

NBER WORKING PAPER SERIES

CRYPTO CAPTURE OF FOREIGN AID

Sumit Agarwal
Peiyi Jin
Eswar S. Prasad
Daniel Rabetti

Working Paper 35655
<http://www.nber.org/papers/w35655>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2026

We are grateful to Lin William Cong, Sean Foley, John Griffin, Campbell Harvey, Shiyang Huang, Bjorn Jorgensen, Jonathan Karpoff, Wayne Landsman, Chen Lin, Evgeny Lyandres, Mei Luo, Michael Sockin, Eugene Soltes, Charles C. Y. Wang, Gloria Yu, and Che Zhang for their invaluable 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). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Sumit Agarwal, Peiyi Jin, Eswar S. Prasad, and Daniel Rabetti. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Crypto Capture of Foreign Aid

Sumit Agarwal, Peiyi Jin, Eswar S. Prasad, and Daniel Rabetti

NBER Working Paper No. 35655

August 2026

JEL No. G15, G18, G29, K29, K42, O16

ABSTRACT

The 2016 Panama Papers leak tightened regulatory enforcement around money laundering and offshore banking. We investigate whether the diversion of foreign aid in developing countries led to a shift to cryptocurrency as an alternative laundering platform. We develop a disbursement-timed forensic measure of cryptocurrency activity, combining on-chain Bitcoin transactions and wallet creation, off-chain exchange records, and IP-linked web traffic, and apply it to World Bank aid disbursements covering \$238 billion across the 93 recipient countries in our estimation sample during 2018-2024. Exploiting the administrative timing of aid tranche arrivals, we find sharp, short-lived surges of crypto activity at the disbursement month, driven mainly by anonymous and newly created wallets on both tax-haven and mainstream exchanges. Blockchain forensics reveal patterns consistent with the placement, layering, and integration sequence of conventional money laundering. We estimate an implied leakage of 2 to 6 cents per aid dollar, which amounts to roughly 1.7 to 4.4 billion dollars of aid diversion across the tranche arrivals we study. Capture carries no funding penalty: the four sectors where we detect it, Transport, Water and Sanitation, Social Protection, and Governance, still absorb half of subsequent World Bank funding. Cryptocurrency facilitates aid diversion, but its transparent ledgers also leave forensic traces that may help detect and recover diverted funds.

Sumit Agarwal
National University of Singapore
ushakri@yahoo.com

Peiyi Jin
Princeton University
jin_peiyi@u.nus.edu

Eswar S. Prasad
Cornell University
and NBER
eswar.prasad@cornell.edu

Daniel Rabetti
National University of Singapore
and Asian Bureau of Finance and
Economic Research
rabetti@nus.edu.sg

I Introduction

Foreign aid is intended to move resources from richer countries to some of the world’s most vulnerable populations. Yet the officials and intermediaries who administer these transfers may also be best placed to divert them. The most direct evidence of such diversion comes from offshore banking: [Andersen et al. \(2022\)](#) show that World Bank disbursements coincide with bank deposits in offshore financial centers known for bank secrecy and estimate that as much as 7.5 cents of every aid dollar is diverted in this way. For decades, this form of concealment relied on a familiar infrastructure of shell companies, nominee owners, and discreet banks in lightly regulated jurisdictions. The 2016 Panama Papers leak ([International Consortium of Investigative Journalists, 2016](#)) brought that infrastructure into public view, and the enforcement wave that followed made traditional tax-haven banking more costly and risky. But greater scrutiny of offshore banks need not have eliminated the demand to conceal diverted funds. It may instead have redirected those funds away from offshore financial centers known for bank secrecy.

Cryptocurrency is a natural candidate for this migration. Pseudonymous wallets can be created quickly, funds can cross borders almost instantaneously, and exchanges operating beyond the effective reach of many regulators provide ready entry and exit points. These features make cryptocurrency potentially attractive to actors seeking to move public money out of sight. They also make the phenomenon difficult to study. The same pseudonymity that facilitates concealment prevents researchers from relying on the leaked ownership documents and country-level bank deposit statistics that underpin the offshore-banking evidence. We therefore ask a simple but consequential question: when foreign aid arrives, does crypto activity shift toward the anonymous and offshore parts of the ecosystem in a manner consistent with diversion and money laundering?

We provide the first systematic evidence on this question. Our analysis centers on a new, disbursement-timed forensic measure of illicit crypto response that combines three complementary layers of data. First, on-chain Bitcoin forensics clusters wallet addresses and tags them to known entities, allowing us to separate anonymous from identified activity for each venue monthly. Second, off-chain trading records from centralized exchanges reveal how activity changes on the platforms themselves. Third, IP-linked web traffic from SimilarWeb refines geographic attribution by linking each platform’s activity to the countries where its users are located. We link these measures to project-level World Bank disbursements from 2018 to 2024, covering \$238 billion disbursed to the 94 recipient countries that enter our estimation sam-

ple. We then examine what happens to a country’s crypto activity in the month an aid tranche arrives. Whereas existing work links cryptocurrency to aggregate capital flight (Graf von Luckner et al., 2023, 2024), our measure ties crypto activity to individual aid projects and, crucially, to the administrative month in which the funds are received. This granularity allows us to distinguish a response to money in hand from a broader response to macroeconomic stress.

We find that the arrival of aid is followed by a sharp reorientation of crypto activity toward precisely the venues and account types that a laundering interpretation would predict. In the arrival month, anonymous transaction volume routed through exchanges domiciled in tax havens rises by roughly 137 percent; identified volume, which is tied to wallets with known owners, moves by barely one-third as much.¹ The same reorientation appears at the extensive margin, although it should be read differently there. New anonymous wallet creation increases by 1.8 to 2.8 IHS units depending on the exchange type, while the corresponding estimates for identified wallets are small and, on tax-haven exchanges, statistically indistinguishable from zero. Because almost all new-wallet creation we observe is unattributable in the first place, we do not read the ratio between the anonymous and identified coefficients as a magnitude.² What the extensive-margin result establishes is that post-disbursement entry into the crypto ecosystem occurs almost entirely through accounts that conceal their owners. The off-chain records corroborate the on-chain evidence: responses on platforms in unregulated jurisdictions are more than ten times larger than those on regulated platforms.

The dynamics are equally revealing. The increase in anonymous tax-haven volume is sharp but short-lived, fading within one to two months. Transaction networks constructed around the five largest disbursements show flows moving toward tax-haven exchanges and mixing services immediately after the funds arrive; by six months, the most salient post-disbursement links are less prominent, although the network has not returned mechanically to the pre-disbursement baseline. Taken together, these patterns resemble the classic placement, layering, and integration sequence of money laundering. Funds enter

¹Outcomes are estimated on the inverse hyperbolic sine (IHS) scale, so a coefficient $\hat{\beta}$ converts to a proportional change as $e^{\hat{\beta}} - 1$; the arrival-month impulse of $\hat{\beta}_0 = 0.864$ gives $e^{0.864} - 1 \approx 1.37$, or 137 percent. We adopt this convention wherever a percentage is quoted. Where the impulse is large, as for new-account creation, the exponentiated figure is unstable in the tails and we report the IHS coefficient itself rather than a percentage.

²Wallet attribution operates at the level of services rather than end users, so a newly observed address inside a named service’s cluster is not a new entity. Newly created wallets are accordingly almost never attributable: anonymous wallets account for 99.9997 percent of new-account mass on tax-haven exchanges and 99.9991 percent on non-tax-haven exchanges, and the identified series is non-zero in only 661 of 5,178 tax-haven country-month cells. This is also why the All and Anonymous columns of Table 2 coincide to three decimals. The anonymity contrast that carries economic content is the one on transaction volume in Table 1, where identified activity accounts for 9 percent of tax-haven and 10 percent of non-tax-haven throughput.

through different types of exchanges, are fragmented across newly created anonymous wallets, are routed toward venues with the weakest oversight, and ultimately move to long-term anonymous storage or offshore off-ramps. No single pattern is dispositive on its own. Their joint occurrence in the narrow window following aid arrival is, however, difficult to reconcile with a general expansion in legitimate crypto use.

The magnitudes are economically meaningful, although their interpretation requires care. Cumulating the implied level change in anonymous tax-haven volume over the post-disbursement window yields a gross leakage ratio of 11.2 percent of disbursed aid. We emphasize that this is not an estimate of net diversion. Cryptocurrency data record gross flows, and the same diverted dollar may appear several times as it moves across wallets. The estimate of 11.2 cents per dollar is therefore an upper bound, attained only if each dollar is observed once. Because our exchange ledgers do not reveal the relevant turnover multiplier, we do not impose one; instead, we report the full mapping from wallet hops to implied net leakage.

For the headline net measure, at the two to five hops that forensic practice treats as typical ([Chainalysis, 2021](#)), the implied leakage is roughly 2 to 6 cents per aid dollar. To fix magnitudes, the 328 tranche arrivals on which the ratio is estimated carry \$78 billion of World Bank disbursements, about one third of all disbursement flow to sample countries over the period, so 2 to 6 cents per dollar corresponds to roughly \$1.7 to \$4.4 billion of diverted aid. We deliberately do not scale this rate to the full disbursement stock: the design identifies the response at large, discrete tranche arrivals, and we find no evidence that the same rate applies to routine disbursements. It falls below the 7.5-cent offshore-banking benchmark of [Andersen et al. \(2022\)](#) once a diverted dollar is observed more than 1.49 times. Under a broad range of plausible layering rates, then, the crypto channel operates on the same order of magnitude as the offshore-banking channel, despite being identified using an entirely independent methodology: blockchain forensics rather than bank deposit records. Our estimates may also be conservative with respect to total crypto-mediated diversion because we observe only Bitcoin, while stablecoins have become a major vehicle for illicit cross-border transfers in developing countries.

