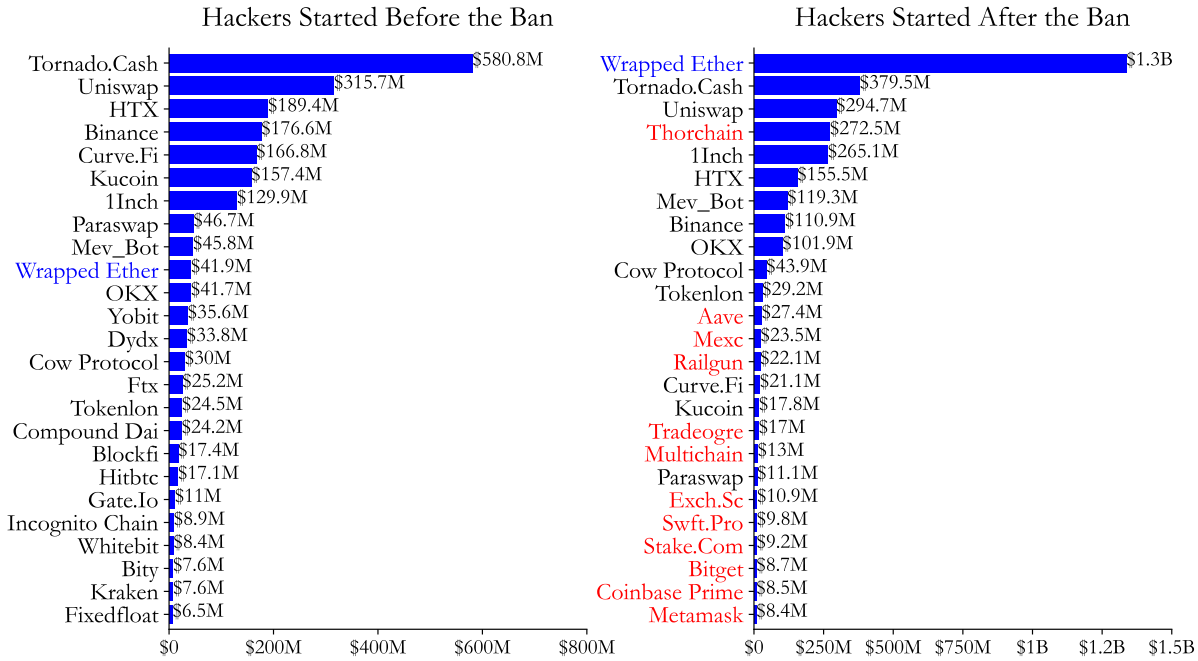

Figure 14: Bybit Hack and Flows to Bridge

This figure shows the flows of proceeds of the $1.4 billion theft from Bybit in February 2025. The left-hand side shows funds exiting Bybit and being forwarded through Ethereum by the hackers to the different services represented in the center. On the right-hand side of Thorchain are Bitcoin addresses, representing addresses downstream of bridge transactions. Concave down edges represent flows from a node in the left to another node in the right, and vice versa for concave up edges. Edges are colored by the total $-amount, and nodes are colored by the corresponding type: service, reported hacker, ETH wallet, BTC wallet, and BTC transaction. On Ethereum we find $1B is bridged through Thorchain, and traced over $42.8M to the OKX Web3 service, $7.1M to Mayachain, $4.6M Lifi, $4.5M to 1inch, and $4.5 to other destinations. On the Bitcoin blockchain we traced $28.7M of the bridged funds to Freebitco.in.

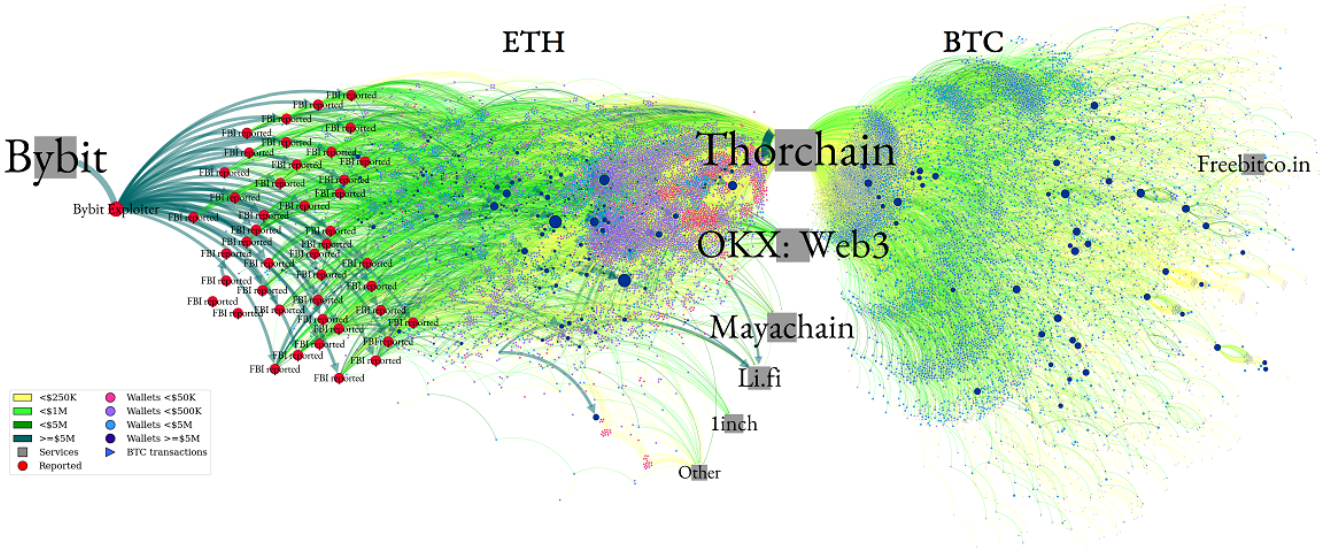

Table 1: Reports Summary

This table presents summary statistics on number of reports and addresses with details provided by scam and by blockchain. The leftmost column splits first by scam, and then represents the same data split by blockchain. Column (1) indicates the total number of reports while Column (2) lists the active addresses within all reports. Column (3) removes duplicates within each category (i.e., it is possible for an address to be named in multiple categories. Column (4) drops addresses with more than 2,000 transactions, which suggests that they may be exchanges. The total inflow to these addresses is displayed in Column (5).

	(1)	(2)	(3)	(4)	(5)
Category	Total Reports	Reports with active address	and remove duplicates within cat.	and remove likely exchanges	Total Inflow
Stolen funds	9,264	9,136	9,136	7,980	$12.91B
Pigbutchering	25,809	21,056	13,172	10,095	$12.44B
Illicit actor	6,885	6,875	6,872	5,158	$11.25B
Contract exploit	938	429	429	277	$5.76B
Phishing	65,897	65,030	62,753	53,966	$4.43B
Scam	30,056	29,872	29,872	22,167	$2.46B
Extortion	192,933	91,975	7,405	6,415	$1.89B
Fake project	4,101	3,854	2,546	1,911	$913.84M
Malware	3,155	2,316	2,141	2,009	$817.02M
Fake returns	4,843	4,407	3,791	2,621	$705.1M
Impersonation	31,496	30,503	28,445	27,440	$145.92M
Airdrop	795	739	494	340	$38.24M
Dark market	833	831	827	577	$28.34M
Sim swap	192	170	152	122	$9.46M
Address poisoning	122,800	122,427	122,427	122,410	$9.26M
Donation scam	639	364	223	162	$2.77M
Bitcoin	244,331	139,383	44,124	37,773	$21.46B
Ethereum	247,985	242,650	229,265	219,185	$31.15B
Tron	8,320	7,951	6,884	1,212	$1.2B
Total	500,636	389,984	280,273	258,170	$53.81B

Table 2: Address Total Inflow Summary

This table presents summary statistics by scam of the total dollar inflow into addresses. The *N* corresponds to Column (6) of Table 1 and total inflow corresponds to Column (7). This table presents mean, standard deviation, and the 25th, 50th, and 75th percentile. Categories are sorted by total inflow.

