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

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

September 23, 2025

Abstract

In a rapidly changing crypto landscape, it is unclear whether anti-money laundering policies accomplish their intended goals or are inexpedient activities. We measure criminal responses to recent enforcement actions in crypto by tracing illicit flows throughout 5.16 TB of Ethereum blockchain transactions. First, the 2022 OFAC sanctions against the Tornado Cash mixer resulted in a 60% decline in monthly volume, 33 basis points in incremental obfuscation costs, twice the probability of detection, and fewer interactions with both Western and less-regulated exchanges. Hackers and related flows switched to swaps and bridges, which are more traceable. Second, OFAC sanctions against other addresses are uncommon but have resulted in over 100,000 BTC and 155,000 ETH stuck on-chain. Third, over $1.34 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 are often necessary. Overall, our results indicate that sanctions, particularly those regarding services, have been effective in seizing funds, moving funds to more traceable and seizable 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 improvement that can help guide recent crypto policy and legislative proposals.

JEL classification: G18, G23, G28, K24, K42
keywords: Anti-Money Laundering Policies, Crypto Flows, DeFi, Tornado Cash, Stablecoins

[*]We thank Zahi Ben-David, Darrell Duffie, Cesare Fracassi, Zhiguo He, Sam Kruger, Yiming Ma, Tom Meling, Andrey Ordin, Aaron Pancost, Alex Pettyjohn, Amin Shams, René Stulz, and Qiang Wang (discussant), as well as conference and seminar participants at the Northern Finance Association (NFA) 2025 Annual Meeting, SITE Financial Regulation, Ohio State University, and the University of Texas at Austin for their helpful comments. We especially thank Juan Antonio Artero Calvo for excellent research assistance. We additionally thank Integra FEC for the 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. What are the economic magnitudes of their impact? Through which channels do they operate? What substitutes emerge in response? 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, particularly following September 11, 2001, as controls to reduce terrorist financing and other 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 AML policies and their enforcement is to increase the probability of detecting both criminals and money launderers, and to punish to the extent that it results in the forfeiture and possible fine of funds (Ferwerda, 2009).

Evaluating the effectiveness of aspects of the AML 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 AML laws matter: A) the conventional view, which holds that AML laws reduce crime by blunting their profitability, and B) the skeptics view, which emphasize that efforts are best spent on targeting the root sources of crime and not on regulating money services.[1] A firm's reputational risk may also serve as a sufficient deterrent to money-laundering activities. 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 cryptocurrency environment, we compile a large amount

[1]Consistent with this view, Cuéllar (2002) argues that there is a tenuous relationship between AML 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 are few 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.

of detailed data on illicit funds that is typically private within banks and not available for researchers or large-scale outside analysis.

In this paper, we empirically evaluate the effectiveness of crypto AML policies and procedures by utilizing the many regulatory and court orders as shocks. We first compile a large dataset of 256,958 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. Using Google Cloud BigQuery, we process 5.16 TB of Ethereum transaction and event data to trace the flow of illicit funds through the blockchain. We then examine how centralized and decentralized exchanges (DEXs) interact with traced addresses before and after government sanctions. Specifically, we use the traced network to analyze two sets of policies: policies targeting *services*, including the high-profile 2022 Tornado Cash ban and DOJ settlements with exchanges, and policies targeting *individual users*, including OFAC sanctions, stablecoin seizures, and the $\$10,000$ reporting threshold.

Our findings show that Tornado Cash, an Ethereum mixer that obfuscates the origin of funds, handles considerable criminal flows from more sophisticated cyber-criminal gangs, with the Lazarus Group from North Korea being the largest identified user. Before its August 2022 U.S. sanction, Western centralized exchanges received 3.99% of Tornado Cash outflows on average, a proportion that declined to 2.14% in the twelve months following the ban, representing a 46.20% reduction. Flows to overseas centralized exchanges did not decline immediately after the ban but began to fall by the end of 2023, after which bridges became the dominant destination for Tornado Cash outflows. 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. The probability of correctly matching transactions from the largest mixer user during a given week also doubles after the ban (from 10.15% before the sanctions to 21.10% afterward), indicating that the decline in volume renders the mixer less useful for obfuscation. Additionally, the reduction in transaction volumes is

more pronounced for Western exchanges relative to overseas exchanges.[2] DeFi swaps and bridges take flows 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, their flows to Tornado Cash decrease by 30%. Rather than using other mixers, they typically move money through DeFi swaps and bridges, outlets that 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 potentially important enforcement mechanism is the freezing of criminal assets. We examine all addresses on the OFAC sanctions list, which includes 380 Bitcoin addresses and 60 Ethereum addresses. Although sanctions are often imposed too late and after funds have left addresses, sanctions nonetheless prove effective in trapping some assets on-chain. We find that approximately 195,000 Bitcoins and 420,000 ETH have passed through these sanctioned addresses, with 51.28% of the Bitcoin and 36.90% of the Ethereum flows remain seemingly stuck on their respective blockchains.

Perhaps more importantly, certain law enforcement agencies and courts have been successful at requesting Tether to freeze assets. In total, we find that over $1.34 billion Tether and $93.12 million USDC have been frozen. Of this, $608.65 million is 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 after freezes increases by approximately 25%, presumably to obfuscate their flows. This indicates that, in addition to costing the criminals their funds, seizures can impose additional costs on criminals in the form of costlier and less regulated services.

We then consider whether DOJ settlements with exchanges deter 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 becomes evident only after the exchange reached settlement agreements during our sample period.

[2]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 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. Additionally, 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 majority of these funds across multiple different wallets.

Finally, we compare the deposit patterns of criminal and other users and find that tainted deposit addresses disproportionately use round-number amounts, being about 12–13 percentage points more likely to use deposit amounts exactly divisible by $500 or $1,000. Tainted users show additional bunching in round-number bins just under the $10,000 Suspicious Activity Report (SAR) threshold, at Western (but not overseas) exchanges. Thus, criminals do appear to have some concern for this reporting threshold, though it is not clear if this threshold is an effective mechanism to detect criminal behavior.

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. Policies targeting services (e.g., mixers and exchanges) appear more effective than actions against users; Actions against individual users are used relatively infrequently and often too late, only after funds are off-ramped. Second, self-policing appears ineffective given the wide range of practices by crypto exchanges in handling criminal flows. Only after settling federal charges did overseas exchanges significantly decrease their flows in OFAC sanctioned addresses. Additionally, overseas exchanges still handle significant

tainted crypto flows. 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 mostly viewed as clean money. Fourth, many large centralized exchanges need to do more monitoring of funds and have more rigorous KYC and KYT procedures. Otherwise, criminal flows can substitute money laundering destinations by using exchanges with more lax AML regulations.

Lastly, we speak to ongoing policy debates. On April 7, 2025, a DOJ order by Deputy Attorney General 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 flows to transparent places and empowering exchanges to monitor their flows. The DOJ's goal of pursuing organized crime, including foreign actors, may be substantially impeded without targeting mixing and tumbling services. Additionally, the judge issuing the recent ruling that OFAC did not have the authority to ban code noted that "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, in the U.S., the GENIUS Act was signed on July 18, 2025, which requires stablecoin issuers to comply with U.S. law enforcement requests for asset seizure.[3] Our analysis shows that seizers can be an effective tool and that 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 AML regime in Italy. Mirenda, Mocetti, and Rizzica (2022) show how organized crime utilizes cash and shell companies to

[3] See Congress website for more details about the the Guiding and Establishing National Innovation for U.S. Stablecoins Act of 2025, or GENIUS Act.

obfuscate transactions entering the banking system. Moore, Clayton, and Anderson (2009) survey the economic structure of online crime and recommend more private data sharing and police enforcement focused on online gangs.[4] 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 AML 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 the variation of policies across countries. 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. 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.[5] Griffin and Mei (2025) map 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 var-

[4] 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 AML enforcement impacts lending by U.S. banks. Dirty money also distort macroeconomic capital allocations (Tanzi, 1996; Quirk, 1996).

[5] 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. They seemingly do not count flows not other closely related addresses.

ious crypto investment scams, Ponzi schemes, ransomware, money laundering, and dark markets.[6] We provide the first examination of AML laws. We also extend the literature by utilizing a comprehensive database of various types of criminal activity, and the analysis also leads to a fuller understanding of which laws might be effective and how criminals evade detection.

II. Data, Background, 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 and Ethereum. Transaction-level data is 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, the background of crypto AML policies, and the tracing methodology.

A. Data

We primarily use two types of data: blockchain transaction data and attribution data. Blockchain transaction data is sourced from the Bitcoin and Ethereum blockchains. Importantly, the fields include the blockchain address of the sender and receiver, which allows us to construct paths of flows between addresses. We study both Bitcoin and Ethereum, with slightly greater emphasis on Ethereum because that is where Tornado Cash, a high-profile mixer, exists. We analyze the Bitcoin blockchain for an analysis of additional OFAC addresses, the Bybit hack, and for additional robustness. We use data from coingecko.com on end-of-day cryptocurrency prices to convert Ethereum and Bitcoin values to dollars.[7]

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

[6] 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, Shin, and Wang (2025), Hamrick, Rouhi, Mukherjee, Feder, Gandal, Moore, and Vasek (2021), and 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).

[7] When converting cryptocurrency value to dollars, we assume prices for the stablecoins Tether, USDC, and DAI are always $1. This subset of currencies constitutes the vast majority of cryptocurrencies we see used in our sample.

or bridged. Swapping is typically when an address uses a service such as a decentralized exchange to 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. On BigQuery, the total Ethereum blockchain amounts to over 5 TB of data.[8]

We use a rich set of sources of attribute data. First, we use data from online sources, such as blockchain.com and etherscan.io, to label addresses that belong to known entities. We focus on service providers such as centralized exchanges, decentralized exchanges, bridges, or mixers. Second, we collect all sanctioned addresses by OFAC, as well as all Ethereum-based addresses where the stablecoin issuers of Tether and USDC have seized assets. Third, we use data on known illegal actors reported by various data collectors. These are mainly reported by victims, discovered by law enforcement, or filed in lawsuits. After dropping reports of insufficient detail to categorize the type of scam, we use 256,958 reported addresses, 37,773 for Bitcoin, and 219,185 for Ethereum. The single largest source of addresses is chainabuse.com, a leading reporting platform where victims and other users describe hacks and scams. To standardize labels across sources, we first hand-label 425 victim reports and fine-tune a GPT-4o-mini large language model (LLM) on this set. We then apply the fine-tuned LLM model to assign scam categories to the remaining reports, ensuring consistent labeling across sources. Table 1 presents summary statistics on reported addresses, tabulated by scam category and blockchain.[9] 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]

Overall, we find that addresses reported as part of stolen funds have the largest total inflow, followed by pig butchering scams, illicit actors, and contract exploits. In Table 2, we present summary statistics for each scam and the averages. We count the flows into these addresses from January 2020

[8] The Ethereum blockchain data consists of two primary tables on Google Cloud BigQuery: the *transactions* table (2.59 TB), which records transaction-level details such as sender, recipient, ETH value transferred, and contract interactions; and the *logs* table (2.57 TB), which stores all smart contract event logs, including token transfers, swaps, and other contract events, emitted during transaction execution.

[9] The Internet Appendix includes additional details about the distribution of and nature of these reports.

[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.

to January 2025. The total inflow to these addresses is a total of $52.61 billion, with $21.46 billion from Bitcoin and $31.15 billion in Ethereum. These beginning addresses should not be used to scope the total amount of activity since these incoming amounts do not capture unreported scam addresses.

B. Crypto AML Policy Background

After the terrorist attacks on September 11, 2001, the U.S. extended the existing anti-money laundering framework of the Bank Secrecy Act and increased cooperation across countries through new international initiatives from the Financial Action Task Force (FATF). Anti-money laundering laws and enforcement may be the main means of deterring activity in important settings such as terrorism and crimes committed by foreign nationals, where authorities may not have the ability to apprehend criminals in non-cooperative nations. Many of the basic principles are being questioned by a crypto industry born as an alternative financial system that is growing in size and political influence.

Cryptocurrency's appeal seemingly lies in its potential as an alternative financial system free from regulation. However, 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. The trade-off is that 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 intermediaries.

The regulatory and enforcement environment is in flux. 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 a precedent would put undue

responsibility on developers to prevent criminals from using their services.[11] 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.[12,13] 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 AML procedures. As part of their settlements, the exchanges agreed to more rigorous monitoring procedures and DOJ oversight. We use these shocks to study how users and exchanges respond to AML enforcement actions.

C. Methodology

We use blockchain data to construct a network of related addresses and analyze their characteristics in order to understand how the network responds to AML regulations. In this subsection, we describe the methodology used to identify related addresses, with *tracing* serving as the primary tool, and gas and deposit address clustering utilized.

C.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 are commonly used by law enforcement, through service providers like Chainalysis and TRM Labs, to follow capital flows. We apply a suite of bulk tracing algorithms used and more fully described by Griffin and Mei (2025) in the context of pig butchering scams.[14] 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

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

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

[13] 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.