Why might cryptocurrency inherit part of the role historically played by offshore banks? Moving funds through a regulated bank generally requires a correspondent relationship and an account attached to a legal identity. Moving funds on-chain requires neither. Value can be fragmented across thousands of newly created wallet addresses and transferred across borders within minutes, outside the banking system and its reporting requirements ([Amiram et al., 2022](#)). Exchanges domiciled in jurisdictions such as the Seychelles, Malta, and several offshore British territories provide entry and exit points where Know-

Your-Customer (KYC) enforcement may be weak, effectively serving as pseudo-offshore banks within the crypto ecosystem ([Financial Action Task Force, 2021](#)). Mixing services offer an additional layer of concealment by pooling transactions from multiple users and redistributing the funds so that the direct on-chain link between sender and receiver is severed.³ Yet this architecture has a feature offshore banking lacks: every transaction leaves an immutable public record. The technology therefore cuts both ways. It facilitates concealment at the identity level, but it also makes the movement of funds observable. We refer to this pattern as the “crypto capture” of aid, a digital descendant of the elite capture that [Reinikka and Svensson \(2004\)](#) documented in Ugandan school grants.

Our identification exploits the administrative calendar of foreign aid. The World Bank releases tranches as projects meet procurement benchmarks, audit checkpoints, and board dates within pre-approved cycles. Accordingly, the month in which a tranche arrives is shaped by bureaucratic milestones rather than by circumstances in the recipient economy ([Kraay, 2012](#); [Andersen et al., 2022](#)). Our identifying variation comes from this wedge between why a country receives aid and when a particular tranche is released. We implement a stacked difference-in-differences design with two-way fixed effects on IHS-transformed outcomes around 328 tranche-arrival spikes. The event studies exhibit flat pre-trends and a discrete increase in the arrival month. This timing is central to our interpretation: the estimated response tracks fund availability, not merely the economic or political conditions that generated the need for aid.

Who, then, operates this channel? We are explicit about the limits of our evidence. We cannot trace a particular aid dollar onto the blockchain, identify the owner of an anonymous wallet, or fully rule out all legitimate economic responses to the arrival of aid. Nevertheless, the evidence consistently points toward actors close to the disbursement chain rather than ordinary beneficiaries. World Bank funds move through government and procurement systems, while intended beneficiaries typically receive schools, roads, health services, and other public goods rather than the disbursement itself ([World Bank, 2024](#)). Using the channel we document requires reliable internet access, exchange accounts, private-key custody, and familiarity with lightly regulated offshore platforms. These capabilities remain scarce among many of the populations that aid is intended to serve: in 2023, only 27 percent of people in low-income countries had internet access ([International Telecommunication Union, 2023](#)), and crypto adoption is strongly

³Cryptocurrency mixers, also called tumblers, deliberately obscure transaction histories by pooling funds from multiple users and redistributing them in a manner that breaks the on-chain link between sending and receiving addresses. [Foley et al. \(2019\)](#) estimate that by the late 2010s roughly 46 percent of Bitcoin transactions involved illegal activity, with about one-quarter of users participating in it; the share has declined as regulated exchanges have grown, but illicit use remains substantial.

associated with financial inclusion and innovation capacity (Auer et al., 2022). The timing reinforces this interpretation. Shock-driven demand should respond to the shock that prompts the aid, such as a disaster, crisis, or election, and may therefore begin well before a tranche is released. Instead, we find the response follows the administrative arrival of the money itself.

We subject this interpretation to several tests. The October 2021 Pandora Papers leak, a shock to offshore financial transparency, attenuates aid-linked activity only among identified transactions tied to document-traceable structures; anonymous channels show no comparable attenuation. This differential response is consistent with pseudonymous wallets being insulated from document-based exposure. A placebo based on natural disasters that are not accompanied by a World Bank disbursement produces a precisely estimated null: without a disbursement, there is no crypto response. The results are unlikely to reflect VPN-induced geographic mismeasurement. Excluding the recipient countries with the heaviest VPN use under three alternative measures leaves the estimates unchanged or slightly larger. We further validate our geographic attribution externally against the Chainalysis Global Crypto Adoption Index, where PPP-weighted correlations increase from 0.689 in 2021 to 0.855 in 2024, and internally using an alternative visit-duration identifier. Within the crypto ecosystem, mixers gain relative to both payment processors and ordinary exchanges after disbursement, isolating increased demand for privacy rather than a broad post-aid expansion in crypto activity. The results are robust to alternative thresholds for defining a tranche-arrival spike and to excluding the 2021 decentralized-finance boom.

The sectoral results help localize both the economic mechanism and the policy problem. Crypto capture is not uniform across aid categories. The arrival-month response is largest in Transport, a sector characterized by large and discretionary construction contracts. It is also significant in Water and Sanitation, Social Protection, and Governance, where disbursement modalities can generate liquid local windfalls. By contrast, in-kind, commodity-based Health aid produces no detectable on-chain response. The exposure is not merely historical: the four sectors in which we find a response account for 48 percent of sector-classified World Bank disbursements and newly approved commitments in the two years after our sample ends. We interpret this as a statement about exposure, not the Bank's allocation rule; our evidence does not imply that sectors with greater leakage receive more aid. Rather, these results may guide efforts to limit aid diversion. Because the capture signal appears when a tranche arrives, our findings indicate that fiduciary supervision is likely most valuable at the point of disbursement, rather than only at project approval or through audits conducted long afterward.

The broader implication is that technological change may alter both the means of corruption and the evidence it leaves behind. Relative to its offshore predecessor, the crypto channel operates at higher frequency, requires no pre-existing banking relationship, and conceals identity through pseudonymous addresses rather than opaque legal entities. But this final feature also creates transparency. Piercing bank secrecy required leaking millions of documents; blockchains publish their ledgers continuously. The same public record that lets us measure crypto capture monthly could let aid agencies and financial regulators monitor disbursements in near real time, flag unusual post-disbursement activity, and, in some cases, trace and recover diverted funds.

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 et al. \(2009\)](#) links aid flows to political-economy distortions in weak-institutional contexts. [Andersen et al. \(2022\)](#) provide direct evidence of elite capture via offshore banking. [Chalendard et al. \(2023\)](#) detect corruption in customs from deviations in administrative transaction records; we apply a related logic to blockchain data. [Graf von Luckner et al. \(2023\)](#) estimate that roughly one-fifth of transactions on global peer-to-peer Bitcoin markets represent international capital flight or flows. [Graf von Luckner et al. \(2024\)](#) 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 et al. \(2018\)](#) and [Foley et al. \(2019\)](#) document the scale of criminal activity in early crypto markets, from drug trades to market manipulation, and show that illicit actors adapt quickly to privacy-enhancing technologies. We build on this work by focusing on a specific mechanism, the laundering of diverted public funds through mixers and loosely regulated offshore platforms; our network analysis supports the “layering” pathways documented in industry reports ([Chainalysis, 2021](#)) and academic tracing work on money laundering

([Ron and Shamir, 2013](#)). A parallel literature traces how illicit value moves through crypto markets: [Griffin and Shams \(2020\)](#) document Tether-driven manipulation of Bitcoin prices, [Griffin and Mei \(2024\)](#) trace on-chain flows from “pig butchering” scams into forced-labor operations, [Griffin et al. \(2025\)](#) evaluate whether crypto anti-money-laundering interventions curtail illicit flows, [Sokolov \(2021\)](#) shows that ransomware payments leave a footprint in blockchain congestion, and [Phua et al. \(2025\)](#) study the economics of scams run through initial coin offerings. We complement this work by exploiting the timing of a specific, quasi-exogenous inflow, World Bank aid disbursements, to identify laundering activity rather than characterizing the aggregate illicit ecosystem.

Finally, our study contributes to the nascent literature that exploits 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 et al. \(2023a\)](#) and [Cong et al. \(2025\)](#) 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 et al., 2019](#); [Budish, 2025](#)). Within this strand, [Londoño-Vélez and Ávila-Mahecha \(2021\)](#) use the Panama Papers leak to show that offshore concealment drives wealth-tax evasion by Colombian elites, [Tørsløv et al. \(2023\)](#) quantify haven-bound profit shifting under explicit bounding assumptions, and [Johannesen et al. \(2020\)](#) and [Guyton et al. \(2021\)](#) document that conventional audits systematically miss offshore evasion, which motivates transaction-level forensic measurement of the kind we develop.

In the next section of the paper, we provide 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 4.7 synthesizes the evidence within the placement, layering, and integration framework. Section 5 presents robustness checks. Section 6 draws out the policy implications. 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 faced pressure to strengthen compliance, and the veil of secrecy around 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 ([Huberman et al., 2021](#)). 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 shows that illicit actors quickly recognized these advantages. [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 breakthroughs in both academia and law enforcement ([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, focusing on World Bank aid. The World Bank is the largest multilateral development lender, disbursing roughly \$30 to \$60 billion annually over our sample period, so we build the sampling frame around it. We use the World Bank’s Project Database (via its API) to obtain detailed information on project loans and grants from the International Development Association (IDA) and the International Bank for Reconstruction and Development (IBRD), the World Bank’s arms.⁴ 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 is the monthly World Bank disbursement flow to a recipient country, constructed by first-differencing loan-level cumulative disbursement records. In the baseline event-study design, a treated episode is a tranche-arrival month in which this flow is at least \$50 million, at least three times the country’s trailing twelve-month median monthly flow, and not preceded by another qualifying spike in

⁴IDA (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. Together, IBRD and IDA disburse approximately \$30–60 billion annually over our sample period (FY2018–2024). 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>).

the prior four months. The identification therefore operates at the monthly frequency; quarterly lagged aid variables are used only in auxiliary off-chain specifications.