Category	N	Total Inflow	Mean	Std	25%	50%	75%
Stolen funds	7,980	$12.91B	$1.62M	$14.61M	$6.69K	$60.32K	$360.33K
Pigbutchering	10,095	$12.44B	$1.23M	$9.98M	$5.7K	$63.58K	$457.83K
Illicit actor	5,158	$11.25B	$2.18M	$62.69M	$150.74	$1.3K	$19.68K
Contract exploit	277	$5.76B	$32.2M	$312.8M	$1.92K	$335.69K	$3.43M
Phishing	53,966	$4.43B	$85.01K	$2.16M	$0	$0	$672.14
Scam	22,167	$2.46B	$117.49K	$1.4M	$536.56	$4.61K	$27.9K
Extortion	6,415	$1.89B	$316.47K	$11.99M	$525.82	$1.33K	$3.28K
Fake project	1,911	$913.84M	$509.38K	$3.16M	$22.85	$10.73K	$165.68K
Malware	2,009	$817.02M	$481.73K	$1.99M	$1.16K	$26.05K	$225.66K
Fake returns	2,621	$705.1M	$358.83K	$10.69M	$722.59	$6.15K	$33.2K
Impersonation	27,440	$145.92M	$5.44K	$105.62K	$5.1	$6.82	$7.87
Airdrop	340	$38.24M	$142.69K	$669.02K	$158.88	$7.26K	$60.88K
Dark market	577	$28.34M	$49.12K	$373.34K	$57.59	$236.16	$2.9K
Sim swap	122	$9.46M	$91.85K	$151.91K	$3.95K	$23.94K	$108.19K
Address poisoning	122,410	$9.26M	$75.65	$12.53K	$0.74	$1.52	$4.1
Donation scam	162	$2.77M	$23.05K	$185.3K	$45.08	$390.7	$2.95K
Total	258,170	$53.81B	$208.42K	$12.76M	$0.74	$3.83	$273.94

Table 3: Effects of the Tornado Cash Ban: Western CEXs vs. Overseas CEXs

This table presents the results of difference-in-differences (DID) regressions examining the effects of the Tornado Cash ban on criminal behaviors and transaction costs. It compares Tornado Cash outflows to western CEXs (treatment group) with those to overseas CEXs (control group). The outcome variables are the number of intermediate hops (columns 1–2), the number of days to exit to a CEX (columns 3–4), and transaction costs, measured in decimals (columns 5–6). Regressions are estimated at the path-exit level, where each "path" corresponds to transfers originating from a single withdrawal of funds that pass through one or more intermediate hops before exiting to a CEX. Each path may include multiple exits, with each exit representing a cash-out event to a CEX and counted separately. The number of days is measured from when the funds initially leave Tornado Cash until they ultimately reach the CEX at each path exit. Transaction cost is expressed in decimal form (0.05 denotes 5%) by summing transaction cost fees paid at each hop and the spread lost from swaps, then dividing by the funds ultimately deposited into CEXs. Tornado Cash was sanctioned on August 8, 2022, and a path exit is assigned a post variable of one if the funds reached a CEX after this date. Western exchanges are Coinbase, Crypto.com, Gemini, and Kraken, and all other exchanges are included as overseas exchanges. All regressions are dollar-weighted by the amount of funds entering CEXs. Exchange fixed effects and year-month fixed effects are included as indicated. Standard errors, clustered by path, are reported in parentheses. The sample period covers March 1, 2022, to December 31, 2023.

Dep. Variable:	Number of Hops		Number of Days		Transaction Cost	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × Post	0.285***	0.285***	117.8***	121.7***	0.00875***	0.00770***
	(0.0992)	(0.108)	(22.27)	(22.14)	(0.00260)	(0.00265)
Treat	-0.157***		32.06***		0.00108	
	(0.0452)		(7.006)		(0.000681)	
Constant	1.749***	1.734***	80.92***	83.79***	0.00512***	0.00525***
	(0.0154)	(0.0137)	(2.296)	(2.142)	(0.000262)	(0.000247)
Exchange FE		✓		✓		✓
Year-Month FE	✓	✓	✓	✓	✓	✓
Observations	132,050	132,040	132,050	132,040	132,050	132,040
Adjusted R^2	0.0868	0.195	0.179	0.276	0.0318	0.0578
Dep. Var. Mean	1.741	1.741	86.93	86.91	0.00545	0.00545
Dep. Var. Std	1.111	1.111	172.7	172.7	0.0262	0.0262

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Effects of the Tornado Cash Ban: Tornado Cash Outflows vs. Criminal-Address Outflows

This table presents the results of difference-in-differences regressions examining the effects of the Tornado Cash ban on its outflows versus criminal flows to CEXs. It compares Tornado Cash outflows to CEXs (treatment group) with flows originating from reported criminal addresses to CEXs (control group). The outcome variables are the number of intermediate hops (columns 1–2), the number of days to exit to a CEX (columns 3–4), and transaction costs, measured in decimals (columns 5–6). Regressions are estimated at the path-exit level, where each "path" corresponds to transfers originating from a single withdrawal of funds that pass through one or more intermediate hops before exiting to a CEX. Each path may include multiple exits, with each exit representing a distinct cash-out event to a CEX and counted separately. The number of days is measured from when funds initially leave the source (Tornado Cash or a criminal address) until they ultimately reach the CEX at that path exit. Transaction cost is expressed in decimal form (0.05 denotes 5%) by summing transaction cost fees paid at each hop and the spread lost from swaps, then dividing by the funds ultimately deposited into CEXs. Tornado Cash was sanctioned on August 8, 2022, and a path exit is assigned a post variable of one if the funds reached a CEX after the ban. All regressions are dollar-weighted by the amount of funds entering CEXs. Exchange fixed effects and year-month fixed effects are included as indicated. Standard errors, clustered by path, are reported in parentheses. The sample period covers March 1, 2022, to December 31, 2023.

Dep. Variable:	Number of Hops		Number of Days		Transaction Cost	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × Post	0.169***	0.122**	131.3***	133.2***	0.00326***	0.00332***
	(0.0464)	(0.0475)	(6.270)	(6.255)	(0.000638)	(0.000639)
Treat	0.168***	0.192***	23.85***	20.95***	0.00384***	0.00349***
	(0.0299)	(0.0300)	(1.942)	(2.044)	(0.000400)	(0.000420)
Constant	1.570***	1.569***	13.41***	13.54***	0.000839***	0.000859***
	(0.0153)	(0.0151)	(0.497)	(0.494)	(0.0000586)	(0.0000613)
Exchange FE		✓		✓		✓
Year-Month FE	✓	✓	✓	✓	✓	✓
Observations	608,713	608,713	608,713	608,713	608,713	608,713
Adjusted R^2	0.0427	0.0549	0.118	0.136	0.00565	0.00712
Dep. Var. Mean	1.583	1.583	17.68	17.68	0.00114	0.00114
Dep. Var. Std	1.527	1.527	69.28	69.28	0.0189	0.0189

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Stablecoin Seizure Summary Statistics

This table summarizes stablecoin seizure statistics for blacklisted addresses overlapping with the traced criminal network. The leftmost column lists each scam category. Column (1) reports the number of origin addresses in that category where the downstream path reaches at least one blacklisted address. Column (2) shows the number of blacklisted addresses downstream of those origins. Column (3) reports the total dollar outflow from the origin addresses. Column (4) reports the dollar amount traced to these blacklisted addresses from origin addresses; this value may exceed the origin outflow in column (3) because a blacklisted address may aggregate flows from multiple origin categories. Column (5) shows the total dollar value frozen in these blacklisted addresses. Column (6) reports the dollar value of stablecoins destroyed by issuers after freezing. The last row presents the aggregate overlap; the columns do not sum because one blacklisted address may appear in multiple paths downstream of reported addresses.

	(1)	(2)	(3)	(4)	(5)	(6)
Category	# associated origins	# blacklisted addresses	$ origin outflow	$ traced to blacklist	$ total blacklisted	$ total destroyed
Impersonation	1589	63	1.52M	2.6M	228.54M	8.44M
Pigbutchering	853	182	4.13B	344.5M	401.57M	52.59M
Address poisoning	331	112	18.71K	11.88K	132.87M	10.34M
Scam	281	134	515.35M	104.57M	301.84M	24.3M
Phishing	58	150	21.04M	3.12M	230.46M	28.91M
Stolen funds	35	37	161.98M	86.03M	95.1M	13.57M
Illicit actor	14	17	52.26M	3.88M	26.4M	3.3M
Fake returns	8	11	26.15M	1.98M	56.01M	0.51M
Fake project	6	11	78.09M	4.88M	59.6M	0
Contract exploit	2	2	180.9K	3.79M	0	0
Donation scam	1	1	9.35K	99.99	0	0
Sim swap	1	1	50.04K	8.2K	25.91K	0
Total	3179	396	4.99B	555.38M	603.14M	109.06M

Table 6: Share of Value to DeFi Services after Asset Seizure

This table examines whether blacklisted addresses and their related addresses increase their use of DeFi services following an asset freeze. It presents the results from a difference-in-differences regression of the DeFi share of outflows, measured as the ratio of funds sent through DeFi services relative to total outflows. Specifically, the following regression is estimated.

$$\textit{DeFi Share}_{g,c,t} = \sum_{t \neq t_{freeze}} \beta_t \times \mathbb{1}(Month = t) \times Treat_g + \mu_c + \gamma_t + \varepsilon_{g,c,t}$$

where $\textit{DeFi Share}_{g,c,t} = \frac{\sum_{i \in g} DeFi\ flows_{i,g,c,t}}{\sum_{i \in g} All\ flows_{i,g,c,t}}$ and it denotes the share of flows sent to DeFi services for the treatment or control group g in cohort c at month t. Each cohort represents a freeze date. The treated group consists of the frozen addresses and their related counterparts. For every treated address, the control group is a random sample of 20 addresses that received inflows of at least $100 within the seven days prior to the freeze. The sample period is 12 months before and after the freeze date of each cohort. Standard errors are clustered by asset seizure cohort and month.

