Crypto Capture of Foreign Aid*

Sumit Agarwal† Peiyi Jin‡ Eswar Prasad§ Daniel Rabetti¶

April 12, 2026

Abstract

Stricter oversight of traditional tax havens has not eliminated illicit financial flows; instead, it has shifted them toward cryptocurrency as a complementary channel. We examine whether foreign aid is diverted to crypto by linking World Bank disbursement data (2018–2024) to granular on-chain Bitcoin forensics, off-chain exchange records, and IP-linked web traffic. Using a stacked difference-in-differences design, we show that aid disbursements trigger a sharp, short-lived reorientation of crypto activity consistent with the classic placement–layering–integration sequence of money laundering. Flows move away from regulated exchanges toward tax-haven platforms and anonymous channels, with new wallet creation driving the extensive-margin response. Effects concentrate immediately after disbursement and dissipate within one to two months. Conservative lower-bound estimates imply leakage of roughly five cents per aid dollar, comparable to offshore banking evidence, and concentrated in countries with pre-existing offshore ties. Exploiting the Pandora Papers as an exogenous shock to concealment incentives, we find further amplification of the crypto-capture channel. Robustness and placebo tests rule out alternative explanations. While crypto enables diversion, its transparent ledgers also create a novel enforcement opportunity by providing real-time forensic traces for detection and potential recovery of misappropriated funds.

JEL classification: G15, G18, G29, K29, K42, O16.
Keywords: Corruption, Foreign aid, cryptocurrency, blockchain forensics, money laundering, tax haven.

---

*We are grateful to Lin William Cong, Micheal Sockin, Charles C. Y. Wang, Sean Foley, Pingyang Gao, John Griffin, Campbell Harvey, Shiyang Huang, Bjørn Jørgensen, Jonathan Karpoff, Evgeny Lyandres, Wayne Landsman, Mei Luo, and Eugene Soltes for their valuable early feedback on this study. We also thank colleagues at Harvard Business School, the National University of Singapore, Cornell University, the University of Hong Kong, and the School of Management and Economics at Tsinghua University for many insightful comments and suggestions. This study further benefited from extensive discussions at the NUS School of Computing Blockchain Labs, the Cornell Fintech Initiative (DEFT Labs), Asian Institute of Digital Finance (AIDF), Tsinghua University Blockchain Research Institute, Midwest Finance Association (MFA), and Asian Bureau of Economics and Finance Research (ABFER).

†National University of Singapore and ABFER, bizagarw@nus.edu.sg
‡National University of Singapore, jin_peiyi@u.nus.edu
§Cornell University and NBER, eswar.prasad@cornell.edu
¶NUS Business School and ABFER, rabetti@nus.edu.sg.

1 Introduction

The international development community channels billions of dollars annually to developing countries through foreign aid disbursements, yet a substantial fraction never reaches its intended beneficiaries. Landmark research documents that aid funds are regularly siphoned off by recipient-country elites and deposited in secret offshore accounts: Andersen, Johannesen and Rijkers (2022) estimate that as much as 7.5 cents of every aid dollar may be diverted in this way. For decades, the offshore banking system provided the infrastructure for this diversion, offering shell companies, nominee ownership, and financial secrecy across jurisdictions with light regulatory touch. The Panama Papers leak of 2016 exposed the scale of this system (International Consortium of Investigative Journalists, 2016) and triggered a wave of enforcement actions that raised the costs and risks of using traditional tax havens. Banks in secrecy jurisdictions came under pressure to strengthen compliance, and the veil surrounding shell corporations was partially lifted. Yet the diversion of foreign aid did not stop. It expanded into new channels. In this paper, we provide the first systematic evidence that corrupt elites have turned to cryptocurrencies as a new conduit for laundering diverted public funds, and that roughly 2 to 5 cents of every World Bank aid dollar disbursed travels through anonymous cryptocurrency channels within months of receipt.

Cryptocurrencies are well suited for this purpose, and their use as a financial concealment vehicle has grown precisely as traditional channels have come under pressure. Unlike wire transfers through regulated banks, cryptocurrency transactions require no correspondent banking relationship and no formal account in the recipient's name. Funds can be fragmented across thousands of freshly created, anonymous wallet addresses and transferred across borders instantly, bypassing the banking system and its reporting requirements (Amiram, Jørgensen and Rabetti, 2022). A proliferating network of exchanges domiciled in lax jurisdictions such as the Seychelles, Malta, and various offshore British territories provides entry and exit points where Know-Your-Customer (KYC) enforcement is minimal, effectively creating pseudo-offshore banks within the crypto ecosystem (Financial Action Task Force, 2021). Mixing services further obscure the provenance of funds by pooling transactions from multiple users and redistributing them in a manner that severs the on-chain link between sender and receiver.[1] At the same time,

[1] Cryptocurrency mixers (also referred to as tumblers) are services that deliberately obfuscate the transaction history of digital assets by pooling funds from multiple users and redistributing them in a manner that breaks the on-chain link between sending and receiving addresses. An estimated 46% of Bitcoin transactions involved illegal trade by the late 2010s, with roughly one-quarter of users participating in illicit activity (Foley, Karlsen and Putniņš, 2019). The share has declined as regulated exchanges have grown, but illicit use remains substantial.

blockchain transactions leave immutable public records, making cryptocurrency a double-edged sword: it attracts illicit actors while simultaneously providing regulators with a forensic tool that has no analogue in traditional offshore banking. We posit that the post-2016 crackdown on conventional tax havens has not eliminated illicit financial flows but has driven the emergence of cryptocurrency as a complementary channel, operated by the same elite networks that maintain offshore banking relationships. We call this "crypto capture" of aid, a digital analog of the classic elite capture documented by Reinikka and Svensson (2004).

We construct a novel dataset linking World Bank project-level disbursements from 2018 to 2024 to forensic on-chain Bitcoin data, off-chain exchange records, and web traffic data from SimilarWeb. The sample spans more than 90 recipient countries and over $300 billion in total disbursements. Wallet addresses are clustered and tagged to known entities, including exchanges, mixers, and payment processors, allowing us to distinguish anonymous from identified transactions at monthly frequency and to attribute crypto activity to specific exchange jurisdictions. We exploit the plausibly exogenous timing of World Bank disbursement tranches, which are governed by administrative milestones within pre-approved project cycles, including procurement benchmarks, auditing checkpoints, and board calendars, rather than by contemporaneous economic conditions (Kraay, 2012; Andersen et al., 2022). This wedge between the motivation for aid and the month in which tranches arrive provides the identifying variation in our stacked difference-in-differences design, estimated by two-way fixed effects (TWFE) on inverse hyperbolic sine (IHS)-transformed outcomes.

The evidence reveals a sharp and short-lived reorientation of cryptocurrency activity following each disbursement, consistent with the classic placement-layering-integration sequence of money laundering. Network graphs constructed around the five largest disbursements show transaction flows surging toward tax-haven exchanges and mixing services immediately after aid receipt and then reverting toward pre-disbursement structure within six months, a timing pattern consistent with front-loaded concealment activity rather than a persistent structural shift in the crypto ecosystem. In the econometric analysis, anonymous transaction volume rises by approximately 116% on tax-haven exchanges in the disbursement month, while identified transactions respond with magnitudes roughly half as large, indicating that the most opaque activity is disproportionately routed through less regulated venues. New anonymous wallet creation, the extensive margin of entry into the crypto ecosystem, increases by 1.5 to 2.7 IHS units across exchange types, while the identified-wallet response is an order of magnitude smaller, confirming that the entire extensive-margin response operates through anonymous channels. Off-chain data from centralized exchanges corroborate these patterns: transaction counts and volume ratios rise selectively on tax-haven and unregulated platforms, with unregulated jurisdictions exhibiting responses more than ten times larger than those in regulated ones. The geographic attribution underlying these estimates is independently validated against the Chainalysis Global Crypto Adoption Index, with Pearson correlations rising from 0.689 in 2021 to 0.855 in 2024, confirming that our web-traffic-based proxy recovers meaningful cross-country variation in cryptocurrency activity.

We quantify the economic magnitude through a leakage ratio. Our primary estimate uses the post-period average treatment effect on anonymous transaction volume on tax-haven exchanges, which captures the layering stage of the laundering cycle, adjusted for turnover using multipliers from blockchain forensics research. As a cross-check, we repeat the calculation using placement-stage anonymous inflows to non-tax-haven exchanges with minimal turnover correction; the two approaches converge on a range of approximately 2 to 5 percent of disbursed aid. These figures are comparable in order of magnitude to the 7.5 percent offshore banking leakage estimated by Andersen et al. (2022) and are derived from an entirely independent methodology based on blockchain forensics rather than banking records. Because our analysis is confined to Bitcoin, and stablecoins—particularly USDT on the Tron blockchain—have emerged as the dominant vehicle for illicit cross-border transfers in developing countries during our sample period, the 2 to 5 percent range represents a lower bound on total crypto-mediated leakage. The aid-to-crypto response is further concentrated in countries with greater pre-existing offshore banking exposure, as measured by the Andersen et al. (2022) offshore-to-GDP index, consistent with two interpretations: either the same elites exploit both channels in parallel, or both channels are driven by common institutional weaknesses that facilitate diversion through whichever vehicle is available.

Several complementary tests strengthen the causal interpretation. The Pandora Papers leak of October 2021 provides a natural experiment on offshore financial transparency: the post-leak decline in aid-linked crypto activity is concentrated entirely in identified transactions, those tied to formal, document-traceable structures, while anonymous channels remain unaffected. This asymmetry is consistent with pseudonymous wallets being insulated from document-based exposure and supports the view that anonymous crypto transactions represent a distinct and more resilient vector for concealing fund flows. A placebo test using disaster shocks that did not trigger World Bank disbursements produces uniformly negative crypto responses, directly ruling out the capital flight interpretation: crisis events alone, without an accompanying liquidity injection, suppress rather than stimulate crypto activity. Finally, significant effects emerge only in higher-governance countries with adequate digital infrastructure, implying that our estimates represent a lower bound on total aid diversion, since countries with weaker institutions and lower internet penetration may divert aid through crypto without generating detectable on-chain signals.

These findings carry direct policy implications. The displacement of illicit flows from offshore banking into cryptocurrency markets suggests that regulatory pressure on traditional tax havens is necessary but not sufficient: tightening one channel redirects rather than reduces diversion. At the same time, the public and auditable nature of blockchain ledgers opens new enforcement possibilities that do not exist in traditional offshore finance. Unlike bank accounts shielded by corporate secrecy and jurisdictional fragmentation, cryptocurrency flows can be monitored in near-real time and linked to specific disbursement events. This creates a tractable window for intervention, potentially enabling aid agencies and financial regulators to flag suspicious post-disbursement activity and, in some cases, trace and recover diverted funds before they are fully integrated into the formal economy.

Together, the evidence suggests a sequence consistent with the classical stages of laundering. In the placement stage, aid funds enter the crypto ecosystem through exchanges of all types. In the layering stage, funds are fragmented across newly created anonymous wallets and routed through tax-haven platforms. In the integration stage, outflows are consistent with funds exiting toward long-term anonymous storage or off-ramping through offshore channels. The crypto channel we document is distinct from traditional offshore banking in three respects: it operates at higher frequency, requires no pre-existing banking relationship, and exploits the pseudonymous architecture of blockchain networks rather than the legal opacity of corporate structures.

Related Literature

Our study bridges several strands of literature in economics and finance. In development economics, a foundational question concerns the extent to which foreign aid benefits intended recipients. Seminal work by Reinikka and Svensson (2004) documents substantial local capture of education aid in Uganda, while Werker, Ahmed and Cohen (2009) links aid flows to political-economy distortions in weak-institutional contexts. Andersen et al. (2022) provide direct evidence of elite capture via offshore banking, estimating that up to 7.5 cents per aid dollar ends up in foreign accounts controlled by recipient-country elites. More recently, Cavallo, Corbacho, Fernandez-Valdovinos, Ocampo and Singh (2024) demonstrate that cryptocurrency markets can function as platforms for capital flight by matching domestic demand for foreign exchange with offshore supply, while Gomez and Zhang (2024) show that trading volumes surge in response to geopolitical events, indicating the growing role of crypto in cross-border capital movements. We extend this literature by identifying a new, real-time digital destination for diverted aid and by demonstrating that blockchain forensics can serve as a complement to more traditional audits and banking records in tracing diversion at high frequency.

Our study also contributes to the literature on cryptocurrencies and illicit finance. Gandal, Hamrick, Moore and Oberman (2018) and Foley et al. (2019) document the scale of criminal activity in early crypto markets, ranging from drug trades to market manipulation. We build on this foundation by focusing on a specific mechanism: the laundering of diverted public funds through the crypto ecosystem. We show that such laundering exploits opaque corners of the market, including mixers and loosely regulated offshore platforms, consistent with Foley et al. (2019)'s finding that illicit actors adapt quickly to privacy-enhancing technologies. Our network analysis offers empirical support for the "layering" pathways documented in industry reports (Chainalysis, 2021) and aligns with academic tracing work on ransomware and money laundering (Ron and Shamir, 2013; Sokolov, 2021).

Finally, our study contributes to the nascent literature that leverages blockchain transparency for economic inference. Makarov and Schoar (2020) uncover cross-border arbitrage flows among exchanges, while Amiram et al. (2022) show that suspicious on-chain activity can predict terrorist attacks. Cong, Grauer, Rabetti and Updegrave (2023a) and Cong, Harvey, Rabetti and Wu (2025) propose frameworks for tracking cybercrime on blockchains. Our work differs from these studies in applying forensic techniques to aid diversion specifically and in strengthening identification through an event-study design that exploits quasi-exogenous variation in disbursement timing. More broadly, we contribute to the study of illicit financial flows, tax evasion, and capital control circumvention (Zucman, 2013; Regner, 2020), and to ongoing policy debates on crypto regulation, echoing calls for combining oversight with technical enforcement capacity (Cumming, Johan and Pant, 2019; Budish, 2025).

The remainder of the paper is organized as follows. Section 2 provides institutional background on foreign aid disbursement and the features of cryptocurrencies that make them attractive for laundering. Section 3 describes the data sources and construction of our on-chain and off-chain samples.

Section 4 presents the empirical analysis, opening with the identification strategy (Section 4.1) and geographic attribution proxy (Section 4.2), then presenting network evidence (Section 4.3), on-chain results (Section 4.4), leakage estimates (Section 4.5), and off-chain results (Section 4.6). Section 5 synthesizes the evidence within the placement-layering-integration framework. Section 6 presents robustness checks. Section 7 concludes.

2 Institutional Background

International aid flows are vulnerable to capture by political and economic elites in recipient countries. Traditional methods for concealing diverted funds involve routing money to offshore bank accounts in secrecy jurisdictions or to shell companies in tax havens. Andersen et al. (2022) find that foreign aid disbursements coincide with significant increases in deposits in offshore financial centers, implying that as much as 7.5 cents of each aid dollar may be diverted. However, the risks of such schemes have increased: the 2016 Panama Papers leak (International Consortium of Investigative Journalists, 2016) and subsequent exposés led to greater enforcement of anti-tax-evasion measures and international cooperation in tracking illicit flows. Banks in traditional havens came under pressure to strengthen compliance, and the veil of secrecy surrounding shell corporations was partially lifted. Consequently, corrupt actors began seeking alternative channels with a lower risk of detection.

2.1 Cryptocurrencies as a Modern Laundering Conduit

Cryptocurrencies offer such a channel. Converting misappropriated funds into cryptocurrency, via exchanges or peer-to-peer markets, transforms identifiable bank transfers into pseudonymous holdings whose provenance can be obscured. Several features make crypto particularly attractive: users can create new addresses at will, fragmenting funds across thousands of wallets to reduce traceability; transfers cross borders instantly without relying on the banking system or its reporting requirements; and no central authority can freeze or scrutinize transactions on the blockchain. Most importantly for our setting, the proliferation of cryptocurrency exchanges in jurisdictions with relaxed regulatory oversight, such as the Seychelles, Malta, and offshore British territories, provides entry and exit points where enforcement of customer identification and anti-money-laundering requirements is often lax, effectively creating pseudo-offshore banks within the crypto ecosystem.

Early evidence confirms that illicit actors recognized these advantages quickly. Foley et al. (2019) show that billions of dollars in illicit transactions flowed through Bitcoin annually by the late 2010s. At the same time, blockchain transactions leave immutable digital trails, and advances in forensic analytics (address clustering, entity tagging, and transaction-pattern identification) have increasingly pierced the veil of pseudonymity, enabling both academic and law-enforcement breakthroughs (Amiram et al., 2022). In our setting, these forensic methods allow us to test whether aid disbursements are followed by patterns of crypto activity consistent with diversion: spikes in anonymous activity, rapid creation of new wallets, and movement of funds from regulated on-ramps to offshore or lightly monitored platforms.

3 Data Acquisition and Processing

Our analysis requires combining data from diverse sources: foreign aid disbursement records, on-chain crypto transaction data (with address identification), and off-chain exchange activity data. Descriptions of specific variables used in our analysis are in Table A.1. Below, we describe each in turn, along with our data processing methods.

3.1 Foreign Aid Data

We construct a high-frequency dataset of foreign aid disbursements, with a focus on World Bank aid. We use the World Bank's Project Database (via its API) to obtain detailed information on project loans and grants provided by the International Development Association (IDA) and the International Bank for Reconstruction and Development (IBRD), the arms of the World Bank.[2] For each project, we extract approval dates and commitment amounts, and, crucially, we use disbursement records that indicate when funds are actually disbursed. Our main aid variable, denoted *Aid*, measures the total amount of World Bank aid disbursed to a given country in the 90-day period leading up to date t. In practice, we aggregate disbursements on a rolling quarterly basis (90 days) to align with our analysis frequency.

[2]IDA (the International Development Association) provides concessional loans and grants to the world's poorest countries, typically those with per capita income below a threshold set by the World Bank and limited access to international capital markets. IBRD (the International Bank for Reconstruction and Development) provides non-concessional loans and guarantees to middle-income and creditworthy low-income countries to support development projects and policy reforms.

Focusing on World Bank aid (as opposed to all Official Development Assistance (ODA)) has a key advantage: the timing of the disbursement data.[3] Most aid datasets (e.g., OECD ODA statistics) are annual and do not reveal precisely when within the year funds were transferred. By contrast, World Bank projects provide more granular disbursement dates. This is crucial for our identification strategy, as we aim to observe crypto flows within narrow windows around aid events. We restrict our sample period to April 2018 through March 2024, which reflects the period for which we have both aid data and reliable crypto transaction data.

Table 1, Panel A reports the top 15 recipients by disbursed amount. Aid is concentrated among a small number of large recipients: India, Indonesia, and Colombia alone account for roughly 16 percent of total disbursements, and the top 15 countries together represent half of all disbursed funds. Recipients span South Asia, Sub-Saharan Africa, Latin America, and Eastern Europe, providing broad geographic variation for identification. Panel B describes the regulatory environment of the exchange jurisdictions in our sample. Approximately 45 percent of these jurisdictions are classified as tax havens (based on exchange registration and regulatory criteria), and 70 percent as crypto-friendly, while the remainder maintain stringent regulatory frameworks. This distribution, which reflects where exchanges are domiciled rather than where aid recipients are located, motivates our heterogeneity analysis by regulatory type.

3.2 On-Chain Data and Blockchain Forensics

The core of our analysis is based on on-chain Bitcoin transaction data. Bitcoin is the dominant cryptocurrency, and we can leverage the public availability of its full transaction ledger. We obtained Bitcoin blockchain data for the period 2018–2024, including information on every transaction and the addresses involved. To synthesize these raw data, we apply identification and clustering techniques to label addresses and group them into wallets (user clusters). It consists of two main steps:

Address Attribution: First, we attribute addresses to known entities using public and forensic sources, including WalletExplorer and open-source blockchain investigations. These sources identify addresses as-

[3] IDA provides concessional financing — interest-free credits and grants — to countries with per capita incomes below the World Bank's operational threshold; IBRD provides market-rate loans to middle-income and creditworthy low-income countries. Together, IBRD and IDA disburse approximately $30–60 billion annually over our sample period (FY2018–2024), making the World Bank the single largest multilateral development lender. Disbursement data are from the World Bank Group Finances portal (`https://financesone.worldbank.org/summaries/ibrd-ida`); for context, total net ODA from DAC members was approximately $215–225 billion annually over this period, as reported by OECD DAC statistics (`https://www.oecd.org/en/topics/official-development-assistance-oda.html`).

sociated with exchanges, mixers, gambling platforms, darknet markets, payment processors, and related services, enabling us to tag transactions involving these services.