Focusing on World Bank aid (as opposed to all Official Development Assistance (ODA)) has a key advantage: disbursement timing. 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 A.2, Panel A reports the top 15 recipients by disbursed amount. Aid is concentrated among a small number of large recipients: India, Indonesia, and Nigeria alone account for roughly 17 percent of total disbursements, and the top 15 countries together represent 54 percent of all disbursed funds. Over April 2018 through March 2024 the 94 recipient countries in our estimation sample receive \$238 billion in World Bank disbursements, of which \$78 billion arrives in the 328 tranche-arrival months that define our treated episodes. 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

Our analysis is based on on-chain Bitcoin transaction data. Bitcoin has historically been the dominant cryptocurrency for illicit cross-border transfers (Foley et al., 2019). 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 the standard clustering-and-tagging pipeline of Meiklejohn et al. (2013), as implemented in WalletExplorer, in two steps: we first group addresses into wallet clusters, then attribute each cluster to a named service when possible. Following the counterparty-classification approach of Foley et al. (2019), we classify each transfer by its counterparty identity and treat as anonymity-seeking the

residual activity whose counterparty cannot be tied to any named service. Crypto trading infrastructure tends to cluster in tax-haven or crypto-friendly jurisdictions, which is why the domicile of the exchange, rather than the location of the trader, organizes the analysis that follows.

Step 1: 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](#)), we cluster addresses that appear jointly as inputs in the same transaction and treat them 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.

Step 2: Service Attribution. We then attribute clusters to known entities using public and forensic sources, including WalletExplorer and open-source blockchain investigations. These sources identify addresses associated with exchanges, mixers, gambling platforms, darknet markets, payment processors, and related services, enabling us to tag transactions involving these services. We classify wallets into one of five categories (Exchange, Mixer, Gambling, Dark Market, or Other Service) using attribution data. A cluster that resolves to a named service is classified as identified; the label is observed at the service level and does not identify the underlying end user. A transfer counts as anonymous when its counterparty is a cluster that no source could match to a named service, even when the other leg is an exchange in a tax-haven jurisdiction. Throughout the analysis, we define anonymity 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 indicate heightened anonymity-seeking behavior. Appendix D provides a worked construction example, tracing a single deposit and withdrawal through the clustering, edge-direction, and change-address conventions to make the anonymous tax-haven flow measure precise.

Our wallet labeling does not cover the full universe of clusters, and the resulting measurement error is one-sided. A cluster belonging to a service WalletExplorer has not attributed is classified as anonymous; the reverse mislabeling is impossible by construction. More complete attribution would likely move some currently-anonymous clusters into the identified bucket, compressing the anonymous-identified contrast

in our estimates. We do not claim the magnitudes are immune to this. The direction is informative, however. Because unattributed identified activity contaminates the anonymous series and never the reverse, the error biases the anonymous coefficient toward the identified one. Since we estimate the anonymous response to be the larger of the two, the estimated anonymous-identified gap is a lower bound on the true gap: incomplete attribution can shrink the contrast we report, but it cannot manufacture it.

The residual concern is whether attribution gaps correlate with disbursement timing. The country-by-platform-by-episode fixed effects in our setting absorb time-invariant differences in attribution coverage. Remaining variation is identified off month-to-month timing of disbursements, governed by administrative schedules unrelated to WalletExplorer’s labeling cadence.

Table A.2 Panels 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. Panel E reports the anonymous share of measured activity by platform type: anonymous flows carry 91.2 (89.8) percent of transaction volume on tax-haven (non-tax-haven) exchanges and essentially all new-account creation on both.

3.3 Off-Chain Exchange Data

While on-chain data capture Bitcoin movements 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.

This distinction is important for wash trading. Wash trading is most likely to affect reported trading volume on centralized exchanges, where the same beneficial owner can trade with itself or coordinate off-

setting orders inside an exchange’s private order book (Cong et al., 2023b; Amiram et al., 2025). For this reason, our core evidence does not rely on self-reported exchange volume. The main outcomes are on-chain Bitcoin transfers, transaction frequency, wallet creation, and flows across identified exchange and service clusters. A wash trade inside a centralized exchange can inflate off-chain volume, but it does not by itself create new on-chain wallets, generate blockchain transfers between clusters, or explain the differential response of anonymous versus identified on-chain activity. We therefore use off-chain exchange data as corroborating evidence on where internal platform activity takes place, while treating the on-chain evidence as the primary basis for inference.

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 most 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 percent of active crypto exchanges in our sample are located in jurisdictions classified as tax havens, about 18 percent are in strongly regulated countries, and about 70 percent are in what can be termed “crypto-friendly” countries (those that actively promote crypto business or lack strict regulations). These categories overlap considerably (for example, many tax haven jurisdictions are also crypto-friendly because of 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. For example, if 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 diverted aid being converted 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, $\text{ihS}(x) = \ln(x + \sqrt{x^2 + 1})$, 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.

We define treated episodes as flow-based tranche-arrival spikes. A recipient country-month is flagged as an event when the World Bank disbursement flow in that month is at least \$50 million and at least three times the country’s trailing twelve-month median flow, with no qualifying spike in the preceding four months. The two screens play distinct roles: the \$50 million floor selects inflows large enough to leave a detectable on-chain footprint, while the three-times-median requirement makes an event mean an inflow that is unusually large for that country, keeping the definition comparable across recipients of very different size; the four-month spacing ensures each event has a clean, non-overlapping pre-period. Because World Bank funds disburse more or less continuously but in tranches of very uneven size, this rule is what converts an uneven disbursement stream into discrete arrival events at which diversion, which requires funds in hand, is actually feasible.

This timing isolates the month in which funds actually arrive rather than an administrative effective date, and yields 328 treated episodes; Section 5.5 shows the findings are insensitive to the \$50 million, three-times, and four-month thresholds. When two or more qualifying spikes to the same country fall within the post-treatment window, we aggregate them into a single episode and set event time zero to the first spike. 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 ($\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$) 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.⁵

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 et al. \(2019\)](#). For each episode e , the treated unit is

⁵This timing-based argument is strongest for instruments whose tranches follow verified project milestones, such as Investment Project Financing, and weaker for policy-based or emergency operations, whose large single-tranche releases can respond to the contemporaneous macroeconomic conditions that also concern us. We do not split the arrival-month impulse by lending instrument, nor decompose announcement timing from arrival timing, because the World Bank records do not carry reliable disbursement-level effective dates. Keying the design on realized monthly flows rather than these administrative dates is precisely the choice that any within-instrument or announcement-based decomposition would have to undo.

the disbursement country $c(e)$, and event time is

$$k \equiv t - t_e,$$

where t_e denotes the tranche-arrival month. We estimate event-time coefficients over the extended window $k \in \{-8, \dots, +5\}$, with the far-pre period $k \in [-12, -9]$ as the omitted reference; anchoring the reference safely outside the mechanically quiet pre-disbursement window, rather than at $k = -1$ adjacent to it, avoids attributing that quiet window to a spurious pre-trend.

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

$$Y_{cpt}^{(e)} = \sum_{k=-8}^{+5} \beta_k \cdot \mathbf{1}\{\text{Treated}_c^{(e)}\} \mathbf{1}\{t - t_e = k\} + \alpha_{cp}^{(e)} + \gamma_t + \varepsilon_{cpt}^{(e)}, \quad (1)$$

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. $\times$ Plat. $\times$ 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 $k \in \{-8, \dots, -5\}$, the clean pre-period lying outside the mechanically quiet window that the tranche-spike definition enforces at $k \in \{-4, \dots, -2\}$, serve as a pre-trend diagnostic: under the parallel-trends assumption, these should be individually and jointly insignificant. Standard errors are clustered at the recipient-country level with the finite-sample correction $G/(G - 1)$, where G is the number of clusters ([Cameron and Miller, 2015](#)).

Omitting the far-pre bin $k \in [-12, -9]$ as the reference, rather than the conventional single period $k = -1$, follows directly from the event definition. The four-month cleanliness screen that isolates each spike mechanically suppresses disbursement flow in the months just before arrival, producing a shallow, localized dip in the outcome at $k \in \{-4, -3\}$ (Appendix Figure A.1). Anchoring at $k = -1$, which sits at the top of the recovery from this dip, would render the sequence $k = -4, \dots, -1$ an upward-sloping “ramp” easily misread as pre-event anticipation, even though it is only the mechanical

dip viewed from its own edge. We therefore reference the nearest window the spike screen leaves unconstrained, $k \in [-12, -9]$, which also keeps the reference disjoint from the pre-trend test window $k \in \{-8, \dots, -5\}$: were the two to coincide, the flatness test would hold by construction. The substantive estimate is insensitive to the choice: the arrival-month impulse for anonymous tax-haven volume is 0.85 under the far-pre reference and 0.77 under $k = -1$, both significant at the one-percent level, and the clean pre-period is flat and jointly insignificant under either normalization. Because calendar-month fixed effects absorb common time variation, the longer reference-to-event gap does not import secular growth in crypto adoption.

To summarize the treatment effect, we report the arrival-month impulse $\hat{\beta}_0$, the coefficient on the $k = 0$ event-time dummy. On the flow-based tranche-arrival panel the response is a sharp impulse that peaks in the arrival month and reverts to baseline within one to two months, so the single-month impulse, rather than a post-period average $\frac{1}{L+1} \sum_{k=0}^L \hat{\beta}_k$ over $k = 0, \dots, 4$, is the object that matches the dynamic path; averaging over the post window mechanically dilutes the impulse against the near-zero post-revert coefficients. Inference on $\hat{\beta}_0$ uses its cluster-robust standard error.