Dep. Variable:	DeFi Share	
	(1)	(2)
Treat × Post	0.266***	0.244***
	(0.0737)	(0.0773)
Treat	-0.169**	-0.164*
	(0.0797)	(0.0873)
Post	-0.0365	
	(0.0333)	
Constant	0.374***	0.356***
	(0.0274)	(0.0157)
Cohort FE	✓	✓
Event Time FE		✓
Observations	10,187	10,187
Adjusted R^2	0.561	0.572
Dep. Var. Mean	0.337	0.337
Dep. Var. Std	0.334	0.334

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Inflow to Tainted Deposit Addresses after the Binance and OKX Settlement

This table presents results from difference-in-differences regressions that test whether tainted inflows to deposit addresses at Binance and OKX declined relative to other exchanges following their respective settlements. Column (1) presents results for the Binance sample, and column (2) presents results for the OKX sample. In each sample, the treatment group consists of tainted deposit addresses at the focal exchange (Binance or OKX), and the control group consists of all other tainted deposit addresses at other exchanges. Specifically, the following regression is estimated.

$$\log\left(1 + Total\ Inflow\right)_{i,e,t} = \beta \times Post_t \times Treat_{i,e} + \mu_i + \gamma_e + \eta_t + \varepsilon_{i,e,t}$$

where $\log\left(1 + Total\ Inflow\right)_{i,e,t}$ is the log of one plus the tainted inflow received by deposit address i on exchange e in month t, and $Post_t$ equals one for months after November 2023 in the Binance sample and for months after February 2025 in the OKX sample. The regression include μ_i for deposit address fixed effects, γ_e for exchange fixed effects, and η_t for month fixed effects. Standard errors are clustered by deposit address and month.

Dep. Variable:	$\log\left(1 + Total\ Inflow\right)$	
	(1)	(2)
Treat × Post	-0.184**	-1.576***
	(0.0792)	(0.153)
Constant	3.160***	4.960***
	(0.0288)	(0.0118)
Address FE	✓	✓
Exchange FE	✓	✓
Year-Month FE	✓	✓
Sample	Binance	OKX
Observations	308,088	29,814
Adjusted R^2	0.229	0.420
Dep. Var. Mean	3.093	4.839
Dep. Var. Std	5.097	5.729

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

For Online Publication

Internet Appendix for: "Are Anti-Money Laundering Laws Effective?"

A. Data from Victim Reports

We collect 311,997 submissions primarily from chainabuse.com and hand-collected from other forum reports. Victims can choose the category that applies to their report or select "other." We use an OpenAI o3-mini model to classify victim reports.[35] We also rely on data partners who have collected addresses directly from scam victims in online groups as well as data vendors like etherscan.com.

We drop 724 invalid cryptocurrency addresses, as determined by being clearly too short (e.g., user submitting website addresses) or of invalid format (e.g., all valid Ethereum addresses start with "0x").

We remove 112,358 inactive addresses, the bulk of which were concentrated in extortion submissions with Bitcoin addresses. This leaves us with more than 389,000 addresses as seen in Table 1. Active addresses may also have been reported more than once and in more than one scam. In Column 3, we remove duplicate addresses within the same scam and find that the vast majority of extortion reports occur with duplicated addresses. Before utilizing our transactions in any blockchain analysis, we also remove any addresses that have been attributed to an identified exchange, service, or a destination with over 2,000 transactions. These may be inadvertently reported by victims if they follow their funds to their end destinations (e.g., reporting an exchange as the scammer). Lastly, to arrive at a unique number of addresses, we sort each address into one scam based on the scam that has the most reports for that address.

Figure IA.2 shows a scatterplot of the average number of transactions on the x-axis and the total inflows into these reported addresses on Bitcoin, Ethereum, and Tron, on the y-axis. The size of the circles is based on the number of active addresses reported. After contract exploit, pig butchering or romance scams are the largest category with more than $20 billion entering these addresses. Notably, these accounts may be understated as there could be general scams or stolen funds with not enough information for our algorithm to classify as pig butchering. Fake projects, phishing, stolen funds, pig butchering, and scams have the highest proportion of addresses on the Ethereum network. Airdrop scams lead to nearly 200 transactions on average but appear to have netted very little activity.

[35] We only keep model suggestions if the API returns a ranked probability of more than 99%. We follow the methodology from this notebook.

B. Additional Figures and Tables

Figure IA.1: Summary statistics on scam sizes

This figure shows the summary statistics of the reported addresses by types of scams. The left panel shows the number of reported addresses in each scam category, and the right panel shows the total dollar inflows received by these addresses. Bars are stacked by blockchain, with orange for Bitcoin (BTC), gray for Ethereum (ETH), and red for TRON (TRX).

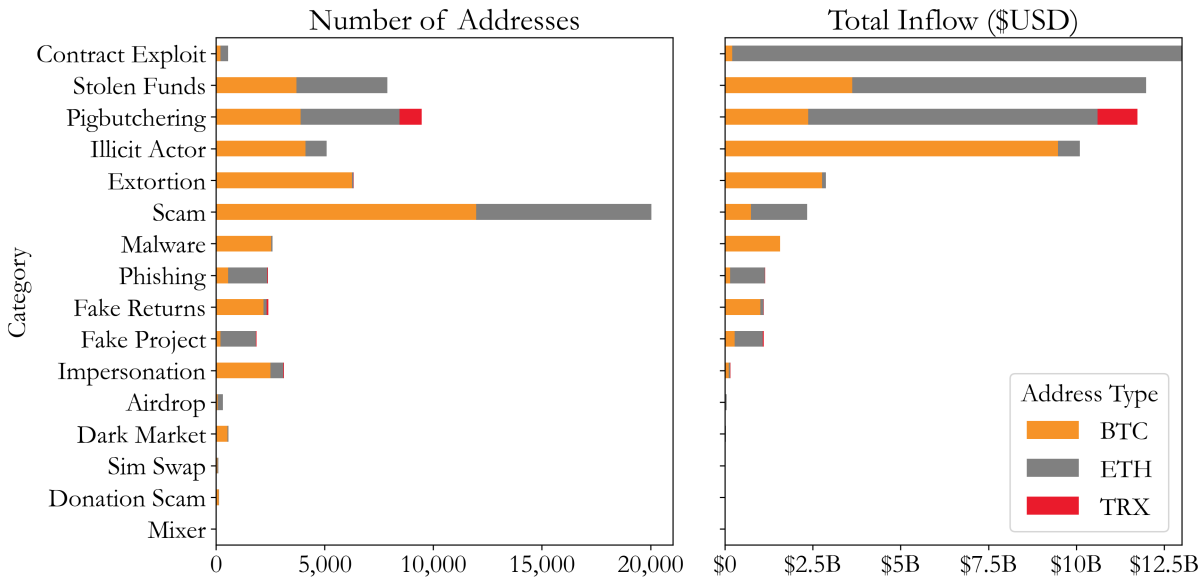

Figure IA.2: Inflows to Reported addresses

This figure shows inflows to reported addresses across different types of scams. Each point represents a specific type of scam. The vertical position (y-axis, logarithmic scale) indicates the total inflow value, and the horizontal position (x-axis) shows the average number of transactions per reported address. The size of each point corresponds to the total number of active reported addresses involved in that specific scam type. Each point contains a pie chart, which breaks down the blockchain composition of the addresses involved, specifically indicating fractions of Bitcoin (BTC), Ethereum (ETH), and TRON (TRX).