[14] These tracing algorithms have been developed and maintained by Integra FEC. The algorithms are essentially a program with a series of steps to follow crypto fund flows to stopping points such as centralized exchange hot wallets. There are separate algorithms per blockchain for tracing funds forward and tracing funds backward.

are commingled in a downstream address that contains flows from other sources, then the forensic researcher must choose how to follow subsequent outflows. We follow commingled funds on a "first-in-first-out" (FIFO) basis, a well-established and 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 and their respective funds. Consequently, traced funds will always follow strictly less than the initial inflow to the originating source.

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.

C.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 currency of Ethereum. The rationale is that all addresses need Ether because blockchain transaction costs can only be paid in Ether. Therefore, we associate the first funder as a way to link potentially 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 it is 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. 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 reduced 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) 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

[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.3 shows destinations cross-tabulated by crime type.

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 discernible. Users appear to funnel more money from decentralized exchanges (i.e., Uniswap, 1inch) than from 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 100,000 ETH ($400 million) per month, but after the ban by October 2022, volume was less than 40,000 ($50 million) per month, or a more than 60% decline. Interestingly, the mixer becomes more attractive for criminals with a 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 weekly tainted criminal address activity as a percentage of the total activity, as shown by the red line in Figure 2. The total flows vary widely, but at times show a sizable fraction of the flows due to criminal activity. In the six months following the sanction, an average of 26.49% of inflows to Tornado Cash are due to criminal flows, more than double the 11.79% observed in the six months prior. This percentage rose further in 2024, reaching an average of 49.62% between March and October. 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 because our sample likely does not capture the entirety of criminal activity. Further, criminal flows are coming from sources that either cannot be 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 originates from criminal activity jumps considerably. Of the identified criminal flows to Tornado Cash, Figure IA.6 plots the top 50 individual senders featuring the North Korean Lazarus Group as first.

C. Is Tornado Cash Less Anonymous Following the Sanctions?

Although Tornado Cash is designed to sever the link between deposits and withdrawals, when liquidity is low, large flows may be more detectable. We investigate how the ban on Tornado Cash affected this aspect of anonymity by estimating, for each week, the probability that a one-ETH withdrawal can be matched to the largest depositor. Perfect attribution is inherently difficult given Tornado Cash's core function of obfuscating transaction flows, so we estimate expected withdrawals by spreading each depositor's inflows across future days according to a declining schedule.

For every depositor, we first observe the daily amount entering Tornado Cash. We predict expected outflows by distributing each depositor's daily inflows into future daily withdrawals using a declining schedule modeled as a Pareto decay function, $f(\Delta) = (1 + \frac{\Delta}{\tau})^{-\alpha}$, where $\alpha = 0.3534$ and $\tau = 30$ days.[17] This declining function assumes that each depositor withdraws a larger fraction of their deposited funds shortly after making the deposit, with the fraction declining progressively for withdrawals made in later periods. By applying this withdrawal projection to each individual depositor's daily inflows, we obtain predicted daily withdrawal amounts for each depositor, which, when aggregated across all depositors on a weekly basis, closely match the actual observed weekly withdrawal totals.

Given that each predicted withdrawal is directly associated with the original depositor, we can calculate the probability of correctly attributing a randomly selected one-ETH withdrawal during week t to the top depositor as follows:

[17] We test six different withdrawal distributions and find that the Pareto distribution with parameters $\alpha = 0.3534$ and $\tau = 30$ days provides predicted weekly withdrawal amounts that most closely match the actual Tornado Cash outflows. Additionally, Béres, Seres, Benczúr, and Quintyne-Collins (2021) document that approximately 70% of Tornado Cash users of the linked deposit-withdraw pairs withdraw their funds within one day of deposit, consistent with the rapid withdrawal schedule implied by our estimated distribution.

$$\text{Prob (Detection)}_t = \frac{\text{predicted withdrawals by the top depositor in } t}{\text{predicted withdrawals in } t}.$$

We calculate this probability weekly and report its ten-week rolling average, as shown by the blue line in Figure 2. Following the August 2022 U.S. Treasury sanctions, the average probability of matching withdrawal transactions to the top depositor significantly increased from 10.15% before the sanctions to 21.10% afterward. This increase primarily resulted from the sharp decline in total inflows after the sanctions, which caused the top depositor's share to become more prominent. Note that there are additional ways to track transactions, such as transaction reporting, timing differences, and transaction clustering that our basic method does not employ, but further detailed analysis of certain transactions might further increase the probability of detection. Thus, the sanction not only reduced Tornado Cash's overall usage but also made it less effective at obscuring fund source identities.

D. 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, as seen in Figure IA.3. We broadly label 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 is the flows from each labeled hack, to each destination, per month.[19] To formally test the effect of the 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

[18] We choose to focus on all hackers because this group has shown more technological sophistication and has 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.

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 log of one plus the amount of ETH 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. After the ban in August, we find that Tornado Cash usage falls significantly and remains approximately 26% lower for about a year.

E. 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 a time series. The bars in each sub-panel represent a monthly distribution that together displays the share of Tornado Cash outflows that terminate in Western central-

ized 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 exchanges 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-statistic = 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-statistic = 5.86, p-value < 0.001). The dollar-weighted share of paths that routed through a swap or other DEX also increased, but it 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.[20] 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 exchanges are

[20] 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.

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 an exchange 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 the characteristics of paths leaving Tornado Cash and entering centralized exchanges to all tainted paths from the traced network to centralized exchanges. The results presented in Table 4 show 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 is 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. These sanctions target individuals linked to international crime varying widely, including North Korea's Lazarus Group, Hezbollah, the Sinaloa cartel, etc. The list further includes cryptocurrency exchanges and other cryptocurrency-related services associated with illicit activity such as mixers or malware vendors. Once an address is sanctioned, U.S.-based entities are legally prohibited from transacting with it. In this section, we evaluate whether these sanctions are effective at inhibiting the flows of these designated addresses.

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

There are 28 days during which OFAC issued sanction orders, and these capture 380 Bitcoin addresses and 60 Ethereum addresses. As shown in Figure IA.7, these addresses belong to many different illicit actors spanning multiple criminal professions related to dark markets, terrorism, ransomware, mixers, etc. The bulk of these sanctions occurred between 2022 and 2024.

In Figure 6 Panel A, we plot the sanctioned Bitcoin (left) and Ethereum (right) addresses. Each bubble represents a sanctioned address, with the bubble size proportional to the address's total lifetime inflow in dollars. The scatterplots illustrate sanction timing by plotting the address age at the time of last transaction, or the number of days between an address's first and last transactions (x-axis), and age at the time of sanctions, or number of days between the first transaction and the sanction date (y-axis). Addresses positioned directly on the 45-degree line were last used to transact exactly on the sanction date, which means users halted all activity after sanctions. In contrast, if an address is far above the 45-degree line, then sanctions are applied significantly after the last active date. The darker red shaded triangular regions highlight addresses that are sanctioned for more than one year and for more than two years after the last transaction date. Bitcoin addresses exhibit significant lags in sanction timing. Specifically, 80.80% of sanctioned Bitcoin addresses have no transactions after they are sanctioned, but 47.47% ceased activity more than one year before sanctions. On average, they are sanctioned 409.13 days after their last transaction. Ethereum addresses, by contrast, tend to be

sanctioned during active use, averaging sanctions 52.14 days before their last recorded transaction. Only 26.53% of Ethereum addresses are sanctioned one year after their last activity. These timing differences suggest that sanctions on Ethereum addresses are more likely to be effective, as they are often imposed while addresses are still active. Since Bitcoin sanctions typically occur after the funds have been moved, the funds only stop criminals from reusing the address, but are generally not actually freezing funds.

In Figure 6 Panel B, we plot the balances of sanctioned addresses and the destinations of their trace outflows with Bitcoin in the left subpanel and Ethereum in the right subpanel. The bars show the balances of these addresses over time and their subsequent destinations, indexed by months since sanctions. In Bitcoin, the total flow that has passed through these transactions totals more than 195,000 Bitcoin. We find that the balance in these addresses seldom stays in the sanctioned wallet.[21] Instead, balances commonly flow to other downstream wallets that reach popular exchanges such as Binance or HTX, and other service providers such as Hydra, a now-defunct Russian dark marketplace. After the sanctions, the eventual destination of these addresses is mostly unchanged, largely because the sanctions occur much later than the last transaction date. In Ethereum, we find that prior to sanctions, these addresses have also established accounts with a few centralized exchanges, such as Bitfinex and Binance. Sanctions are more likely to occur when there is still a balance in the wallet, as shown by the light blue bars. After being sanctioned, around 150,000 Ethereum were sent to Tornado Cash, providing a direct channel for money launderers to plausibly avoid sanctions. A similar amount was sent to other Ethereum wallets and remained on-chain.

In summary, the sanctioned addresses received over 420,000 Ethereum and 195,000 Bitcoin. Valued at prices at the time funds entered the sanctioned addresses, these amounts correspond to about $1.55 billion and $4.25 billion, respectively. Of these total inflows, close to 155,000 ETH (36.90%) and 100,000 BTC (51.28%) remain on-chain. Using prices at the time of deposit, sanctioned entities off-ramped funds to centralized exchanges and mixers: $524.62 million to Binance, $456.61 million to Tornado Cash, $323.17 million to Hydra, $147.22 million to HTX, and the remainder $253.02 million

[21] The light blue stacked bars at the top of each bar indicate the BTC still held in sanctioned wallets, i.e., the running balance. These bars are often barely visible because the BTC is typically transferred out quickly after arrival.

to other exchanges. Of these, only the flows to Tornado Cash occur after sanctions; all other funds exited the Ethereum and Bitcoin blockchains before sanctions through well-established centralized exchanges.[22]

The assets that remained on-chain are effectively frozen, as sanctioned entities have been unable to transfer these funds through centralized exchanges and convert them into fiat currency. Thus, sanctions, though used relatively infrequently, appear effective in that they sever these actors from the traditional financial system. The remainder have either already entered centralized exchanges or a mixer service such that the on-chain path can no longer be traced. Naturally, sanctions would be more effective if they could be used more often, enforced sooner, or before a centralized exchange accepts the funds. Discouraging the use of mixers would also help make sanctions more effective since mixers are a preferred method to avoid sanctions.

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 increase due to related freezes.

Figure 7 presents data on 1,798 instances where Ethereum addresses were prohibited from future Tether interactions. In total, almost $1.34 billion in Tether is held in frozen addresses.[23] Of the $1.34

[22] Figure IA.8 and Figure IA.9 show the total inflow to sanctioned addresses and their services of choice, as well as the dollarized value of the respective amounts stuck on-chain in event time. Panel A of both figures plots the Bitcoin and Ethereum outflows over calendar time, respectively. Importantly, the large inflows to Tornado Cash during the middle of 2022 correspond to a series of hacks targeting DeFi platforms in the Ethereum network. The largest of these was the Axie Infinity Ronin bridge hack perpetrated by the Lazarus Group. See here for further details on this hack.

[23] 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.

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, scams, impersonation, and phishing addresses. This may indicate that stablecoin issuers are more likely to respond to freeze requests from law enforcement in investigations tied 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 is received by the addresses that have been frozen, compared to the exact $603.14 million frozen.[24] 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 is currently seized. In total, we find that $1.34 billion Tether and $93.12 million USDC have been frozen. Of this, $608.65 million is 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 seizures have a wide variety in the 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.[25] 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. Bars above zero indicate inflows and bars below zero indicate outflows. Hatched bars denote flows that arise from token swaps rather than simple transfers. In Panel A, we see that frozen addresses react immediately: within the first hour, they swap USDC into DAI, and in the following hours, they move out ETH and DAI, in

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

[25] We also examine USDC freeze events in Figure IA.10, and find that they are relatively less common than Tether freezes. Nevertheless, similar patterns emerge where wallets with frozen USDC largely cease all other activity, with a few exceptions. Table IA.2 presents the statistics on the frozen USDC addresses found in our traced network.

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 documents a similar response over the next 24 hours, with both USDC and TUSD swapped into DAI and a rise in transaction activity in the hours after the freeze.[26] These observations are consistent with a rapid shift away from assets that can be frozen. USDC and TUSD, like Tether, can be frozen by their issuers, whereas DAI does not have a freeze function.[27] The evidence suggests that freezes are effective and that coordination by stablecoin issuers is important for enforcement. We formally test this behavior of increased transaction volume using a difference-in-differences framework. 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. The estimated coefficients are plotted in Figure IA.11, where the outcome variables are transaction value and transaction count. Table IA.3 tabulates results and finds that the related addresses transferred an average of 466% more dollar flow in 1.36 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 consists 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:

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

[26] TUSD refers to TrueUSD, a U.S. dollar–backed stablecoin that can also be frozen by its issuer.

[27] 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.