Three sets of specification checks support these choices, all reported in the appendix. We do not adopt Poisson pseudo-maximum-likelihood (PPML) as our primary estimator because the linear IHS specification is computationally transparent and directly comparable across the large number of outcome–platform–specification combinations we report; PPML instead serves as a scale-invariance check. Because the IHS coefficient is not scale-invariant when the outcome is predominantly zero ([Chen and Roth, 2024](#)), Appendix E.2 re-estimates the identical stacked $\text{Treat} \times \text{Post}$ with a scale-equivariant PPML estimator and with separate extensive- and intensive-margin decompositions, and shows that the baseline result is corroborated on the scale-free extensive margin; on the sparse spike panel the levels-based PPML impulse is imprecise and so is uninformative rather than contradictory. Appendix E.2 also confirms that the main results are qualitatively unchanged when the interacted fixed effects are replaced with additive country and month fixed effects. Appendix E.3 verifies that the baseline findings are robust to alternative control group compositions, including restricting the control pool to developing economies and excluding neighboring countries.

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.⁶

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 et al. \(2023b\)](#) 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:

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

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

⁶Examples include tax-haven-domiciled exchanges such as HTX (<https://www.htx.com/>) and regulated exchanges such as Kraken (<https://www.kraken.com/>).

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.

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 $Y_{cpt} = \text{VisitsRatio}_{cpt} \times G_t$ 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. Colors distinguish service categories, and the plotted labels identify each category directly.

Figure 1 shows the network for the five largest aid disbursements in our sample (summary statistics in Table A.3), plotting address-level snapshots one month before, one month after, and six months after each event. Before disbursement, the network is dominated by non-tax-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 is more muted, with non-tax-haven exchanges regaining relative prominence and obfuscation-related flows subsiding. The six-month network should therefore be read as attenuation rather than a full return to the pre-disbursement baseline, 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. Panel A is a difference network comparing the week before each of the five largest aid disbursements with the week after it. The most visible change is a link from the category labeled “Non Haven Exchange” (non-tax-haven exchanges) to tax-haven exchanges, with additional flows involving mixers and payment platforms. This pattern is consistent with funds moving out of regulated exchange channels and toward venues with weaker oversight, while mixing services become more active in the same window. Panel B compares the same pre-disbursement week with six months after disbursement. The tax-haven node is no longer the central feature of the graph; the remaining differences are concentrated around non-tax-haven exchanges, payment platforms, mixers, and gambling services. We therefore read Figure 2 as qualitative evidence that the strongest tax-haven reallocation is front-loaded, not as proof that the network fully returns to baseline. This attenuation is consistent with laundering activity being 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 1 and A.7 report the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) of aid disbursements on transaction volume and transaction frequency, estimated using a stacked event-study difference-in-differences design on the flow-based tranche-arrival panel (328 episodes) 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 4.2), the country-level variation in the tables below traces to the same geographic signal as the web-traffic attribution in Section 4.2. 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 4.6) and the network graphs (Section 4.3).

We begin with transaction volume in Table 1. On tax-haven exchanges, the response is systematically stronger for anonymous than for identified transactions. In Panel A, the arrival-month impulse for anonymous transaction volume is 0.864, compared with 0.301 for identified transactions. The same pattern holds for both inflows and outflows. For inflows (Panel B), anonymous volume rises by 0.773 versus 0.220 for identified transactions. For outflows (Panel C), anonymous volume increases by 0.776 versus 0.245. 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 using fragmented transfers and intermediary wallets to reduce traceability.

Table 1 also reveals that aid disbursements 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 1.387, compared with 0.864 on tax-haven exchanges. Identified transaction volume also rises substantially, with a coefficient of 0.622 on non-tax-haven exchanges versus 0.301 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 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.7 shows that transaction frequency follows the same pattern. On tax-haven exchanges, the arrival-month impulse for anonymous transaction frequency is 0.888, compared with 0.327 for identified

transactions in Panel A. For inflows (Panel B), anonymous frequency rises by 0.799 versus 0.266 for identified transactions. For outflows (Panel C), anonymous frequency increases by 0.771 versus 0.239. On non-tax-haven exchanges, anonymous frequency is again larger in absolute magnitude, reaching 1.306 in Panel A, compared with 0.628 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 Table A.10 shows 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 0.864, compared with 0.301 for identified transactions, while the corresponding frequency coefficients are 0.888 versus 0.327. The baseline transaction-volume and transaction-frequency coefficients are reported in Tables 1 and A.7. Likewise, when neighboring countries are excluded from the control group, the corresponding estimates remain very similar: 0.863 versus 0.300 for volume and 0.887 versus 0.326 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.

The inference is not sensitive to how standard errors are computed. Treatment is assigned at the recipient-country level, so we cluster there throughout; clustering two-way by recipient country and calendar month instead, to allow for arbitrary within-month correlation across countries, raises the standard error on the anonymous arrival-month impulse only slightly, by about four percent on tax-haven and eleven percent on non-tax-haven exchanges, and the effect remains significant at the one-percent level. As a design-based check that does not rely on an asymptotic variance at all, we also run a randomization-

inference test: holding the panel fixed, we permute each treated country’s arrival month within its own event window and re-estimate the arrival-month impulse, repeating this 2,000 times to build the distribution of impulses expected when the timing of disbursements carries no information. The observed anonymous impulse lies beyond every one of the 2,000 placebo draws on both platform types (the largest tax-haven placebo impulse is 0.54, compared with the observed tax-haven anonymous impulse of 0.864 in Table 1), for a randomization p -value of 0.0005.

4.4.2 New Wallet Creation

We next examine whether aid disbursements are followed by new wallet creation, 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 2 reports the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) by exchange jurisdiction and anonymity status using a stacked event-study difference-in-differences on the flow-based tranche-arrival panel with linear TWFE and IHS-transformed outcomes.

The most striking feature of Table 2 is the near-complete concentration of new wallet creation in anonymous accounts. On tax-haven exchanges, the arrival-month impulse for anonymous new accounts is 1.775 IHS units in Panel A and highly significant across all three panels, while the corresponding effect for identified accounts is much smaller and never statistically significant: 0.044 in Panel A, 0.005 in Panel B, and 0.033 in Panel C. The anonymous and all-account coefficients are numerically identical on tax-haven exchanges (1.775 in Panel A, 1.685 in Panel B, 1.636 in Panel C), and the reason is worth stating plainly: anonymous wallets are 99.9997 percent of new-account mass on tax-haven exchanges and 99.9991 percent on non-tax-haven exchanges, so the All series is numerically the Anonymous series and the two regressions return the same coefficient and the same standard error. This reflects how wallets are attributed rather than an estimation artifact. Attribution resolves clusters to named services, so a newly observed address inside a service’s cluster is not a new entity, and newly created wallets are almost never attributable; the identified new-account series is non-zero in only 661 of 5,178 tax-haven country-month cells. We therefore read Table 2 as establishing that essentially all post-disbursement entry into the crypto ecosystem runs through wallets that resolve to no named service, and we do not interpret the ratio of the anonymous to the identified coefficient as a magnitude. The anonymity contrast that carries economic content is the one in Table 1, where identified activity accounts for 9 percent of tax-haven and 10 percent

of non-tax-haven transaction volume and both columns are estimated on well-populated series.

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 2.832 for aggregate anonymous accounts (Panel A), 2.739 for anonymous inflow accounts (Panel B), and 2.643 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 comparatively small: 0.260 in Panel A, 0.198 in Panel B, and a statistically insignificant 0.073 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 2.6 to 2.8 IHS units versus 1.6 to 1.8 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 1 and A.7, 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.

4.5 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 percent 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, $\text{ihs}(y) = \ln(y + \sqrt{y^2 + 1})$, 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 use anonymous transaction volume on tax-haven exchanges (Table 1, Panel A), which captures funds as they pass through lightly regulated platforms during the layering stage of the laundering cycle documented in Section 4.7. Because a disbursed tranche is a one-time stock, we adopt a cumulative accounting convention: the numerator sums the implied level change over the full post-disbursement window $k = 0, \dots, 4$ rather than counting a single month, and we compute both the event-study coefficient path and the baseline on the same stacked estimation panel, so that the magnitude is internally consistent with the regression in Table 1. The pre-disbursement baseline, the average geo-adjusted anonymous tax-haven volume over $k \in \{-4, -3, -2\}$ for treated cells on the estimation panel, is $\bar{y}^0 \approx 1,068$ BTC. The response is sharply concentrated at impact ($\hat{\beta}_0 = 0.864$) and decays over the following months, so the cumulative level change is $\sum_{k=0}^4 (e^{\hat{\beta}_k} - 1) \bar{y}^0 \approx 1,237$ BTC. At the sample-period average Bitcoin price of \$21,550, this is approximately \$26.7 million in additional anonymous tax-haven throughput per treated tranche.⁷ Against the mean true tranche flow of \$238 mil-

⁷The sample-period average is the simple mean of monthly average Bitcoin prices, each month's average taken over daily closing prices from CoinMarketCap, computed over April 2018 through August 2023. The daily series terminates on 25 August 2023, so the average excludes the final seven months of the sample period (September 2023 through March 2024), over which Bitcoin traded well above this level. Valuing the response at a higher average price would raise the dollar figure and hence the leakage ratio, so this convention is the conservative one. Valuing each treated episode at its own arrival-month price instead leaves the ratio essentially unchanged: extending the CoinMarketCap series over its missing tail (September 2023 through February 2024) with Yahoo Finance BTC-USD daily closes, which agree with CoinMarketCap to about 0.01 percent on overlapping months, the baseline-weighted mean arrival price is \$22,328, within 4 percent of the flat constant because only about 5 percent of the pre-period baseline weight falls in the post-August-2023 high-price months, and the gross ratio becomes 11.6 percent. The magnitude is therefore not driven by the timing of tranches relative to the Bitcoin price cycle.