Addresses attributed to known services are classified as non-anonymous at the service level, meaning the platform is observable on-chain. This does not identify the underlying end user. Wallets interacting with identified services remain classified as anonymous unless they themselves can be attributed to a known service, even when transacting with exchanges in tax-haven jurisdictions.

Wallet Clustering: Individual blockchain addresses are highly fragmented, as users can generate many addresses. To recover user-level entities, we apply a union–find clustering algorithm to the transaction graph. Following standard practice (Ron and Shamir, 2013), addresses that appear jointly as inputs in the same transaction are clustered and treated as controlled by a single entity. Applying this and related heuristics yields wallet clusters corresponding to services, and prior work shows that aggregate exchange and flow patterns are robust to alternative clustering rules.

Wallets are classified into one of five categories—Exchange, Mixer, Gambling, Dark Market, or Other Service—using attribution data. Wallets that cannot be matched to known services are labeled as anonymous. Throughout the analysis, anonymity is defined on-chain: anonymous wallets are those not attributable to identifiable services, rather than unidentified individuals. We compare activity involving anonymous versus identified wallets, where post-disbursement increases in anonymous activity are indicative of heightened anonymity-seeking behavior.

Identification Under Incomplete Wallet Labeling: We acknowledge that our labeling does not cover the full universe of wallets. This limitation introduces classification noise rather than systematic mislabeling: anonymous wallets are not incorrectly assigned.

Identification relies on a stacked difference-in-differences design that exploits within-entity changes in the timing of aid disbursements. As long as labeling is unrelated to the timing of aid events, incomplete coverage does not bias the estimates; rather, it attenuates the estimated effects toward zero.

Table 1 Panel C and D summarize the cryptocurrency transaction platforms in our on-chain data. Panel C reports account counts and 24-hour trading volumes across 15 platforms. The sample spans major centralized exchanges (Binance, Kraken, Huobi), privacy-enhancing services (CoinJoinMess), and smaller or specialized platforms. CoinJoinMess accounts for 77 percent of all new accounts but only 2 percent of trading volume, consistent with its role as an anonymization tool rather than a trading venue. By contrast, Binance, Kraken, and Huobi collectively represent 82 percent of total volume. Panel D reports the geographic distribution of exchange activity by domicile. Trading volume is heavily concentrated in exchanges based in jurisdictions with favorable regulatory environments: Malta, Singapore, and the Seychelles together account for 72 percent of total 24-hour volume, while the United States contributes only 3 percent. This pattern underscores the tendency of crypto trading infrastructure to cluster in tax-haven or crypto-friendly jurisdictions.

3.3 Off-Chain Exchange Data

While on-chain data capture the movement of Bitcoin between addresses, they do not directly reveal trades or transactions occurring within exchanges (for example, two customers trading Bitcoin on Coinbase's order book leave no trace on the public blockchain aside from their deposit/withdrawal transactions). To complement our analysis, we collect off-chain data on exchange activity from CoinMarketCap's API and other aggregate sources.

The off-chain dataset includes quarterly figures on trading volume and transaction counts for a large set of cryptocurrency exchanges worldwide. We compile statistics for exchanges that account for the majority of global trading volume and categorize them by geography and regulatory environment. Specifically, we group exchanges into the same jurisdictional categories as above: tax haven-based vs. non-haven, regulated vs. unregulated, and crypto-friendly vs. not friendly. These designations draw on exchange registration information and indices of crypto regulatory openness.

Some summary statistics from our exchange classification highlight why this distinction matters: roughly 45% of active crypto exchanges in our sample are located in jurisdictions classified as tax havens, about 18% are in strongly regulated countries, and about 70% are in what can be termed "crypto-friendly" countries (those that actively promote crypto business or lack strict regulations). There is considerable overlap among these categories (for example, many tax haven jurisdictions are also crypto-friendly due to low regulation). The key point is that a significant portion of crypto trading infrastructure is based in places that could serve as new "offshore" havens for illicit flows.

Using CoinMarketCap data, we focus on off-chain transaction counts (trades or transfers) and exchange volumes around aid events. If, for example, a large aid disbursement to Country X is followed by a spike in trading volume on exchanges in the Cayman Islands, this would be consistent with elites in Country X converting aid money into crypto via those exchanges. We also examine the activity ratio between different groups of exchanges (e.g., tax haven vs. non-haven) around aid events to assess whether relative shifts occur.

4 Empirical Analysis

We open with the identification design and geographic attribution strategy that underpin all subsequent econometric results. We then present three bodies of empirical evidence: network graphs that visualize cryptocurrency flows around major disbursements, on-chain transaction analysis using a stacked difference-in-differences design, and off-chain exchange-level analysis.

4.1 Identification Strategy

Foreign aid disbursements are staggered across countries and may recur within the same country over time. To accommodate this treatment structure, we employ a stacked difference-in-differences (DiD) design at the country-platform-month level, estimated using two-way fixed effects (TWFE) with outcomes transformed via the inverse hyperbolic sine (IHS). We choose linear TWFE with IHS as the primary specification for two reasons. First, country-level crypto transaction data contain a large mass of zero observations, many country-platform-month cells record no Bitcoin activity whatsoever, which renders log-linear models infeasible without ad hoc adjustments. The IHS transformation, $\text{ihs}(x) = \ln\big(x + \sqrt{x^2+1}\big)$, accommodates zeros naturally while preserving a semi-elasticity interpretation analogous to log-linear models for large values of x (Bellemare and Wichman, 2020). Second, linear TWFE on IHS-transformed outcomes is computationally transparent, facilitating replication across the large number of outcome-platform-specification combinations in our analysis. As a robustness check, Appendix D.4 confirms that the main results are qualitatively unchanged when estimated using Poisson pseudo-maximum-likelihood (PPML). Appendix D.5 verifies that the baseline findings are robust to alternative control group compositions, including restricting the control pool to developing economies and excluding neighboring countries.

We define episodes as follows. When two or more disbursements to the same country fall within the post-treatment window, we aggregate them into a single episode and set event time zero to the month of the initial disbursement. This prevents post-treatment periods from being mechanically re-treated and guards against bias from overlapping episodes. For each treated episode, the control group consists of countries that receive no World Bank disbursement during the entire analysis window (a "no-overlap" control group). To guard against spillovers through trade or migration channels, we also exclude from the control group any country that shares a land border with the treated country, and any country whose own treatment window overlaps with the current episode. As a robustness check, we additionally restrict the control group to developing countries only, ensuring comparability in institutional environment and baseline crypto adoption.

Aid allocation may respond to observable economic conditions, raising concerns about endogeneity. Our design includes country-by-platform-by-episode (Ctry.$\times$Plat.$\times$Ep.) and calendar-month fixed effects, and event-study plots confirm the absence of differential pre-trends in crypto outcomes. Following Kraay (2012) and Andersen et al. (2022), we further exploit the fact that World Bank disbursements are governed by multi-year project cycles and predetermined release schedules. While the allocation of aid across countries is endogenous to economic conditions, the precise timing of disbursement tranches within approved projects is driven by administrative milestones such as procurement benchmarks, auditing cycles, and board calendars that are plausibly orthogonal to contemporaneous local shocks. This wedge between the motivation for aid and the month in which funds arrive is the source of identifying variation in our design.

Let c index countries, p index platforms (tax-haven or non-tax-haven crypto exchanges), and t index months. We define an "episode" e as a treated country-month pair corresponding to an aid disbursement event, and stack multiple episodes following Cengiz, Dube, Lindner and Zipperer (2019). For each episode e, the treated unit is the disbursement country $c(e)$, and event time is

$$k \equiv t - t_e,$$

where t_e denotes the month of the disbursement. We estimate event-time coefficients over the window $k \in \{-4, \ldots, 4\}$ and omit $k = -1$ as the reference period.

For each outcome Y_{cpt}, we estimate the following specification separately for tax-haven and non-tax-haven platforms:

$$Y_{cpt}^{(e)} \;=\; \sum_{\substack{k=-4 \\ k\neq -1}}^{4} \beta_k \cdot \mathbf{1}\big\{\text{Treated}_c^{(e)}\big\}\,\mathbf{1}\big\{t - t_e = k\big\} \;+\; \alpha_{cp}^{(e)} \;+\; \gamma_t \;+\; \varepsilon_{cpt}^{(e)}, \tag{1}$$

where the superscript (e) indexes the stacked episode. Each observation is weighted by the inverse of the number of episodes in which it appears, following the stacking-weights approach of Cengiz et al. (2019), ensuring that no unit receives disproportionate influence from appearing in multiple cohorts.

Here, $\alpha_{cp}^{(e)}$ denotes Ctry.×Plat.×Ep. fixed effects that absorb all time-invariant differences in baseline platform usage and crypto adoption across countries within each episode. Calendar-month fixed effects γ_t absorb global shocks common to all countries and platforms, such as Bitcoin price movements, market-wide traffic cycles, and major news events. The pre-treatment coefficients β_k for $k \in \{-4, \ldots, -2\}$ serve as a pre-trend diagnostic: under the parallel-trends assumption, these should be individually and jointly insignificant. Standard errors are clustered at the Ctry.×Plat.×Ep. level with the finite-sample correction $G/(G-1)$, where G is the number of clusters (Cameron and Miller, 2015).

To summarize the post-disbursement treatment effect, we report the post-period average:

$$\overline{\text{ATT}} \;=\; \frac{1}{L+1}\sum_{k=0}^{L} \hat{\beta}_k, \qquad L = 4,$$

with standard error $\text{SE}(\overline{\text{ATT}}) = \sqrt{\mathbf{w}'\,\widehat{V}\,\mathbf{w}}$, where $\mathbf{w} = \big(\frac{1}{L+1}, \ldots, \frac{1}{L+1}\big)'$ selects the post-period positions in the cluster-robust variance-covariance matrix $\widehat{V}$.

4.2 Geographic Attribution via IP-Linked Web Traffic

A central empirical challenge in studying cryptocurrency activity is the absence of reliable geographic identifiers in on-chain data. Blockchain transactions do not reveal users' country of origin, and exchange-level data typically aggregate activity across multiple jurisdictions. To address this, we use web traffic to cryptocurrency exchange platforms as a proxy for country-level engagement with cryptocurrency trading.[4]

Web visits and time spent on crypto platforms capture information acquisition, account access, and transaction execution behavior that typically precedes or accompanies trading activity. While web traffic does not directly measure realized on-chain transactions, it provides a high-frequency, geographically attributable signal of trading intent and participation. Prior work has adopted related strategies: Foley et al., 2019 combine on-chain forensics with digital footprint data to attribute Bitcoin transactions across jurisdictions, and Cong, Li, Tang and Yang, 2023b use SimilarWeb traffic rankings to benchmark legitimate trading activity. Our approach extends this methodology by linking geographically attributed trading activity to specific disbursement events rather than estimating static cross-sectional patterns.

We construct three country-month measures of web traffic. Total visits captures the absolute intensity of engagement with crypto platforms. Total visit duration measures depth of engagement; where duration data are missing for a country-platform pair, we approximate it by multiplying the observed visit count by the global average time spent per visit in that month. Finally, the visits ratio, defined as a country's share of global visits to crypto-related websites in a given month, captures relative position in the global distribution of crypto traffic. We use this ratio to scale global Bitcoin transaction volume and transaction frequency, constructing implied country-month measures of on-chain activity:

$$\widetilde{Y}_{cpt} \;=\; \text{VisitsRatio}_{cpt} \times G_t,$$

where G_t denotes the relevant global on-chain aggregate. All outcomes are IHS-transformed before estimation.

We focus on a short monthly event window for two reasons. First, it guards against confounding from slow-moving structural changes such as improvements in internet infrastructure, which require substantial planning time and are unlikely to generate the sharp, quickly dissipating response we observe. Second, daily regressions are infeasible due to the high mass of zero disbursement observations at that horizon. Our first specification estimates the effect of aid disbursements on web-traffic flows directly. Our second specification rescales visit flows by the global Bitcoin price, which is a common shock that does not vary by country and preserves the exogeneity of the treatment variable.

[4] Examples include tax-haven-domiciled exchanges such as HTX (`https://www.htx.com/`) and regulated exchanges such as Kraken (`https://www.kraken.com/`).

An important caveat applies to the interpretation of our geo-adjusted on-chain outcomes. Because these measures are constructed by scaling global Bitcoin aggregates by each country's web-traffic share, variation in $Y_{ecpt} = \text{VisitsRatio}_{cpt} \times G_t$ is driven entirely by cross-country differences in the visits ratio, since G_t does not vary across countries within a given month. The geo-adjusted on-chain results therefore do not constitute an independent body of evidence from the web-traffic analysis; rather, they decompose the web-traffic response along dimensions—anonymity status, exchange jurisdiction, and activity margin—that raw traffic data cannot distinguish. This decomposition reveals that the post-disbursement surge is disproportionately concentrated in anonymous transactions and on tax-haven platforms, but the country-level identification in both sets of results traces to the same source of geographic variation. Independent corroboration comes from two other pillars: the off-chain exchange data, which use exchange-reported transaction counts not derived from web traffic, and the network graphs, which trace address-level Bitcoin flows without relying on geographic attribution. The on-chain tables should therefore be read as a detailed anatomy of the web-traffic response rather than as a separate confirmation of it.

4.3 Network Evidence

We construct directed network graphs to visualize cryptocurrency flows before and after major aid disbursements. Each node represents a service cluster such as an exchange or mixer, with size proportional to total volume handled, and directed edges represent Bitcoin flows with thickness proportional to flow volume. Nodes are color-coded by service type: non-tax-haven exchanges (purple), tax-haven exchanges (orange), payment platforms (blue), mixers (pink), and gambling or other services (green).

Figure 1 shows the network for the five largest aid disbursements in our sample (summary statistics in Table A.2), plotting address-level snapshots one month before, one month after, and six months after each event. Before disbursement, the network is dominated by non-haven exchanges, reflecting routine trading and remittance activity; tax-haven platforms and mixers are present but peripheral. In the month following disbursement, the network reconfigures sharply: flows surge toward tax-haven exchanges, which become central redistribution hubs, while mixers intensify their role as intermediaries channeling funds away from regulated on-ramps. Six months later, the pattern partially reverts, with non-haven exchanges regaining relative prominence and obfuscation-related flows subsiding, suggesting that concealment activity is front-loaded rather than persistent.

Figure 2 presents the same dynamics at the category level, aggregating platforms of the same service type into single nodes to clarify inter-category flow patterns. Within one week of a major disbursement, inflows to tax-haven exchanges such as HTX intensify while outflows from regulated non-haven platforms rise simultaneously, consistent with funds being withdrawn from compliant exchanges and redirected toward jurisdictions with weaker oversight. Flows to mixing services also thicken visibly, suggesting that the placement and layering stages of the laundering cycle activate in concert. By six months post-disbursement, the network topology closely resembles the pre-event baseline: elevated inflows to tax-haven platforms dissipate, outflows from regulated exchanges normalize, and the transient connections to privacy-enhancing services fade. This rapid reversion reinforces the interpretation that laundering activity is triggered by discrete disbursement events rather than a persistent shift in the structure of illicit financial flows.

Overall, the network evidence provides a visual depiction of how aid disbursements are followed by a rapid, temporary reorientation of crypto flows from transparent to opaque venues, consistent with concealment behavior documented in law enforcement and forensic studies.

4.4 On-Chain Results

4.4.1 Anonymous vs. Identified Transactions

To quantify heterogeneity by anonymity status and exchange jurisdiction, Tables 2 and A.4 report post-period average treatment effects ($k = 0, \ldots, 4$) of aid disbursements on transaction volume and transaction frequency, estimated using a stacked difference-in-differences design with linear TWFE and a no-overlap control group. All outcomes are IHS-transformed, so coefficients can be interpreted as semi-elasticities. Figure 3 plots the full event-study profiles using geo-adjusted measures that scale global Bitcoin activity by each country's web-traffic share. Because these geo-adjusted outcomes are constructed by multiplying a global on-chain aggregate by each country's web-traffic share (Section **??**), the country-level variation in the tables below traces to the same geographic signal as the web-traffic analysis in Section 4.4.3. The decomposition by anonymity status, exchange jurisdiction, and activity margin that follows is therefore best read as an anatomy of the web-traffic response rather than as independent confirmation of it; independent corroboration comes from the off-chain exchange data (Section **??**) and the network graphs (Section **??**).

We begin with transaction volume in Table 2. On tax-haven exchanges, the response is systematically stronger for anonymous than for identified transactions. In Panel A, the post-period average treatment effect for anonymous transaction volume is 0.770, compared with 0.431 for identified transactions. The same pattern holds for both inflows and outflows. For inflows (Panel B), anonymous volume rises by 0.695 versus 0.376 for identified transactions. For outflows (Panel C), anonymous volume increases by 0.706 versus 0.334. These results indicate that the post-aid increase in on-chain transaction value on tax-haven exchanges is disproportionately concentrated in anonymous channels, consistent with the use of fragmented transfers and intermediary wallets to reduce traceability.

A second feature of Table 2 is that aid disbursements also generate large and statistically significant increases on non-tax-haven exchanges. In Panel A, the coefficient for anonymous transaction volume on non-tax-haven exchanges is 1.226, compared with 0.770 on tax-haven exchanges. Identified transaction volume also rises substantially, with a coefficient of 0.647 on non-tax-haven exchanges versus 0.431 on tax-haven exchanges. The larger absolute magnitudes on non-tax-haven exchanges are not surprising from a placement-stage perspective: regulated, non-tax-haven platforms are the primary fiat-to-crypto on-ramps, since they have the banking relationships and liquidity necessary to absorb large initial conversions. Funds must enter the ecosystem somewhere, and the most accessible entry points are mainstream exchanges. What distinguishes the tax-haven response is not its absolute size but its composition: the anonymous-to-identified gap is systematically larger on tax-haven platforms, reflecting their function as layering hubs where funds are reshuffled across anonymous wallets and routed onward rather than simply converted. The two findings are therefore complementary rather than contradictory: non-tax-haven exchanges absorb the initial placement flow, while tax-haven exchanges concentrate the anonymity-oriented layering activity.

Table A.4 shows that transaction frequency follows the same pattern. On tax-haven exchanges, the post-period average treatment effect for anonymous transaction frequency is 0.757, compared with 0.356 for identified transactions in Panel A. For inflows (Panel B), anonymous frequency rises by 0.678 versus 0.310 for identified transactions. For outflows (Panel C), anonymous frequency increases by 0.687 versus 0.267. On non-tax-haven exchanges, anonymous frequency is again larger in absolute magnitude, reaching 1.158 in Panel A, compared with 0.579 for identified transactions. The frequency results therefore reinforce the interpretation from transaction volume: aid disbursements increase activity throughout the exchange ecosystem, but the increase is disproportionately concentrated in anonymous channels, especially on tax-haven platforms.

Appendix tables show that this pattern is robust to alternative control-group constructions. When the control group is restricted to developing countries that did not receive World Bank disbursements during the estimation window, the estimated effects remain nearly unchanged. For example, for all transactions, the coefficient on anonymous transaction volume in tax-haven exchanges is 0.776, compared with 0.435 for identified transactions, while the corresponding frequency coefficients are 0.762 versus 0.359. Likewise, when neighboring countries are excluded from the control group, the corresponding estimates remain very similar: 0.768 versus 0.430 for volume and 0.755 versus 0.355 for frequency. The same stability holds on non-tax-haven exchanges and for new-account creation. These appendix results indicate that the main findings are not driven by comparisons with high-income non-recipients or by geographic spillovers from nearby countries.

The evidence reveals two complementary patterns. First, aid disbursements generate a large and significant increase in on-chain Bitcoin activity on both tax-haven and non-tax-haven exchanges, with anonymous transactions consistently exhibiting the strongest responses. Second, the anonymous–identified gap is systematically larger on tax-haven platforms, suggesting that while aid-linked crypto activity is widespread, the most opaque segment of the transaction network is disproportionately routed through exchanges that offer greater anonymity and weaker regulatory oversight.

4.4.2 New Wallet Creation

We next examine whether aid disbursements are followed by the creation of new wallets, a salient extensive-margin response commonly associated with laundering strategies that rely on fresh identifiers to reduce traceability. Panels (c) and (d) of Figure 3 plot event-study estimates for the geo-adjusted number of newly created Bitcoin accounts, and Table 3 reports the post-period average treatment effects ($k = 0, \ldots, 4$) by exchange jurisdiction and anonymity status using a stacked difference-in-differences with linear TWFE and IHS-transformed outcomes.