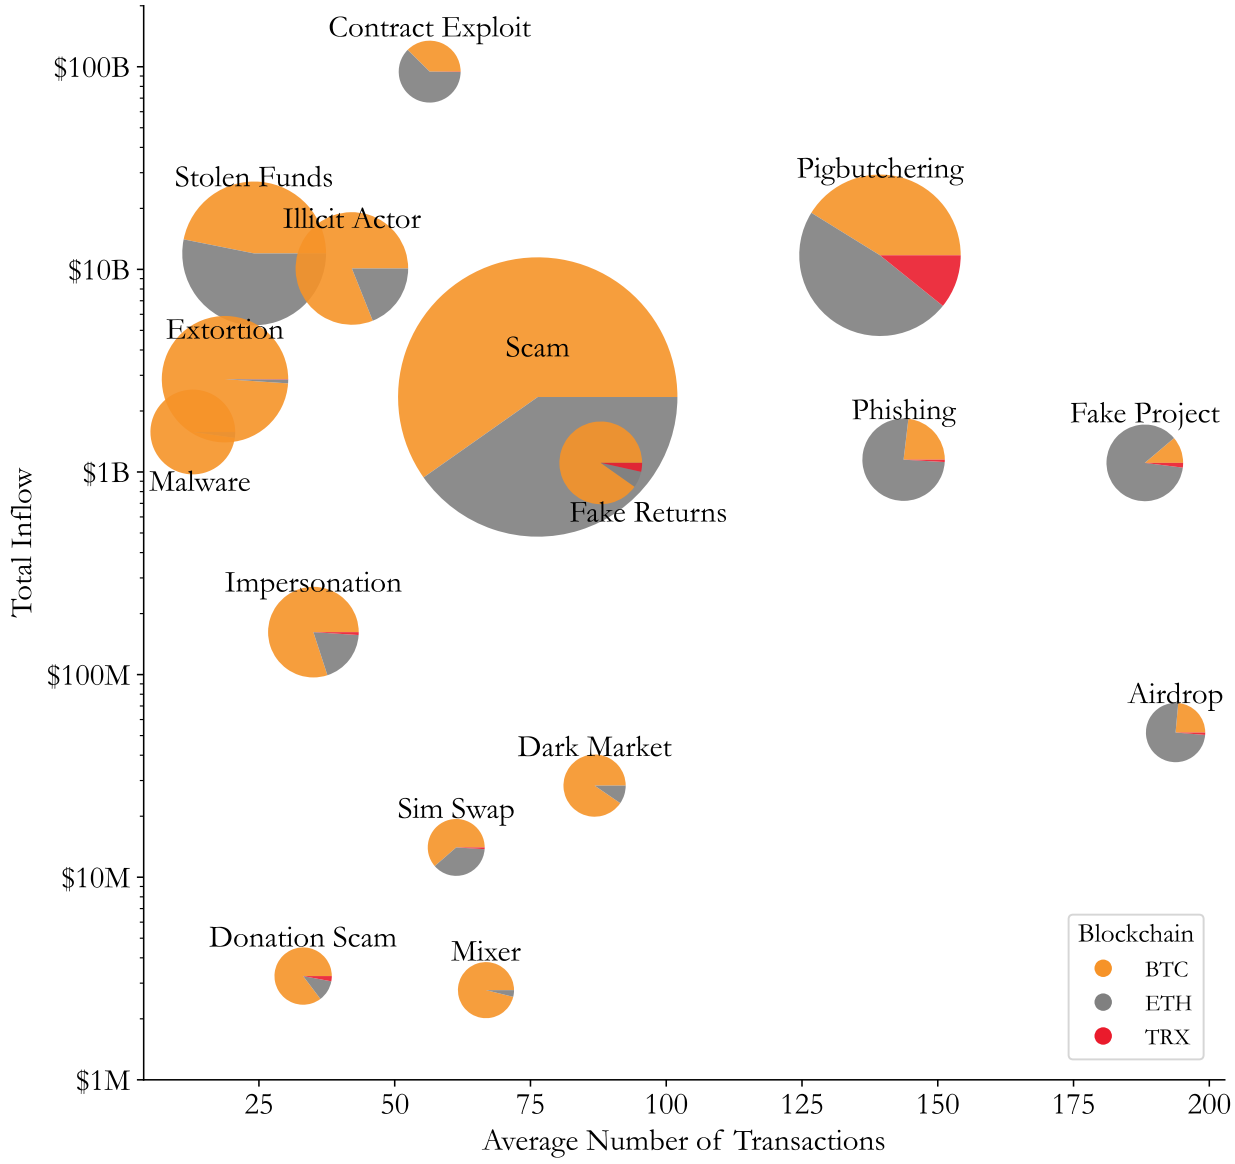

Figure IA.3: Hacker Flows to Tornado Cash vs. Other Services

This figure compares the monthly traced hacker flows to Tornado Cash versus other services over time. The sample consists of 5,867 unique hackers, identified from addresses labeled as contract exploits, stolen funds, or illicit actors. Each hacker's flows are tracked on Ethereum and split between Tornado Cash (red line) and all other services (blue line) to form a balanced hacker-month-destination panel. Points display the average monthly flows in $Log(1 + TracedFlows)$, and the error bars show 90% confidence intervals. Vertical dashed lines mark the May 6, 2022 sanction of Blender.io and the August 8, 2022 sanction of Tornado Cash by OFAC.

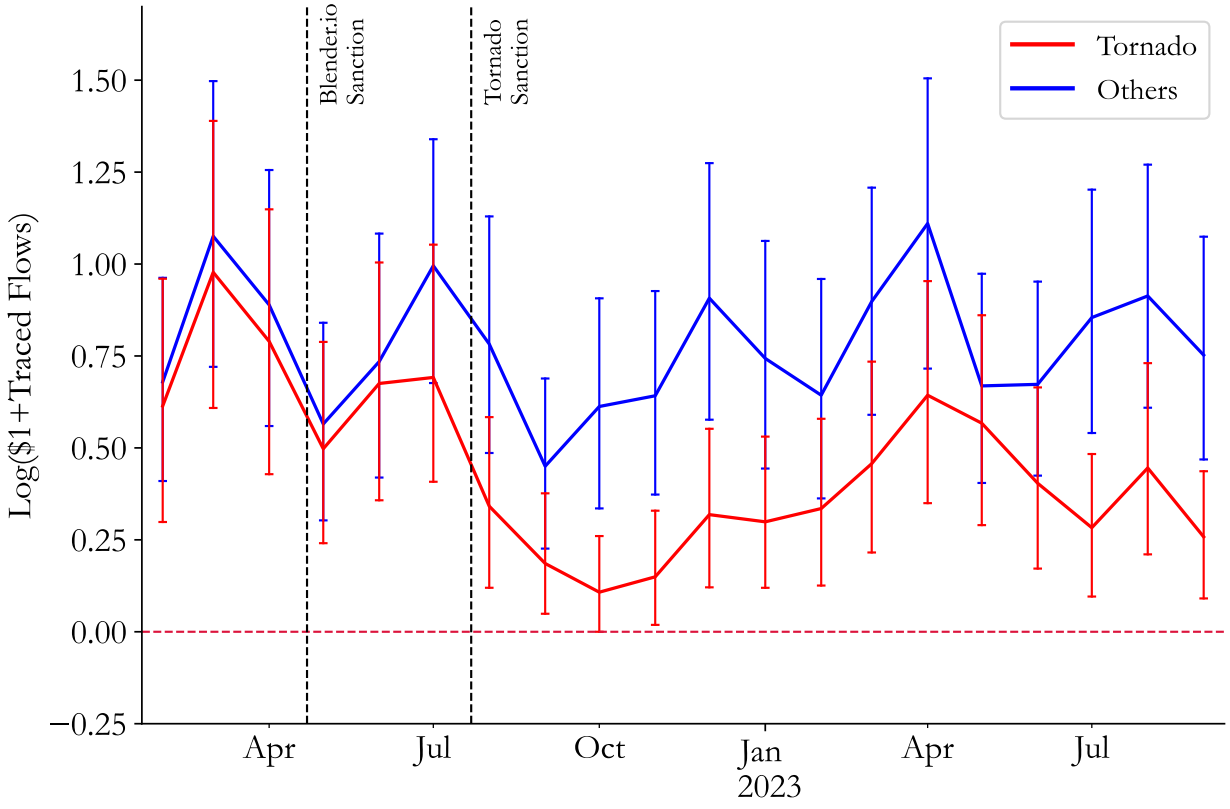

Figure IA.4: Destinations of Tornado Cash Outflows

This figure shows Tornado Cash inflows (Panel A) and outflows (Panel B) by hop. The y-axis lists the various entities found in each corresponding trace, and the hatches stand for the hops in each trace. The left panels are before the ban, the center panels are during the ban, and the right panels are after the ban was lifted. Hatches indicate hop distance from Tornado Cash.