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, and γ_t are month fixed effects. Figure 9 plots the coefficients by event time, and Table 6 tabulates the coefficients. Overall, we find that 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 the asset seizures as part of AML 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 IA.12, 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.[28] The spread is calculated based on the dollar value of crypto of input compared to the amount received as output. 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

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

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 is services where on-chain funds can 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 the best-positioned players 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.[29] 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.[30] We use these announcements to evaluate whether customer flows to these exchanges changed following the pleas.[31] We focus on changes around the settlement of two tainted flows: Tornado Cash and victim-reported illicit flows.

[29] See here for the DOJ press release on Binance settlement.

[30] See here for the DOJ press release on OKX settlement.

[31] 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.

A. Do 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 10 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% reduction, which is statistically significant at the 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 began. 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 outflows from 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 activities. 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 that directly appear in our traced network, as discussed in the methodology section.[32]

To investigate this shift, Panel A of Figure 11 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 consists of all tainted deposit addresses at other exchanges. Panel B plots differences-in-difference coefficients. While there is no immediate impact in the months just

[32] 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.

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 11.

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 design. Column (2) of Table 7 reports the regression results, and Figure IA.13 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 in 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 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. Deposit Patterns by Tainted Deposit Addresses

Data collected as part of the Bank Secrecy Act is a cornerstone of the anti-money laundering policies of the U.S. Treasury's Financial Crimes Enforcement Network (FinCEN) (Cuéllar, 2002; Gao, Pacelli, Schneemeier, and Wu, 2023). In 2023, more than 15% of all FBI investigations were directly linked to reports of suspicious transactions.[33] Under these policies, any deposit of $10,000 or more must be reported under a Currency Transaction Report (CTR), and any suspicious transactions, such as repeated deposits, deposits just under the threshold, or funds of questionable origin, must be reported under a Suspicious Activity Report (SAR). Crypto exchanges are also required to report suspicious

[33] Source: see FinCEN Year in Review

transactions as part of their money service business license. If criminals are concerned with such reporting, they might structure their deposits to avoid detection. In this section, we test whether tainted deposit addresses are more likely to use round-numbered transaction amounts or to bunch just below the $10,000 threshold. We categorize deposit addresses by whether they are tainted by traced criminal flows and study their transaction patterns relative to untainted addresses.

A. Use of Round Numbers

A natural starting point is to examine whether illicit actors favor round-numbered deposit amounts, which could reflect deliberate behavioral structuring or simple heuristics (Nigrini, 2018). Figure 12 plots two distributions: the distribution of transaction counts for illicit flows in red and the distribution of all other flows in blue. The other flows likely contain considerable illicit flows as well, though mixed in with more non-illicit activity. Note that the y-axis is on a logarithmic scale. Bins divisible by $1,000 are represented by squares, and bins divisible by $10,000 are shown as triangles. Round-number bins, depicted by these shapes, are noticeably elevated compared to neighboring bins, reflecting the general tendency to use round numbers. When comparing illicit actors to others, the red illicit markers are almost always above their corresponding marker for other flows after $1,000 because illicit actors generally use larger transaction sizes. However, the separation between illicit and other flow markers becomes especially pronounced at round-number values, suggesting tainted addresses use round numbers more frequently.

To formally test whether tainted deposit addresses are more likely to use round numbers, we estimate regressions where the dependent variable is an indicator equal to one if the deposit amount is exactly divisible by 500 or 1000. The independent variable, $Tainted$, is an indicator equal to one if the deposit address is tainted by traced criminal flows. Columns (1) and (3) of Table 8 show that tainted deposits are significantly more likely to use round numbers. Specifically, tainted addresses are 13 percentage points more likely to deposit amounts divisible by 500, and 12 percentage points more likely to deposit amounts divisible by 1000. Columns (2) and (4) add an interaction term between the tainted indicator and a Western exchange dummy to test for differential behavior across jurisdictions. The

interaction terms are small and statistically insignificant, indicating that the use of round-numbered deposit amounts does not differ systematically between Western and overseas exchanges.

To account for the fact that tainted deposit addresses tend to use larger transaction sizes, we repeat the analysis using only deposits greater than $100. As shown in Table IA.4, the results are similar. As an additional robustness test, we examine whether the effect is specific to round-number values by testing deposit amounts that are close but not exactly divisible by round numbers. If the results are truly driven by round-number behavior, we should not observe similar effects when the dependent variable captures non-round-number amounts. In Table IA.5, we redefine the dependent variable to capture deposits that leave a remainder of 1, 499, or 999 when divided by 500 or 1000. Across these specifications, the coefficients on the tainted indicator are either statistically insignificant or extremely small in magnitude. These findings support the interpretation that tainted addresses specifically favor round-numbered deposit values.

B. Bunching around the Reporting Threshold

Beyond round-number effects, we also study whether tainted addresses exhibit strategic behavior around the $10,000 reporting threshold. We begin by plotting the distribution of deposit transaction sizes within a narrow window around the $10,000 threshold in Panel A of Figure 13. The figure displays the density of deposits for tainted (red) and untainted (blue) addresses, using solid circles for flows divisible by $100 and hollow circles otherwise. High-order polynomials are fitted separately for each group, with dashed lines corresponding to bins not divisible by $100. The distribution is shown separately for Western and overseas exchanges.

The results reveal several important patterns. First, there is a clear spike in the frequency of deposits exactly at the $10,000 threshold, as shown in Figure 12. Second, there is more bunching right below the threshold than right above it, as supported by the discontinuity between the dashed lines to the left and right of the $10,000 threshold. This pattern holds for both tainted and untainted flows, as indicated by the overlapping blue and red dashed lines. It is also observed in both Western and overseas exchanges. Third, when focusing on round-number bins below the threshold, tainted flows

bunch at these numbers more than untainted flows, though mainly at Western exchanges. This is most evident at deposit amounts such as $9,900, $9,800, and $9,500. To formally test the illicit actors' behavior of bunching below the threshold, we estimate the following regression form separately for every $10 deposit bin k in the $9,500–$10,500 window:

$$\mathbb{1}(d \in k)_{d,k,i,e,t} = \alpha + \beta^k \times Tainted_i + \gamma_e + \eta_t + \varepsilon_{d,k,i,e,t}$$

where $\mathbb{1}(d \in k)$ is an indicator variable equal to one if deposit d falls into bin k, and $Tainted_i$ is an indicator for whether deposit address i is tainted by illicit flows. The coefficient β^k captures the difference in the probability that a tainted deposit address uses a specific deposit amount (bin k) relative to an untainted address. All regressions include exchange fixed effects to account for time-invariant exchange-specific characteristics, and year-month fixed effects to control for common temporal shocks. The regressions are run separately for Western and overseas exchanges.

The regression results are presented in Panel B of Figure 13, with green markers corresponding to Western exchanges and purple markers to overseas exchanges. The estimates reveal several patterns. First, at the $10,000 threshold, the coefficients are positive for both Western and overseas exchanges, indicating that tainted deposit addresses are more likely to use this exact amount compared to untainted ones. However, the effect is notably smaller in Western exchanges, suggesting that bunching at the threshold is mitigated in Western exchanges. Second, at deposit amounts immediately below the threshold, such as the $9,990 and $9,980 bins, the coefficients are negative, implying that tainted addresses are less likely to use these slightly sub-threshold amounts relative to untainted users. This effect is again weaker in Western exchanges, indicating that illicit actors are more likely to use these near-threshold values when interacting with Western exchanges than with overseas ones. Third, within Western exchanges, the coefficients are positive and statistically significant at round-number bins below the threshold, including $9,900, $9,800, and $9,500. Together, these results are consistent with the patterns in Panel A that illicit actors favor round-numbered deposit sizes and bunch just below the reporting threshold at Western exchanges.

Critics of CTR and SAR reporting argue that these requirements generate unnecessary busywork

and produce a volume of reports too large to be meaningfully analyzed (Cuéllar, 2002). Overall, while the requirements appear to influence deposit patterns, their ability to effectively discriminate between tainted activity does not appear as an extremely strong pattern. Effort spent on the $10,000 threshold might be more effectively applied to monitoring criminal flows.

VIII. 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?

Figure 14 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. Furthermore, hackers do not seem to widely use identified non-KYC exchanges, as these are relatively short-lived and only handle relatively small crypto deposits, as shown in Figure IA.15. 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 sanc-

tions.[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 15 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 $42.8 million sent to OKX, but the majority of funds, or $968.15 million, we traced 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 million (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 $98.7 million has 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.[35]

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 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

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

[35] Figure IA.16 plots the cumulative tainted BTC associated with said hack, and the corresponding inflows to Freebitco.in deposits.

sophisticated attacks like this are potentially traceable.

IX. 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 costly to criminals, with enforcement actions against services, such as mixers and exchanges, being the most consequential.

Nevertheless, our analysis indicates many areas for improvement. Even though OFAC sanctions keep funds on-chain, few addresses and dollar amounts are sanctioned, and often, too late. Additionally, while overseas exchanges see reduced sanction-related flows after enforcement actions, they continue to receive 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 weak spots. Therefore, reputational risk and the goodwill of exchanges to stop crime do not appear to be effective deterrents in the crypto space. Reducing flows to the Tornado Cash mixers led to greater 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 be automated for reasonable costs. We hope additional research will further analyze crypto crime and enforcement to help enact reasonable policies that both facilitate legitimate capital formation and 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.

Béres, Ferenc, István A Seres, András A Benczúr, and Mikerah Quintyne-Collins, 2021, Blockchain is watching you: Profiling and deanonymizing ethereum users, in *2021 IEEE international conference on decentralized applications and infrastructures (DAPPS)*, 69–78, IEEE.

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.

Gao, Janet, Joseph Pacelli, Jan Schneemeier, and Yufeng Wu, 2023, Dirty money: How banks influence financial crime, *Working Paper* .

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* .

Nigrini, Mark, 2018, A fingerprint of fraud, *Journal of Accountancy* .

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.

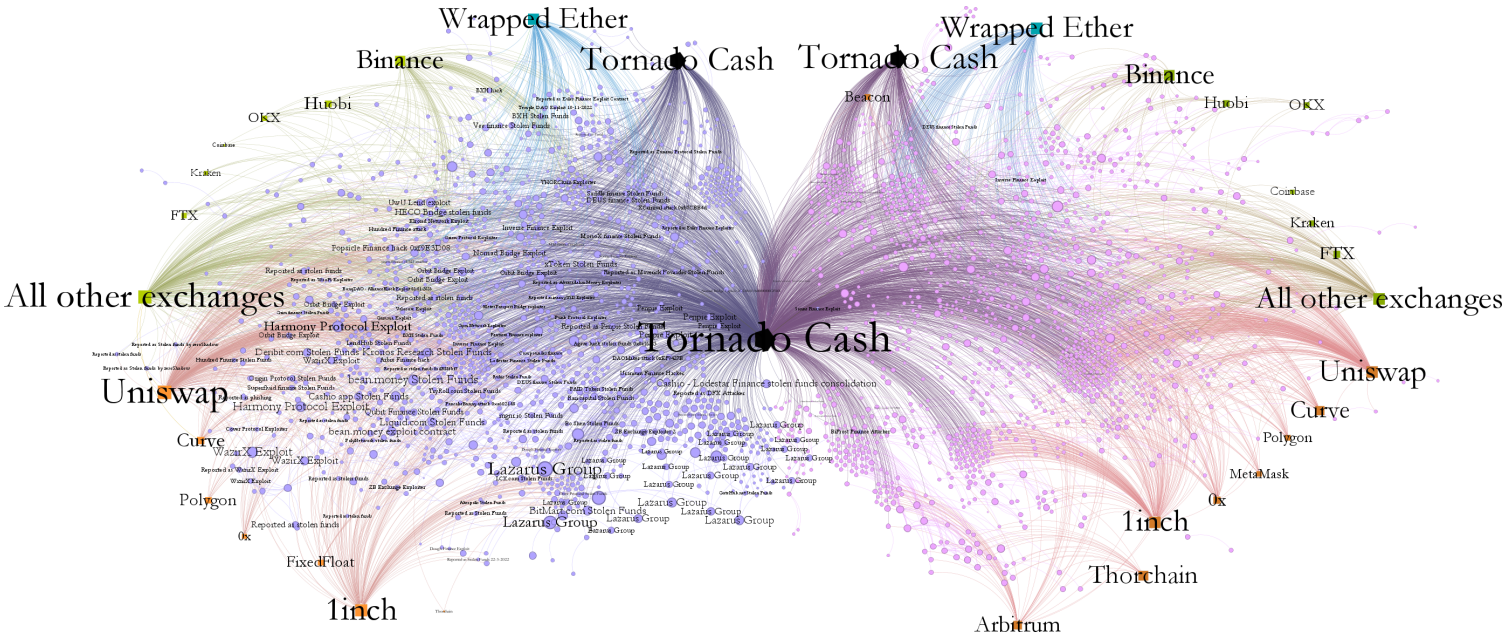

Figure 2: Inflows to Tornado Cash

This figure shows the total inflows and traced criminal inflows to Tornado Cash from 2021 to 2024. The bars indicate the monthly total inflows to Tornado Cash. The different shades in the bars represent the inflows from the top one depositor (dark blue), the top two to five depositors (blue), and all other depositors (light blue). The red line represents the percentage of traced criminal inflows from all scam types. The blue line shows the probability of matching the withdrawal transaction to the top depositor. This probability is estimated using a forward projection approach that allocates each depositor's inflows to future days according to a Pareto decay function; the weekly predicted outflows are then used to calculate the share attributable to the top depositor. While the bars are based on monthly totals, the percentages and probabilities shown in the lines are computed at weekly frequency and plotted as ten-week rolling averages. 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.