lion over the same treated episodes, the implied gross leakage ratio is

$$\lambda_{\text{gross}} = \frac{\Delta \text{CryptoValue}}{\text{Aid}} \approx 11.2\%.$$

The split between flat and per-episode pricing is immaterial across every baseline and accumulation convention. Under the headline dip baseline ($k \in [-4, -2]$), where the baseline-weighted mean arrival price is $\bar{P}^w = \$22,328$, the gross leakage ratio is 11.2% under the flat constant versus 11.6% per episode on the cumulative $k=0 \dots 4$ response, and 13.3% versus 13.7% on the arrival-month impulse ($k=0$). Anchoring instead on a cleaner, more distant pre-window ($k \in [-8, -5]$, $\bar{P}^w = \$21,107$) raises the level of the ratio but leaves the flat/per-episode gap negligible: 15.4% versus 15.1% cumulatively and 18.2% versus 17.8% on impulse. The two pricing conventions never diverge by more than half a percentage point, confirming that the estimated leakage is not an artifact of when tranches fall in the Bitcoin price cycle. Prices are CoinMarketCap daily-close monthly means through August 2023—the series that defines the flat constant—extended over September 2023 through February 2024 with Yahoo Finance BTC-USD daily-close monthly means, the two agreeing to about 0.01% on overlapping full months.

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. Writing $m \geq 1$ for the number of times a diverted dollar is observed, net leakage is

$$\lambda_{\text{net}}(m) = \frac{\lambda_{\text{gross}}}{m},$$

a rectangular hyperbola with constant elasticity -1 . We do not estimate m : the exchange ledgers we observe are disconnected around each entity and carry no edges along the laundering path, so in-sample turnover is not recoverable. Figure 5 therefore reports the entire curve rather than a point estimate.

Two features of the curve carry the magnitude claim. First, because layering can only inflate observed volume relative to diverted principal, $m \geq 1$ holds by construction, so the gross ratio of 11.2 percent is a sharp upper bound on net leakage, attained at the no-layering anchor $m = 1$ and requiring no assumption about turnover. Second, net leakage falls below the 7.5 percent offshore-banking benchmark of [Andersen et al. \(2022\)](#) once m exceeds $m^* = 1.49$, so crypto is the larger channel only under near-zero layering. Figure 5 plots the curve under both conventions for λ_{gross} , the cumulative response over

$k = 0, \dots, 4$ used throughout and the arrival-month impulse at $k = 0$, whose gross ratio is 13.3 percent; the band between them is narrow relative to the vertical distance either curve travels in m , so turnover, not the summation window, governs the magnitude. The Chainalysis *Crypto Crime Report* describes illicit funds passing through two to five intermediary wallets before exiting (Chainalysis, 2021), which read off the cumulative curve would place net leakage between 5.6 percent and 2.2 percent. The 328 treated episodes carry \$78.2 billion in disbursements, so that range corresponds to approximately \$4.4 billion at $m = 2$ and \$1.7 billion at $m = 5$; at the assumption-free anchor $m = 1$ the bound is \$8.7 billion. These figures apply to the tranche arrivals on which the ratio is estimated, which are one third of disbursement flow to sample countries, and we do not extrapolate the rate to the remaining two thirds. That is forensic guidance rather than a moment of the turnover distribution, so we impose no value of m and report the curve rather than any point on it. What the design delivers without any turnover assumption is the bound: net crypto-mediated leakage is at most approximately 11 percent of disbursed aid. This bound is comparable in order to the 7.5 percent 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.

Finally, our leakage estimates capture only the Bitcoin channel and therefore represent a lower bound on total crypto-mediated diversion. 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.

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 pro-

vide independent confirmation that observed blockchain flows are tied to real economic activity. Because the off-chain data are reported only at quarterly frequency, the monthly stacked event-study design used throughout the paper cannot be applied at this frequency; this section accordingly serves as descriptive corroboration that the on-chain flows correspond to real exchange activity, not as an independent causal estimate.

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 3, 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 percent increase in foreign aid is associated with a significant increase in both volume growth, of 2.39 pp, and the volume ratio, of 0.39 pp, each at the 1 percent level. Non-haven exchanges show an insignificant volume response and a near-zero volume ratio, with the cross-group difference (Δ) significant at the 1 percent level for both measures. The regulated-versus-unregulated and crypto-friendly-versus-unfriendly splits yield nearly identical patterns: the volume-ratio response is concentrated in unregulated jurisdictions, more than ten times the regulated-country effect, and crypto-friendly jurisdictions show a significant 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 3, 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 effect in low-disclosure countries is significant at the 1 percent level and an order of magnitude larger than the point estimate in high-disclosure countries, which is insignificant. For corruption control, the below-median-group coefficient is significant while the above-median estimate 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 4 isolates this effect by comparing transaction activity across five offshore financial centers (Hong Kong, Singapore, Seychelles, the Cayman 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 is positive for transaction counts and positive and significant at the 5 percent level for the volume ratio, 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 Huobi case also points to a concrete enforcement margin. Although blockchain transfers can pass through anonymous wallets, fiat-to-crypto entry and crypto-to-fiat exit usually require exchange accounts, banking relationships, and Know-Your-Customer records ([Financial Action Task Force, 2021](#)). On- and off-ramping activity is therefore the most detectable part of the sequence. Stricter monitoring of large post-disbursement deposits, withdrawals, and internal transfers at major exchanges, especially exchanges that relocate across regulatory regimes or operate in crypto-friendly offshore jurisdictions, would target the point where identities are most likely to be recorded. The Huobi relocation illustrates why enforcement should follow exchange entities as well as exchange domicile: the same platform can move across jurisdictions while continuing to intermediate aid-linked 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.

4.7 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 wallets that resolve to no named service 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.

The one-to-two-month persistence of the spike is consistent with the operational frictions of laundering at scale. A large and unexpected inflow cannot be anonymized instantaneously without creating a visible transaction cluster. Actors therefore have incentives to split funds across wallets, wait between

transfers, and route balances through mixers, intermediary wallets, and lightly regulated platforms in batches. These services also face finite liquidity and processing capacity: when a large transfer wave arrives, mixing and redistribution take time because funds must be pooled with other flows before they can be withdrawn with a weaker on-chain link to the source. This batching logic matches the event-study pattern: the response begins immediately after disbursement, remains elevated while funds are fragmented and relayed through intermediaries, and fades once the layering sequence is completed ([Financial Action Task Force, 2021](#); [Chainalysis, 2021](#)).

Integration (Realization or Storage). The final stage involves either converting crypto holdings back into fiat currency to re-enter the formal financial system or retaining them as digital assets for long-term storage. Although our data does not directly capture the ultimate disposition of funds, 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 evidence points to attenuation rather than a permanent reconfiguration of flows.

5 Robustness Checks and Extensions

This section presents a series of robustness checks and mechanism tests organized into three tiers. We begin with three tests of identification validity: the Pandora Papers natural experiment, which tests whether an exogenous transparency shock differentially affects anonymous versus identified channels; a placebo built on disaster shocks without disbursements, which rules out demand-driven explanations; and a sensitivity analysis over the thresholds that define tranche-arrival spikes. We then present three

checks of measurement validity: an analysis of VPN-mediated geographic mismeasurement, which signs the direction of any plausible bias and bounds its magnitude by excluding high-VPN recipients and interacting the response with VPN intensity; 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 probe the mechanism and specification: a service-level mixer crowd-out test showing that the response concentrates in privacy-oriented activity rather than a broad post-aid crypto boom, and a specification check excluding the May to September 2021 DeFi boom window. Additional robustness results, including alternative control group compositions, on-chain outcomes without web-traffic adjustment, and governance heterogeneity, 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.

5.1 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. Leaked offshore records have served as identifying variation for hidden elite wealth before ([Londoño-Vélez and Ávila-Mahecha, 2021](#)); we extend this logic from banking structures to on-chain anonymity.

To test whether the baseline aid-crypto relationship changed after this transparency shock, we estimate a pooled stacked specification $y \sim \text{DisbPost} \times \text{PandoraPost} \mid \alpha_i + \gamma_t$, where DisbPost equals one in post-disbursement months for treated episodes and PandoraPost indicates calendar months on or after October 2021. The interaction coefficient Δ captures the change in the aid-crypto relationship following the leak.

Table 5 reports the results. The most striking finding is that the post-Pandora attenuation is concentrated in identified rather than anonymous transactions. For transaction volume, the interaction term Δ for identified activity is negative on both platform types and significant on tax-haven exchanges, at -0.531 , though the non-tax-haven counterpart of -0.469 does not reach conventional significance,

while the corresponding coefficients for anonymous volume are small and statistically insignificant. New-account creation shows the same asymmetry even more sharply: the identified interaction is again negative and significant on both platform types, at -0.275 and -0.340 respectively, whereas anonymous new accounts show no significant change. For transaction frequency the identified interaction is likewise negative on both platform types but does not reach conventional significance, and anonymous frequency again shows no change. Across all three outcomes the sign pattern is identical (identified activity attenuates after the leak while anonymous activity does not), and it is statistically significant for the two outcomes, volume and new accounts, that directly capture value movement and entry. Figure 4 illustrates this asymmetry by overlaying the pre- and post-Pandora event-study paths within each panel: while identified transaction volume and new-account creation exhibit a pronounced post-disbursement response that attenuates after the Pandora leak (the post-Pandora arrival-month impulse collapses toward zero), 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.