The most striking feature of Table 3 is the near-complete concentration of new wallet creation in anonymous accounts. On tax-haven exchanges, the post-period average for anonymous new accounts is 1.457 IHS units and highly significant across all three panels, while the corresponding effect for identified accounts is an order of magnitude smaller: 0.114 in Panel A, 0.074 in Panel B, and a statistically insignificant 0.037 in Panel C. Notably, anonymous and all-transaction coefficients are numerically identical on tax-haven exchanges (1.457 in Panel A, 1.376 in Panel B, 1.374 in Panel C), indicating that the entire extensive-margin response on these platforms operates through anonymous wallets.

Non-tax-haven exchanges display a qualitatively similar but amplified pattern. Anonymous new-account creation is positive and strongly significant across all panels, with considerably larger magnitudes than on tax-haven platforms: the coefficient reaches 2.706 for aggregate anonymous accounts (Panel A), 2.629 for anonymous inflow accounts (Panel B), and 2.478 for anonymous outflow accounts (Panel C). As on tax-haven exchanges, the anonymous and all-transaction coefficients are virtually identical, confirming that new wallet creation is driven almost entirely by anonymous accounts regardless of exchange jurisdiction. Identified new-account creation on non-tax-haven exchanges is positive but very small: 0.170 in Panel A, 0.065 in Panel B, and 0.105 in Panel C, preserving the sharp anonymity gap observed on tax-haven platforms.

Two features of these results merit emphasis. First, the anonymous-to-identified gap in new wallet creation is far more extreme than the corresponding gap in transaction frequency and volume documented above. Rather than scaling up activity through pre-existing identified accounts, aid episodes are associated with rapid entry into the crypto ecosystem via newly created anonymous wallets, consistent with intentional reshuffling aimed at obscuring provenance from the outset. This matters for the layering interpretation: it suggests that anonymity-seeking is not a secondary step taken after an initial identified conversion, but a deliberate strategy employed from the first point of contact with the crypto ecosystem. Actors appear to create unattributed entry points for funds rather than converting via identified accounts and anonymizing later, a pattern that is significantly harder to trace forensically.

Second, the larger anonymous new-account effects on non-tax-haven exchanges (approximately 2.5 to 2.7 IHS units versus 1.4 to 1.5 on tax-haven platforms) reflect the placement logic described above. Non-tax-haven exchanges serve as the initial on-ramp, and the near-zero identified new-account response on both exchange types confirms that the new participants entering the ecosystem around disbursement events are choosing anonymous accounts from the start, regardless of platform jurisdiction. Combined with the transaction-level results in Tables 2 and A.4, which show significant activity increases on both exchange types, this pattern indicates that aid-linked crypto adoption operates through anonymous channels as a general phenomenon, with the most opaque segment of the transaction network concentrated on tax-haven exchanges where regulatory oversight is weakest.

4.4.3 Platform Engagement Evidence

The on-chain results document sharp post-disbursement increases in anonymous transaction volume, frequency, and new wallet creation, but they do not directly reveal the behavioral channel through which aid recipients enter the cryptocurrency ecosystem. Web traffic to exchange platforms provides this evidence. Visits and time spent on crypto exchanges capture information acquisition, account access, and transaction execution—activities that necessarily precede or accompany the on-chain flows documented above. By examining platform engagement directly, we can distinguish between an intentional demand-side response to aid receipt and passive variation driven by slow-moving factors such as internet penetration or general cryptocurrency adoption trends.

Figure 4 plots stacked DiD event-study estimates of web traffic to cryptocurrency-exchange websites around foreign-aid disbursement events. Panel (a) reports total visits and Panel (b) reports total visit duration, both measured at the country-month level. In both panels, the response rises sharply in the disbursement month ($k = 0$), reaches its peak on impact, and then declines gradually over the following months while remaining positive for several post-treatment periods. Pre-treatment coefficients are relatively modest and do not exhibit a pronounced pre-trend immediately prior to treatment. Table A.3 in the appendix confirms these findings across linear TWFE and PPML specifications: the TWFE coefficient on IHS-transformed total visits is approximately 2.28, stable across all three control-group definitions, with corresponding PPML incidence-rate ratios of 15.5–20.9.

Three features of this response identify it as evidence of a deliberate behavioral evidence. First, the response is immediate: it peaks in the disbursement month itself, ruling out slow-moving drivers such as infrastructure expansion or gradual cryptocurrency adoption. Second, it attenuates within a few months, matching the temporal profile of the on-chain transaction and new-wallet results in Sections 4.4.1 and 4.4.2 and consistent with a discrete conversion event rather than a permanent shift in internet behavior. Third, the traffic is directed specifically at exchange platforms—venues where fiat-to-crypto conversion occurs—rather than at general cryptocurrency information sites. Together, these dynamics establish that aid disbursements trigger active, purposeful engagement with cryptocurrency trading infrastructure, providing the demand-side behavioral link between the receipt of aid funds and the on-chain patterns documented above.

4.5 Leakage Ratio

To quantify the economic magnitude of aid-linked cryptocurrency activity, we translate the estimated semi-elasticities into dollar-valued transaction volumes and express them as a share of foreign-aid disbursements. A useful benchmark is Andersen et al. (2022), who estimate offshore banking leakage of approximately 7.5% of disbursed aid using BIS cross-border deposit data. Our approach uses an entirely independent methodology based on blockchain forensics, making the two estimates genuinely comparable rather than mechanically related.

Our regression outcomes are IHS-transformed, $\text{ihs}(y) = \ln\bigl(y + \sqrt{y^2+1}\bigr)$, where y denotes geo-adjusted transaction volume in Bitcoin. For large y, $\text{ihs}(y) \approx \ln(2y)$, so coefficients admit a semi-elasticity interpretation (Bellemare and Wichman, 2020). We construct a pre-disbursement baseline $\bar{y}^0 = 1{,}271$ BTC using the average anonymous transaction volume on tax-haven exchanges in the window $k \in \{-4,-3,-2\}$ for treated-country cells. We use the post-period average treatment effect ($\widehat{\text{ATT}} = 0.770$), which averages the event-time coefficients over $k = 0,\ldots,4$, rather than the impact coefficient at $k = 0$ alone. This choice captures the cumulative footprint of laundering activity across the full post-disbursement window: because the placement-layering-integration cycle unfolds over several months, restricting the calculation to the impact month would understate the total volume of aid-linked transactions that pass through the crypto ecosystem, while coefficients at later horizons reflect the layering and integration stages documented in Section 5. The implied level change is $\Delta y \approx (e^{0.770}-1) \times 1{,}271 \approx 1{,}474$ BTC, corresponding to a 116% increase relative to the pre-disbursement baseline. At the sample-period average Bitcoin price of $21,763, this translates to approximately $32.1 million in additional anonymous transaction value per month over the five-month post-disbursement window. Against an average disbursement of $305 million, the implied *gross* leakage ratio is

$$\lambda_{\text{gross}} = \frac{\Delta\text{CryptoValue}}{\text{Aid}} \approx 10.5\%.$$

This gross ratio measures transaction-value exposure, not net diversion, because cryptocurrency volumes reflect gross flows: each diverted dollar may be observed multiple times as funds pass through successive wallets or exchanges before exiting the crypto ecosystem. Blockchain forensics research suggests that illicit funds typically traverse two to five wallet hops before off-ramping into fiat currency, implying a turnover multiplier in the range $m \in \{2,3,5\}$ (Chainalysis, 2021). We report the turnover-adjusted ratio $\lambda_{\text{net}} \approx \lambda_{\text{gross}}/m$ for each value, yielding adjusted estimates of approximately 5.3%, 3.5%, and 2.1%, respectively. We stress that m is not directly observed in our data; these values should be interpreted as a transparent sensitivity range rather than point estimates of net diversion. The central estimate of approximately 2 to 5% is comparable in order of magnitude to the 7.5% offshore banking leakage in Andersen et al. (2022), derived here from a channel that operates at higher frequency, requires no pre-existing banking relationship, and exploits pseudonymous blockchain architecture rather than the legal opacity of corporate structures.

As a cross-check on both the coefficient choice and the turnover assumption, we repeat the calculation using anonymous inflows to non-tax-haven exchanges ($\widehat{\text{ATT}} = 1.101$, Table 2, Panel B), which capture the initial placement of fiat-denominated aid into the crypto ecosystem before subsequent layering across wallets and tax-haven platforms inflates observed volumes. The pre-disbursement baseline for treated-country cells on non-tax-haven exchanges is $\bar{y}^0_{\text{NTH}} = 200$ BTC, substantially smaller than the tax-haven baseline, consistent with volumes accumulating as funds are fragmented during the layering stage. Applying the same turnover multiplier range $m \in \{2, 3, 5\}$ yields adjusted estimates of approximately 1.4%, 1.0%, and 0.6%, respectively. These placement-based estimates are smaller than the layering-based range of 2.1–5.3%, as expected: the non-tax-haven baseline captures only the initial on-ramp volume before funds are recycled across multiple wallets and exchanges. The two approaches nonetheless confirm that crypto-mediated leakage is on the order of 1 to 5 percent of disbursed aid, with the tax-haven layering estimate providing an upper bound and the non-tax-haven placement estimate providing a lower bound.

The offshore heterogeneity results reported in Table 6 help discipline the comparison between the crypto and offshore banking channels. Countries with greater pre-existing offshore banking exposure, as measured by the Alstadsæter–Johannesen–Zucman offshore-to-GDP index (Alstadsæter, Johannesen and Zucman, 2018), following Andersen et al. (2022), exhibit significantly larger crypto responses to aid disbursements: a one-standard-deviation increase in offshore exposure amplifies the aid-induced response by 0.223 IHS units for transaction volume, 0.232 for frequency, and 0.479 for new accounts on tax-haven exchanges, with somewhat smaller but significant magnitudes on non-tax-haven platforms. This complementarity—whereby the crypto response is strongest precisely where offshore banking is most accessible—is consistent with two distinct interpretations. Under the first, the same actors exploit both channels in parallel, drawing on offshore-banking expertise and intermediary networks to lower the cost of engaging with crypto markets. Under the second, both channels are driven by a common underlying factor—the quality of elite financial networks, concentrated access to aid flows, or weak domestic institutions—that facilitates diversion through whichever vehicle is available, without any individual necessarily using both channels simultaneously. Our data cannot definitively distinguish between these mechanisms, which would require individual-level linking of on-chain wallets to offshore account holders.[5]

The distinction carries practical consequences. If the same actors operate across both channels simultaneously, single-channel enforcement leaves the parallel infrastructure intact, requiring coordinated, multi-channel oversight. If instead the two channels serve similar but distinct populations connected by common institutional weaknesses, the reallocation elasticity is less clear, and institutional reform may prove more effective than channel-specific regulation.

Finally, our leakage estimates capture only the Bitcoin channel and therefore represent a lower bound on total crypto-mediated diversion for a reason distinct from the geographic selection discussed above. Stablecoins—particularly USDT on the Tron blockchain—have emerged during our sample period as the dominant vehicle for cross-border illicit transfers in developing countries, offering dollar denomination, near-zero fees, and settlement times that surpass Bitcoin, while operating on blockchains where forensic attribution remains less developed. This shift accelerated after 2021, meaning that the Bitcoin-based estimates in the later years of our sample may understate the total crypto footprint of aid diversion even in countries with adequate digital infrastructure. Privacy-enhancing cryptocurrencies and decentralized finance protocols represent further unobserved channels. The 2 to 5 percent range should accordingly be interpreted as a lower bound on crypto-mediated leakage, not as an estimate of total diversion through digital assets.

4.6 Off-Chain Transaction Analysis

To complement the on-chain analysis, we examine off-chain exchange activities recorded by centralized platforms across jurisdictions. Off-chain data capture transactional behavior that occurs within exchange platforms (order matching, deposits, and internal transfers) that does not appear on public blockchains. While on-chain data trace the movement of crypto assets across wallets, they cannot distinguish between speculative transfers, custodial holdings, and active trading. Off-chain records thus provide independent confirmation that observed blockchain flows are tied to real economic activity.

[5] See Appendix D.3 for details on the offshore banking complementarity analysis.

Our empirical strategy adopts a lead–lag framework at the quarterly frequency. We construct two dependent variables. The first is the logarithmic level change in exchange volume between adjacent quarters, capturing the absolute post-aid shift. The second is a logarithmic relative ratio comparing volumes on suspect platforms to a benchmark group of regulated exchanges (Coinbase, Kraken, and Gemini): $\log(EX\_Haven_t/Benchmark_t) - \log(EX\_Haven_{t-1}/Benchmark_{t-1})$. This benchmark-adjusted measure filters out common macro trends and isolates jurisdiction-specific behavioral shifts following aid disbursements.

4.6.1 Platform Heterogeneity

Table 4, Panel A reports the effect of lagged foreign aid on off-chain activity across three platform classifications: tax-haven versus non-haven, unregulated versus regulated, and crypto-friendly versus not crypto-friendly.

The results reveal a consistent pattern. On tax-haven exchanges, a 1% increase in foreign aid is associated with a 2.39 percentage point increase in transaction count growth and a 0.39 percentage point increase in the volume ratio, both significant at the 1% level. Non-haven exchanges show an insignificant transaction-count response and a near-zero volume ratio (-0.03), with the cross-group difference (Δ) significant at the 1% level for both measures. The unregulated–regulated and crypto-friendly–unfriendly splits yield nearly identical patterns: the volume-ratio response in unregulated jurisdictions (0.37) is more than ten times the regulated-country effect (0.03), and crypto-friendly jurisdictions show a significant 0.39 response against a null effect in unfriendly jurisdictions. These results indicate that the regulatory environment, rather than cryptocurrency per se, governs whether exchanges serve as conduits for rapid fund movement following aid receipt.

4.6.2 Institutional Quality Heterogeneity

Table 4, Panel B examines whether the aid–crypto relationship varies with recipient-country institutional quality by interacting aid with the World Bank's CPIA governance score, domestic credit depth, control of corruption, and disclosure standards. In each case, the effect of aid on off-chain crypto flows is larger and more precisely estimated in the below-median group. The contrast is sharpest for disclosure: the coefficient in low-disclosure countries (0.12, significant at the 1% level) is twelve times the point estimate in high-disclosure countries (0.01, insignificant). For corruption control, the low-group coefficient

(0.03) is significant while the high-group estimate (0.02) is smaller and only marginally significant. These patterns suggest that weaker institutions, particularly those with limited transparency requirements, are more susceptible to the diversion of aid into crypto channels.

4.6.3 The Huobi Effect

Huobi's relocation from Seychelles to Singapore in May 2021 provides a natural case study in how a major exchange's shift in regulatory domicile affects the geographic transmission of aid-linked flows. Seychelles is classified as a tax-haven jurisdiction in our sample, while Singapore is not, meaning that Huobi's transaction volume migrated from the tax-haven to the non-haven category at a known point in time. This relocation allows us to isolate the exchange-specific contribution to Singapore's off-chain response, separately from any pre-existing jurisdiction-level trend.

Singapore, a crypto-friendly jurisdiction with moderate regulatory oversight, exhibits a particularly strong off-chain response to foreign aid after May 2021, the month Huobi shifted its operational headquarters from Seychelles to Singapore. Table 5 isolates this effect by comparing transaction activity across five offshore jurisdictions (Hong Kong, Singapore, Seychelles, Islands, and Malta), with a final column reporting the Huobi Effect: the difference between Singapore aggregates that include Huobi and those that exclude it.

The Huobi Effect coefficients are 0.04 for transaction counts and 0.02 for the volume ratio (the latter significant at the 5% level), modest relative to the cross-jurisdiction effects but statistically distinguishable from zero. Huobi's presence contributes measurably, though not predominantly, to Singapore's post-aid activity. Rather than driving the entire Singapore effect, the exchange appears to amplify existing flows, consistent with its role as a semi-offshore platform that absorbed activity that might otherwise have been routed to more opaque jurisdictions. This case study illustrates how the presence and relocation decisions of major exchange entities can shape the geographic transmission of aid-linked financial flows.

The off-chain results corroborate the on-chain findings and address an important identification concern: whether observed on-chain spikes reflect speculative wallet rebalancing rather than actual exchange usage. The fact that exchange-reported transaction counts increase precisely in the jurisdictions and periods where on-chain activity rises indicates that blockchain surges are accompanied by real deposit and trading behavior. Together with the institutional-quality heterogeneity and the Huobi case study, these results reinforce the interpretation that aid disbursements generate measurable behavioral responses at both the platform and jurisdiction level, with tax-haven and semi-offshore venues serving as the primary conduits.

5 The Laundering Pattern

The on-chain and off-chain evidence is consistent with a layered laundering pattern operating through cryptocurrency markets in the wake of foreign aid disbursements. We interpret the observed dynamics within the classical framework of placement, layering, and integration, while emphasizing the distinct informational roles of on-chain and off-chain data. On-chain data capture the movement and routing of assets across wallets and exchanges, whereas off-chain data identify where economically meaningful trading or internal reallocation occurs within exchange platforms. This distinction is important: it separates transient transit flows from the settlement points where funds are actively redistributed. The three data sources in our analysis map naturally onto the three stages of the laundering cycle, providing converging evidence for each.

Placement (Entry into Crypto). The placement stage involves the initial conversion of fiat-denominated aid funds into cryptocurrency. Large-scale fiat-to-crypto conversion requires access to regulated banking infrastructure and liquid on-ramps, which are most readily available on compliant, non-tax-haven exchanges. Consistent with this mechanism, the on-chain analysis in Section 4.4 reveals statistically significant increases in both transaction volume and frequency on non-tax-haven exchanges during the disbursement month, with the largest absolute magnitudes observed at this early stage. The web-traffic evidence reinforces this interpretation: visits and time spent on exchange platforms surge sharply in the disbursement month and begin to attenuate shortly thereafter, reflecting a rapid wave of platform engagement consistent with fiat conversion rather than slow-moving changes in digital infrastructure. These contemporaneous inflows reflect initial entry into the crypto ecosystem via platforms that facilitate fiat on-ramping, rather than through opaque venues where the barriers to large-scale conversion are higher.

Layering (Obfuscation via Transfers). Once funds enter the crypto system, subsequent transfers are used to reduce traceability and obscure provenance. The layering stage typically involves withdrawals from regulated exchanges to self-hosted wallets, fragmentation across multiple addresses, and routing through offshore or lightly regulated platforms. Our evidence points to tax-haven exchanges as central nodes in this process. The on-chain analysis documents sharp increases in transaction volume, transaction frequency, and new anonymous wallet creation on tax-haven exchanges immediately following disbursements, indicating intensive address-level reshuffling. The near-complete concentration of new wallet creation in anonymous accounts, with the identified-wallet response an order of magnitude smaller, is consistent with intentional fragmentation across fresh identifiers to obscure fund provenance.

The off-chain evidence sharpens this interpretation. While on-chain activity rises broadly across exchange types, reflecting funds moving through the network, off-chain activity responds selectively on lightly regulated, tax-haven platforms. Because off-chain volume is recorded only when assets are actively used within an exchange's private ledger, this pattern indicates that tax-haven venues are not merely transit points but sites where internal reallocation and redistribution take place. The response is also short-lived, consistent with rapid completion of the layering cycle rather than prolonged accumulation.

Integration (Realization or Storage). The final stage involves either converting crypto holdings back into fiat currency for re-entry into the formal financial system, or retaining them as digital assets for long-term storage. Although the ultimate disposition of funds is not directly observable in our data, we document significant on-chain outflows from tax-haven exchanges following disbursements, indicating that funds leave these platforms after the layering stage rather than accumulating indefinitely. The off-chain evidence confirms that these outflows are preceded by active within-platform trading and internal transfers on tax-haven exchanges, consistent with funds being consolidated before final disposition rather than merely transiting through the platform. The concentration of both on-chain outflows and off-chain activity in the same narrow post-disbursement window as the placement and layering responses suggests that integration is executed as part of a single, coordinated sequence rather than deferred indefinitely. The rapid decline in transaction volume within one to two months of disbursement reinforces this interpretation: funds are withdrawn from active circulation rather than retained for ongoing use, consistent with conversion to fiat or transfer to long-term storage outside the exchange ecosystem. After this window closes, the network topology reverts to its pre-disbursement structure.