Panel A: Tornado Cash Inflows

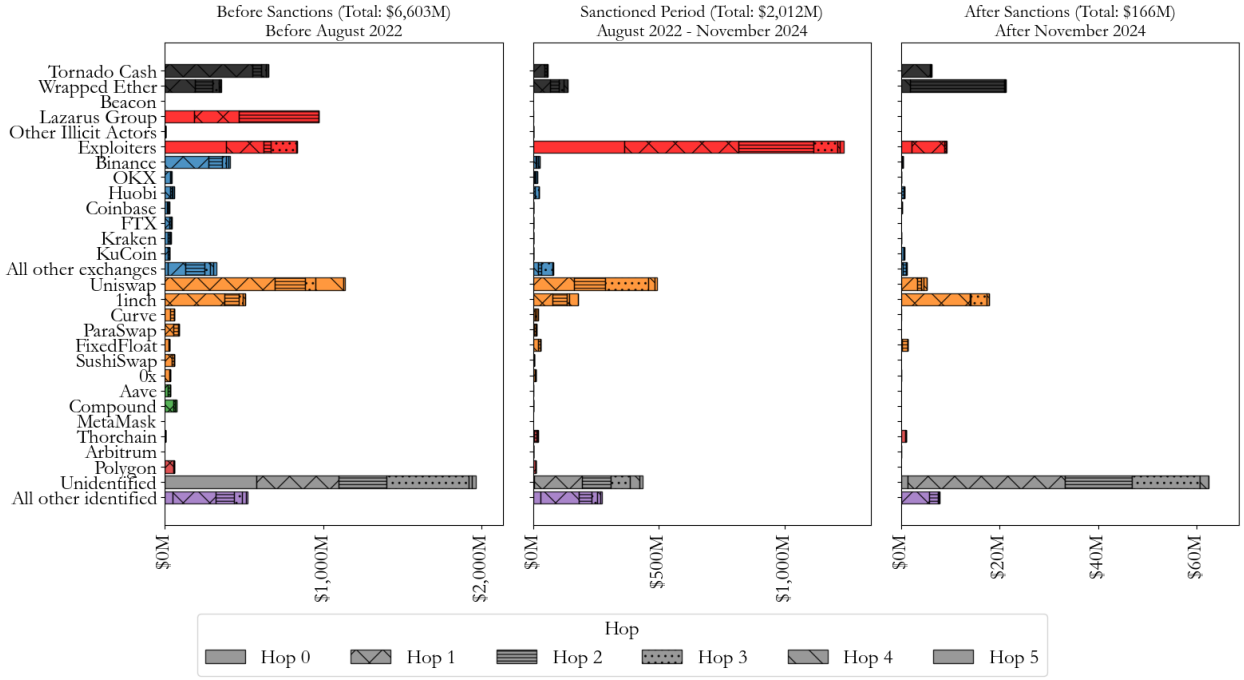

Panel B: Tornado Cash Outflows

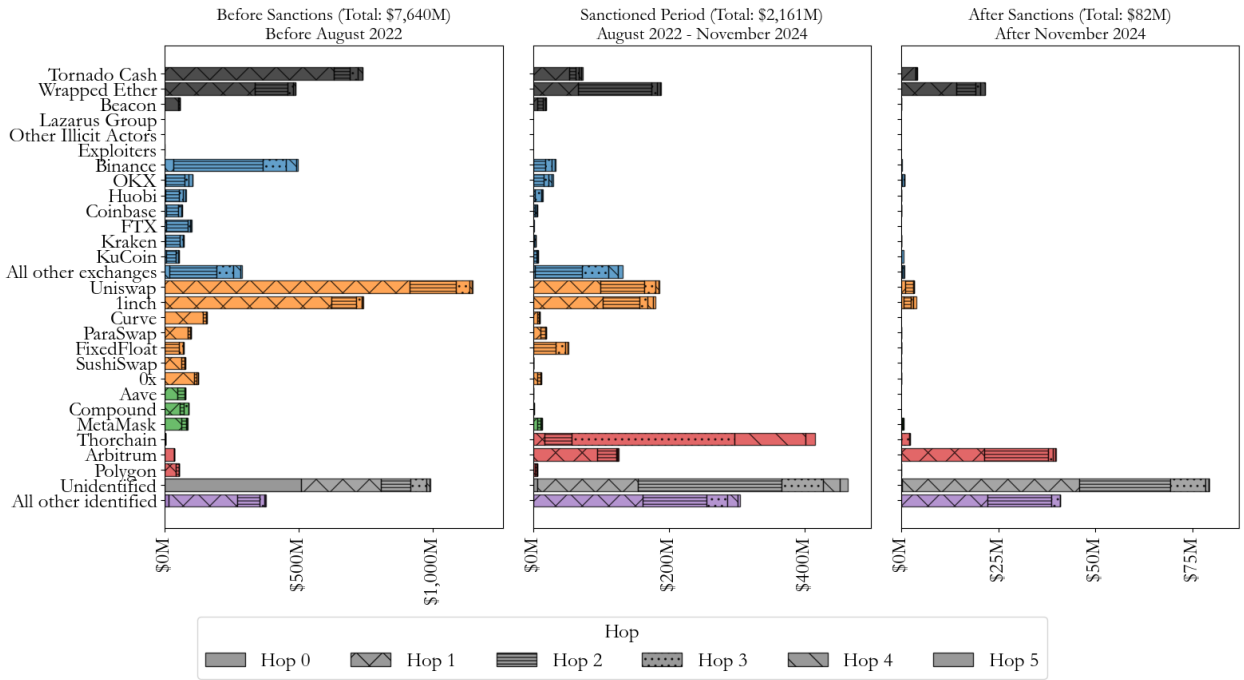

Figure IA.5: Transaction Costs of Moving Funds from Tornado Cash to CEXs

This figure compares the transaction costs of moving funds from Tornado Cash to domestic and overseas CEXs before and after the Tornado Cash ban, which was imposed on August 8, 2022. Transaction costs are first calculated for each individual path from Tornado Cash to the destination, and then aggregated by weighting each path's cost by the amount of funds moved into the respective CEXs. The classification of a path as pre-ban or post-ban is based on the end date of the path. Blue bars represent domestic CEXs and orange bars represent overseas CEXs. The 95% confidence intervals are shown.

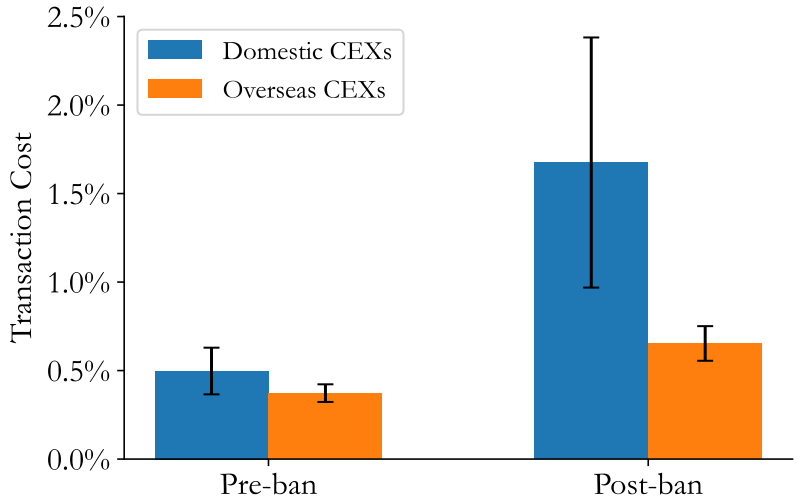

Figure IA.6: Top Tornado Cash Related Scams

This figure shows the total amounts deposited by the top 50 reported addresses into Tornado Cash. Each horizontal bar represents an individual reported Ethereum address and is colored according to its scam category. Hatches indicate in which hop of the backward trace the reported addresses forwarded the funds for depositing to Tornado Cash.