One might object that this asymmetry is a mechanical artifact of how the anonymous and identified series are constructed rather than genuine evidence about anonymity-seeking behavior. We meet this objection with a level-split placebo. We discard the true anonymous and identified labels and instead partition the same global on-chain total into two synthetic series by a fixed fraction that matches the observed anonymous share but contains no genuine anonymity information, then re-estimate the identical Pandora triple difference. Were the observed asymmetry an artifact of the split, this placebo would reproduce it. It does not. In a specification that permits arbitrary relabeling of the split, the true labels deliver an identified interaction of -0.531 on tax-haven exchanges and -0.469 on non-tax-haven exchanges, as reported in Table 5. The tax-haven estimate is significant at the five-percent level, while the non-tax-haven

estimate is negative but not statistically significant at conventional levels. Against anonymous interactions of -0.095 and 0.275 , neither significant, this implies an identified-minus-anonymous asymmetry of -0.436 and -0.744 that matches the table’s core pattern: identified activity falls more after Pandora, while anonymous activity does not show a comparable decline. The level-matched placebo instead yields an asymmetry of $+0.04$ and $+0.01$, statistically zero and if anything of the opposite sign, so the mechanical split reproduces essentially none of the effect. The finding is stable across a ladder of fixed fractions from 0.09 to 0.91: the placebo interaction is flat in the split fraction and never approaches the identified coefficient. The anonymity asymmetry is a property of the data, not of the construction.

A second concern is that October 2021 sits amid a turbulent period for crypto markets: China’s ban on Bitcoin mining and the associated hash-rate migration peaked in mid-2021, the market reached its cyclical price high that November, and the Terra and FTX collapses followed through 2022. Could the post-Pandora break simply capture one of these confounds? It cannot, for two reasons. First, the estimand is a triple difference with month fixed effects, so any shock common to all countries and platforms in a given calendar month, including the mining ban, the price peak, and the subsequent crashes, is absorbed and cannot load on the interaction; only variation that is both differential between post-disbursement and control episodes and differential across the October 2021 break survives. Second, and decisively, these macro shocks are anonymity-neutral.⁸ A hash-rate migration relocates where blocks are produced and a price crash rescales aggregate volume, but neither raises the cost of routing funds through traceable, KYC-linked channels relative to pseudonymous ones, and neither does so selectively in the country-months that follow an aid tranche. To mimic our result a confound would have to attenuate identified post-disbursement flows while leaving anonymous post-disbursement flows intact, precisely the differential that a transparency leak targeting named offshore structures produces and that a market-wide repricing or a recomposition of mining does not. The mining ban in particular changes the geography of supply, not the traceability of demand, so it generates no disbursement-interacted anonymity asymmetry.

This finding also provides indirect support for the causal interpretation of the baseline results. The

⁸Bitcoin transaction fees also vary substantially over our sample, but they are levied per transaction on the size of the transaction in bytes rather than on the value transferred, so they are negligible relative to a large transfer even at fee peaks: the median anonymous tax-haven transfer in our panel is about 1.1 BTC (roughly \$24,438 at the \$21,550 sample-period average price used in Section 4.5), and the peak monthly-average fee over the sample (\$31 in April 2021) is about 0.13 percent of that median; a \$1 million transfer moved as a single transaction would pay the same \$31 (0.003 percent). They are also common to all countries within a calendar month and therefore absorbed by the month fixed effects, and they do not differ between traceable and pseudonymous counterparties, so they cannot produce the anonymity-differential response we document. Fees would matter only if diversion were fragmented into transfers small enough for a per-transaction charge to bind.

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.

5.2 Disaster-Without-Disbursement Placebo

A natural concern is that the observed cryptocurrency activity following aid disbursements reflects a demand-side response to the shocks that can prompt aid, rather than diversion of aid funds. Under this alternative, households and firms convert local assets into cryptocurrency to hedge against disaster-induced macroeconomic instability or political risk, behavior that may coincide with disaster shocks but operates independently of aid inflows.

The timing profile of the baseline results provides initial evidence against this interpretation. For the majority of disbursements in our sample, which are driven by development policy reforms and investment project milestones, the lag between the triggering event and the arrival of funds is typically several months to over a year. If hedging behavior were driving the crypto response, activity should rise during this extended pre-disbursement window as elites anticipate the associated policy changes. Instead, the event-study estimates show flat pre-trends and a sharp spike at the disbursement month, consistent with diversion, which requires the physical arrival of funds before they can be captured, and inconsistent with shock-driven demand, which responds to the uncertainty shock rather than the liquidity event.

For disaster-related disbursements, where the lag between shock and disbursement is shorter, we provide a direct placebo test. 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. The placebo cohort is, if anything, more crypto-capable than the baseline recipients rather than less: its members are richer (median GDP per capita of roughly \$17,700 against \$3,800), more digitally connected (median internet penetration of 85 against 61 percent), and show higher baseline crypto engagement (median share of web traffic directed to exchanges of 0.040 against 0.025 percent), so any disaster-induced shift into cryptocurrency should surface more strongly here, and the flat placebo is a conservative test rather than an artifact of a less crypto-active sample. Table 6 reports the stacked TWFE estimates.

The results are unambiguous. The placebo estimates do not replicate the positive, significant baseline effects: coefficients are small and insignificant across every outcome and exchange type. Arrival-month impulses on anonymous transaction volume are only 0.087 and 0.153 IHS units on tax-haven and non-tax-haven exchanges, and the corresponding figures for transaction frequency and new account creation are similarly small; none of the eighteen coefficients attains even ten-percent significance. Disaster exposure unaccompanied by a World Bank disbursement produces no detectable change in on-chain crypto activity.

Statistical insignificance alone would leave the placebo open to the charge of being underpowered, so we make the power claim directly. If anything the placebo is better powered than the baseline: it draws on 129,381 country $\times$ platform $\times$ episode $\times$ month cells against 18,499 in the baseline stack, because it requires only the absence of a disbursement and so admits roughly sixty control countries per episode against the baseline's three, tightening standard errors in a direction that works against the null. Yet the placebo intervals still exclude the baseline effects they are meant to reproduce. For anonymous transaction volume the impulse of 0.087 has a 95 percent interval from -0.14 to 0.31 on tax-haven exchanges and from -0.15 to 0.46 off them, both below the baseline estimates of 0.864 and 1.387, and new accounts tell the same story, with intervals well short of the baseline impulses of 1.775 and 2.832. Wherever the baseline effect is nonzero, the upper limit of the placebo interval is between a quarter and a third of it, so the placebo would have detected an effect a third the size of the one we report. The null is therefore informative rather than inconclusive.

This clean null directly contradicts the demand-side interpretation. If households and firms 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 leave crypto activity unchanged, implying that the disbursement itself, rather than disaster-induced uncertainty, is the necessary condition for the observed crypto response. The placebo also addresses related macroeconomic channels: currency devaluation and broader economic deterioration plausibly follow natural disasters, yet these conditions are equally present in the placebo sample, where crypto activity does not move. 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.

Finally, the null placebo confirms that the baseline estimates are not confounded by disaster exposure. In the baseline specification, the control group includes country-episodes that experienced natural disasters without subsequent disbursements. Because disasters without aid generate no detectable crypto response, these control-group observations neither inflate nor deflate the counterfactual, so the estimated treatment effect isolates the disbursement itself rather than any disaster-induced shift in crypto activity. This reinforces the interpretation that the estimated leakage reflects aid diversion specifically, not a generic response to crisis conditions.

5.3 VPN-Mediated Geographic Mismeasurement

A further concern with our geographic attribution is that recipient-country users may mask their location through commercial VPN services, causing crypto-exchange traffic from country C to be re-attributed to other jurisdictions. We address the concern in two parts: a structural argument about the direction of any plausible bias, and direct empirical tests of its magnitude.

Two routing scenarios are theoretically conceivable. In the first, treated-country users route through other recipient countries that serve as controls for that episode; the leakage contaminates the control mean upward at the treatment month, and the contrast attenuates from both sides. In the second, users in non-recipient countries route through exit nodes located in the treated country, which would inflate C 's measured activity at event time and bias $\hat{\beta}$ upward. The first scenario attenuates $\hat{\beta}$ toward zero and is therefore conservative; the second is the one we must rule out.

The reverse-routing scenario is not supported by the empirical literature on commercial VPN services. [Khan et al. \(2018\)](#) document, in a study of 62 commercial VPN providers, that exit nodes cluster heavily in the United States and Western Europe, with limited and uneven presence in low- and middle-income countries that comprise the bulk of our aid-recipient sample. On the demand side, [Dutkowska-Zuk et al. \(2022\)](#) survey 729 U.S. VPN users and find that the dominant motivations (privacy, security, and circumvention of geographic content restrictions) generate traffic that routes away from the user's home toward developed-country or content-rich jurisdictions, not toward aid recipients. The multi-perspective study by [Ramesh et al. \(2023\)](#) on 1,252 U.S. VPN users reaches the same conclusion. Even granting some baseline rate of reverse routing, a stationary routing share is absorbed by the country and time fixed effects; for $\hat{\beta}$ to be biased the share would have to spike specifically at the disbursement month,

which would require control-country users to coordinate their VPN routing with World Bank disbursement calendars. There is no documented behavioral mechanism that would produce such a correlation. By contrast, the first scenario, in which a treatment-induced surge in C leaks via VPN to other panel countries, is mechanically consistent with our specification and operates in the conservative direction.