6 Robustness Checks and Extensions

This section presents a series of robustness checks organized into three tiers. We begin with two tests of identification validity: the Pandora Papers natural experiment, which tests whether an exogenous transparency shock differentially affects anonymous versus identified channels, and a capital flight placebo, which uses disaster shocks without disbursements to rule out demand-driven explanations. We then present two checks of measurement validity: external validation of the geographic attribution proxy against the Chainalysis Global Crypto Adoption Index, and re-estimation using an alternative visit-duration-based geographic identifier. Finally, we report a specification check excluding the May–September 2021 DeFi boom window. Additional robustness results—including PPML estimation, alternative control group compositions, on-chain outcomes without web-traffic adjustment, governance heterogeneity, and offshore banking complementarity—are reported in the appendix. Results across all checks are consistent with the baseline and reinforce the interpretation that aid disbursements trigger deliberate, anonymity-seeking cryptocurrency activity.

6.1 Pandora Papers

The Pandora Papers, published in October 2021 by the International Consortium of Investigative Journalists, exposed the offshore financial holdings of hundreds of politicians, public officials, and their associates across more than 90 countries. The leak represents a large, plausibly exogenous shock to the perceived risks of using offshore financial structures for illicit purposes: it dramatically increased public and regulatory scrutiny of cross-border fund flows and demonstrated that records previously thought to be confidential could be exposed at any time.

To test whether the baseline aid-crypto relationship changed after this transparency shock, we estimate the same DiD specification used in the baseline analysis, $y \sim \text{DisbPost} \mid \alpha_i + \gamma_t$, separately on the pre-Pandora subsample (months before October 2021) and the post-Pandora subsample (months from October 2021 onward), and report the difference $\Delta = \hat{\beta}_{\text{post}} - \hat{\beta}_{\text{pre}}$. This split-sample approach avoids the collinearity that arises when estimating a triple interaction in the no-overlap panel and provides a transparent test of whether the aid-crypto relationship changed after the leak.

Table 7 reports the results. The most striking finding is that the post-Pandora decline is concentrated in identified rather than anonymous transactions, a pattern that reverses the conventional deterrence prediction. For transaction volume on tax-haven exchanges, the pre-Pandora coefficient for identified transactions is positive and significant, but the post-Pandora coefficient drops sharply, yielding a negative and significant Δ. Transaction frequency tells the same story: identified activity falls markedly after the leak. The pattern is consistent across both tax-haven and non-tax-haven platforms and across all three outcome families (volume, frequency, and new accounts).

By contrast, anonymous transaction activity is largely unaffected by the Pandora Papers disclosure. The pre- and post-Pandora coefficients for anonymous volume and frequency remain similar in magnitude, and the corresponding Δ estimates are small and statistically insignificant. Anonymous new-account creation follows the same pattern, showing no meaningful change across the two subperiods. Figure 5 illustrates this asymmetry directly: while identified transaction volume and new-account creation exhibit a pronounced post-disbursement response that attenuates after the Pandora leak, the dynamic patterns for anonymous wallets remain broadly stable across the two subperiods.

This asymmetry has a clear interpretation. The Pandora Papers exposed the identities behind formally structured offshore vehicles, including shell companies, trusts, and named bank accounts, precisely the infrastructure that underpins identified crypto transactions conducted through KYC-compliant channels. The leak raised the expected cost of routing funds through traceable pathways, inducing a selective withdrawal from identified activity. Anonymous transactions, by contrast, operate outside this identity-linked infrastructure: pseudonymous wallets and privacy-enhancing techniques are not vulnerable to document-based exposure. The persistence of anonymous activity in the face of heightened scrutiny suggests that actors using these channels had already optimized for untraceability and were therefore insulated from a shock that targeted formal record-keeping structures.

This finding also provides indirect support for the causal interpretation of the baseline results. The timing of the structural break coincides with an independently documented shock to offshore financial transparency rather than any change in aid policy or disbursement procedures. The fact that only the identified channel responds, while anonymous activity continues unabated, reinforces the view that the two transaction types serve fundamentally different purposes in the post-disbursement crypto ecosystem.

6.2 Capital Flight Placebo

A natural concern is that the observed cryptocurrency activity following aid disbursements reflects capital flight by wealthy elites rather than diversion of aid funds. Under this alternative, individuals with portable wealth convert local assets into cryptocurrency to hedge against macroeconomic instability or political risk, behavior that may coincide with disaster shocks but operates independently of aid inflows. We address this concern through a placebo exercise using disaster shocks without disbursements.

We define placebo episodes as country-years in which a natural disaster was recorded in EM-DAT (2018–2024) but no World Bank disbursement followed. If the baseline patterns reflect disaster-induced macroeconomic disruption rather than aid disbursements, similar increases should appear in this placebo sample. Table 8 reports the stacked TWFE estimates.

The results are unambiguous. Far from replicating the positive effects in the baseline, the Treat$\times$Post coefficients are negative and highly significant across nearly all outcomes and exchange types. Anonymous transaction volume declines by 0.985 IHS units on tax-haven exchanges and by 1.065 on non-tax-haven exchanges; anonymous transaction frequency falls by 1.007 and 1.052, respectively. New anonymous account creation exhibits the largest negative effects, with coefficients of -1.750 on tax-haven and -1.904 on non-tax-haven exchanges. The negative effects extend to identified transactions as well, with one narrow exception: identified new-account creation on non-tax-haven exchanges shows a small positive coefficient (0.026), though the magnitude is economically negligible relative to the large negative effects on anonymous accounts.

A natural follow-up question is whether these negative coefficients reflect disaster-induced destruction of digital infrastructure rather than the absence of a liquidity injection. The two alternatives generate opposing cross-sectional predictions. Under infrastructure destruction, the suppression should be largest in low-governance countries, where digital networks are most fragile.[6] Under the liquidity-injection interpretation, the suppression should be broadly similar across governance levels, since the absence of aid affects all countries regardless of infrastructure quality.

The governance-split analysis in Table 9 discriminates between these alternatives. The baseline positive aid-to-crypto effect is concentrated entirely in higher-governance countries—precisely those with the

[6] Cross-country correlations between the WGI Control of Corruption score and standard measures of digital infrastructure (e.g., ITU fixed broadband subscriptions per capita, World Bank WDI internet penetration rates) exceed 0.6 in our sample period.

most resilient digital infrastructure—while below-median governance countries show near-zero effects. If infrastructure destruction were driving the negative placebo coefficients, the suppression should instead be concentrated in low-governance countries with fragile networks. Two additional observations reinforce this conclusion: most disaster episodes in EM-DAT involve regional rather than nationwide disruption, with national-level connectivity typically remaining operational outside the affected area (Aker and Mbiti, 2010; Blumenstock, Eagle and Fafchamps, 2016); and the baseline results themselves require functioning infrastructure to generate detectable on-chain signals—if disasters routinely destroyed infrastructure at the national level, the baseline estimates in higher-governance countries would also be attenuated, which they are not.

The concentration of baseline effects in higher-governance countries reflects detection capacity rather than diversion propensity. Crypto-based laundering requires reliable internet connectivity, exchange access, and financial literacy—infrastructure that is more prevalent where governance is stronger. Higher governance does not imply the absence of elite corruption; it implies the presence of digital ecosystems through which corruption can be executed via crypto and detected in our data. In weaker-institution countries, aid diversion may be equally or more prevalent, but is more likely to operate through cash-based channels, informal value transfer systems, or peer-to-peer transactions invisible to blockchain forensics. Our baseline estimates therefore represent a lower bound on total aid diversion.

This pattern directly contradicts the capital flight interpretation. If wealthy elites were converting assets into cryptocurrency in response to disaster-induced instability, we would expect positive coefficients in precisely this placebo sample. Instead, disaster shocks unaccompanied by disbursements suppress rather than stimulate crypto activity, and the governance split confirms that this suppression cannot be attributed to infrastructure loss. Together with the Pandora Papers result and the concentration of effects in anonymous channels on lightly regulated exchanges, this constellation of findings points toward deliberate routing of diverted aid funds through the crypto ecosystem rather than any passive market response.

6.3 Validation of Geo-Adjusted Measures

6.3.1 External Validation

To validate the geographic attribution proxy, we correlate our country-year measures with the independently constructed Chainalysis Global Crypto Adoption Index (2020–2024), which combines proprietary on-chain data, exchange-reported volumes, and web-traffic patterns from sources entirely distinct from our SimilarWeb-based methodology.[7] We collapse our proxy to country-year means and merge with the publicly released Chainalysis rankings, yielding 81 matched observations across approximately 30 countries. Spearman rank correlations between our IHS-transformed proxy and the Chainalysis score range from 0.371 to 0.439 in the pooled specification, and reach -0.543 against the Chainalysis rank on non-tax-haven exchanges. The correspondence is stronger on non-tax-haven platforms, consistent with Chainalysis measuring mainstream adoption, and remains significant on tax-haven exchanges ($\rho \approx 0.39$–0.45, $p < 0.001$). These correlations confirm that the web-traffic-based geographic attribution recovers meaningful cross-country variation in cryptocurrency activity independently of our identification strategy.

6.3.2 Alternative Geographic Identifier

Our baseline geo-adjusted outcomes allocate global on-chain activity to countries using each country's share of total visits to cryptocurrency exchanges. A potential concern is that raw visit counts may overweight brief or incidental page views, attributing on-chain activity to countries whose users engage only superficially with exchange platforms. Table 10 re-estimates the baseline specifications replacing the visit-share ratio with a visit-duration ratio, which weights platform traffic by the relative time users spend on each exchange and is arguably more informative about actual trading intent.

The results are qualitatively robust. Point estimates remain positive and strongly significant across all three outcome categories and both exchange types. For transaction volume, the coefficient is 0.487 for anonymous transactions on tax-haven exchanges and 0.509 on non-tax-haven exchanges; for transaction frequency, 0.480 and 0.502, respectively. New anonymous account creation shows the largest effects at

[7] Details on the Chainalysis index methodology, the rank-to-score conversion, and year-by-year results are provided in Appendix C.

0.575 on tax-haven and 0.601 on non-tax-haven exchanges. Identified transactions are positive and significant for volume and frequency but notably smaller for new accounts, preserving the sharp anonymity gap observed in the baseline. The magnitudes are somewhat attenuated relative to the visit-share specification, as expected: the duration ratio assigns less weight to countries with high visit counts but short session times. The qualitative pattern is fully preserved, confirming that the baseline findings are not an artifact of the specific web-traffic proxy used for geographic attribution.

6.4 DeFi Boom Exclusion

The period from May to September 2021 witnessed an unprecedented surge in decentralized finance activity that reshaped global cryptocurrency markets and could, in principle, confound our estimates if the associated increase in on-chain volume coincided with aid disbursement events. To ensure that the baseline results are not driven by this transient market-wide shock, we re-estimate all specifications after dropping the May–September 2021 window from both the treatment and control samples. Table A.7 in Appendix D.6 reports the results. The point estimates remain positive, strongly significant, and quantitatively similar to the baseline across all outcome families, exchange types, and anonymity categories—for example, anonymous transaction volume on tax-haven exchanges is 0.730 compared with 0.770 in the baseline—confirming that the 2021 DeFi boom does not drive our findings.

7 Conclusion

This study provides systematic evidence that foreign-aid disbursements trigger short-lived but substantial laundering activity in cryptocurrency markets. Aid flows to developing countries produce sharp post-disbursement surges in anonymous Bitcoin transactions consistent with the classic placement-layering-integration sequence: funds enter through exchanges of all types, are fragmented across newly created anonymous wallets, and are routed through tax-haven platforms where regulatory oversight is weakest. The entire extensive-margin response operates through anonymous accounts, new anonymous wallet creation increases by 1.5 to 2.7 IHS units across exchange types in the disbursement month, and turnover-adjusted leakage amounts to approximately 2 to 5 percent of disbursed aid. The Pandora Papers shock selectively suppressed identified transactions while leaving anonymous channels unaffected; placebo tests using disasters without disbursements yield uniformly negative responses; and significant effects emerge only in countries with adequate digital infrastructure, implying that our estimates represent a lower bound on total diversion.

Our findings carry three direct policy implications. First, the offshore heterogeneity analysis reveals that the crypto response to aid disbursements is stronger in countries where traditional offshore banking is more accessible, indicating that the two channels serve overlapping populations of elites rather than wholly distinct ones. This complementarity means that the same actors already operate across both channels simultaneously, so enforcement targeting a single channel leaves the parallel infrastructure, and the networks behind it, intact. Effective containment requires coordinated, multi-channel oversight: tightening know-your-customer standards on crypto exchanges in tax-haven jurisdictions while simultaneously maintaining pressure on traditional offshore banking infrastructure.

Second, the transparency of public blockchains creates an enforcement opportunity that has no analogue in traditional offshore finance. Unlike offshore bank accounts, which are protected by layers of corporate secrecy and jurisdictional fragmentation, Bitcoin flows are recorded immutably on a public ledger and can be monitored in near-real time using the same forensic techniques we employ. Aid-disbursing institutions could, in principle, implement post-disbursement surveillance windows, flagging unusual surges in anonymous crypto activity in recipient countries and conditioning subsequent tranches on the absence of suspicious patterns. International financial intelligence units with access to blockchain analytics could deploy these tools to trace placement flows back to their origin and, where exchange records are obtained through legal channels, identify the accounts involved. Such a system would be more tractable to build than the equivalent for offshore banking, precisely because the underlying ledger is public.

Third, our leakage estimates, while economically significant, almost certainly understate total diversion for two distinct reasons. The first is geographic: significant effects emerge only in higher-governance countries with sufficient digital infrastructure, implying that countries with weaker institutions and lower internet penetration may divert aid through crypto without generating detectable on-chain signals. The second is technological: our analysis is confined to Bitcoin, yet the landscape of illicit cryptocurrency flows has shifted markedly during our sample period. Stablecoins—particularly USDT on the Tron blockchain—have emerged as the dominant vehicle for cross-border illicit transfers in developing countries, offering dollar denomination, near-zero fees, and faster settlement than Bitcoin while operating on blockchains where forensic tooling remains less mature. This shift accelerated after 2021, meaning that our estimates may capture the historically important channel while understating the channel that has grown most rapidly in the later years of our sample. Privacy-enhancing cryptocurrencies such as Monero and Zcash, as well as decentralized finance protocols that facilitate cross-chain transfers, represent additional conduits beyond our current data. Extending the forensic toolkit to cover stablecoin flows on Tron and Ethereum, cross-chain bridges, wrapped tokens, and atomic swap protocols is a pressing priority for both researchers and regulators.

Several important questions remain open. Our identification strategy exploits variation in the timing of disbursements but cannot directly observe the identities of diverting actors or link specific on-chain wallets to named officials. Future work linking blockchain forensics to leaked administrative datasets, audit reports, or legal proceedings could strengthen the causal chain from disbursement to individual culpability. It would also be valuable to exploit the reallocation margin more directly: our offshore complementarity results suggest that the same networks operate across both channels simultaneously, but the elasticity of reallocation between them when one becomes costlier remains unknown. More broadly, the approach developed here, combining event-study identification with multi-source blockchain forensics, opens new avenues for studying illicit financial flows, tax evasion, and regulatory arbitrage across a rapidly digitizing global financial system.

References

Aker, Jenny C. and Isaac M. Mbiti, "Mobile Phones and Economic Development in Africa," *Journal of Economic Perspectives*, 2010, *24* (3), 207–232.

Alstadsæter, Annette, Niels Johannesen, and Gabriel Zucman, "Who Owns the Wealth in Tax Havens? Macro Evidence and Implications for Global Inequality," *Journal of Public Economics*, 2018, *162*, 89–100.

Amiram, D., B. N. Jørgensen, and D. Rabetti, "Coins for Bombs: The Predictability of On-chain Transfers for Terrorist Attacks," *Journal of Accounting Research*, 2022, *60* (2), 427–466.

Andersen, Jørgen Juel, Niels Johannesen, and Bob Rijkers, "Elite Capture of Foreign Aid: Evidence from Offshore Bank Accounts," *Journal of Political Economy*, 2022, *130* (2), 388–425.

Bellemare, Marc F. and Casey J. Wichman, "Elasticities and the Inverse Hyperbolic Sine Transformation," *Oxford Bulletin of Economics and Statistics*, 2020, *82* (1), 50–61.

Blumenstock, Joshua E., Nathan Eagle, and Marcel Fafchamps, "Airtime Transfers and Mobile Communications: Evidence in the Aftermath of Natural Disasters," *Journal of Development Economics*, 2016, *120*, 157–181.

Budish, Eric, "Trust at scale: The economic limits of cryptocurrencies and blockchains," *The Quarterly Journal of Economics*, 2025, *140* (1), 1–62.

Callaway, Brantly and Pedro H. C. Sant'Anna, "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Cameron, A. Colin and Douglas L. Miller, "A Practitioner's Guide to Cluster-Robust Inference," *Journal of Human Resources*, 2015, *50* (2), 317–372.

Cavallo, Eduardo, Ana Corbacho, Carlos Fernandez-Valdovinos, Juan Jose Ocampo, and Ranjit Singh, "Crypto as a Marketplace for Capital Flight," Technical Report WP/24/124, International Monetary Fund 2024.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *The Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Chainalysis, "Crypto Crime Report. Available at https://go.chainalysis.com/2021-Crypto-Crime-Report.html," 2021.

Cong, Lin W., Kim Grauer, Daniel Rabetti, and Henry Updegrave, "Blockchain Forensics and Crypto-Related Cybercrimes. Book Chapters. Available at http://dx.doi.org/10.2139/ssrn.4358561," 2023.

Cong, Lin William, Xi Li, Ke Tang, and Yang Yang, "Crypto Wash Trading," *Management Science*, 2023, *69* (11), 6427–6454.

Cong, Will, Campbell Harvey, Daniel Rabetti, and Zong-Yu Wu, "An anatomy of crypto-enabled cybercrimes," *Management Science*, 2025, *71* (4), 3622–3633.

Cumming, Douglas J, Sofia Johan, and Anshum Pant, "Regulation of the crypto-economy: Managing risks, challenges, and regulatory uncertainty," *Journal of Risk and Financial Management*, 2019, *12* (3), 126.

Financial Action Task Force, "Updated Guidance for a Risk-Based Approach to Virtual Assets and Virtual Asset Service Providers," Technical Report, Financial Action Task Force, Paris October 2021.

Foley, Sean, Jonathan R. Karlsen, and Tālis J. Putniņš, "Sex, Drugs, and Bitcoin: How Much Illegal Activity Is Financed through Cryptocurrencies?," *Review of Financial Studies*, 2019, *32* (5), 1798–1853.

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

Gomez, Luciana and Wei Zhang, "Cryptocurrencies as a Vehicle for Capital Exodus: Evidence from the Russia–Ukraine Conflict," *Journal of International Money and Finance*, 2024, *135*, 102827.

International Consortium of Investigative Journalists, "The Panama Papers," 2016. Leak exposing offshore holdings of elites. Available at https://www.icij.org/investigations/panama-papers/.

Kraay, Aart, "How Large Is the Government Spending Multiplier? Evidence from World Bank Lending," *Quarterly Journal of Economics*, 2012, *127* (2), 829–887.

Makarov, Igor and Antoinette Schoar, "Trading and Arbitrage in Cryptocurrency Markets," *Journal of Financial Economics*, 2020, *135* (2), 293–319.

Regner, Tobias, "Capital Controls and Crypto: Evading Capital Flight," *Journal of Economic Perspectives*, 2020, *34* (3), 187–204.

Reinikka, Ritva and Jakob Svensson, "Local Capture: Evidence from a Central Government Transfer Program in Uganda," *Quarterly Journal of Economics*, 2004, *119* (2), 679–705.

Ron, Dorit and Adi Shamir, "Quantitative analysis of the full bitcoin transaction graph," in "Financial Cryptography and Data Security: 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers 17" Springer 2013, pp. 6–24.

Sokolov, Konstantin, "Ransomware Activity and Blockchain Congestion," *Journal of Financial Economics*, 2021, *141* (2), 771–782.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Werker, Eric, Faisal Z. Ahmed, and Charles Cohen, "How Is Foreign Aid Spent? Evidence from a Natural Experiment," *American Economic Journal: Macroeconomics*, 2009, *1* (2), 225–244.

Zucman, Gabriel, "The Missing Wealth of Nations: Are Europe and the U.S. Net Debtors or Net Creditors?," *Quarterly Journal of Economics*, 2013, *128* (3), 1321–1364.