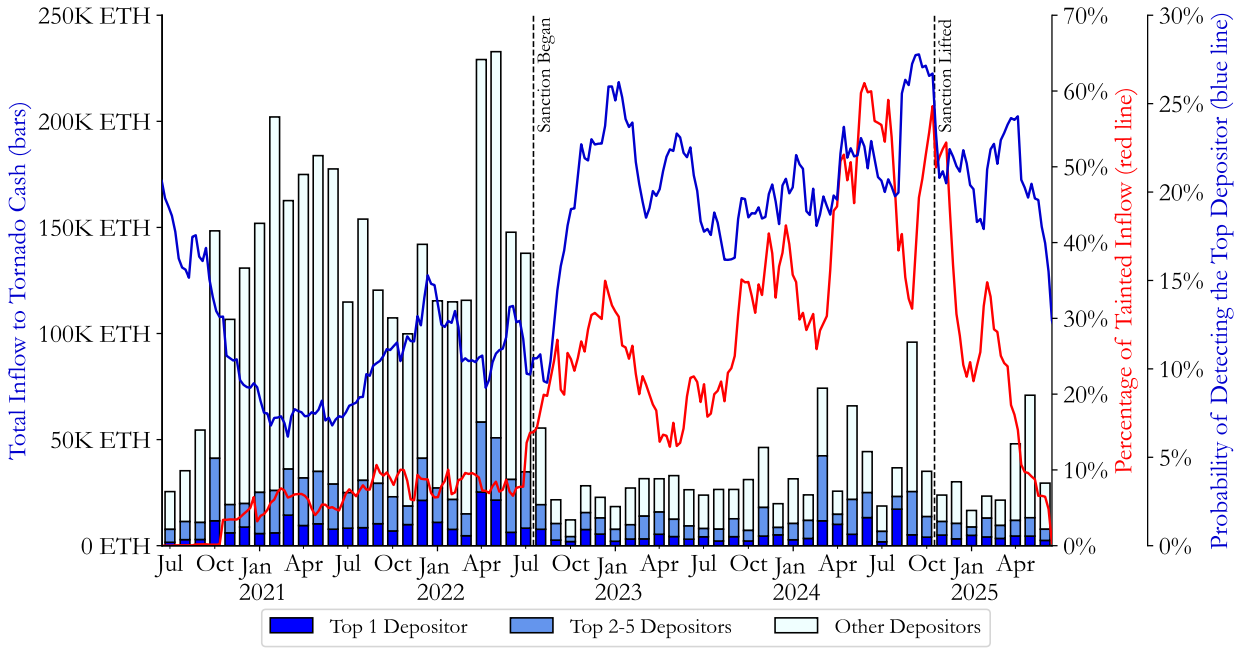

Figure 3: Effects of U.S. Treasury Sanctions Against Tornado Cash on Hacker Flows

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

$$\log(1 + TaintedFlows)_{i,d,t} = \sum_{t \neq July2022} \beta_t \times \mathbb{1}(Month = t) \times Tornado_d + \delta \times Tornado_d + \mu_i + \gamma_t + \varepsilon_{i,d,t}$$

where $\log(1+TaintedFlows)_{i,d,t}$ is the natural logarithm of one plus the amount of ETH 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. The vertical dashed line marks the August 8, 2022 sanction of Tornado Cash.

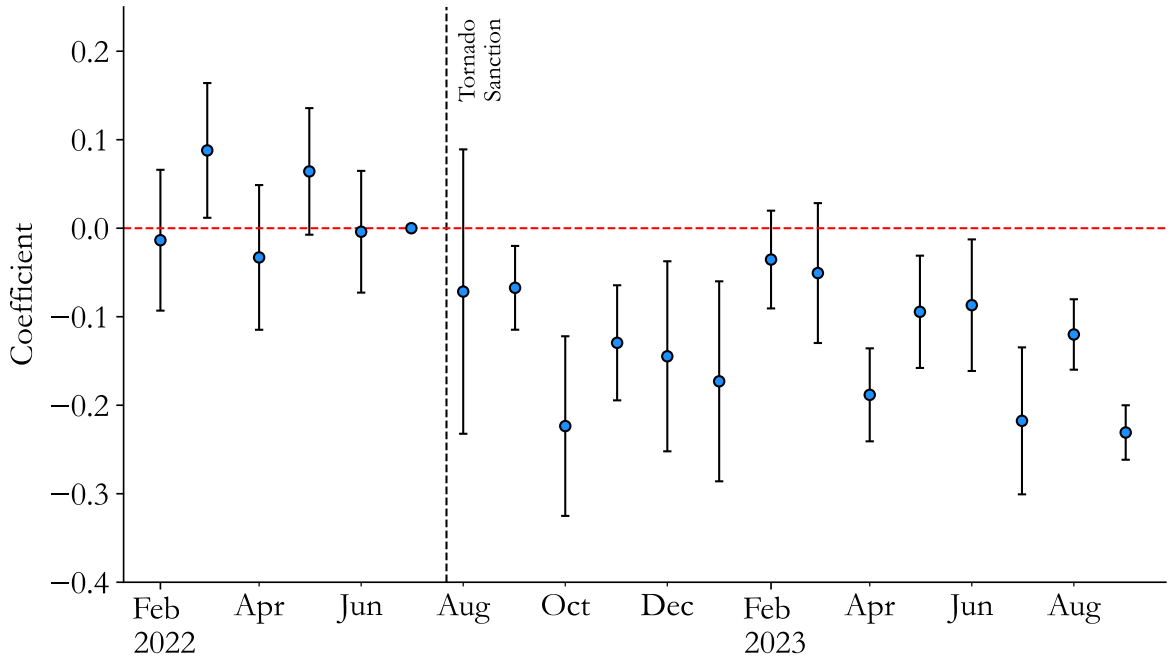

Figure 4: Tornado Cash Outflows to Destinations by Intermediate Hop

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 ss) 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

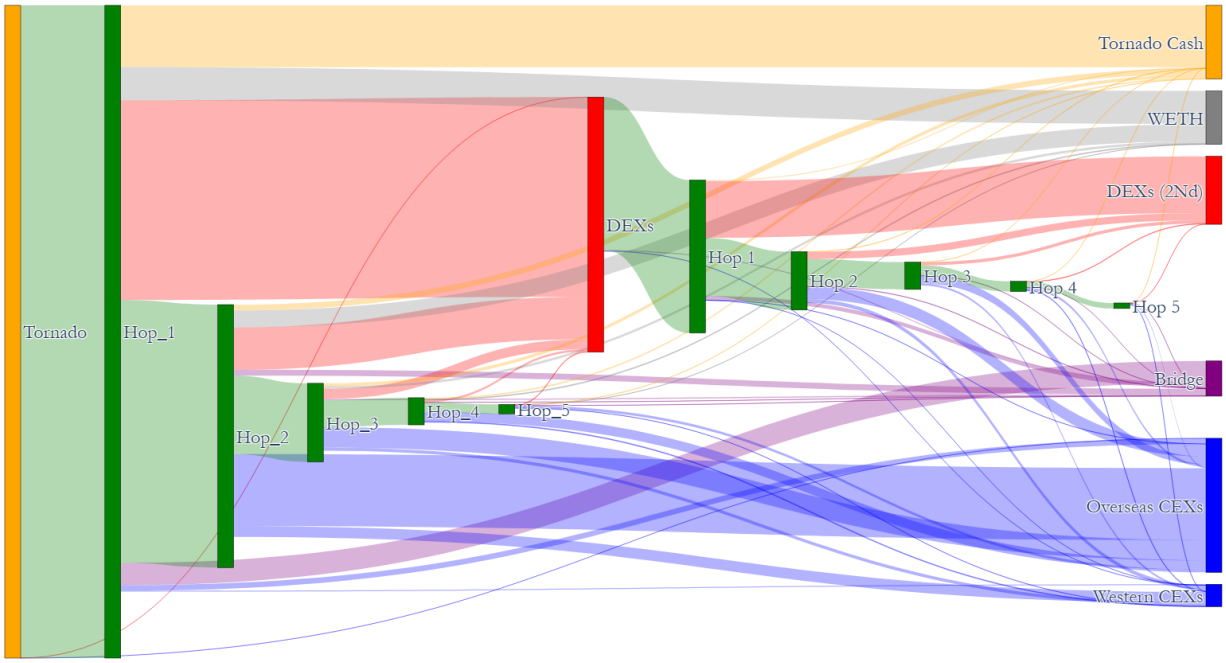

Panel B: After Tornado Cash Ban

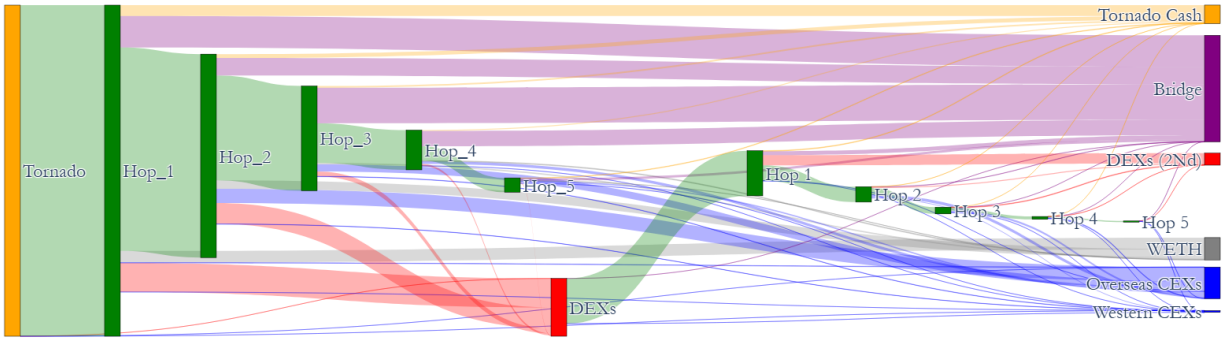

Figure 5: Effects of Tornado Cash Ban on Its Outflows to Different Destinations

This figure examines the effects of the Tornado Cash ban on its outflows to different destinations. It shows Tornado Cash outflows to Western exchanges in Panel A, overseas exchanges 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 exchanges 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.