We provide three complementary exclusion-based tests of how much VPN-affected traffic could plausibly be driving the result. Table 7 reports all three. In the first, we drop Indonesia, Turkey, and Pakistan, three countries consistently flagged in the policy and industry literature as having the highest commercial VPN penetration among our recipient sample. In the second, we drop the top decile of recipient countries by Google Trends search intensity for the term “VPN” over the sample period (90th-percentile cutoff of 14.8 on the 0–100 scale; six countries excluded). The descriptive ranking underlying this construction is reported in Appendix Table A.5.

The manual list aggregates heterogeneous industry and survey figures, and Google Trends search intensity is a relative attention index rather than an adoption rate. To complement them with a consistent survey-based measure, we use GlobalWebIndex (GWI) data on the share of internet users aged 16–64 who report using a VPN or proxy server, as compiled in the *Digital 2026 Global Overview Report* by DataReportal and We Are Social (Q2 2025 survey wave). The GWI series applies the same instrument and definition across covered countries. Our third rule drops the top decile of recipient countries by GWI VPN adoption. GWI covers 20 recipient countries; we compute the cutoff over that set and retain uncovered recipients, which are typically smaller economies.

If VPN-driven mismeasurement were generating the baseline result, all three exclusions should attenuate the estimates substantially. The opposite occurs. Relative to the baseline arrival-month impulse on anonymous tax-haven transaction volume, each of the three exclusions (the manual VPN list, the Google Trends top decile, and the top decile by GWI adoption) leaves the $\text{Treat} \times \text{Post}$ coefficient essentially unchanged or slightly larger. The same pattern holds across transaction frequency and new-account creation, and across both tax-haven and non-tax-haven platforms. Dropping high-VPN countries does not weaken the aid-crypto response; if anything, it slightly strengthens it.

A third specification, reported in Appendix Table A.6, interacts $\text{Treat} \times \text{Post}$ with VPN intensity directly, both as a discrete top-half indicator (HighVPN) and as a continuous standardized Google Trends score (VPNScore). Neither interaction is statistically significant for any outcome or platform type: the point estimates are small relative to the arrival-month impulse and imprecisely estimated. Thus, there is

no evidence that higher recipient-country VPN intensity strengthens the aid-crypto response, which is the positive gradient that VPN-mediated routing would have to generate to bias the baseline estimate.

5.4 Mixer Crowd-Out

The post-disbursement evidence shows that the mixer response is not a broad crypto-service boom. After aid is disbursed, mixer activity rises relative to payment services, centralized exchanges (CEXs), decentralized exchanges (DEXs), and gambling services, suggesting the response is concentrated in privacy-oriented activity rather than in general on-chain use.

Appendix Table A.11 compares CoinJoinMess mixer transaction counts with those of other service types using no-overlap stacked TWFE, with standard errors clustered at the calendar month level. If aid disbursement simply raised all types of crypto activity, mixer transactions should increase similarly to payments, CEXs, DEXs, or gambling services. The opposite occurs. Mixer activity rises sharply relative to payment services: the gap is positive in the disbursement month, grows further by $k = 1$, and remains positive and significant at the 1 percent level over the full post window, where the gap reaches 0.642 IHS units. The corresponding gaps are also positive and significant at the 1 percent level, at 0.410 IHS units for CEXs and 0.622 for DEXs. Mixer activity does not rise detectably relative to gambling. Gambling platforms are themselves widely used as informal mixing venues (bettors deposit and withdraw to break the on-chain link between funds), so the absence of a significant mixer-minus-gambling gap after disbursement is expected: both categories serve the same anonymity-seeking demand. The broader pattern shows that the response is concentrated in privacy-oriented activity rather than in general on-chain use.

This pattern addresses several alternative interpretations. First, the service comparison rules out the view that the result is only a broad post-aid crypto boom: mixer activity rises more than ordinary crypto services. Second, the payment-service comparison speaks to payment adoption or general liquidity shocks. If aid disbursement simply increased recipients' use of crypto payment rails, payment services should rise similarly; they do not. Third, the stronger mixer response relative to CEXs and DEXs suggests that the pattern is not only users moving funds through exchanges after receiving aid.

5.5 Sensitivity to the Tranche-Spike Thresholds

Because treated episodes are defined by a threshold rule (a monthly disbursement flow of at least \$50 million and at least three times the country’s trailing twelve-month median, with no qualifying spike in the preceding four months), a natural concern is that these cutoffs are arbitrary and that the results could hinge on them. Table 8 addresses this directly. It re-defines episodes across a 3×3 grid of absolute (\$25M, \$50M, \$100M) and relative ($2\times$, $3\times$, $5\times$) thresholds, spanning 207 to 461 treated episodes, and re-estimates the baseline arrival-month impulse on each resulting panel. The anonymous impulse is positive and significant at the one-percent level in every cell, for both transaction volume and new-account creation and on both exchange types, and the anonymous-to-identified gap is preserved throughout (identified impulses, omitted for space, remain small and mostly insignificant). The clean pre-period ($k \in \{-8, \dots, -5\}$) stays jointly insignificant across the grid. The headline results are therefore not an artifact of the particular thresholds used to isolate tranche-arrival spikes; if anything, stricter thresholds, which retain only the largest and most clearly anomalous inflows, yield larger point estimates.

5.6 Validation of Geo-Adjusted Measures

5.6.1 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.⁹ We collapse our proxy to country-year means and merge with the publicly released Chainalysis rankings, yielding 81 matched observations across approximately 30 countries. We quote a single family of statistics throughout: PPP-weighted Pearson correlations estimated year by year. These are 0.689 in 2021, 0.518 in 2022, 0.591 in 2023, and 0.855 in 2024, significant in every year after 2020. The correspondence is stronger on non-tax-haven platforms, consistent with Chainalysis measuring mainstream adoption rather than the anonymity-seeking activity that predominates offshore. The pooled specification, which blends the noisy inaugural 2020 index with the stronger later years and is correspondingly more conservative, is reported in Appendix C.2. These correlations

⁹Details on the Chainalysis index methodology, the rank-to-score conversion, and year-by-year results are in Appendix C.2.

confirm that the web-traffic-based geographic attribution recovers meaningful cross-country variation in cryptocurrency activity independently of our identification strategy.

5.6.2 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 9 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 arrival-month impulse is 1.606 for anonymous transactions on tax-haven exchanges and 1.913 on non-tax-haven exchanges; for transaction frequency, 1.629 and 1.829, respectively. New anonymous account creation shows the largest effects at 2.909 on tax-haven and 3.338 on non-tax-haven exchanges. Identified transactions are positive and significant for volume and frequency but far smaller for new accounts, where the tax-haven impulse of 0.202 is statistically insignificant, preserving the sharp anonymity gap observed in the baseline. If anything, the duration-ratio magnitudes are somewhat larger than the visit-share impulses rather than attenuated, but 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.

5.7 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 to September 2021 window from both the treatment and control samples. Table A.12 in Appendix E.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 0.931 compared with 0.864 in the baseline), confirming that the 2021 DeFi boom does not drive our findings.

6 Policy Implications

This section examines whether the crypto-capture signal is concentrated in particular sectors, documents whether those sectors continue to receive World Bank money after the sample window closes, and then discusses the implications for fiduciary supervision.

6.1 Which Sectors Capture More

The baseline result identifies an average aid–crypto response across all disbursements, but the underlying mechanism predicts different responses for different kinds of aid. Aid that flows to beneficiary households, public-sector salaries, or community-level grants creates distributed liquidity that can be converted into cryptocurrency by many small actors. Aid that flows through procurement contracts to a small number of firms creates concentrated liquidity, in a form less suited to immediate on-chain conversion. The classical aid-capture literature emphasizes the latter channel, in which elite capture is most visible in large infrastructure and energy contracts. The project-type breakdown in Table A.13 tests which mechanism the on-chain response is consistent with.

We classify the universe of World Bank projects approved in fiscal years 2018–2024 into eleven sector categories on the basis of the project title, and tag each of the 328 spike episodes by its dominant project sector, the sector receiving the largest IBRD/IDA disbursement flow in the episode’s arrival month. Appendix F sets out the classification procedure and reports the full set of category-by-outcome estimates. Under this tagging, Environment, Climate, and Disaster contains no spike episodes, and Digital and Communications rests on only three episodes and should be read with caution.

Table A.13 reports the $k = 0$ arrival-month impulse, estimated separately within each stratum on that category’s subset of spike episodes. On the paper’s headline outcome, anonymous transaction volume on tax-haven exchanges, the capture ranking is led by Transport, whose arrival-month impulse of 2.676 is the largest of any sector and strongly significant, followed by Water and Sanitation at 1.528, Social Protection and Jobs at 1.518, and Governance and Public Sector at 0.743, each statistically significant. The coefficients on Finance and Private Sector, Energy, Agriculture and Rural Development, Education, and Digital and

Communications are not distinguishable from zero, and Health is weakly negative. New-account creation tracks the same ordering, with Transport again largest at 5.491, and the non-tax-haven columns are broadly similar, with Digital and Communications and Energy turning positive and significant on non-tax-haven platforms. The response is therefore concentrated in a handful of categories rather than declining monotonically across the full ranking. A natural concern with cross-category comparisons is that categories receiving larger tranches may produce mechanically larger on-chain responses. Two facts weigh against this reading. First, category explains little of the variation in tranche size: regressing the log arrival-month disbursement flow on the category indicators yields an R^2 of 0.097 (adjusted 0.071), so more than ninety percent of that variance is within rather than between categories. Second, the Spearman rank correlation between a category's mean arrival-month flow and its impulse is $\rho = 0.21$ ($p = 0.56$), indistinguishable from zero. The extremes make the point: Transport has the largest impulse, of 2.676, and the smallest mean tranche, of \$97.6 million, while Finance and Private Sector has the largest mean tranche, of \$328.2 million, and an impulse of 0.913 that is not distinguishable from zero.