Figures

Figure 1. Network Analysis Around Top Five Aid Disbursements. Larger nodes represent higher centrality, and darker edges indicate greater volume. The panels show the networks one month before, one month after, and six months after the major aid disbursements. Colors distinguish service categories: non-tax-haven exchanges (purple), tax-haven exchanges (orange), payment platforms (blue), mixers (pink), and gambling/others (green).

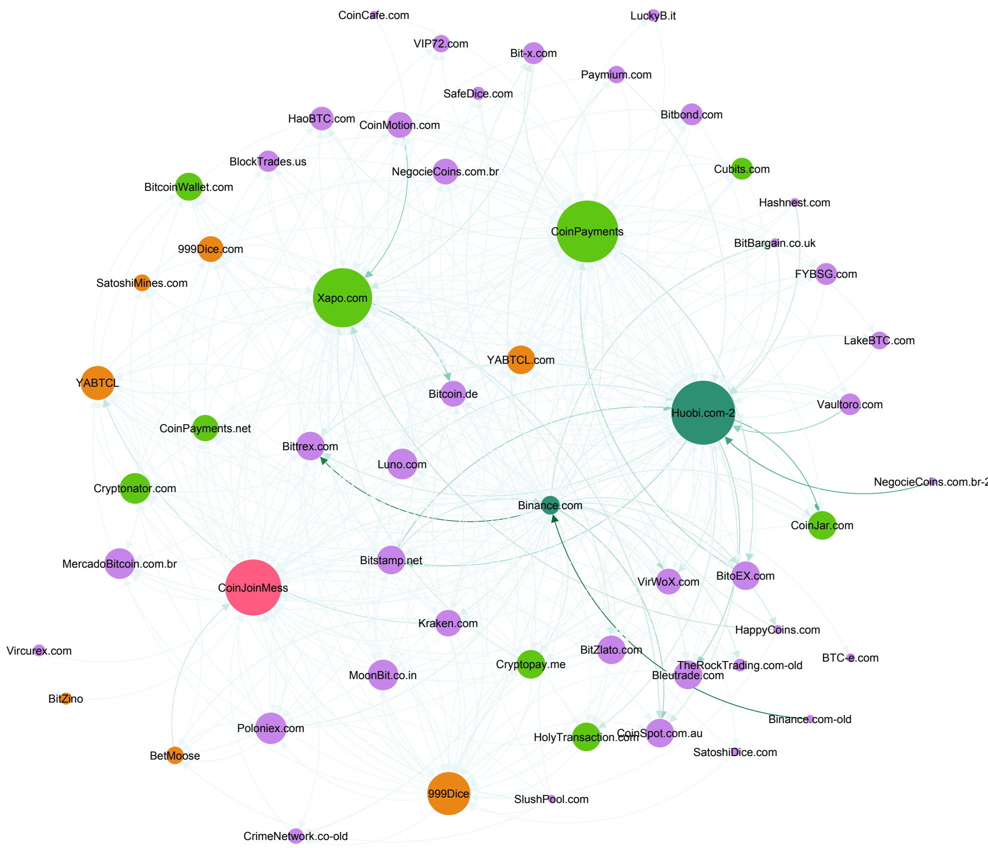

(a) Network One Month Before Aid

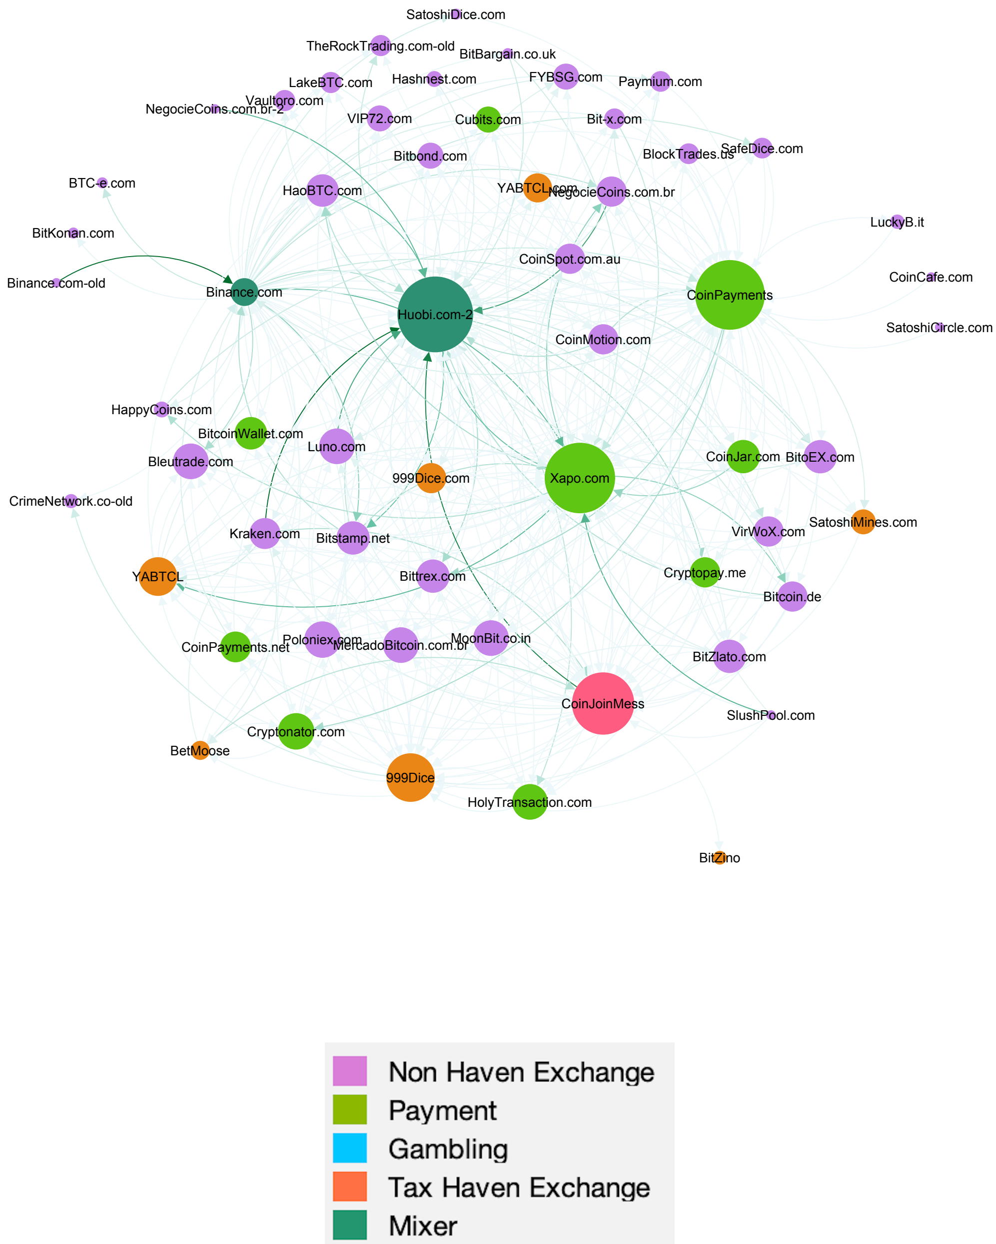
SatoshiDice.com
TheRockTrading.com-old
BitBargain.co.uk
LakeBTC.com
Hashnest.com
FYBSG.com
Paymium.com
Vaultoro.com
NegocieCoins.com.br-2
VIP72.com
Cubits.com
Bit-x.com
BlockTrades.us
SafeDice.com
Bitbond.com
BTC-e.com
HaoBTC.com
YABTCL.com
NegocieCoins.com.br
LuckyB.it
BitKonan.com
CoinSpot.com.au
CoinCafe.com
Binance.com-old
Binance.com
Huobi.com-2
CoinPayments
SatoshiCircle.com
CoinMotion.com
HappyCoins.com
BitcoinWallet.com
Luno.com
CoinJar.com
BitoEX.com
Bleutrade.com
999Dice.com
Xapo.com
CrimeNetwork.co-old
SatoshiMines.com
Kraken.com
Bitstamp.net
VirWoX.com
YABTCL
Bittrex.com
Cryptopay.me
Bitcoin.de
Poloniex.com
MercadoBitcoin.com.br
MoonBit.co.in
CoinPayments.net
BitZlato.com
CoinJoinMess
SlushPool.com
Cryptonator.com
BetMoose
999Dice
HolyTransaction.com
BitZino
Non Haven Exchange
Payment
Gambling
Tax Haven Exchange
Mixer

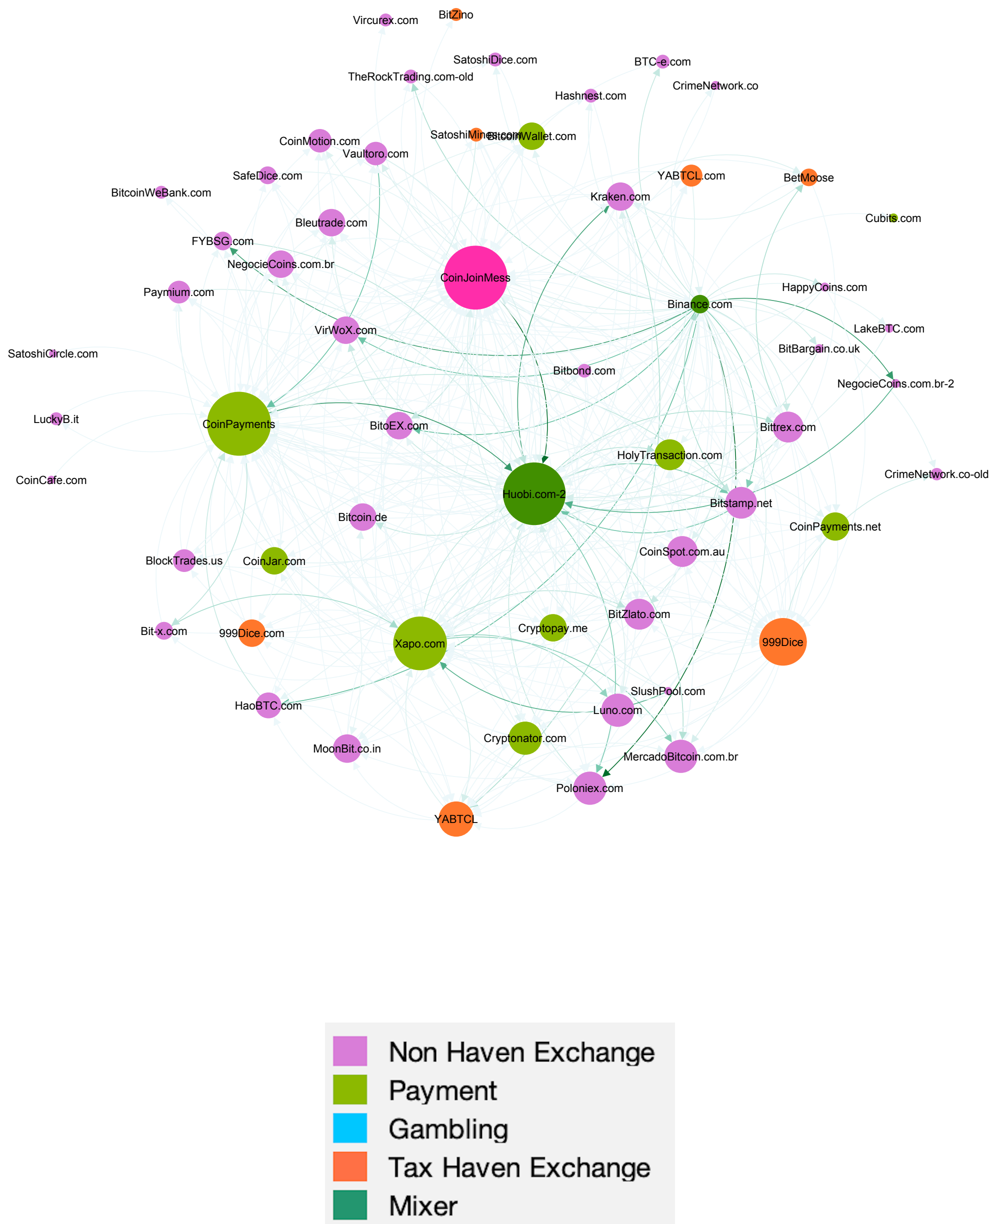
Vircurex.com
BitZino
SatoshiDice.com
BTC-e.com
TheRockTrading.com-old
CrimeNetwork.co
Hashnest.com
SatoshiMines.com
BitcoinWallet.com
CoinMotion.com
Vaultoro.com
SafeDice.com
YABTCL.com
BetMoose
BitcoinWeBank.com
Kraken.com
Bleutrade.com
Cubits.com
FYBSG.com
NegocieCoins.com.br
CoinJoinMess
Paymium.com
HappyCoins.com
Binance.com
VirWoX.com
LakeBTC.com
BitBargain.co.uk
SatoshiCircle.com
Bitbond.com
NegocieCoins.com.br-2
LuckyB.it
CoinPayments
BitoEX.com
Bittrex.com
HolyTransaction.com
CoinCafe.com
CrimeNetwork.co-old
Huobi.com-2
Bitstamp.net
Bitcoin.de
CoinPayments.net
CoinSpot.com.au
BlockTrades.us
CoinJar.com
BitZlato.com
Bit-x.com
999Dice.com
Cryptopay.me
Xapo.com
999Dice
SlushPool.com
HaoBTC.com
Luno.com
Cryptonator.com
MoonBit.co.in
MercadoBitcoin.com.br
Poloniex.com
YABTCL
Non Haven Exchange
Payment
Gambling
Tax Haven Exchange
Mixer

Figure 2. Clustered Network Analysis Around Top Five Aid Disbursements. Larger nodes represent higher centrality, and darker edges indicate greater volume. The figures show changes in network structure around the largest aid disbursements, comparing one week before vs. one week after, and one week before vs. six months after. Colors distinguish service categories: non-tax-haven exchanges (purple), tax-haven exchanges (orange), payment platforms (blue), mixers (pink), and gambling/others (green).

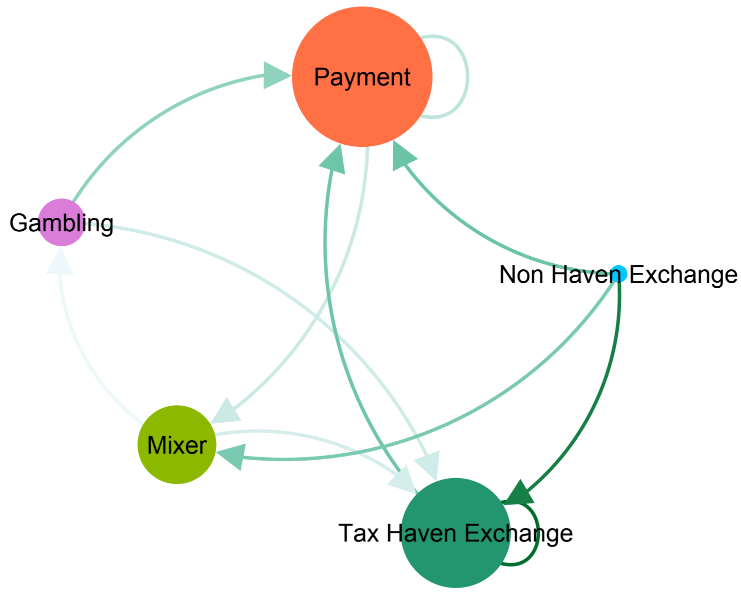

(a) Network Difference One Week Before and After the Top Five Aids

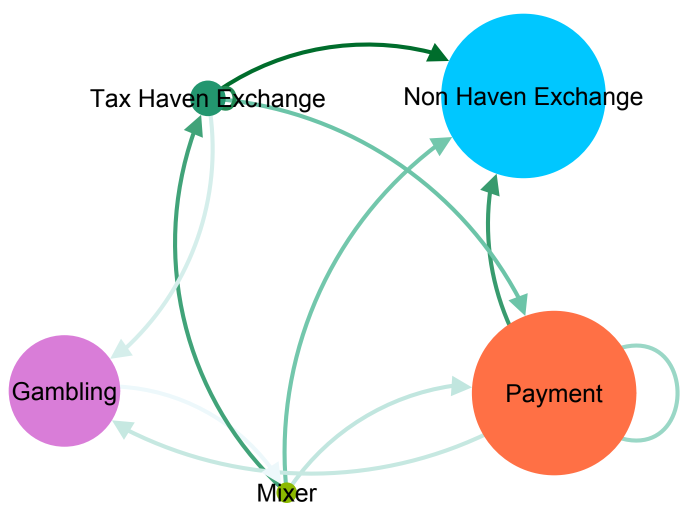

(b) Network Difference One Week Before and Six Months After the Top Five Aids

Figure 3. Stacked DiD: Foreign Aid and Geo-Adjusted Bitcoin Activity by Wallet Type. These figures plot stacked difference-in-differences estimates of the dynamic response of geo-adjusted Bitcoin transaction volume, transaction frequency, and number of newly created accounts around foreign-aid disbursement events, separately for anonymous and identified wallets. Event time k is measured in months relative to the disbursement month ($k = 0$), with $k = -1$ omitted as the reference period. Each outcome is constructed by interacting a country's visits ratio with the corresponding global measure of on-chain activity, yielding an implied country-level measure. Panels (a) and (b) plot transaction volume; Panels (c) and (d) plot the number of newly created accounts. Left column corresponds to anonymous wallets and the right column to identified wallets. The specification includes country fixed effects and month fixed effects, and standard errors are clustered at the country level. Error bars denote 95% confidence intervals.

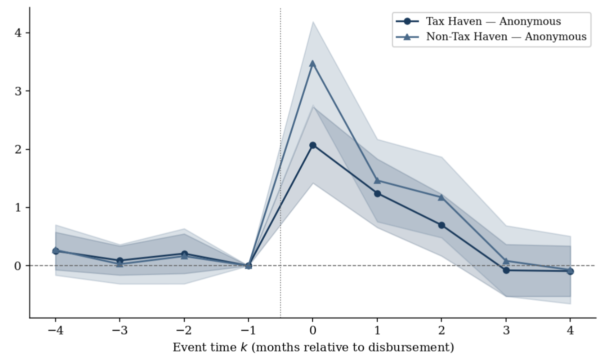

(a) Transaction volume: anonymous wallets

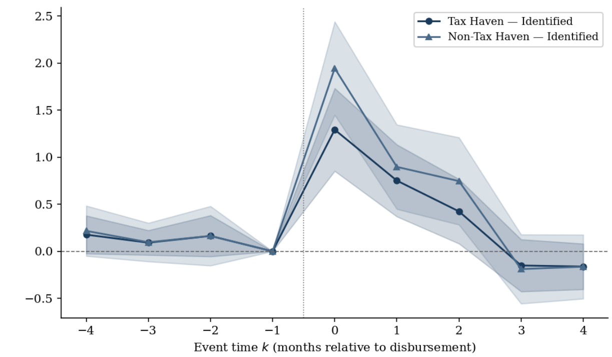

(b) Transaction volume: identified wallets

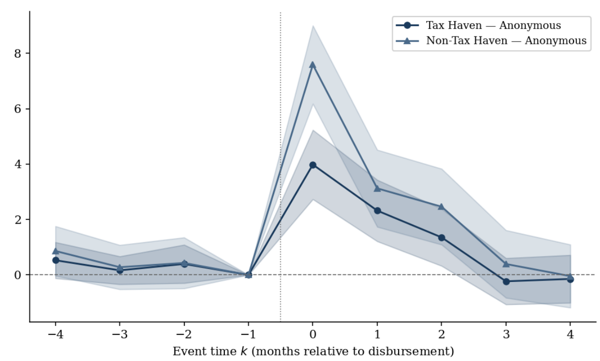

(c) New accounts: anonymous wallets

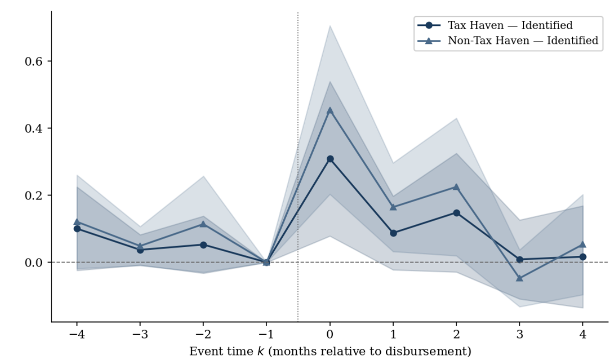

(d) New accounts: identified wallets

Figure 4. Stacked DiD: Foreign Aid and Web Traffic to Crypto-Exchange Websites. These figures plot stacked difference-in-differences estimates of the dynamic response of web traffic to cryptocurrency-related websites around foreign-aid disbursement events. Event time k is measured in months relative to the disbursement month ($k = 0$), with $k = -1$ omitted as the reference period. Panel (a) plots total visits and Panel (b) plots total visit duration (in seconds), both aggregated at the country-month level. The specification includes country fixed effects and month fixed effects, and standard errors are clustered at the country level. Error bars denote 95% confidence intervals.