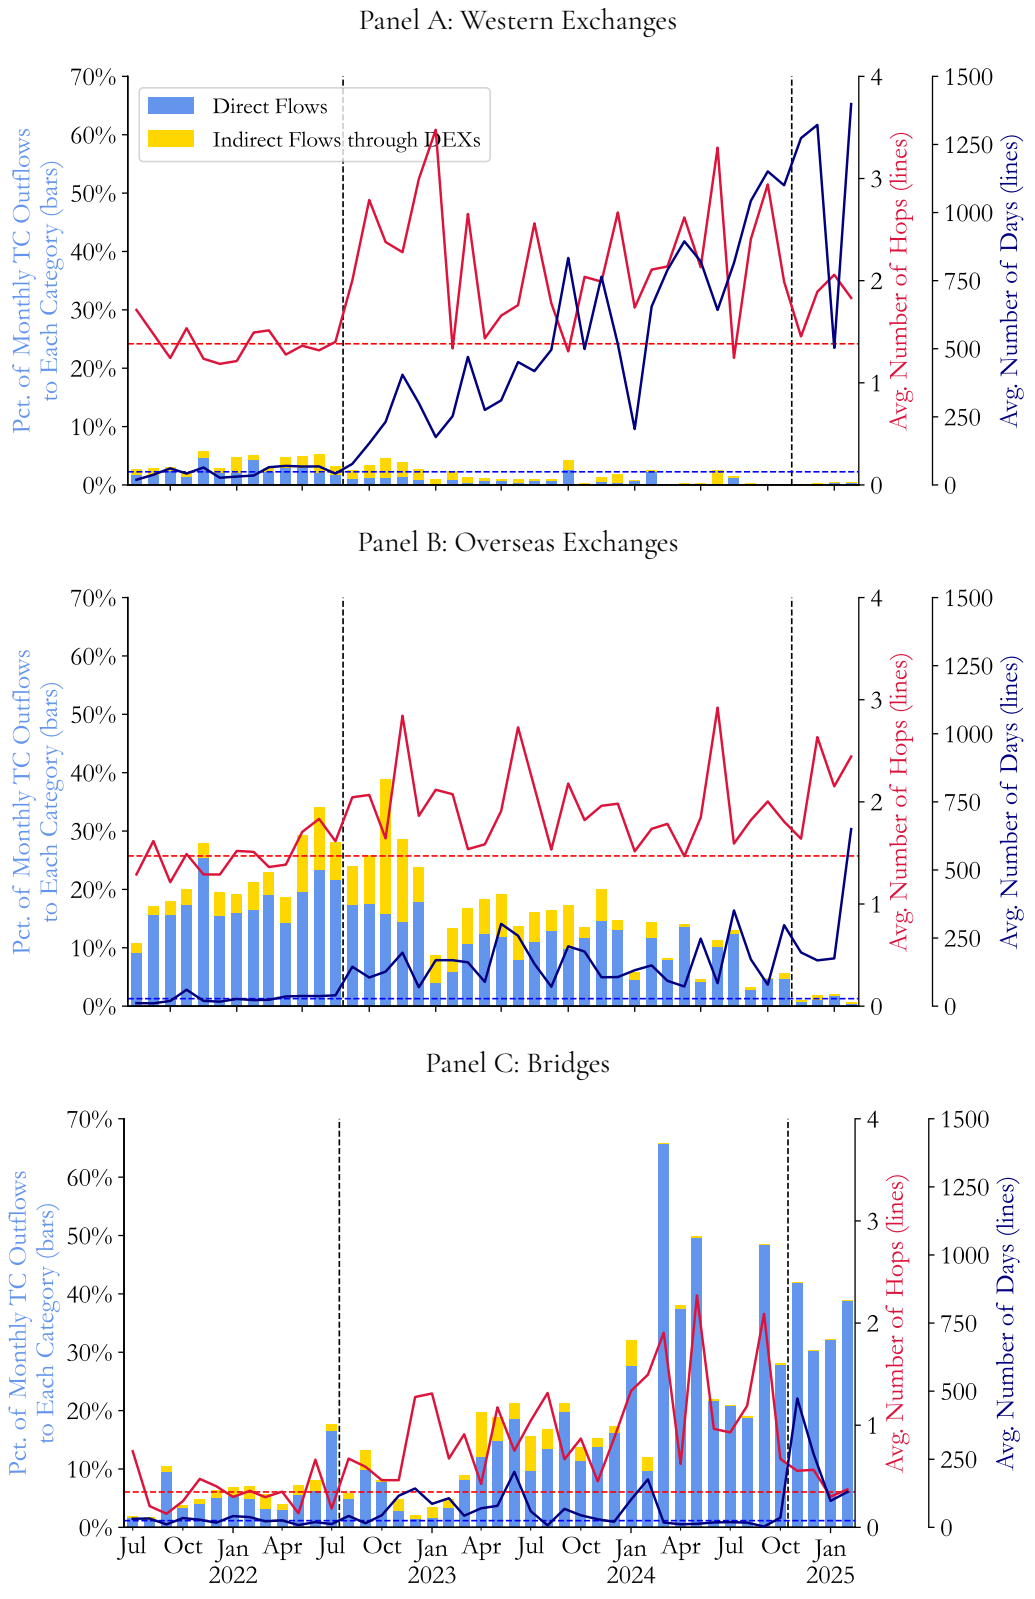

Figure 6: Effect of OFAC Sanctions

This figure plots OFAC sanction activity. Panel A shows scatterplots for BTC and ETH sanctioned addresses, with the x-axis indicating the number of days between the first and last transaction, and the y-axis showing the number of days between the first transaction and the sanction date. Dot size reflects total lifetime inflow, and the red line (rightmost axis) shows cumulative inflow. Histograms along the top and right display the percentage of addresses by last transaction day and sanction day, respectively. The shaded triangles show the timing of sanctions: addresses above the 45-degree line were sanctioned after their last transaction, while those below were sanctioned before. Darker red triangles highlight addresses sanctioned with a delay of more than 1 year and more than 2 years after their last transaction, respectively. Annotated percentages indicate the share of addresses in each group. Panel B plots destinations of flows leaving addresses sanctioned by OFAC, indexed by months since sanctions were imposed. Assets in the wallet correspond to those assets in the sanctioned address, whereas the assets on-chain are those that have been forwarded in the corresponding blockchain, but have yet to reach a known service.