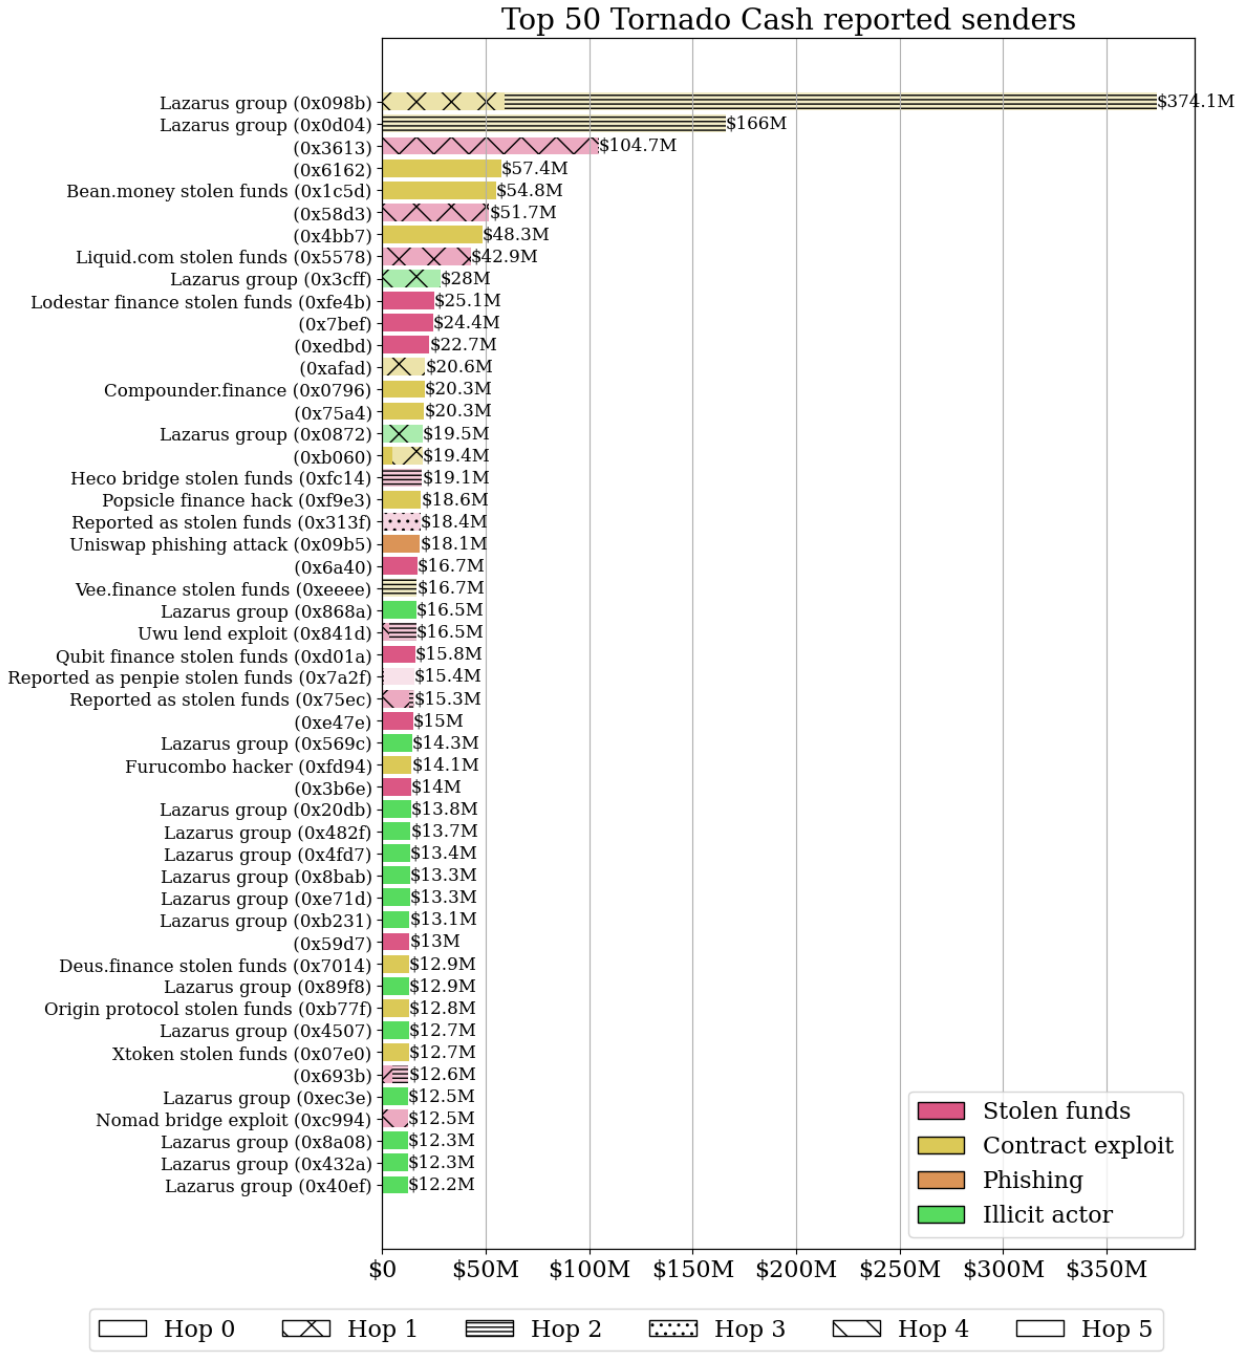

Figure IA.7: Effects of Tether Seizure on Transaction Volume by Hours

This figure presents results from a dynamic difference-in-differences regression that tests whether frozen addresses and their related addresses increase transaction volume after the account has been frozen by Tether. Related addresses to the frozen ones are found through gas clustering. The treated group consists of the frozen addresses and their related counterparts, while for every treated address, the control group is a random sample of 20 addresses that also received Tether in the preceding three days. Specifically, the following regression is estimated: plots the coefficients from the following regression:

$$y_{i,t} = \sum_{t \in [-24,24]} \beta_t \times \mathbf{1}(Hour = t) \times Treat_i + \mu_i + \gamma_t + \varepsilon_{i,t}$$

where $y_{i,t}$ denotes the transaction volume or transaction count of address i in hour t, $Post_t$ is a binary variable that equals to one for hours after the freeze, μ_i are address fixed effects, and γ_t are event time fixed effects. Standard errors are clustered by cohort, address, and event time. The top panel plots the regression coefficients when transaction dollar volume is the dependent variable and the bottom panel presents number of transactions.