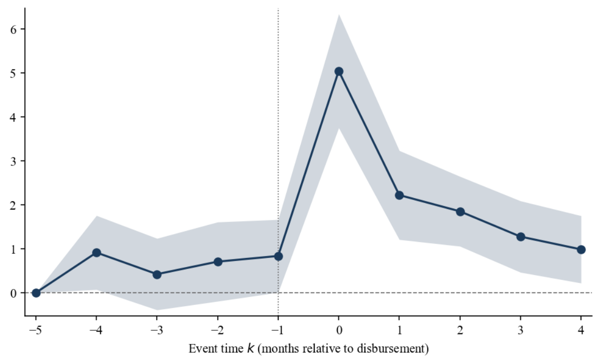

(a) Total visits

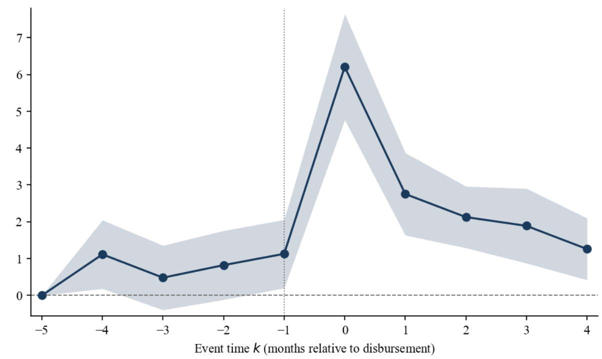

(b) Total time spent (seconds)

Figure 5. Stacked DiD: Foreign Aid and Geo-Adjusted Bitcoin Activity Before and After the Pandora Papers. These figures plot stacked difference-in-differences estimates of the dynamic response of geo-adjusted Bitcoin transaction volume and number of newly created accounts around foreign-aid disbursement events, separately for anonymous and identified wallets, comparing the pre-Pandora (before October 2021) and post-Pandora (October 2021 onward) subperiods. Event time k is measured in months relative to the disbursement month ($k = 0$), with $k = -1$ omitted as the reference period. Each outcome is constructed by interacting a country's visits ratio with the corresponding global measure of on-chain activity, yielding an implied country-level measure. Panels (a) and (b) plot transaction volume; Panels (c) and (d) plot the number of newly created accounts. Left column corresponds to anonymous wallets and right column to identified wallets. The specification includes country fixed effects and month fixed effects, and standard errors are clustered at the country level. Shaded areas denote 95% confidence intervals.

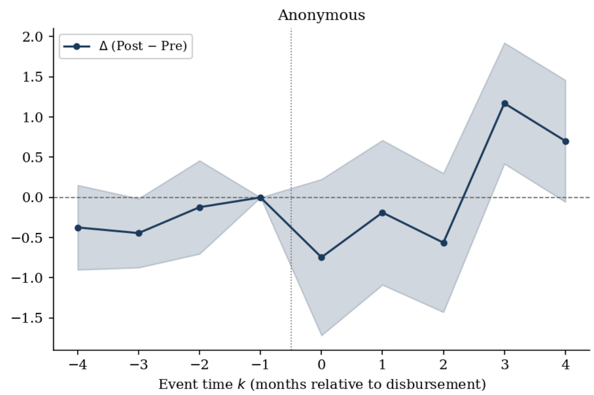

(a) Transaction volume: anonymous wallets

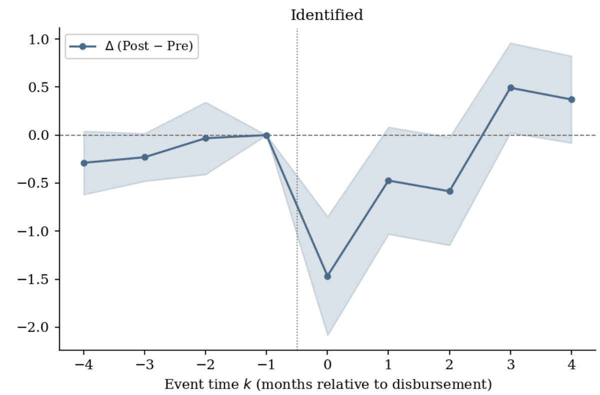

(b) Transaction volume: identified wallets

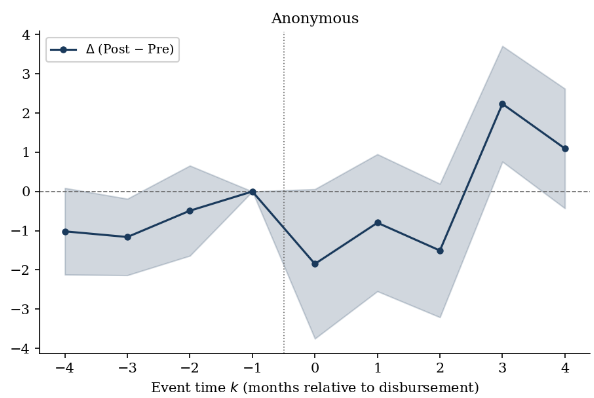

(c) New accounts: anonymous wallets

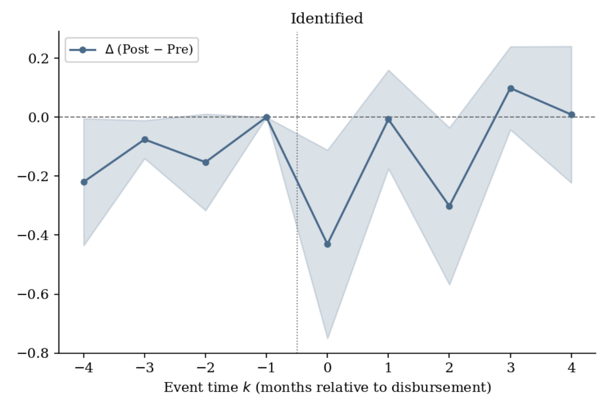

(d) New accounts: identified wallets

Tables

Table 1. Overview of Aid Recipients and Crypto Transaction Platforms. This table summarizes key statistics on aid distribution and crypto exchange activity, organized in four panels. Panel A reports the top 15 recipient countries by total disbursed aid (USD millions) and their share of the global sample. Panel B shows the regulatory-type distribution of crypto exchanges in the sample. Panel C describes platform-level account and volume statistics. Panel D reports the top countries by 24-hour crypto exchange volume (USD thousands).

Panel A. Top 15 Aid Recipient Countries by Disbursed Amount

Country	**Disbursed Amount** *(USD millions)*	**Fraction**
India	8,968	0.06
Indonesia	7,726	0.05
Colombia	7,116	0.05
Ukraine	6,750	0.04
Philippines	6,541	0.04
Nigeria	5,206	0.03
Bangladesh	4,892	0.03
Morocco	4,584	0.03
Pakistan	4,345	0.03
Ethiopia	3,740	0.02
Argentina	3,506	0.02
Turkiye	3,416	0.02
Ecuador	3,373	0.02
Egypt, Arab Republic of	3,326	0.02
Kenya	3,198	0.02
Total (top 15)	**76,694**	**0.50**

Panel B. Regulatory-Type Distribution of Crypto Exchanges

	Tax Haven	**Regulated**	**Crypto Friendly**
Proportion	0.45	0.18	0.70

Table 1 (continued) — Overview of Aid Recipients and Crypto Transaction Platforms.

Panel C. Platform Distribution by New Accounts and 24h Volume

Platform	**# Accounts** *(thousands)*	**% Accounts**	**Volume 24h** *(thousands)*	**% Volume**
CoinJoinMess	31,404.07	0.77	1,579.36	0.02
Binance	2,135.84	0.05	28,254.06	0.34
Kraken	2,050.89	0.05	20,808.49	0.25
CoinPayments	1,855.92	0.05	539.19	0.01
Luno	1,633.95	0.04	1,164.35	0.01
Huobi	1,229.52	0.03	19,029.29	0.23
Bittrex	495.58	0.01	3,482.29	0.04
YABTCL	16.92	0.00	1.75	0.00
999Dice	13.77	0.00	4.21	0.00
Bitstamp	13.17	0.00	5,639.09	0.07
Xapo	13.12	0.00	670.07	0.01
BetMoose	1.14	0.00	1.20	0.00
HitBTC	0.10	0.00	2.99	0.00
Cex	0.02	0.00	7.42	0.00
BitZino	0.01	0.00	0.06	0.00
Total	**40,864.02**	**1.00**	**83,300.68**	**1.00**

Panel D. Top Countries by Crypto Exchange Volume

Country	**Volume 24h** *(USD thousands)*	**Fraction**
Malta	71,954.47	0.33
Singapore	60,677.21	0.27
Seychelles	26,541.83	0.12
South Korea	12,887.99	0.06
Hong Kong	11,673.39	0.05
UAE	10,125.61	0.05
USA	6,096.03	0.03
Australia	5,757.90	0.03
Estonia	3,793.66	0.02
Cayman Islands	2,792.98	0.01
Total (top 10)	**212,200.07**	**0.97**

Table 2. Impact of Foreign Aid on Transaction Volume. This table reports the effect of foreign-aid disbursements on on-chain Bitcoin transaction volume across tax-haven and non-tax-haven exchanges, separately for all, anonymous, and identified transactions. Panels distinguish aggregate activity (Panel A), inflows (Panel B), and outflows (Panel C). The reported coefficient is the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE with a no-overlap control group. All specifications include country-by-platform-by-episode (Ctry.×Plat.×Ep.) and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Crypto Exchanges in Tax Haven			Crypto Exchanges in Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: All Transactions						
Treat×Post	0.777***	0.770***	0.431***	1.235***	1.226***	0.647***
	(0.124)	(0.123)	(0.079)	(0.153)	(0.151)	(0.098)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.01	0.04	0.04	0.03
Panel B: Inflows						
Treat×Post	0.706***	0.695***	0.376***	1.116***	1.101***	0.602***
	(0.113)	(0.111)	(0.073)	(0.139)	(0.137)	(0.092)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.01	0.04	0.04	0.03
Panel C: Outflows						
Treat×Post	0.709***	0.706***	0.334***	1.120***	1.117***	0.465***
	(0.113)	(0.112)	(0.062)	(0.139)	(0.139)	(0.077)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.01	0.04	0.04	0.03

Table 3. Impact of Foreign Aid on New Accounts. This table reports the effect of foreign-aid disbursements on new Bitcoin account creation across tax-haven and non-tax-haven exchanges, separately for all, anonymous, and identified accounts. Panels distinguish aggregate activity (Panel A), inflows (Panel B), and outflows (Panel C). The reported coefficient is the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE with a no-overlap control group. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Crypto Exchanges in Tax Haven			Crypto Exchanges in Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: All Transactions						
Treat×Post	1.457***	1.457***	0.114***	2.706***	2.706***	0.170***
	(0.235)	(0.235)	(0.037)	(0.299)	(0.299)	(0.040)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.00	0.04	0.04	0.01
Panel B: Inflows						
Treat×Post	1.376***	1.376***	0.074**	2.628***	2.629***	0.065**
	(0.220)	(0.220)	(0.030)	(0.290)	(0.290)	(0.031)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.00	0.04	0.04	0.00
Panel C: Outflows						
Treat×Post	1.374***	1.374***	0.037	2.478***	2.478***	0.105***
	(0.225)	(0.225)	(0.024)	(0.276)	(0.276)	(0.025)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.00	0.04	0.04	0.01

Table 4. Impact of Foreign Aid on Off-Chain Transactions: Platform Heterogeneity and Country Characteristics. Panel A reports the effect of lagged foreign aid on the quarterly difference in off-chain volume (columns 1–3) and volume ratios (columns 4–6), each measured relative to benchmark platforms (Coinbase, Kraken, and Gemini). Within each row group the first two columns report separate estimates for the indicated platform classification and its complement; the third column (Δ) reports the difference between the two. Coefficients are multiplied by 100. Panel B splits the sample at the country-level median of each institutional characteristic: CPIA score, domestic credit to GDP, control of corruption (World Bank WGI), and financial disclosure. All specifications include two-quarter lagged foreign aid and one-quarter lagged BTC volume as controls, with country and time fixed effects. Standard errors in parentheses, clustered by aided country. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel A. Off-Chain Transactions and Volume Ratios by Platform Type

	Lead − Lag			**Lead − Lag Ratio**		
	Tax Haven	Non-Haven	**Δ**	Tax Haven	Non-Haven	**Δ**
Foreign Aid$_{\text{lagged}}$	2.39***	1.51	0.88***	0.39***	−0.03	0.42***
	(0.55)	(1.16)	(0.31)	(0.11)	(0.14)	(0.02)
Controls	✓	✓		✓	✓	
Country FE	✓	✓		✓	✓	
Time FE	✓	✓		✓	✓	
Obs.	1,727	1,727		1,727	1,727	
Adj. R^2	0.29	0.21		0.39	0.14	
	Unregulated	Regulated	**Δ**	Unregulated	Regulated	**Δ**
Foreign Aid$_{\text{lagged}}$	2.38***	1.97***	0.41***	0.37**	0.03***	0.34***
	(0.64)	(0.55)	(0.05)	(0.16)	(0.01)	(0.08)
Controls	✓	✓		✓	✓	
Country FE	✓	✓		✓	✓	
Time FE	✓	✓		✓	✓	
Obs.	1,727	1,727		1,727	1,727	
Adj. R^2	0.25	0.19		0.29	0.21	
	Crypto Friendly	Not Friendly	**Δ**	Crypto Friendly	Not Friendly	**Δ**
Foreign Aid$_{\text{lagged}}$	2.41***	1.47**	0.94***	0.39**	−0.00	0.39***
	(0.64)	(0.58)	(0.03)	(0.16)	(0.00)	(0.08)
Controls	✓	✓		✓	✓	
Country FE	✓	✓		✓	✓	
Time FE	✓	✓		✓	✓	
Obs.	1,727	1,727		1,727	1,727	
Adj. R^2	0.25	0.37		0.29	0.50	

Panel B. Heterogeneity by Country Characteristics

	CPIA		**Domestic Credit**		**Control for Corruption**		**Disclosure**	
	High	*Low*	*High*	*Low*	*High*	*Low*	*High*	*Low*
Foreign Aid$_{\text{lagged}}$	0.03	0.03***	0.02	0.03***	0.02**	0.03***	0.01	0.12***
	(0.02)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.04)	(0.04)
Controls	✓	✓	✓	✓	✓	✓	✓	✓
Country FE	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	267	701	736	605	696	691	53	152
Adj. R^2	0.35	0.29	0.30	0.28	0.46	0.29	0.40	0.45

Table 5. Huobi Effect. This table estimates the Huobi Effect by comparing transaction differences between all Singapore-based exchange platforms and those excluding Huobi, from May 2021 to April 2024. The results are benchmarked against tax-haven countries. May 2021 marks Huobi's headquarters relocation from Seychelles to Singapore. The regressions control for the two-quarter lagged foreign aid amount and one-quarter lagged BTC volume, include time and country fixed effects, and cluster standard errors at the aided-country level. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Hong Kong	Singapore	Seychelles	Islands	Malta	Huobi Effect
	(coefficients multiplied by hundreds)					
Panel A: Lead - lag						
Foreign Aid $_{lagged}$	3.35***	1.71***	2.80***	3.87***	1.90***	0.04
	(0.704)	(0.62)	(0.57)	(0.85)	(0.52)	(0.13)
Controls	✓	✓	✓	✓	✓	✓
Country FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
# Obs	1,727	1,727	1,727	1,727	1,727	867
Adj. R^2	0.22	0.39	0.33	0.31	0.38	0.87
Panel B: Lead - lag ratio						
Foreign Aid $_{lagged}$	0.25***	0.01	0.16***	0.04**	0.18**	0.02**
	(0.06)	(0.03)	(0.04)	(0.02)	(0.08)	(0.01)
Controls	✓	✓	✓	✓	✓	✓
Country FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
# Obs	1,727	1,727	1,727	1,727	1,727	867
Adj. R^2	0.47	0.27	0.29	0.56	0.15	0.87

Table 6. Offshore Banking Exposure and Aid-Induced Crypto Responses. This table tests whether the aid–cryptocurrency relationship varies with countries' pre-existing exposure to offshore banking. The specification augments the baseline stacked difference-in-differences with an interaction between Treat×Post and a country-level moderator z, defined as the Alstadsæter–Johannesen–Zucman (AJZ) offshore-to-GDP index for 2007, which measures the ratio of cross-border deposits in offshore financial centres to GDP. The moderator is standardised (mean zero, unit standard deviation) over the estimation sample. A positive interaction coefficient indicates that the aid-induced crypto response is larger in countries with greater offshore banking exposure, consistent with complementarity between the two channels. Outcomes are IHS-transformed geo-adjusted measures of anonymous on-chain activity. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Crypto Exchanges in Tax Haven	Crypto Exchanges in Non-Tax-Haven
Panel A: Transaction Volume		
Treat×Post×z	0.223***	0.129**
	(0.043)	(0.057)
Ctry.×Plat.×Ep. FE	✓	✓
Time FE	✓	✓
Obs.	5,351	5,351
Adj. R^2	0.07	0.08
Panel B: Transaction Frequency		
Treat×Post×z	0.232***	0.138**
	(0.058)	(0.062)
Ctry.×Plat.×Ep. FE	✓	✓
Time FE	✓	✓
Obs.	5,351	5,351
Adj. R^2	0.07	0.08
Panel C: New Accounts		
Treat×Post×z	0.479***	0.333***
	(0.113)	(0.125)
Ctry.×Plat.×Ep. FE	✓	✓
Time FE	✓	✓
Obs.	5,351	5,351
Adj. R^2	0.07	0.08

Table 7. Crypto Laundering after the Pandora Papers. This table tests whether the aid–crypto relationship intensified after the Pandora Papers leak (October 2021). Rather than estimating a triple interaction in a single pooled regression, we run the baseline stacked difference-in-differences specification separately on the pre-Pandora and post-Pandora subsamples and report the difference $\Delta = \hat{\beta}_{\text{post}} - \hat{\beta}_{\text{pre}}$. The Pandora indicator equals one for months $\geq$ 2021-10. The reported coefficient in each subsample is the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE with a no-overlap control group. Three sections report results for transaction volume, transaction frequency, and new accounts; columns distinguish all, anonymous, and identified transactions. Columns 1–3 restrict to exchanges domiciled in tax havens; columns 4–6 to non-tax-haven exchanges. Receiver and sender splits are excluded. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level; the standard error on Δ is $\sqrt{SE^2_{\text{pre}} + SE^2_{\text{post}}}$ (conservative, assumes independence across non-overlapping windows). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Tax Haven			Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: Transaction Volume						
$\hat{\beta}_{\text{post}}$ (Post-Pandora)	1.105***	1.104***	0.354***	0.986***	0.984***	0.269***
	(0.187)	(0.187)	(0.077)	(0.169)	(0.169)	(0.063)
$\hat{\beta}_{\text{pre}}$ (Pre-Pandora)	1.367***	1.348***	1.003***	1.253***	1.222***	0.964***
	(0.273)	(0.269)	(0.198)	(0.250)	(0.245)	(0.188)
Δ (Post − Pre)	-0.262	-0.244	-0.650***	-0.267	-0.238	-0.695***
	(0.331)	(0.328)	(0.212)	(0.302)	(0.297)	(0.199)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs. (pre)	9,432	9,432	9,432	9,432	9,432	9,432
Obs. (post)	8,760	8,760	8,760	8,760	8,760	8,760
Panel B: Transaction Frequency						
$\hat{\beta}_{\text{post}}$ (Post-Pandora)	0.995***	0.993***	0.339***	0.928***	0.926***	0.269***
	(0.171)	(0.170)	(0.070)	(0.160)	(0.160)	(0.058)
$\hat{\beta}_{\text{pre}}$ (Pre-Pandora)	1.373***	1.366***	0.844***	1.233***	1.222***	0.783***
	(0.276)	(0.275)	(0.166)	(0.247)	(0.245)	(0.154)
Δ (Post − Pre)	-0.378	-0.373	-0.505***	-0.305	-0.296	-0.514***
	(0.325)	(0.323)	(0.180)	(0.295)	(0.293)	(0.164)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs. (pre)	9,432	9,432	9,432	9,432	9,432	9,432
Obs. (post)	8,760	8,760	8,760	8,760	8,760	8,760
Panel C: New Accounts						
$\hat{\beta}_{\text{post}}$ (Post-Pandora)	2.160***	2.160***	0.076***	2.116***	2.116***	0.012
	(0.349)	(0.349)	(0.028)	(0.342)	(0.342)	(0.010)
$\hat{\beta}_{\text{pre}}$ (Pre-Pandora)	2.851***	2.851***	0.206**	2.696***	2.696***	0.115*
	(0.542)	(0.542)	(0.084)	(0.513)	(0.513)	(0.065)
Δ (Post − Pre)	-0.690	-0.690	-0.130	-0.580	-0.580	-0.103
	(0.645)	(0.645)	(0.088)	(0.616)	(0.616)	(0.066)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs. (pre)	9,432	9,432	9,432	9,432	9,432	9,432
Obs. (post)	8,760	8,760	8,760	8,760	8,760	8,760

Table 8. Placebo Test of Natural Disasters Without World Bank Disbursement. This table assesses whether crisis exposure alone, rather than aid inflows, drives the baseline crypto-activity results. Placebo treatments are defined as country–years in which a natural disaster was recorded in EM-DAT (2018–2024) but no World Bank disbursement followed. The reported coefficient is the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on all-transaction aggregates (IHS-transformed), estimated by linear TWFE with a no-overlap control group. If the baseline results reflect aid inflows rather than disaster shocks, coefficients here should be near zero. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. *$p < 0.10$, **$p < 0.05$, ***$p < 0.01$.