Panel A: Sanctioned Day versus Last Transaction Day

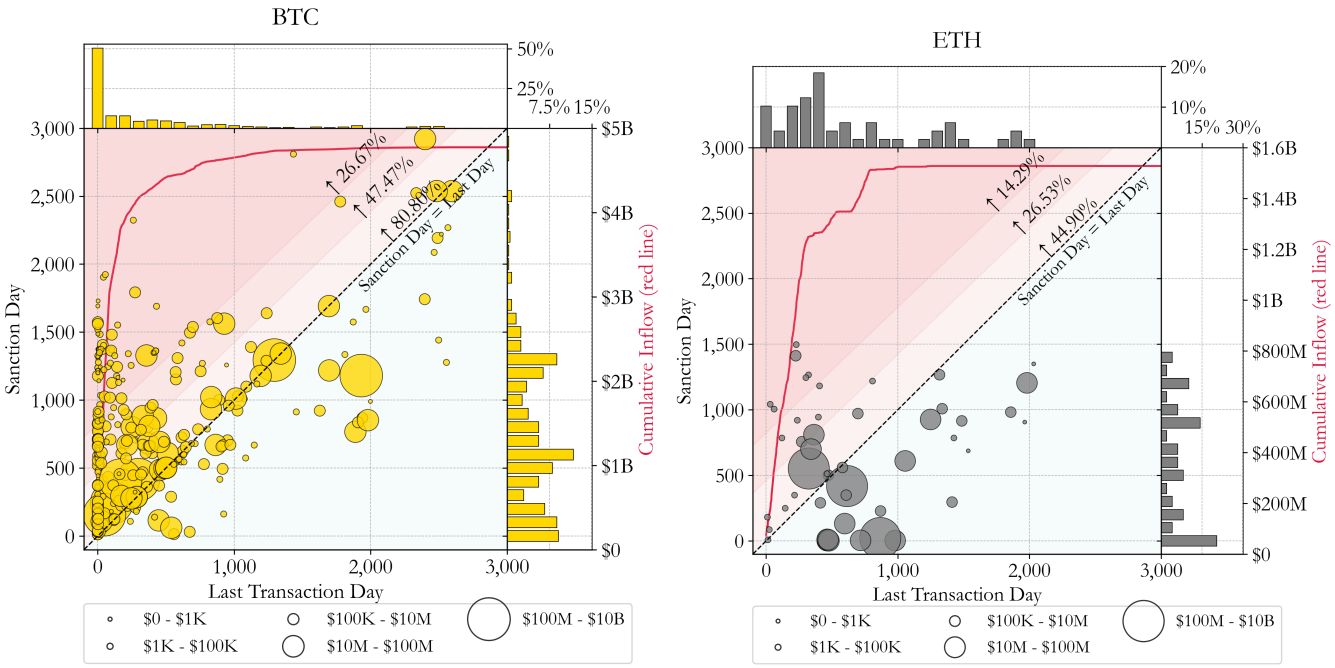

Panel B: Balance Traced from Sanctioned Addresses to End Destinations

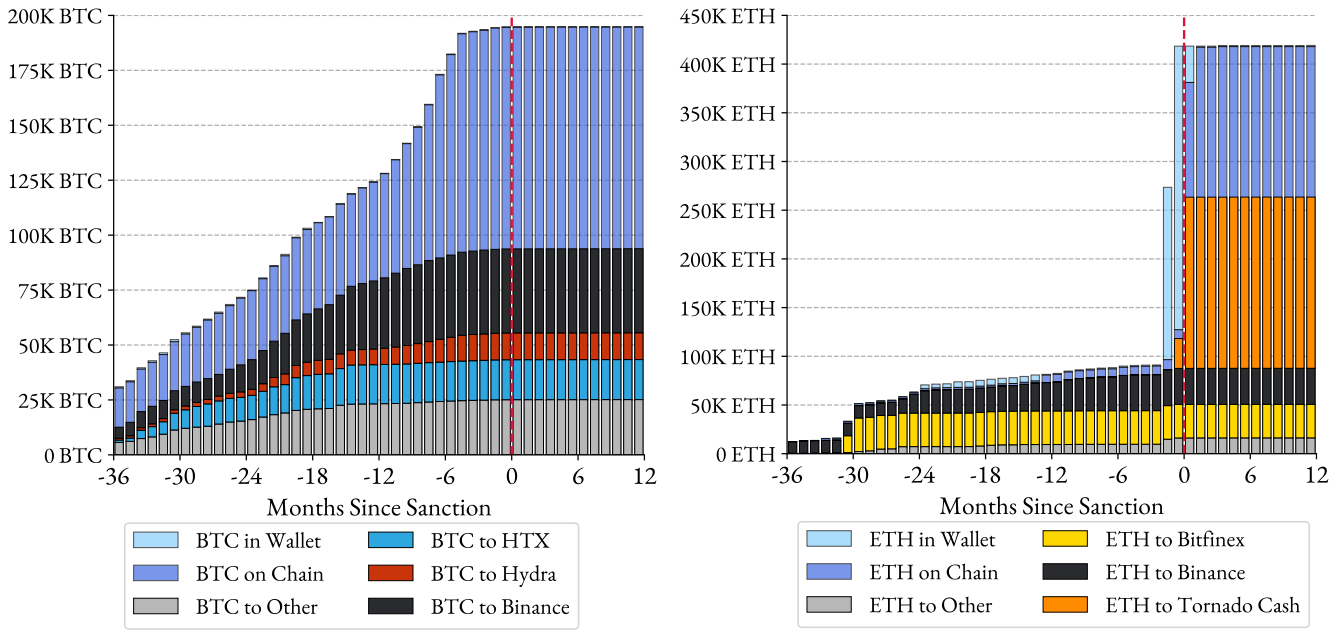

Figure 7: Tether Freezes

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

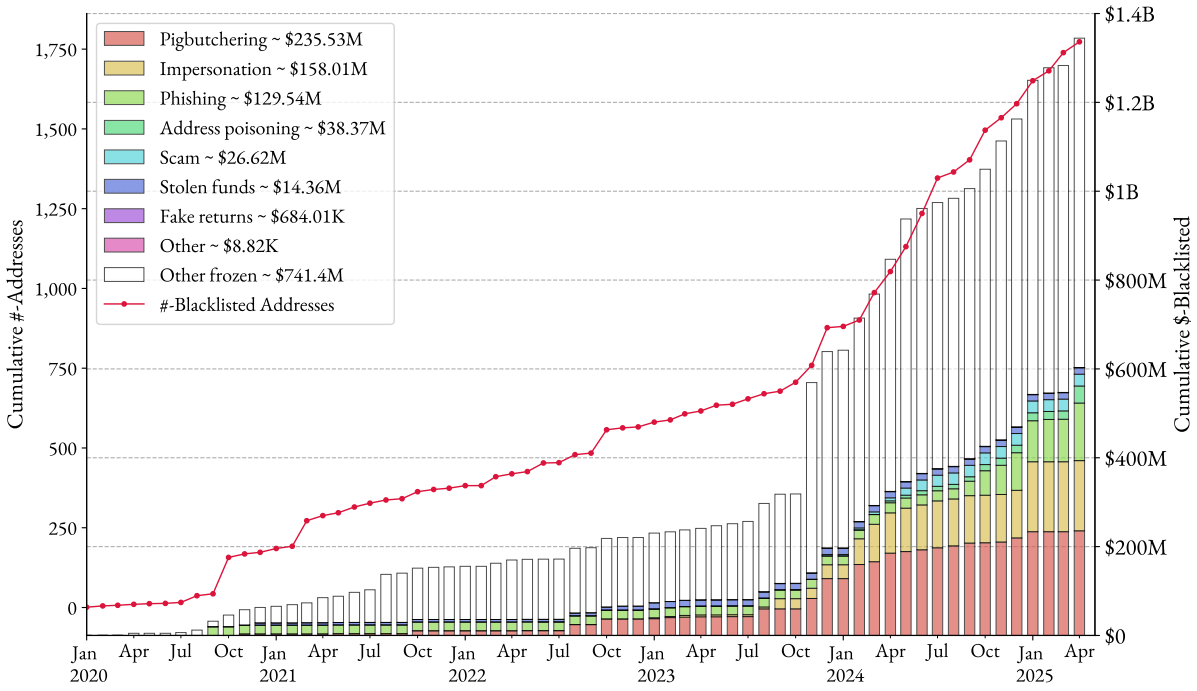

Panel B: Activity of Frozen Addresses Prior to Seizure

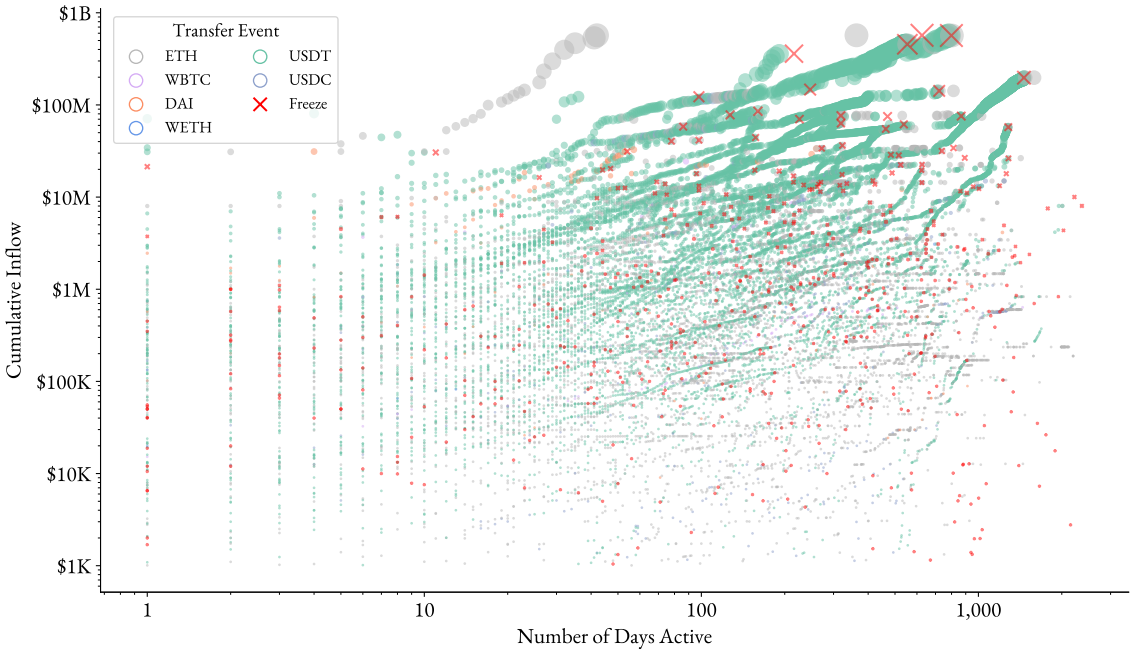

Figure 8: Tether Seizure: Immediate Effects

This figure presents inflows and outflows indexed by the hours prior to and after the Tether freeze for the targeted addresses and the addresses clustered around these. For both panels, inflows and outflows are reflected in the left-hand y-axis, and these flows are broken into their corresponding cryptocurrency, here denoted by the corresponding colors. Inflows are represented by upward-pointing bars, and outflows are likewise represented by downward-pointing bars. Hatched sections correspond to amounts swapped. On the opposite y-axis, the running balance is represented. Panel A shows the flows corresponding to the frozen addresses, which cannot forward Tether after the frozen date, marked here by the dotted vertical line at the zero hour. Panel B shows the outflows of addresses related to but excluding the addresses that were frozen. These additional addresses were clustered using a conservative common-funder heuristic.

Panel A: Frozen addresses

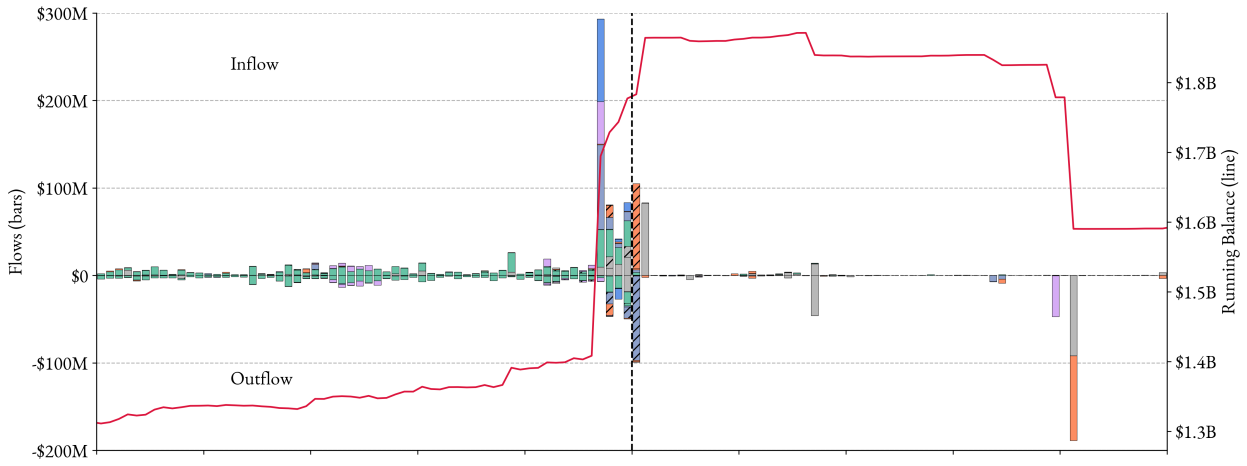

Panel B: Clustered addresses

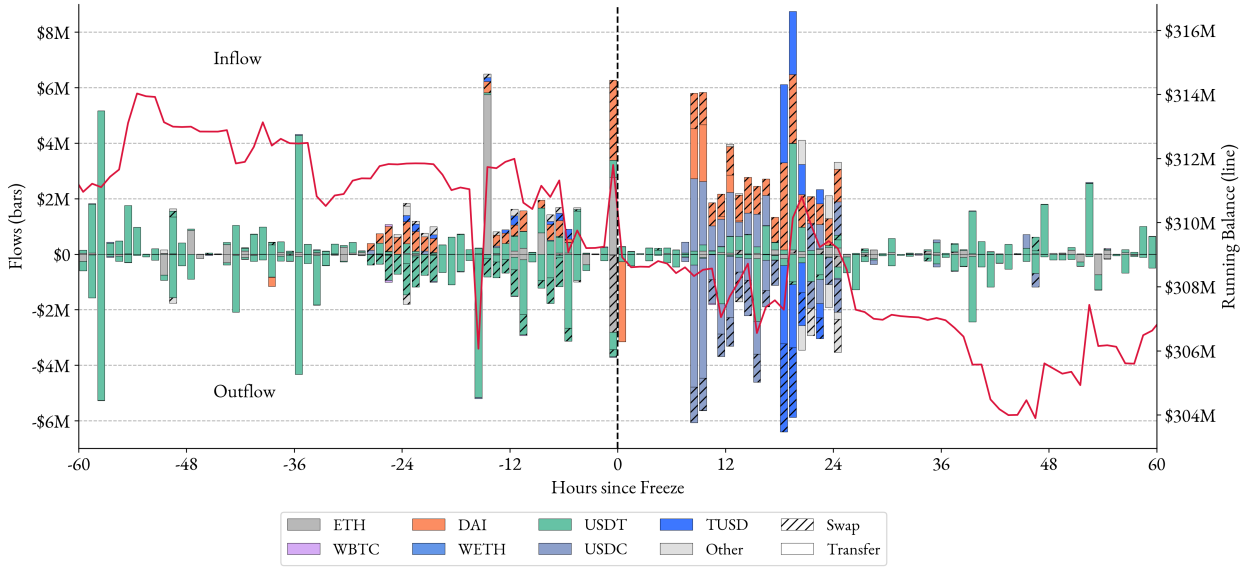

Figure 9: Tether Seizure: Long-Term Effects

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 consists 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 and is weighted by total transaction value. 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.

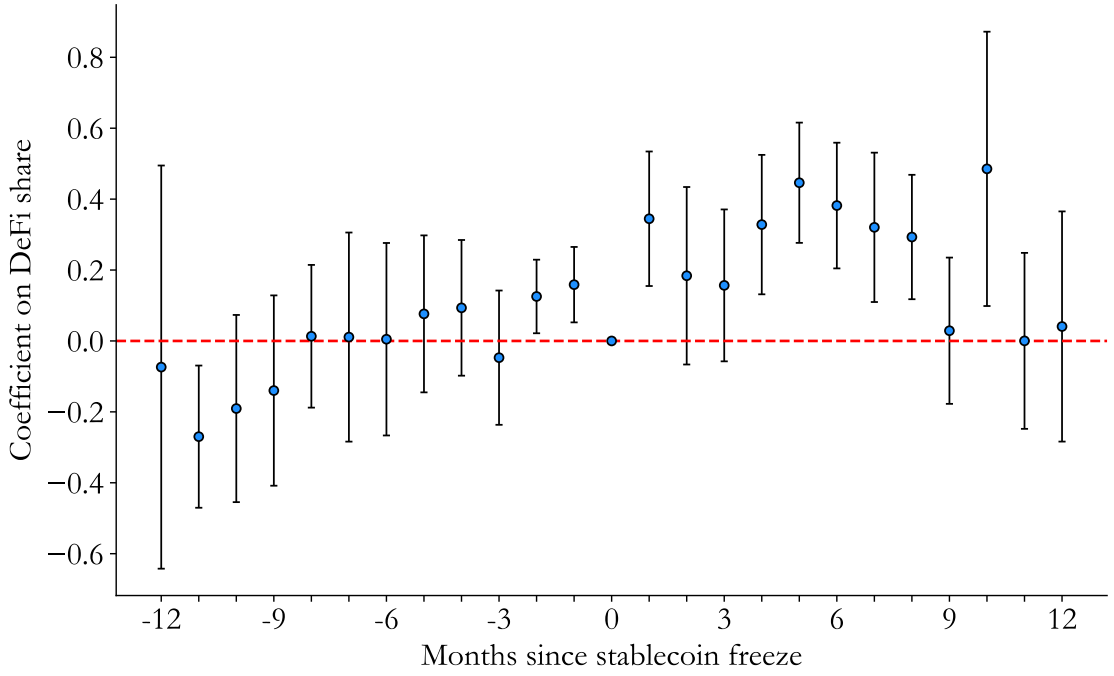

Figure 10: Tornado Cash Outflows to Binance around Binance Settlement

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 exchanges (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 exchanges 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 exchanges using a difference-in-differences (DID) regression, with the estimated coefficients plotted over time. Year-month fixed effects and exchange fixed effects are included. Standard errors are clustered by exchanges. 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

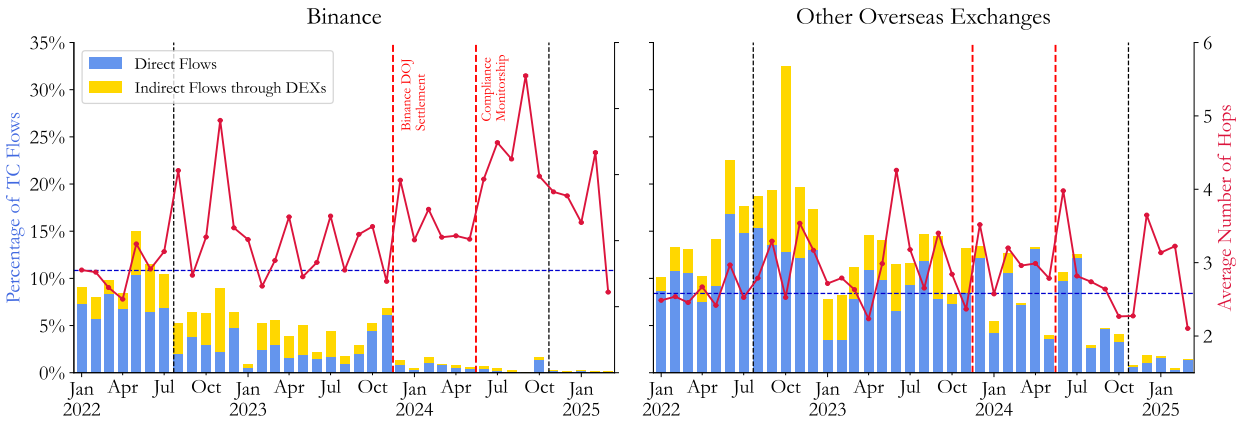

Panel B: DID Analysis on Number of Hops

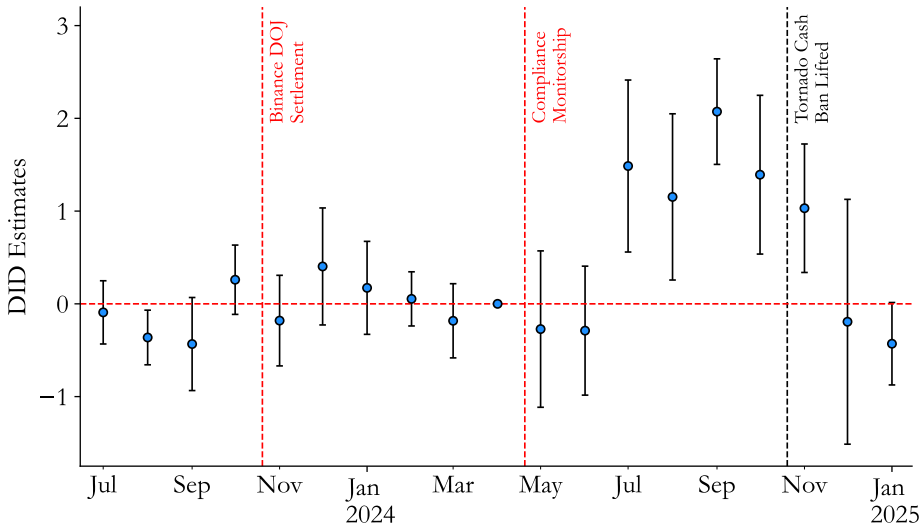

Figure 11: Tainted Flows to Binance around Binance Settlement

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.