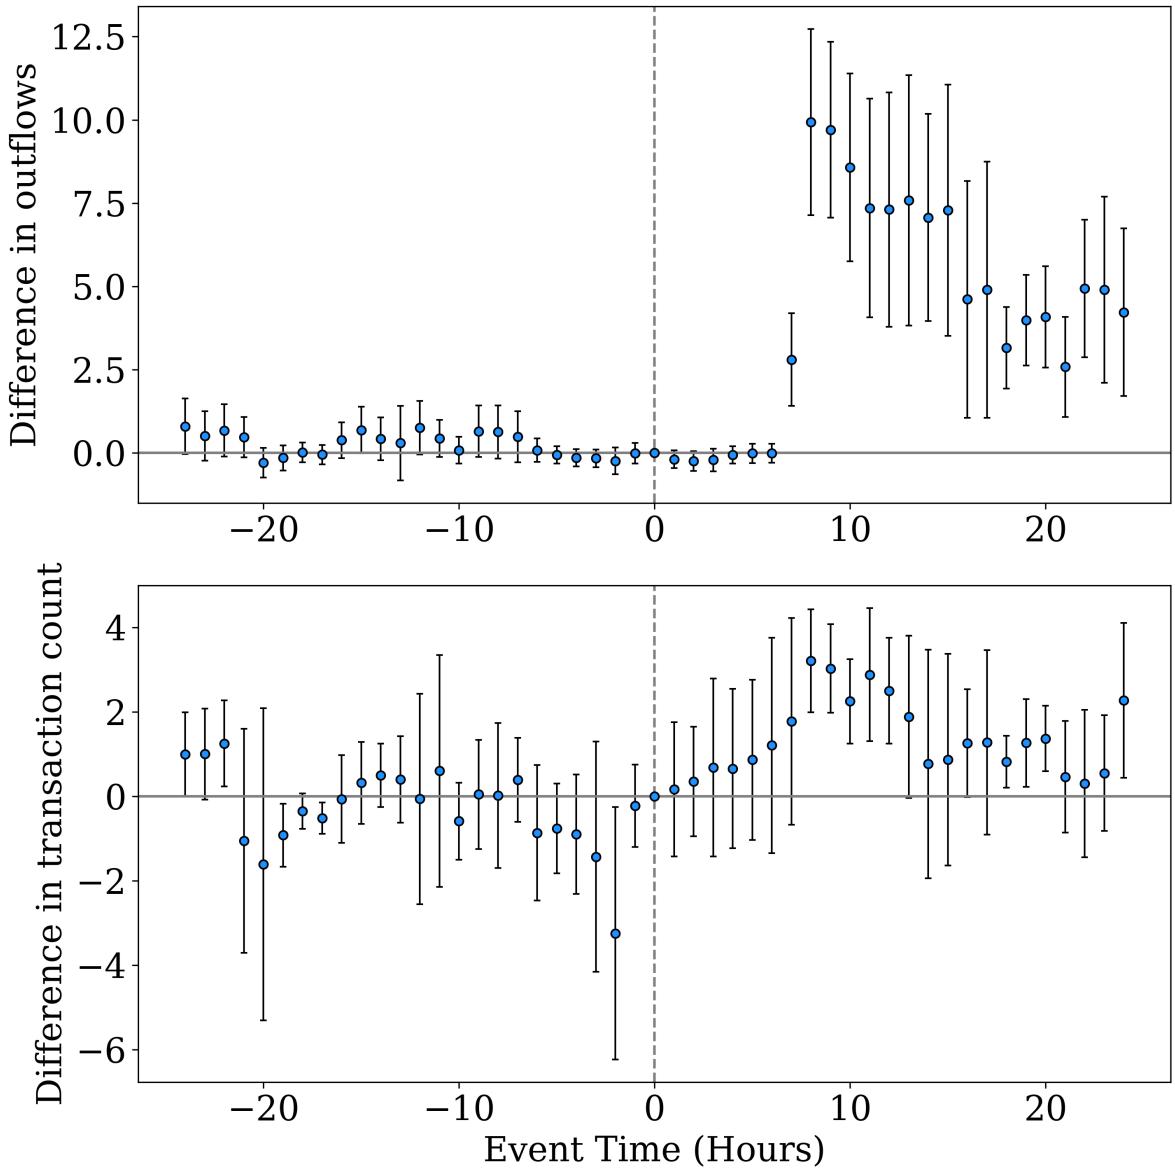

Figure IA.8: Tainted Flows to OKX around OKX Settlement

This figure examines whether tainted flows to OKX declined following the OKX settlement in February 2025. Panel A shows monthly inflows to OKX tainted deposit addresses (dashed line) compared to inflows to all other tainted deposit addresses (solid line). Panel B presents the estimated coefficients from a difference-in-differences regression at the deposit address-month level. The treatment group is composed of tainted OKX deposit addresses, and the control group consists of tainted deposit addresses at all other exchanges. The regression is of the form:

$$\log\left(1 + Total\ Inflow_{i,t}\right) = \sum_{t \neq t_{Nov2023}} \beta_t \times \mathbb{1}(Month = t) \times Treat_i + \mu_i + \gamma_e + \eta_t + \varepsilon_{i,t}$$

where $\log\left(1 + Total\ Inflow_{i,t}\right)$ denotes the log of one plus the total tainted inflow received by deposit address i in exchange e in month t. The regression includes deposit address fixed effects, exchange fixed effects, and month fixed effects. The red dashed line corresponds to the OKX settlement date. Standard errors are clustered by deposit address and month.

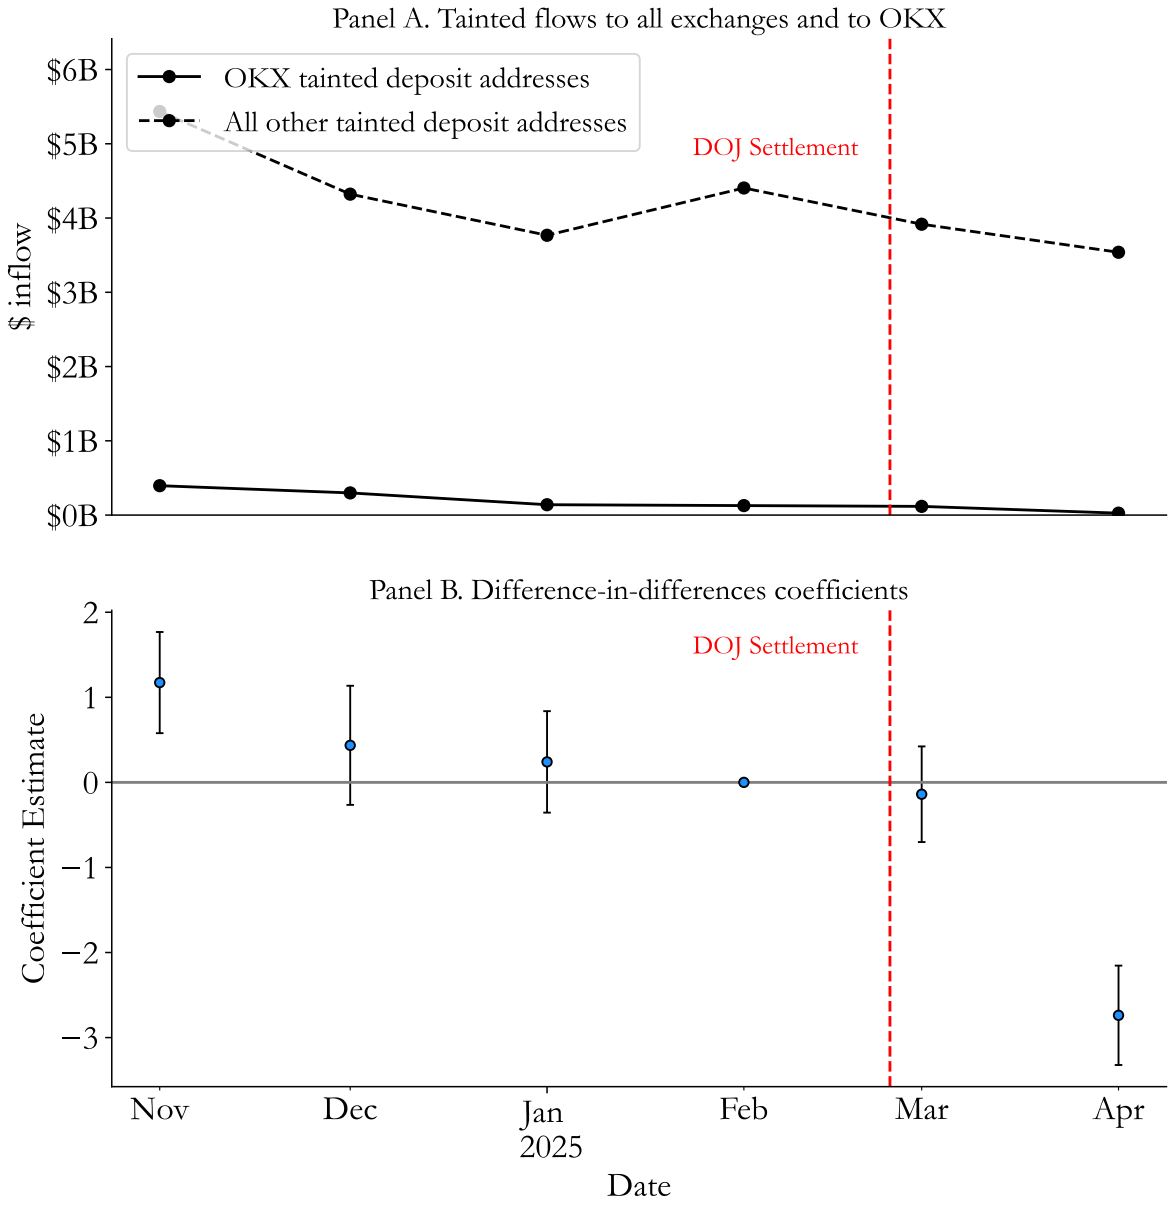

Figure IA.9: Traced Flows to Destinations

This figure shows destinations for each traced scam on the Ethereum blockchain. Panel A shows the totals for each scam. Each horizontal bar is broken by the percentage of the total amount traced to each category, which is reflected in the lower-x axis. The corresponding red dots to each bar indicate the total amount traced for each scam, which are represented in the upper x-axis. Panel B shows the category flows as a percentage of the total monthly flows.