	Tax Haven			Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: Transaction Volume						
Treat×Post	-0.994***	-0.985***	-0.796***	-1.082***	-1.065***	-0.923***
	(0.059)	(0.058)	(0.047)	(0.067)	(0.066)	(0.057)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	13,399	13,399	13,399	13,399	13,399	13,399
Adj. R^2	0.00	0.00	0.01	0.00	0.00	0.01
Panel B: Transaction Frequency						
Treat×Post	-1.008***	-1.007***	-0.646***	-1.057***	-1.052***	-0.789***
	(0.060)	(0.060)	(0.037)	(0.066)	(0.065)	(0.048)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	13,399	13,399	13,399	13,399	13,399	13,399
Adj. R^2	0.00	0.00	0.01	0.00	0.00	0.01
Panel C: New Accounts						
Treat×Post	-1.750***	-1.750***	-0.894***	-1.904***	-1.904***	0.026***
	(0.105)	(0.105)	(0.061)	(0.121)	(0.121)	(0.006)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	13,399	13,399	13,399	13,399	13,399	13,399
Adj. R^2	0.00	0.00	0.02	0.00	0.00	0.00

Table 9. Governance Heterogeneity: Impact of Foreign Aid on Transaction Volume. This table tests whether the aid–cryptocurrency relationship varies with recipient-country institutional quality by splitting the sample at the median of the World Bank's Worldwide Governance Indicators Control of Corruption score (WGI CC). Panel A reports results for below-median governance countries; Panel B for countries at or above the median. The reported coefficient is the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE with a no-overlap control group. Columns distinguish all, anonymous, and identified transactions across tax-haven and non-tax-haven exchanges. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Tax Haven			Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: Low governance (below median WGI CC)						
Treat×Post	-0.048	-0.051	-0.004	0.225	0.222	0.073
	(0.129)	(0.127)	(0.092)	(0.211)	(0.208)	(0.137)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	11,583	11,583	11,583	11,583	11,583	11,583
Adj. R^2	0.00	0.00	0.00	0.00	0.00	0.00
Panel B: High governance (at/above median WGI CC)						
Treat×Post	0.562^{**}	0.553^{**}	0.338^{**}	0.898^{***}	0.887^{***}	0.508^{***}
	(0.234)	(0.232)	(0.164)	(0.269)	(0.266)	(0.179)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	10,663	10,663	10,663	10,663	10,663	10,663
Adj. R^2	0.00	0.00	0.00	0.00	0.00	0.00

Table 10. Robustness Check: Alternative Geographic Identifier (Visit-Duration Ratio). This table replaces the baseline geographic identifier with a visit-duration ratio to verify that the main results are not driven by the choice of geo-attribution method. The reported coefficient is the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE with a no-overlap control group. Three sections report results for transaction volume, transaction frequency, and new accounts; columns distinguish all, anonymous, and identified transactions by platform type. Outcomes are IHS-transformed after the duration-ratio adjustment. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Tax Haven			Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: Transaction volume						
Treat×Post	0.488***	0.487***	0.394***	0.509***	0.509***	0.436***
	(0.060)	(0.060)	(0.049)	(0.062)	(0.062)	(0.054)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Panel B: Transaction frequency						
Treat×Post	0.481***	0.480***	0.375***	0.503***	0.502***	0.426***
	(0.059)	(0.059)	(0.047)	(0.061)	(0.061)	(0.052)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Panel C: New accounts						
Treat×Post	0.575***	0.575***	0.040	0.601***	0.601***	0.092***
	(0.070)	(0.070)	(0.025)	(0.073)	(0.073)	(0.032)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096

Appendix: Crypto Capture of Foreign Aid

Agarwal, Jin, Prasad, and Rabetti

Appendix

A Variable Descriptions

This appendix summarizes the key variables used in the empirical analysis and provides definitions for the main outcomes, treatments, and classifications employed throughout the paper. Table A.1 reports variable names and descriptions, covering foreign aid disbursements, cryptocurrency transaction outcomes, wallet classifications, exchange characteristics, and country-level controls.

The primary explanatory variable is foreign aid disbursement, measured at the quarterly level and lagged by one quarter to align with the timing of on-chain responses. The main outcome variables capture different dimensions of cryptocurrency activity, including transaction volume, transaction frequency (transaction times), and the number of newly created wallet addresses. These outcomes are analyzed in logarithmic form to reduce sensitivity to extreme values and to facilitate interpretation of dynamic responses in the stacked event-study framework.

Wallets are classified as anonymous or identified based on whether they are linked to a platform or verified entity, allowing us to distinguish opaque from traceable transaction channels. Exchanges are further categorized by tax-haven status, regulatory environment, and jurisdictional attitudes toward cryptocurrency adoption. These classifications support heterogeneity analyses that contrast activity on offshore versus regulated platforms.

To benchmark exchange-specific activity, we construct comparison measures using major regulated exchanges (Coinbase, Kraken, and Gemini). Lead–lag transformations of both level and ratio outcomes are used to study short-run changes around aid disbursement events, capturing dynamic adjustments rather than long-run differences.

Finally, the table lists country-level covariates drawn from external sources, including governance quality (CPIA and Control of Corruption), financial development (Domestic Credit), and disclosure practices. These controls are used in auxiliary specifications and robustness checks to account for cross-country differences in institutional quality and financial infrastructure.

Table A.1. Variable Definitions and Descriptions

Variable	Description
Foreign Aid	Total foreign aid disbursement (lagged by one quarter) used as the main independent variable.
Transaction Volume	Value of cryptocurrency transactions on a given platform or within a region (log-transformed).
Transaction Frequency	Number of cryptocurrency transactions (log-transformed).
Number of New Accounts	Number of newly created wallet addresses (log-transformed).
Anonymous / Identified	Classification of wallets based on ID verification status: anonymous or platform-verified.
Tax Haven / Non-Haven	Exchange classification based on tax jurisdiction status.
Regulated / Unregulated	Classification according to regulation by major financial authorities.
Crypto-Friendly / Not Friendly	Classification based on the jurisdiction's policy attitude toward cryptocurrency.
Benchmark	Transaction volume on benchmark exchanges (Coinbase, Kraken, and Gemini).
Lead-Lag Level Change	One-quarter difference in outcome variable (log-transformed).
Lead-Lag Ratio Change	Difference in log transaction ratios: $\log\left(\frac{EX_t}{BM_t}\right) - \log\left(\frac{EX_{t-1}}{BM_{t-1}}\right)$.
CPIA	Country Policy and Institutional Assessment (CPIA) score from the World Bank; higher values indicate stronger governance.
Domestic Credit	Domestic credit to the private sector as a percentage of GDP; proxy for financial development.
Control of Corruption	World Governance Indicator reflecting the strength of anti-corruption mechanisms; higher values indicate stronger control.
Disclosure	Index measuring public financial disclosure and transparency practices.

We examine the aggregate network dynamics of on-chain Bitcoin transactions around the largest foreign-aid disbursement events. We identify the five dates with the highest total disbursed amounts between April 2018 and March 2024 and track how network-level activity—including transaction volume, the number of active addresses, and the flow of funds across exchanges—evolves in the days surrounding each event. Table A.2 lists these five disbursement dates, the corresponding amounts, and the receiving countries. The events range from concentrated disbursements to one or two countries (e.g., Tajikistan and Ukraine in July 2023) to broad distributions spanning ten or more recipients (e.g., December 2020), allowing us to assess whether the network response varies with the geographic scope of aid flows.

Table A.2. Top Five Disbursements by Date and Amount. We present the aggregate network dynamics for the five dates with the largest foreign aid distributions between April 2018 and March 2024. This analysis captures the significant impact of these top aid distribution events on the overall network.

Date	Disbursed Amount (in millions)	Receiving Countries
2023-07-25	1,507.40	Tajikistan, Ukraine
2023-11-28	1,443.43	Turkiye, India
2022-12-05	1,386.16	Indonesia, Romania
2020-12-17	1,347.77	Kiribati, Cote d'Ivoire, Nigeria, Uganda, Malawi, Papua New Guinea, Dominican Republic, Uzbekistan, Pakistan, Bosnia and Herzegovina
2023-12-18	1,301.26	Senegal, Sierra Leone, Burkina Faso, India, Albania, North Macedonia, Ukraine, Turkiye

B Network Analysis: Nigeria

This appendix provides a visual, case-based illustration of how crypto transaction networks evolve around large foreign-aid disbursements, focusing on Nigeria as a representative high-volume recipient country. Nigeria experienced several large World Bank disbursement episodes during our sample period and features a diverse ecosystem of regulated exchanges, offshore platforms, and intermediary services, making it well suited for network-based analysis.

The figures construct directed transaction networks using wallet- and platform-level on-chain data aggregated over short event windows around the five largest aid disbursements to Nigeria. Nodes represent crypto platforms or services (exchanges, payment platforms, mixers, gambling services, and others), and directed edges represent aggregated transaction flows between nodes within the event window. Node size reflects platform centrality in the transaction network, while edge width is proportional to transfer volume. Platforms are color-coded by service category and jurisdictional status, with particular emphasis on tax haven versus non-tax haven exchanges.

Figure A.1 plots the average network structure one month before, one month after, and six months after the largest aid disbursement events. Relative to the pre-disbursement period, the post-disbursement network exhibits a marked increase in transaction volume and connectivity, particularly toward tax haven exchanges and among exchange-to-exchange links. This pattern is consistent with a temporary intensification of intermediary activity following aid receipt. Six months after disbursement, overall network density and flow intensity decline, with the structure reverting toward pre-disbursement levels.

Figure A.2 presents network-difference graphs that directly contrast transaction flows before and after the largest aid episode. The left panel highlights changes between one week before and one week after the disbursement, while the right panel compares one week before to six months after. These difference plots emphasize that the post-aid surge is concentrated in short-run increases in flows involving exchanges—especially offshore platforms—rather than persistent, long-run reconfiguration of the network.

Figure A.1. Network Analysis: Largest Aid (Nigeria). The figures show the average network one month before, one month after, and six months after the top five disbursement dates. The width of arrows represents transfer magnitudes; node sizes represent platform centrality. Non-tax haven exchanges are purple, tax haven exchanges are orange, payment platforms are blue, mixers are pink, and gambling/others are green. Compared to the pre-disbursement period, post-disbursement networks exhibit greater transfer volume toward tax havens and intensified flows among exchanges, while six months later, the network reverts to lower activity.

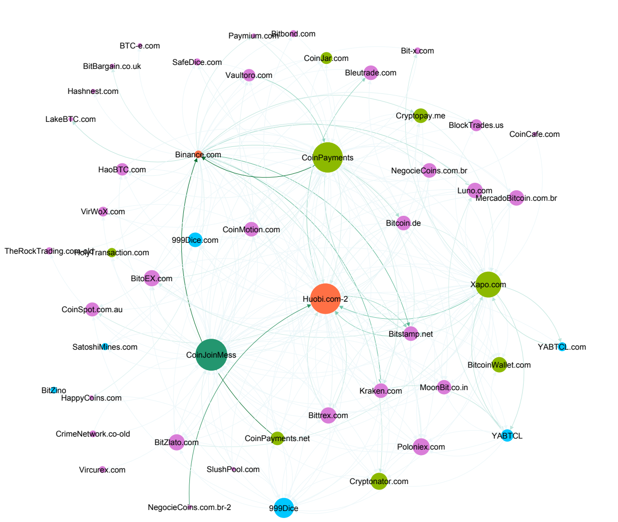

(a) One month before the aid

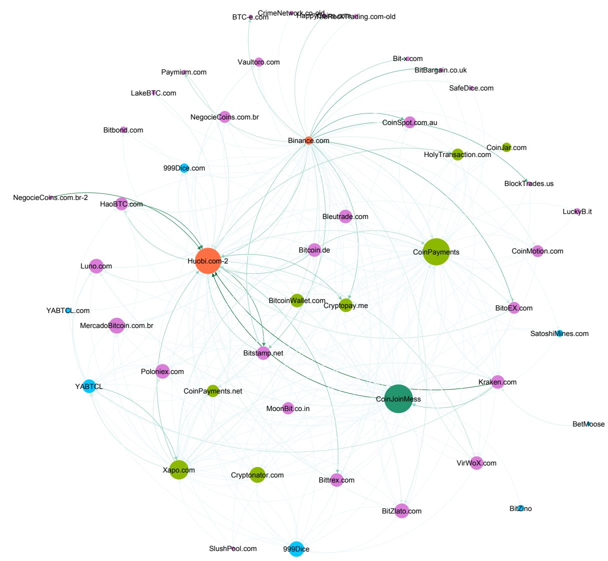

(b) One month after the aid

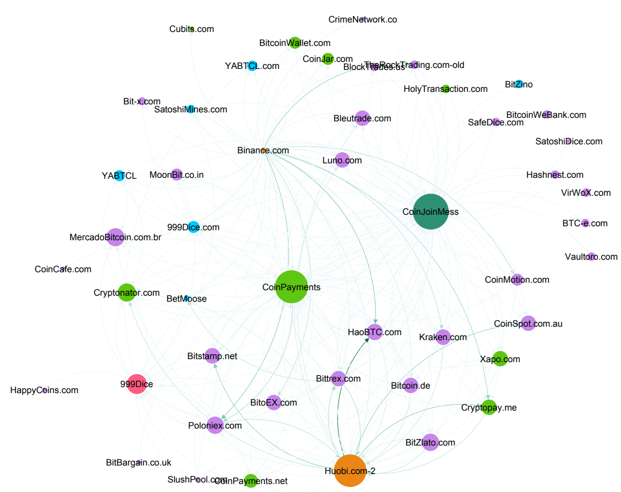

(c) Six months after the aid

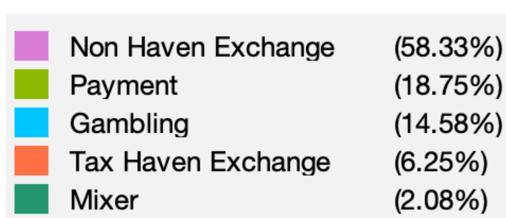

Service Category Color Legend

Figure A.2. Network Differences: Largest Aid (Nigeria). These plots show the differences in network structures before and after major aid disbursements. The left panel compares one week before and one week after the largest aid event, while the right panel compares one week before and six months after. Color coding follows the same convention: non-tax haven exchanges (purple), tax haven exchanges (orange), payment platforms (blue), mixers (pink), gambling and others (green).

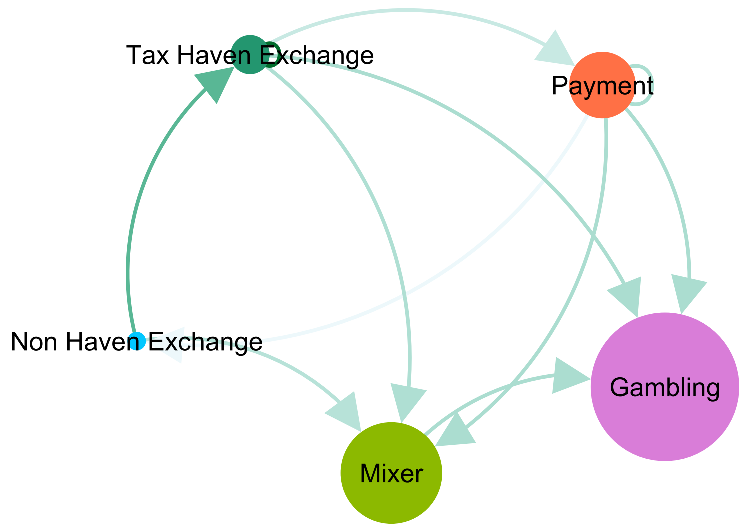

Network Difference One Week Before and After the Largest Aid

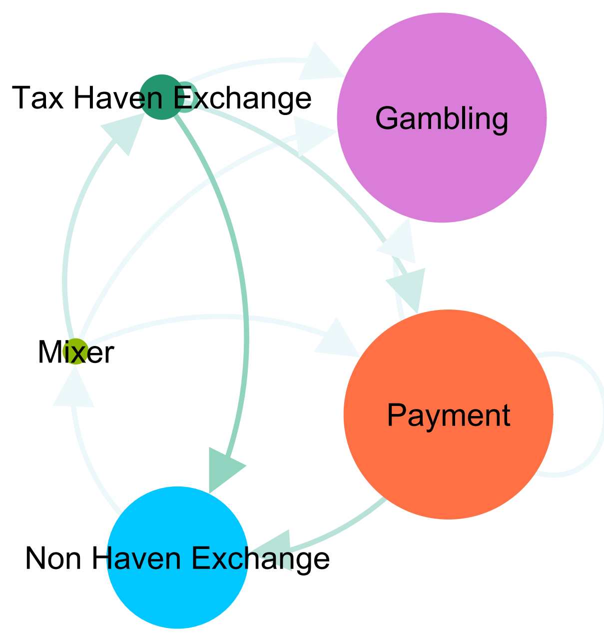

Network Difference One Week Before and Six Months After the Largest Aid

C External Validation of Geo-Adjusted Measures

We validate our geo-adjusted on-chain proxy against the independently constructed Chainalysis Global Crypto Adoption Index (2020–2024), which ranks approximately 150 countries annually on a normalised 0–1 scale using on-chain transaction data, exchange-reported volumes, and web-traffic patterns from sources entirely distinct from our SimilarWeb-based methodology.[1]

We collapse our IHS-transformed geo-adjusted measures to country–year means and merge with the top-20 countries from each annual Chainalysis report, the publicly verified portion of the rankings. Where Chainalysis published normalised scores directly (e.g., India = 1.00, Nigeria = 0.80, Vietnam = 0.60 in the 2023 index), we use those values; remaining entries are imputed via the log-rank transformation $s = 1 - \log(r)/\log(N)$, which better approximates the heavy-tailed shape of the actual score distribution than a linear mapping. The merged sample contains 65 country–year observations across 20 countries and five years.

A key methodological consideration is that the Chainalysis index normalises raw on-chain activity by PPP per capita, explicitly down-weighting high-income countries where large absolute transaction volumes reflect economic size rather than adoption intensity. Our raw proxy does not apply this adjustment, so a naive correlation would be attenuated by scale differences between rich and poor countries: the United States and China generate enormous absolute volumes but rank only moderately on the Chainalysis index. To place the two measures on a conceptually comparable footing, we weight each country–year observation by the inverse of its PPP per capita (World Bank WDI, `NY.GDP.PCAP.PP.CD`), matching the Chainalysis weighting philosophy.