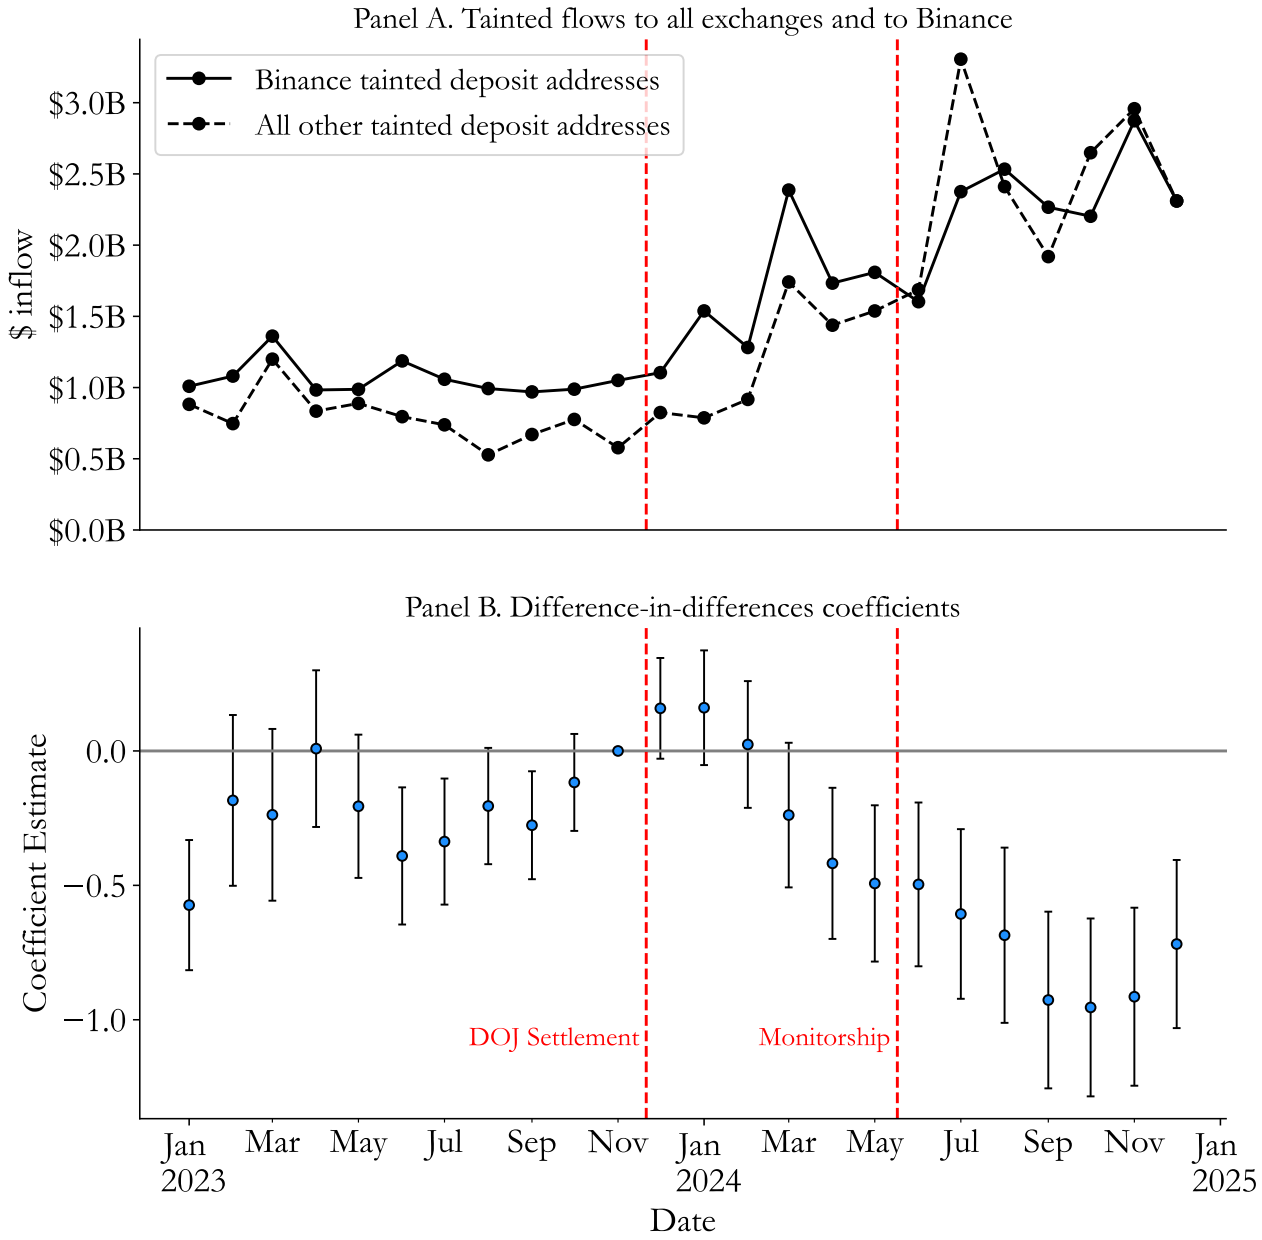

Figure 12: Round Number Deposits by Tainted Deposit Addresses in Exchanges

This figure plots the distribution of transaction sizes for all exchanges split into bins of $100. Red dots show illicit flows, and blue dots show all other flows. Right-bound bins divisible by $10,000 are triangles, $1,000 are squares, and all others are circles.

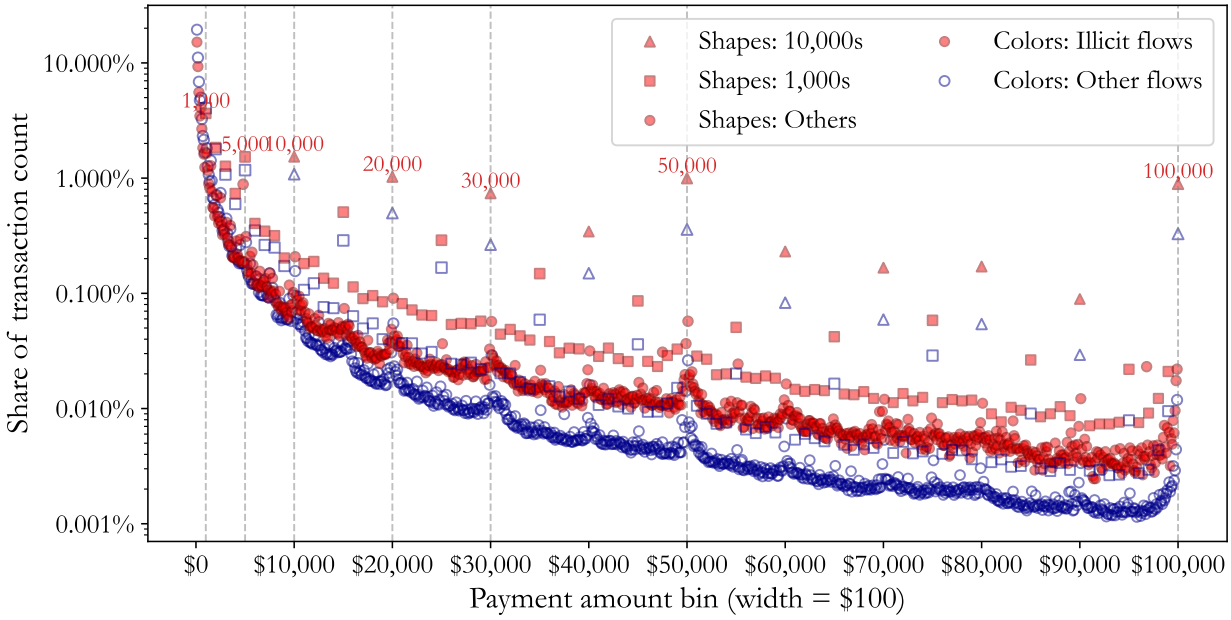

Figure 13: Deposit Bunching in Western Exchanges

This figure investigates deposit bunching around the $10,000 threshold. Panel A plots the density of deposit transactions within bins ranging from $9,500 to $10,500, with the y-axis representing the percentage of total transactions in this range. The $10,000 bin precisely captures the round number; bins to its left are left-inclusive (e.g., $9,990-$9,999.99), while bins to its right are right-inclusive (e.g., $10,000.01-$10,010). The left subpanel displays data for Western exchanges, and the right subpanel for Overseas exchanges. Solid red (blue) circles represent tainted flows (other flows) divisible by $100, with hollow red (blue) circles denoting tainted flows (other flows) indivisible by $100. Panel B presents the estimated coefficient β^k from the regression

$$\mathbb{1}(d \in k)_{d,k,i,e,t} = \alpha + \beta^k \times Tainted_i + \gamma_e + \eta_t + \varepsilon_{d,k,i,e,t},$$

run for each deposit bin k. Here, $\mathbb{1}(d \in k)$ is an indicator for a deposit d falling into bin k, and $Tainted_i$ is an indicator for whether deposit address i is tainted. β^k thus captures the differential probability of a tainted deposit occurring in bin k relative to a non-tainted deposit, after controlling for exchange fixed effects and year-month fixed effects. The regressions are run separately for Western (green) and Overseas (purple) exchanges, with vertical bars indicating 95% confidence intervals. Standard errors are clustered by exchange-time.

Panel A: Deposit Frequency around $10,000 Threshold

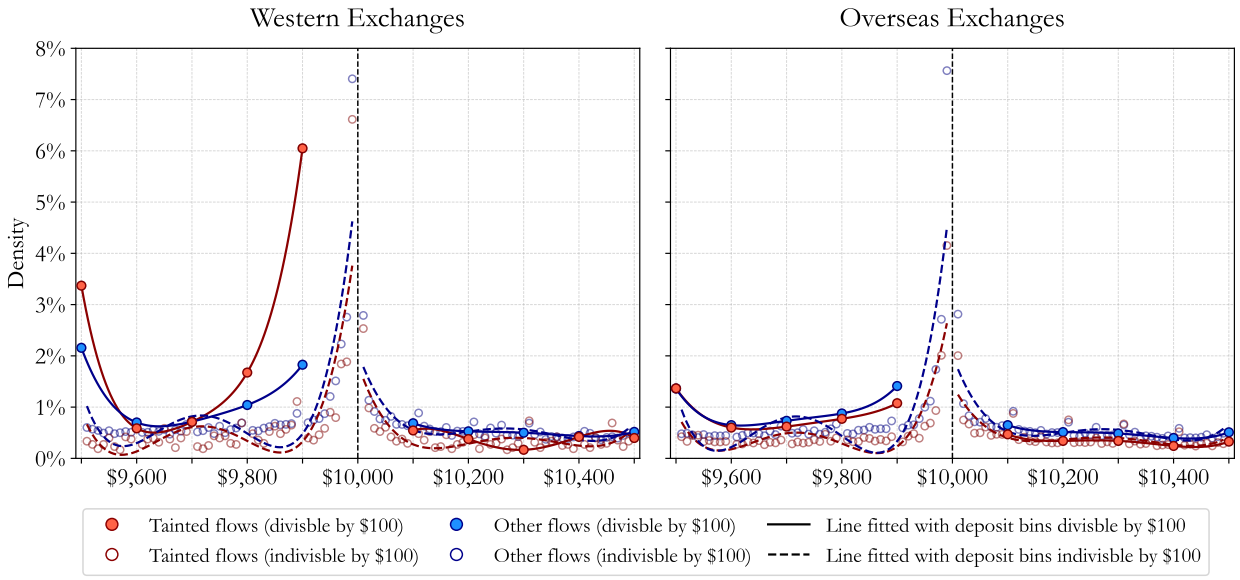

Panel B: Deposit Bunching Regressions

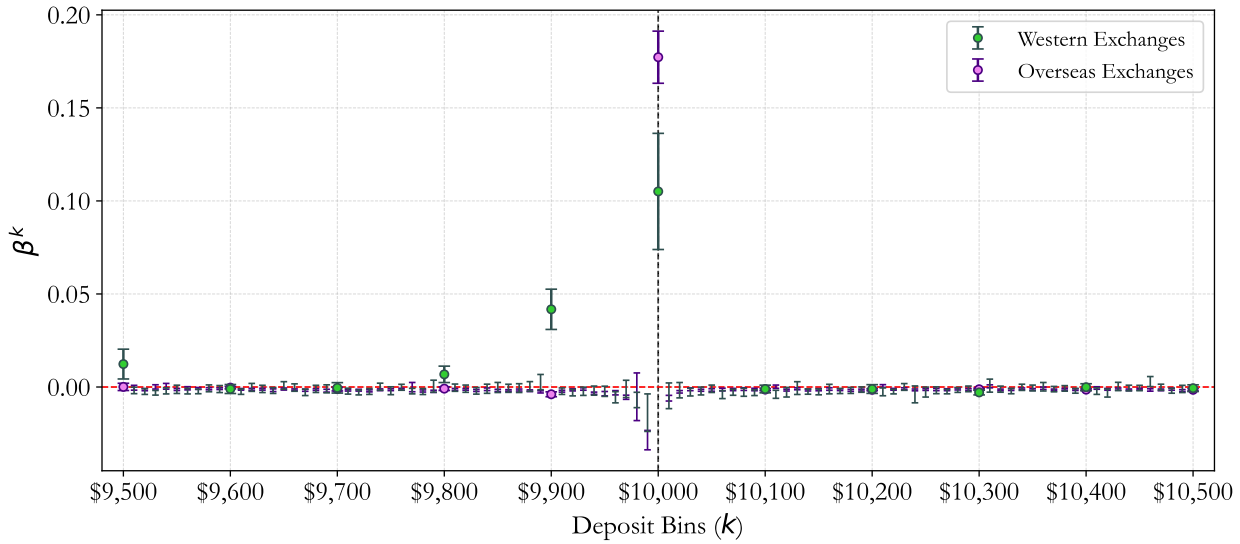

Figure 14: Effects of Tornado Cash Ban on Hacker Flows Destinations

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.