Panel A: Traced Flows by Scams

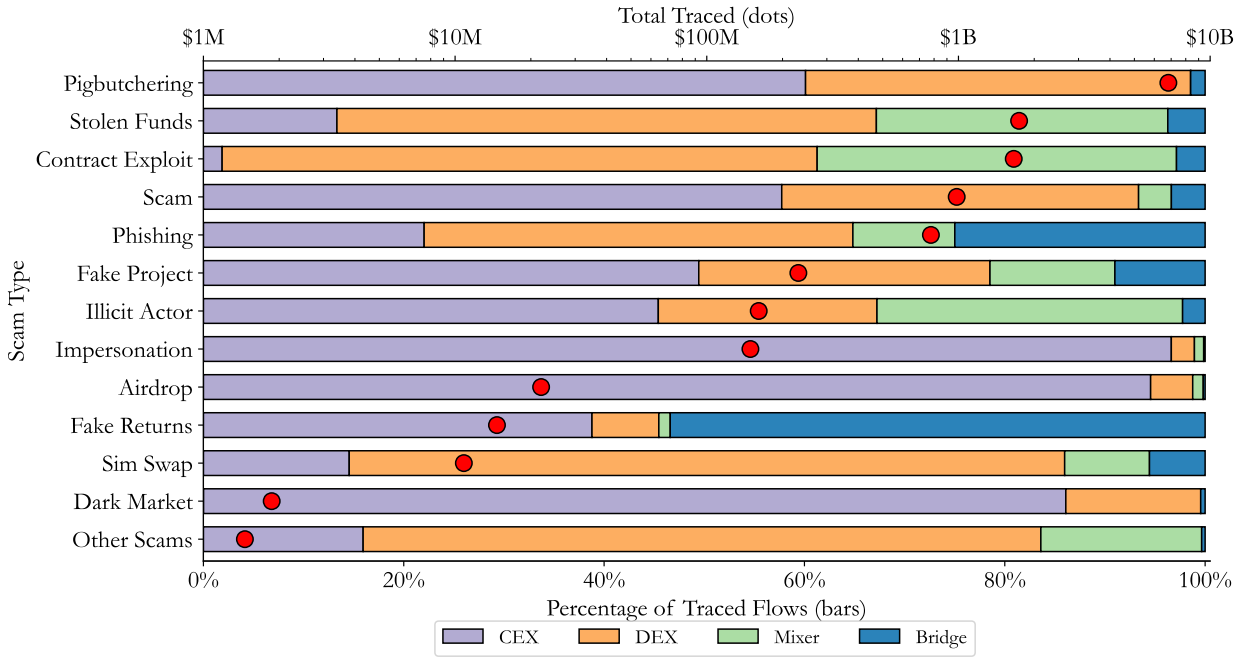

Panel B: Traced Flows by Months

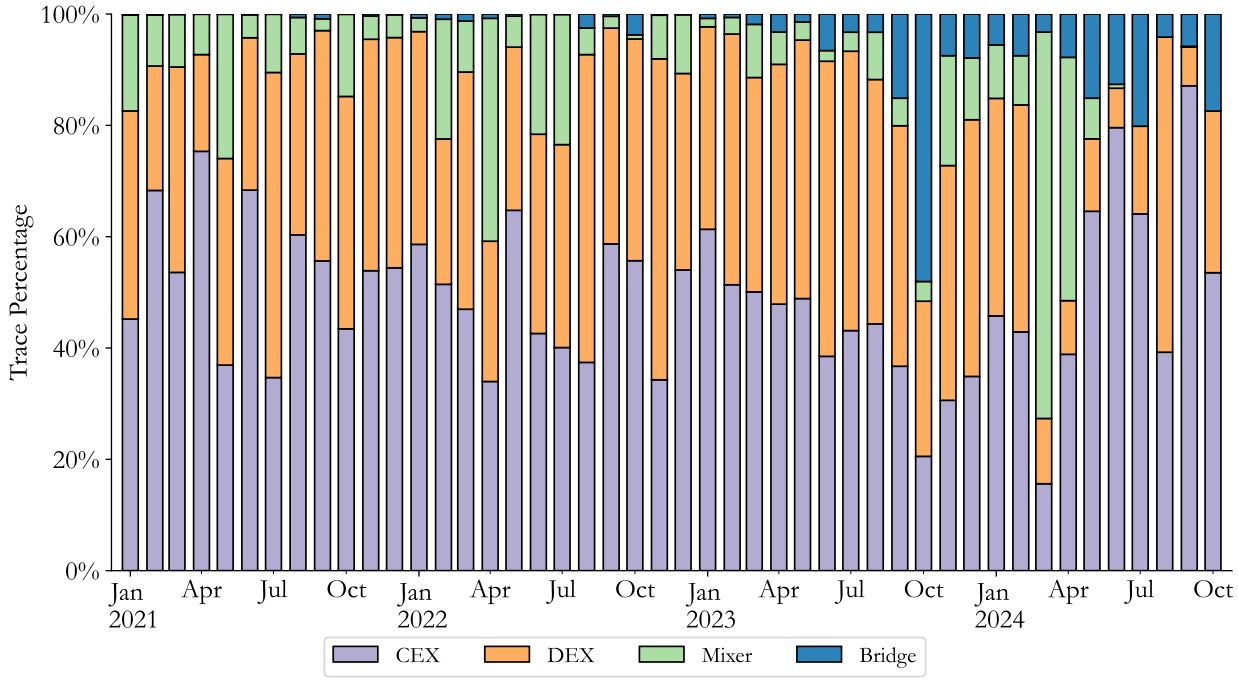

Table IA.I: Scam definitions

Type	Description
Extortion	Using intimidation to force someone to pay money under the threat of releasing sensitive information
Phishing	Tricking individuals to providing personal information, such as login information, often by driving traffic to a malicious website to gain access to wallets and secret keys
Impersonation	Imitating a famous person, brand, or organization for fraudulent purposes, often to ask for money or drive traffic to a fraudulent website to steal sensitive data, including "Nigerian Prince-type" scams where a rich foreigner promises to share the profits from a fake investment opportunity
Pig butchering	Confidence scam that lures victims into fabricated cryptocurrency investment schemes often by gaining a victim's affection, spoofing exchanges, and manipulating victims with fake returns or taxes
Donation scam	Posing as charities to exploit donors
Airdrop	Unsolicited mass distribution of a cryptocurrency token to direct investors to a malicious website
Fake returns	Promising fake returns to convince victims to invest through: promising fake payments - payout scam or load up scam faking returns by paying profits to earlier investors with funds from more recent investors - Ponzi scheme fraudulently and artificially inflating the price of a cryptocurrency, such as during a pump and dump scam - using misleading information or other techniques
Fake returns	Pretending to build a fake project only to abandon the project after attracting investors, such as rug pulls or exit scams
Sim swap	Convincing a mobile carrier to port a victim's phone number to their SIM card to gain access to sensitive accounts
Contract exploit	Exploiting vulnerabilities in a smart contract or other protocols to drain funds
Malware	Inserting malicious software into victim systems that can be used to either extract private information or to extort victims for ransoms
Stolen funds	All other unidentified means of pilfering an online entity's wallet, including through social engineering or gaining access through hacks
Illicit actors	Wallets of known sanctioned entities, including state and non-state actors, for uses like terrorism financing or software exploits
Dark market	Wallets linked to platforms that exchange CSAM or illicit goods
Scam	Broad category for other reports of scams, as categorized by a data vendor, but provide no further details

Table IA.II: Effects of Tether Seizure on Transaction Volume

This table presents results from a difference-in-differences regression that tests whether frozen addresses and their related addresses increase transaction volume after the account has been frozen by Tether. Related addresses to the frozen ones are found through gas clustering. The treated group consists of the frozen addresses and their related counterparts, while for every treated address, the control group is a random sample of 20 addresses that also received Tether in the preceding three days. Specifically, the following regression is estimated:

$$y_{i,t} = \sum_{t \in [-24,24]} \beta_t \times Treat_i \times Post_t + \mu_i + \gamma_t + \varepsilon_{i,t}$$

where $y_{i,t}$ denotes the transaction volume or transaction count of address i in hour t, $Post_t$ is a binary variable that equals to one for hours after the freeze, μ_i are address fixed effects, and γ_t are event time fixed effects. Standard errors are clustered by cohort, address, and event time. This table correspond to results in Figure IA.7.

	(1) log(Transaction Volume)	(2) Transaction Count
Treat × Post	4.666***	1.360***
	(0.474)	(0.252)
Address FE	✓	✓
Event Time FE	✓	✓
Observations	9161730	9161730
R-squared	0.699	0.501
Number of Clusters	49	49

Standard errors in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$