The year-by-year results provide the clearest evidence that our proxy recovers meaningful cross-country variation in cryptocurrency activity. PPP-weighted Pearson correlations rise from near zero in 2020 to 0.689 in 2021, 0.518 in 2022, 0.591 in 2023, and 0.855 in 2024 (all $p < 0.05$ except 2020). The weak 2020 correlation is expected rather than troubling: SimilarWeb's clickstream coverage of developing-country web traffic, the primary source for our geographic attribution, was substantially thinner in the early years

[1] See `https://www.chainalysis.com/blog/2024-global-crypto-adoption-index/` for details. The index covers the period ending approximately June of each report year and ranks between 151 and 155 countries.

of our sample, and the Chainalysis index methodology itself was less refined in its inaugural edition. As both data sources matured and geographic coverage deepened, the correspondence strengthened, reaching 0.855 by 2024. This pattern is precisely what one would expect of a valid proxy whose underlying data quality improves over time, and it directly contradicts the hypothesis that the correlation is a statistical artefact.

The pooled specification, which averages across all five years and therefore blends the noisy 2020 data with the stronger later years, yields a more conservative but still highly significant PPP-weighted Pearson correlation of 0.495 ($p < 0.001$) for transaction volume and 0.513 ($p < 0.001$) on non-tax-haven exchanges. Correlations are systematically stronger on non-tax-haven exchanges than on tax-haven exchanges, consistent with the Chainalysis index capturing mainstream adoption at regulated venues rather than the anonymity-seeking activity that predominates on offshore platforms, precisely the distinction our empirical design exploits. Together, the year-by-year and pooled results confirm that the web-traffic-based geographic attribution recovers meaningful cross-country variation in cryptocurrency activity, validating its use as a geographic allocation device in our main analysis.

D Additional On-chain Analysis

This section reports supplementary results that extend and validate the main findings along six dimensions. First, we compare the web-traffic results across linear TWFE and PPML estimators, confirming robustness to functional form. Second, we show that the on-chain transaction-frequency results extend consistently across all three activity margins (aggregate, inflows, outflows). Third, we split the sample by recipient-country governance quality, revealing that detectable crypto responses are concentrated in higher-governance countries with adequate digital infrastructure. Fourth, we re-estimate all baseline specifications using PPML on level outcomes, confirming that the findings are not driven by the IHS transformation. Fifth, we restrict the identifying variation to administratively determined disbursement timing, addressing residual endogeneity concerns. Sixth, we test whether the aid–crypto relationship is amplified in countries with greater pre-existing offshore banking exposure.

D.1 Web Traffic

Table A.3 reports the post-period average treatment effect of foreign-aid disbursements on crypto-exchange web traffic from tax-haven platforms. Columns 1–3 present TWFE estimates on IHS-transformed outcomes across three control-group definitions; columns 4–6 report the corresponding PPML incidence-rate ratios (IRR), averaged over post periods $k = 0, \ldots, 4$, with $H_0\colon \text{IRR} = 1$. The PPML specifications include event-time dummies only, without unit or calendar-time fixed effects, because the large mass of zero observations in country-platform-episode cells induces separation under high-dimensional fixed effects; the PPML columns therefore serve as a directional check on functional form rather than a specification of equal standing with the fully saturated TWFE model.

Panel A shows that foreign-aid receipt is associated with a large and highly significant increase in total visits to tax-haven crypto exchanges. The TWFE coefficient is approximately 2.28 ($p < 0.01$) and is virtually unchanged across control specifications. The PPML estimates imply a fifteen- to twentyfold increase in visit counts (IRR ≈ 15.5–20.9, $p < 0.01$), consistent with the IHS results in direction and significance, though the larger PPML magnitudes reflect both the multiplicative interpretation of the

IRR and the absence of fixed effects that would otherwise absorb level differences across countries and global trends in web traffic. Panel B documents a similarly strong response in total visit duration, with a TWFE coefficient of approximately 2.85 ($p < 0.01$) and PPML IRRs of 6.6–6.8. The stability of the TWFE estimates across the no-overlap, no-overlap+developing, and no-overlap+no-neighbors samples indicates that the results are not sensitive to the composition of the control group.

Table A.3. Impact of Foreign Aid on Crypto Exchange Web Traffic. This table reports the effect of foreign-aid disbursements on web traffic to cryptocurrency-exchange websites hosted in tax-haven jurisdictions. Columns 1–3 report the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE across three control-group definitions: baseline no-overlap, no-overlap restricted to developing countries, and no-overlap excluding neighboring countries. Columns 4–6 report the corresponding PPML incidence-rate ratios (IRR), computed as $\overline{\text{IRR}} = \exp\left(\frac{1}{L}\sum_{k\geq 0}\hat{\beta}_k\right)$ with delta-method standard errors and $H_0: \text{IRR} = 1$; the PPML specifications include event-time dummies only (no unit or calendar-time fixed effects). Panel A reports total visits; Panel B reports total visit duration (in seconds). All TWFE specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	TWFE			PPML		
	Baseline	+ Developing	+ No neighbors	Baseline	+ Developing	+ No neighbors
Panel A: Total visits						
Treat×Post	2.278***	2.278***	2.278***	15.536***	15.536***	20.909***
	(0.330)	(0.330)	(0.331)	(8.954)	(8.954)	(13.567)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✗	✗	✗
Time FE	✓	✓	✓	✗	✗	✗
Obs.	4,462	4,462	4,363	4,462	4,462	4,363
Panel B: Total visit duration (seconds)						
Treat×Post	2.847***	2.847***	2.849***	6.620***	6.620***	6.804***
	(0.346)	(0.346)	(0.346)	(2.129)	(2.129)	(2.473)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✗	✗	✗
Time FE	✓	✓	✓	✗	✗	✗
Obs.	4,462	4,462	4,363	4,462	4,462	4,363

D.2 On-Chain Transaction Frequency

Table A.4 reports post-period average ATTs for on-chain transaction frequency using a stacked event-study difference-in-differences design with a no-overlap control group. Results are presented for tax-haven exchanges (columns 1–3) and non-tax-haven exchanges (columns 4–6), separately for all, anonymous, and identified transactions.

Across all three panels (aggregate activity, inflows, and outflows), the treatment effect is positive and statistically significant for both tax-haven and non-tax-haven exchanges. The magnitudes are consistently larger for non-tax-haven platforms (e.g., 1.163 versus 0.759 for all transactions in Panel A), and within each exchange type the effect is driven predominantly by anonymous transactions: anonymous-transaction coefficients closely track the aggregate estimates (e.g., 0.757 versus 0.759 on tax-haven exchanges), while identified transactions show substantially smaller responses (0.356). This pattern is stable across activity margins: inflows (Panel B) and outflows (Panel C) mirror the aggregate results in Panel A, suggesting that aid inflows are associated with a broad-based increase in on-chain Bitcoin activity, with a disproportionate share routed through anonymous channels.

Table A.4. Impact of Foreign Aid on Transaction Frequency. This table reports the effect of foreign-aid disbursements on on-chain Bitcoin transaction frequency across tax-haven and non-tax-haven exchanges, separately for all, anonymous, and identified transactions. Panels distinguish aggregate activity (Panel A), inflows (Panel B), and outflows (Panel C). The reported coefficient is the post-period average of Treat×Post ($k = 0, \dots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE with a no-overlap control group. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Crypto Exchanges in Tax Haven			Crypto Exchanges in Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: All Transactions						
Treat×Post	0.759***	0.757***	0.356***	1.163***	1.158***	0.579***
	(0.122)	(0.121)	(0.063)	(0.147)	(0.147)	(0.085)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.02	0.04	0.04	0.03
Panel B: Inflows						
Treat×Post	0.681***	0.678***	0.310***	1.077***	1.070***	0.521***
	(0.108)	(0.108)	(0.056)	(0.136)	(0.135)	(0.079)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.01	0.04	0.04	0.03
Panel C: Outflows						
Treat×Post	0.689***	0.687***	0.267***	0.997***	0.993***	0.433***
	(0.113)	(0.112)	(0.049)	(0.131)	(0.131)	(0.067)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	9,096	9,096	9,096	9,096	9,096	9,096
Adj. R^2	0.02	0.02	0.01	0.03	0.03	0.03

D.3 Offshore Banking Complementarity

Table 6 tests whether the aid–cryptocurrency relationship varies with countries' pre-existing exposure to offshore banking. We augment the baseline stacked difference-in-differences specification with an interaction between Treat$\times$Post and a country-level moderator z measuring offshore banking intensity. Following Andersen et al. (2022), we use the Alstadsæter–Johannesen–Zucman (AJZ) offshore-to-GDP index for 2007, which captures the ratio of a country's cross-border deposits in offshore financial centres to its GDP. The moderator is standardised to mean zero and unit standard deviation over the estimation sample. A positive interaction coefficient implies that aid-induced crypto activity is amplified in countries where traditional offshore channels are already well established.

The interaction coefficients are positive and statistically significant across all six specifications. On tax-haven exchanges, a one-standard-deviation increase in offshore exposure amplifies the aid-induced response by 0.223 IHS units for transaction volume, 0.232 for transaction frequency, and 0.479 for new accounts. Non-tax-haven exchanges display the same pattern with somewhat smaller but significant magnitudes (0.129, 0.138, and 0.333 respectively). Overall, the interaction adds roughly 30–50% on top of the baseline Treat$\times$Post effect for countries one standard deviation above the mean.

The positive interaction has two possible interpretations that carry different implications for measuring aggregate aid diversion. Under the first, the same elites exploit both channels in parallel: offshore banking expertise—familiarity with cross-border transfers, access to intermediaries, pre-existing structures for moving funds outside domestic oversight—lowers the cost of engaging with crypto markets. If this interpretation is correct, a portion of the diverted funds captured by our crypto estimates also appears in the offshore deposit flows measured by Andersen et al. (2022), and summing the two leakage figures would overcount total diversion.

Under the second interpretation, both channels are driven by a common underlying factor—weak domestic institutions, concentrated elite access to aid flows, or the quality of elite financial networks—that facilitates diversion regardless of the specific vehicle used. In this case, different individuals may use different channels, meaning the two leakage estimates could be closer to additive. Our data cannot definitively distinguish between these mechanisms: doing so would require individual-level linking of on-chain wallets to offshore account holders, which is beyond the scope of blockchain forensics absent leaked administrative records.

The distinction carries practical consequences. If the same actors operate across both channels simultaneously, single-channel enforcement leaves the parallel infrastructure intact, requiring coordinated, multi-channel oversight. If instead the two channels serve similar but distinct populations connected by common institutional weaknesses, the reallocation elasticity is less clear, and institutional reform may prove more effective than channel-specific regulation. In either case, aggregate aid diversion is bounded above by the sum of our crypto estimate (2–5% under turnover adjustment) and the 7.5% offshore banking estimate of Andersen et al., but the true figure likely falls below this sum given the institutional linkages documented here. More broadly, the finding that offshore-banking exposure amplifies rather than crowds out the crypto response underscores the difficulty of containing illicit flows through single-channel regulation.

D.4 Alternative Specification

Table A.5 re-estimates the baseline treatment effects using Poisson pseudo-maximum-likelihood (PPML) on level outcomes. The PPML specification includes event-time dummies only, without unit or calendar-time fixed effects. This restricted specification reflects a computational constraint rather than a modelling choice: the high-dimensional Ctry.×Plat.×Ep. fixed effects used in the linear TWFE cannot be included in the PPML estimator because the large mass of zero observations in country-platform-episode cells induces separation, causing many fixed-effect groups to be dropped entirely during estimation. The PPML results should therefore be interpreted as a directional robustness check on the functional form—confirming that the findings are not driven by the IHS transformation—rather than as a specification of equal standing with the fully saturated TWFE model. The reported statistic is the average incidence-rate ratio (IRR) over post periods $k = 0, \ldots, 4$, where IRR > 1 indicates a treatment-induced increase.

Across all three outcome families (transaction volume in Panel A, transaction frequency in Panel B, and new accounts in Panel C), the IRRs are large, positive, and statistically significant at the 1% level, corroborating the direction and significance of the TWFE results. On tax-haven exchanges, the IRRs for all and anonymous transactions range from approximately 11 to 13, implying that treated countries experience roughly a tenfold increase in on-chain activity relative to controls. Non-tax-haven exchanges exhibit even larger IRRs (≈ 20–24), consistent with the pattern observed in the linear specifications. The larger PPML magnitudes relative to the TWFE estimates reflect both the multiplicative interpretation of the IRR and the absence of unit and time fixed effects that would otherwise absorb level differences across countries and global trends in crypto activity.

We suppress one degenerate cell: identified new-account creation on tax-haven exchanges (Panel C), where near-zero baseline counts produce an IRR of approximately 38,794 with a standard error of similar magnitude, driven entirely by the exponential transformation of a large coefficient off a negligible base. The remaining PPML results broadly confirm that the baseline findings are not an artefact of the IHS transformation or the linear functional form.

Table A.5. Impact of Foreign Aid on Crypto Activity (PPML). This table re-estimates the baseline treatment effects using Poisson pseudo-maximum-likelihood (PPML) on level outcomes with a no-overlap stacked control group. The specification includes event-time dummies only (no unit or calendar-time fixed effects). The reported statistic is the average incidence-rate ratio (IRR) over post-treatment periods $k = 0, \dots, 4$, computed as $\overline{\text{IRR}} = \exp\left(\frac{1}{L}\sum_{k\geq 0}\hat{\beta}_k\right)$; an IRR greater than one indicates a treatment-induced increase. Standard errors (in parentheses) are obtained via the delta method and clustered at the Ctry.×Plat.×Ep. level; H_0: IRR $= 1$. $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Tax Haven			Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: Transaction Volume						
Treat×Post	11.691***	11.597***	12.842***	22.816***	22.749***	22.302***
	(7.772)	(7.719)	(8.505)	(14.171)	(14.142)	(14.112)
Obs.	4,903	4,903	4,903	4,903	4,903	4,903
Panel B: Transaction Frequency						
Treat×Post	12.200***	12.235***	10.114***	23.783***	23.889***	20.806***
	(9.072)	(9.097)	(7.594)	(16.301)	(16.369)	(14.406)
Obs.	4,903	4,903	4,903	4,903	4,903	4,903
Panel C: New Accounts						
Treat×Post	12.982***	12.981***	38794.195***	23.644***	23.641***	56.455***
	(10.297)	(10.297)	(42550.847)	(17.398)	(17.395)	(29.574)
Obs.	4,903	4,903	4,903	4,903	4,903	4,903

D.5 Alternative Control Composition

A remaining concern is that the baseline control pool may include countries whose cryptocurrency-adoption trajectories differ systematically from those of aid recipients. Table A.6 addresses this through two complementary restrictions: limiting the control group to developing economies that received no disbursement during the estimation window, and excluding countries that share a land border with any treated country. The results are closely aligned with the baseline across both specifications. Anonymous transaction volume on tax-haven exchanges is 0.776 and 0.768 under the two restrictions, compared with 0.770 in the baseline. Transaction frequency and new anonymous account creation show the same stability, and the anonymous-to-identified gap is preserved throughout. The observation counts shift modestly, confirming that few control units are affected by either filter. These results indicate that the main findings are not driven by comparisons with dissimilar non-recipients or by geographic spillovers from neighboring countries.

Table A.6. Impact of Foreign Aid on Crypto Activity across Alternative Control Compositions. This table combines two robustness checks: (i) restricting the control group to developing countries that did not receive World Bank disbursements during the estimation window, and (ii) excluding neighboring countries from the control group under the no-overlap design. The estimator in the first specification is the Sun and Abraham (2021) saturated cohort×relative-time DiD, numerically equivalent to Callaway and Sant'Anna (2021), with a no-overlap control group. The second specification uses stacked event-study DiD (linear TWFE). Three sections report results for transaction volume, transaction frequency, and new accounts, all for aggregate transaction measures only. Outcomes are IHS-transformed. The reported coefficient is the equal-weight average treatment effect over post-treatment periods $k = 0, \ldots, 4$. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Crypto Exchanges in Tax Haven						Crypto Exchanges in Non-Tax Haven					
	All		Anonymous		Identified		All		Anonymous		Identified	
	Dev. Ctrl	No Nbr	Dev. Ctrl	No Nbr	Dev. Ctrl	No Nbr	Dev. Ctrl	No Nbr	Dev. Ctrl	No Nbr	Dev. Ctrl	No Nbr
Panel A: Transaction Volume												
Treat×Post	0.783***	0.775***	0.776***	0.768***	0.435***	0.430***	1.244***	1.234***	1.234***	1.225***	0.654***	0.646***
	(0.124)	(0.124)	(0.123)	(0.123)	(0.079)	(0.079)	(0.153)	(0.152)	(0.152)	(0.151)	(0.098)	(0.098)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010
Adj. R^2	0.02	0.02	0.02	0.02	0.01	0.01	0.04	0.04	0.04	0.04	0.03	0.03
Panel B: Transaction Frequency												
Treat×Post	0.764***	0.757***	0.762***	0.755***	0.359***	0.355***	1.172***	1.163***	1.167***	1.157***	0.585***	0.578***
	(0.122)	(0.122)	(0.122)	(0.121)	(0.063)	(0.063)	(0.148)	(0.147)	(0.147)	(0.147)	(0.086)	(0.085)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010
Adj. R^2	0.02	0.02	0.02	0.02	0.02	0.02	0.04	0.04	0.04	0.04	0.03	0.03
Panel C: New Accounts												
Treat×Post	1.468***	1.453***	1.468***	1.453***	0.115***	0.114***	2.723***	2.705***	2.723***	2.705***	0.172***	0.170***
	(0.236)	(0.235)	(0.236)	(0.235)	(0.037)	(0.037)	(0.301)	(0.299)	(0.301)	(0.299)	(0.040)	(0.040)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010	9,092	9,010
Adj. R^2	0.02	0.02	0.02	0.02	0.00	0.00	0.04	0.04	0.04	0.04	0.01	0.01

D.6 DeFi Boom Exclusion

The period from May to September 2021 witnessed an unprecedented surge in decentralized finance activity that reshaped global cryptocurrency markets and could, in principle, confound our estimates if the associated increase in on-chain volume coincided with aid disbursement events. To ensure that the baseline results are not driven by this transient market-wide shock, we re-estimate all specifications after dropping the May–September 2021 window from both the treatment and control samples.

Table A.7 reports the results. The point estimates remain positive, strongly significant, and quantitatively similar to the baseline across all outcome families, exchange types, and anonymity categories. Anonymous transaction volume on tax-haven exchanges is 0.730 (compared with 0.770 in the baseline), and anonymous transaction frequency is 0.728 (versus 0.757). New anonymous account creation is 1.353 on tax-haven and 2.332 on non-tax-haven exchanges, closely tracking the baseline magnitudes of 1.457 and 2.706. The anonymous-to-identified gap is preserved throughout, and the modest reduction in observation counts (from 9,096 to 8,918) confirms that few episodes fall within the excluded window. These results confirm that the 2021 DeFi boom does not drive our findings.

Table A.7. Robustness Check: DeFi Boom Exclusion. This table re-estimates the baseline treatment effects after excluding the May–September 2021 period from both the treatment and control samples to ensure that the unprecedented surge in decentralized finance activity during this window does not confound the estimates. The reported coefficient is the post-period average of Treat×Post ($k = 0, \ldots, 4$) from a stacked difference-in-differences on IHS-transformed outcomes, estimated by linear TWFE with a no-overlap control group. Three sections report results for transaction volume, transaction frequency, and new accounts; columns distinguish all, anonymous, and identified transactions by platform type. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses, clustered at the Ctry.×Plat.×Ep. level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

	Crypto Exchanges in Tax Haven			Crypto Exchanges in Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Transaction Volume						
Treat×Post	0.734***	0.730***	0.389***	1.042***	1.037***	0.493***
	(0.135)	(0.134)	(0.085)	(0.165)	(0.163)	(0.104)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	8,918	8,918	8,918	8,918	8,918	8,918
Adj. R^2	0.01	0.01	0.01	0.03	0.03	0.02
Transaction Frequency						
Treat×Post	0.729***	0.728***	0.339***	0.983***	0.980***	0.456***
	(0.134)	(0.134)	(0.069)	(0.160)	(0.159)	(0.092)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	8,918	8,918	8,918	8,918	8,918	8,918
Adj. R^2	0.01	0.01	0.01	0.03	0.03	0.02
New Accounts						
Treat×Post	1.353***	1.353***	0.133***	2.332***	2.332***	0.198***
	(0.257)	(0.257)	(0.042)	(0.322)	(0.322)	(0.045)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	8,918	8,918	8,918	8,918	8,918	8,918
Adj. R^2	0.01	0.01	0.00	0.03	0.03	0.01