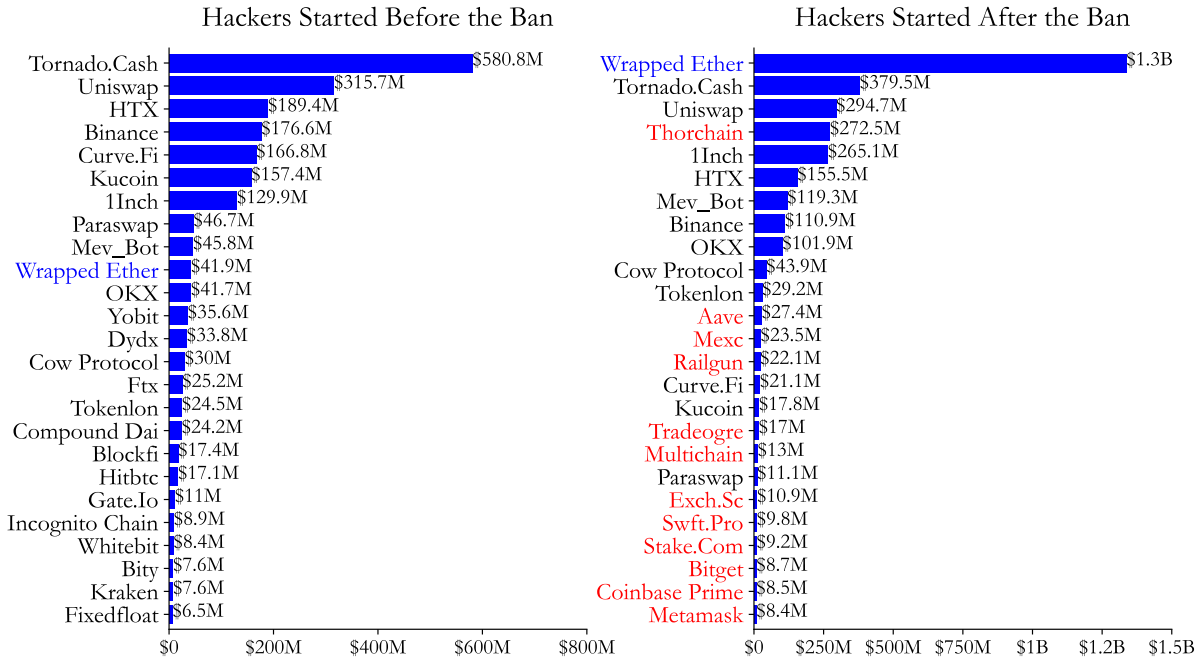

Figure 15: Bybit Hack and Flows to Bridge

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 on the left to another node on 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 about $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 $98.7M of the bridged funds to Freebitco.in, and $1.9M to other exchanges. Last updated in August 2025.

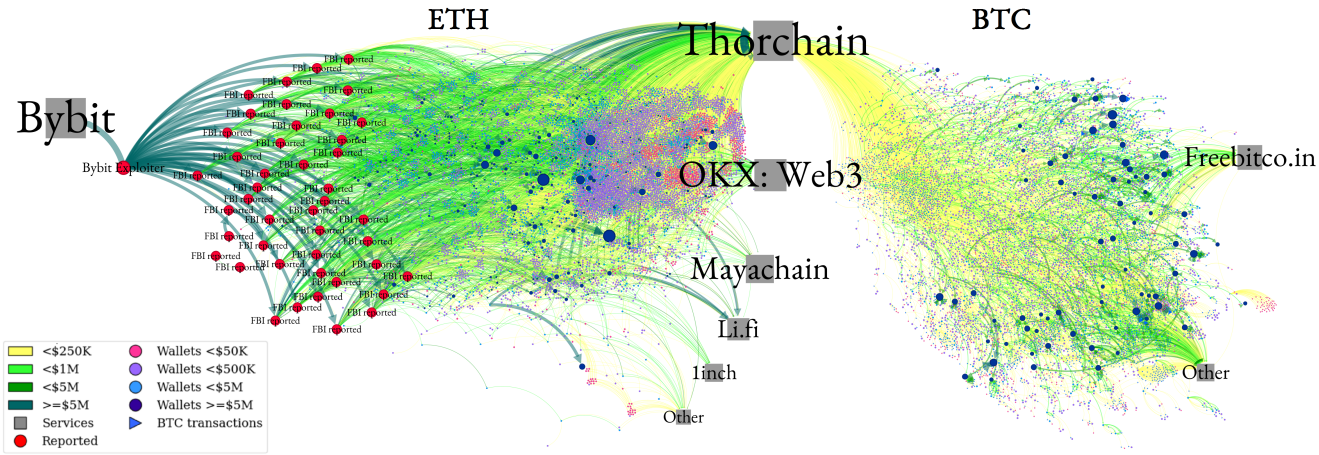

Table 1: Reports Summary

This table presents summary statistics on the 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	8,244	8,204	8,204	7,980	$12.91B
Pigbutchering	23,613	19,018	11,725	9,071	$11.30B
Illicit actor	5,370	5,367	5,364	5,158	$11.25B
Contract exploit	938	429	429	277	$5.76B
Phishing	65,852	64,992	62,724	53,942	$4.42B
Scam	26,942	26,796	26,796	22,167	$2.46B
Extortion	192,918	91,966	7,399	6,409	$1.89B
Fake project	4,043	3,798	2,502	1,873	$880.33M
Malware	3,143	2,304	2,129	2,008	$817.01M
Fake returns	4,697	4,277	3,698	2,535	$696.78M
Impersonation	31,406	30,446	28,404	27,408	$138.04M
Airdrop	788	733	490	336	$37.55M
Dark market	755	753	749	577	$28.34M
Sim swap	191	169	151	121	$9.39M
Address poisoning	122,800	122,427	122,427	122,410	$9.26M
Donation scam	616	354	218	158	$2.75M
Bitcoin	244,331	139,383	44,124	37,773	$21.46B
Ethereum	247,985	242,650	229,265	219,185	$31.15B
Total	492,316	382,033	273,389	256,958	$52.61B

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 (4) of Table 1 and total inflow corresponds to Column (5). This table presents the mean, standard deviation, and the 25th, 50th, and 75th percentiles. 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	9,071	$11.30B	$1.25M	$9.65M	$5.20K	$64.19K	$502.15K
Illicit actor	5,158	$11.25B	$2.18M	$62.69M	$150.74	$1.30K	$19.68K
Contract exploit	277	$5.76B	$32.20M	$312.80M	$1.92K	$335.69K	$3.43M
Phishing	53,942	$4.42B	$84.82K	$2.16M	$0	$0	$666.74
Scam	22,167	$2.46B	$117.49K	$1.40M	$536.56	$4.61K	$27.90K
Extortion	6,409	$1.89B	$316.78K	$11.99M	$525.61	$1.33K	$3.28K
Fake project	1,873	$880.33M	$501.04K	$3.19M	$22.65	$10.30K	$155.36K
Malware	2,008	$817.01M	$482.01K	$1.99M	$1.16K	$26.19K	$225.73K
Fake returns	2,535	$696.78M	$369.84K	$10.92M	$671.91	$6.09K	$33.21K
Impersonation	27,408	$138.04M	$5.15K	$99.57K	$5.10	$6.81	$7.84
Airdrop	336	$37.55M	$142.25K	$673.10K	$158.88	$7.44K	$60.88K
Dark market	577	$28.34M	$49.12K	$373.34K	$57.59	$236.16	$2.90K
Sim swap	121	$9.39M	$92.10K	$152.64K	$3.92K	$22.33K	$110.95K
Address poisoning	122,410	$9.26M	$75.65	$12.53K	$0.74	$1.52	$4.10
Donation scam	158	$2.75M	$23.68K	$188.45K	$45.75	$405.59	$2.95K
Total	256,958	$52.61B	$204.73K	$12.76M	$0.73	$3.78	$230.85

Table 3: Effects of the Tornado Cash Ban: Western versus Overseas Exchanges

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 (treatment group) with those to overseas exchanges (control group). The outcome variables are the number of intermediate hops (columns 1–2), the number of days to exit to an exchange (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 an exchange. Each path may include multiple exits, with each exit representing a cash-out event to an exchange and counted separately. The number of days is measured from when the funds initially leave Tornado Cash until they ultimately reach the exchange 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 exchanges. Tornado Cash was sanctioned on August 8, 2022, and a path exit is assigned a post variable of one if the funds reached an exchange 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 exchanges. 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)	
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 exchanges. It compares Tornado Cash outflows to exchanges (treatment group) with flows originating from reported criminal addresses to exchanges (control group). The outcome variables are the number of intermediate hops (columns 1–2), the number of days to exit to an exchange (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 an exchange. Each path may include multiple exits, with each exit representing a distinct cash-out event to an exchange 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 exchange 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 exchanges. Tornado Cash was sanctioned on August 8, 2022, and a path exit is assigned a post variable of one if the funds reached an exchange after the ban. All regressions are dollar-weighted by the amount of funds entering exchanges. 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)
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: USDT Seizure Summary Statistics

This table summarizes USDT 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	1,589	63	1.52M	2.60M	228.54M	8.44M
Pigbutchering	853	182	4.13B	344.50M	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	24.30M
Phishing	58	150	21.04M	3.12M	230.46M	28.91M
Stolen funds	35	37	161.98M	86.03M	95.10M	13.57M
Illicit actor	14	17	52.26M	3.88M	26.40M	3.30M
Fake returns	8	11	26.15M	1.98M	56.01M	0.51M
Fake project	6	11	78.09M	4.88M	59.60M	0
Contract exploit	2	2	180.90K	3.79M	0	0
Donation scam	1	1	9.35K	99.99	0	0
Sim swap	1	1	50.04K	8.20K	25.91K	0
Total	3,179	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 first regression includes the cohort fixed effects, and the second adds event time fixed effects, and both are weighted by total transaction value. 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)	
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. All regressions 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)
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

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

Table 8: Use of Round Number Deposits by Tainted Deposit Addresses

This table presents regression results examining the relationship between tainted deposits and the likelihood of transactions being exact multiples of 500 or 1000. Specifically, the following regression is estimated.

$$\mathbb{1}(\text{Divisible by 500})_{i,e,t} = \alpha + \beta_0 \times Tainted_i + \beta_1 \times Tainted_i \times Western_{i,e} + \gamma_e + \eta_t + \varepsilon_{i,e,t}$$

where d represents the deposit transaction value; $\mathbb{1}(\text{Divisible by 500})$ is a binary indicator that takes a value of 1 if the transaction value is divisible by 500; $Tainted_i$ is a binary indicator equal to one if deposit address i is identified as tainted; and $Western_{i,e}$ is a binary indicator equal to one if deposit address i is on an exchange e that is classified as a Western exchange. All regressions control for exchange fixed effects (γ_e) and year-month fixed effects (η_t). Standard errors are clustered by exchange.

Dep. Variable:	$\mathbb{1}$(Divisible by 500)		$\mathbb{1}$(Divisible by 1000)	
	(1)	(2)	(3)	(4)
Tainted	0.129***	0.111***	0.121***	0.104***
	(0.008)	(0.015)	(0.007)	(0.015)
Tainted × Western		0.020		0.019
		(0.017)		(0.016)
Exchange FE	✓	✓	✓	✓
Year-Month FE	✓	✓	✓	✓
Observations	88,580,957	88,580,957	88,580,957	88,580,957
Adjusted R^2	0.040	0.040	0.028	0.028

Standard errors in parentheses

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