The two ends of the ranking correspond to the two canonical leakage channels. On the procurement side, where [Olken \(2007\)](#) measures roughly 24 percent missing expenditures in Indonesian road projects and [Bandiera et al. \(2009\)](#) document similar waste in public purchasing, Transport, a category of large and discretionary construction contracts, tops the ranking. On the distributive side, where [Reinikka and Svensson \(2004\)](#) find only 13 percent of a Ugandan capitation grant reaching schools absent a hardened payment rail ([Muralidharan et al., 2016](#)), Social Protection, Water and Sanitation, and wage-heavy Governance fill ranks two through four, echoing the household transfers of the cash-transfer general-equilibrium literature ([Egger et al., 2022](#)). Where that literature measures leakage one program and channel at a time in program-specific units, the arrival-month impulse gives a single comparable metric across all sectors, and it locates the exception where the mechanism predicts: Health, disbursed heavily in-kind, creates no local liquid windfall and shows no on-chain response. The ranking does not adjudicate cleanly between recipient-side and procurement-side diversion; both channels leave an on-chain footprint, whereas capital-intensive Energy, where capture may remain in domestic cash and physical assets, does not.

6.2 Continuing World Bank Exposure

Table 10 joins each sector's capture intensity to three post-window funding measures. Legacy-book disbursement (April 2024 to June 2026) is gross cumulative disbursement from the live IBRD/IDA Statement of Loans and Credits, net of the March 2024 baseline; it is money that actually moved, but it is largely predetermined by pre-2024 commitments. Newly approved commitments (board approvals from April 2024 to June 2026, taken from the World Bank Projects API) are the Bank's forward-looking allocation choice after the study window closed. Approval-rate acceleration divides annualized new commitments after the window by the annualized in-sample rate, removing sector scale. The four sectors whose capture impulse is significant at the 10 percent level absorb 48 percent of sector-classified post-window disbursement (\$260 billion of \$538 billion) and 48 percent of sector-classified newly approved commitments (\$108 billion of \$224 billion), across 546 of 1,136 classified new board approvals. Their share of annualized commitments falls from 55 percent in sample to 48 percent after the window. We report this decline as a level rather than test it, and we draw no inference from its sign: on either side of it, close to half of the Bank's continuing money flows to sectors where capture is measurably present. The four-sector composition is moreover robust to tightening the significance threshold to 5 percent: all four impulses remain significant at that level, including Water and Sanitation, the marginal sector at $p = 0.048$, so the disbursement and commitment shares are unchanged. The sector map covers 98 percent of post-window disbursed dollars, 95 percent of newly committed dollars, and 92 percent of new approvals, and recomputing the two shares on unclassified-inclusive denominators gives 48 percent and 46 percent, so nothing rests on the choice of denominator. Figure 6 visualises the join: the upper-right quadrant of Panel B, with high measured capture and high continuing funding, contains Transport and Governance and Public Sector.

We do not claim that the Bank rewards leaky sectors, and we place no weight on the cross-sector correlation between capture intensity and subsequent funding. The supervisory case rests on exposure: how many dollars continue to flow into sectors where capture is measurably present, not on any claim about the Bank's allocation rule. That exposure has a country-level precedent: corrupt governments do not receive less aid ([Alesina and Weder, 2002](#)) and disbursement conditionality is enforced unevenly ([Kilby, 2009](#)), so supervisory intensity is a choice variable rather than a constant. And supervision works when applied: raising the audit probability from about 4 percent to 100 percent cut missing expenditures by

roughly eight percentage points (Olken, 2007), project-level supervision quality is a robust micro correlate of World Bank project outcomes (Denizer et al., 2013), and hidden offshore balances relocate when scrutiny tightens (Johannessen and Zucman, 2014).

6.3 Implications for Supervision

World Bank fiduciary supervision is organized around project characteristics and country/institutional risk assessments, including procurement and financial-management risk (World Bank, 2024). The heterogeneity results supply a third dimension that neither captures: the measured propensity of a disbursement in a given sector to surface as anonymous, tax-haven, on-chain volume in the month it arrives. Three implications follow.

First, Transport is the priority case. It carries the largest capture impulse of any sector on both volume and new accounts, and it is the only sector whose post-window commitments (\$35.8 billion) exceed its entire six-year in-sample total (\$29.9 billion); on an annualized basis, its approval rate accelerates by a factor of 3.2, the steepest of any category. The highest measured capture therefore coincides with high and rising exposure, in precisely the contract-intensive category where the procurement-leakage literature locates diversion.

Second, supervision should bind at disbursement rather than at approval. The capture signal is an arrival-month impulse: it appears at $k = 0$, in the month the tranche lands, a window that ex-ante appraisal and ex-post audit both miss. The instruments this implies are disbursement-linked: staged tranche release, contemporaneous transaction monitoring, and beneficiary verification timed to the arrival month.

Third, distributive and procurement aid call for different instruments. Capture in Social Protection and in Water and Sanitation is consistent with many small actors converting distributed liquidity, which calls for payment-rail integrity, the digital-identity and beneficiary-reconciliation margin that Muralidharan et al. (2016) show to be effective. Capture in Transport is consistent with concentrated, contract-level diversion, which calls for procurement-side responses such as beneficial-ownership disclosure for contractors and contract-level disbursement audits. A single fiduciary instrument applied uniformly across sectors would be mistargeted in both directions.

None of this establishes that supervision is currently misallocated relative to capture risk: we cannot observe the Bank’s existing sector-level supervisory intensity, and the reward correlation is underpowered

by construction. The claim is narrower. A large, identified share of continuing World Bank money flows to sectors where aid arrival demonstrably coincides with anonymous offshore crypto activity; that coincidence is observable in near-real time; and supervision that ignores it leaves an identified signal unused.

7 Conclusion

This paper asks where and how foreign aid is diverted now that enforcement has raised the cost of parking it in offshore banks, and finds part of the answer on the blockchain. Linking World Bank project disbursements to on-chain Bitcoin forensics, off-chain exchange records, and IP-linked web traffic, and exploiting the administrative timing of 328 tranche-arrival spikes across the 94 recipient countries in our estimation sample, we document a sharp and short-lived reorientation of crypto activity in the month aid arrives. Anonymous transaction volume on tax-haven exchanges rises by roughly 137 percent, entry into the crypto ecosystem runs almost entirely through newly created anonymous wallets, mixers gain relative to ordinary venues, and the surge dissipates within one to two months: the placement, layering, and integration sequence of money laundering compressed into the weeks after arrival. Cumulated over the post-disbursement window, the response implies gross leakage of 11.2 cents per disbursed dollar, an upper bound attained only if no diverted dollar is counted twice; at the layering rates forensic practice treats as typical, net leakage is 2 to 6 cents. Applied to the \$78 billion of World Bank money that arrives in those 328 tranche months, that is roughly \$1.7 to \$4.4 billion of diverted aid. Flat pre-trends, a placebo of disasters without disbursements, and the Pandora Papers' selective suppression of identified but not anonymous activity all point in the same direction: toward diversion of money in hand rather than shock-driven demand.

The results tell donors not only that capture occurs, but where and when. It concentrates in Transport, Water and Sanitation, Social Protection, and Governance, sectors whose disbursements create liquid windfalls, in contrast to in-kind Health aid, which shows none; and those four sectors still absorb 48 percent of sector-classified World Bank disbursements and new commitments in the two years after our sample closes. Because the signal appears in the month the tranche lands, supervision should bind at disbursement rather than at approval or ex-post audit, through staged tranche release, transaction monitoring timed to the arrival month, and scrutiny of the on- and off-ramps where laundered funds must

touch the regulated financial system. And because blockchains are public, such monitoring can operate in close to real time using the same forensics we deploy here. We are equally clear about what the evidence cannot establish. It cannot name the actors behind anonymous wallets; it cannot say whether crypto substitutes for or complements older laundering channels; and because it tracks only Bitcoin while stablecoins now carry much of the illicit cross-border volume in developing countries, it almost certainly understates total crypto-mediated diversion. Extending the forensic toolkit to stablecoins and cross-chain protocols, and linking on-chain traces to audit reports and legal proceedings, are the natural next steps. More broadly, the approach developed here, event-study identification joined to multi-source blockchain forensics, offers a template for studying illicit finance as it digitizes: the technologies that move money out of sight now generate the public records by which it can be found.

References

- Alesina, Alberto and Beatrice Weder**, “Do Corrupt Governments Receive Less Foreign Aid?,” *American Economic Review*, 2002, 92 (4), 1126–1137.
- 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.
- Amiram, Dan, Evgeny Lyandres, and Daniel Rabetti**, “Trading Volume Manipulation and Competition Among Centralized Crypto Exchanges,” *Management Science*, 2025, 0 (0).
- 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.
- Auer, Raphael, Marc Farag, Ulf Lewrick, Lev Menand Orazem, and Martin Zoss**, “Banking in the Shadow of Bitcoin? The Institutional Adoption of Cryptocurrencies,” BIS Working Papers 1013, Bank for International Settlements 2022.
- Bandiera, Oriana, Andrea Prat, and Tommaso Valletti**, “Active and Passive Waste in Government Spending: Evidence from a Policy Experiment,” *American Economic Review*, 2009, 99 (4), 1278–1308.
- 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.
- Budish, Eric**, “Trust at scale: The economic limits of cryptocurrencies and blockchains,” *The Quarterly Journal of Economics*, 2025, 140 (1), 1–62.
- Cameron, A. Colin and Douglas L. Miller**, “A Practitioner’s Guide to Cluster-Robust Inference,” *Journal of Human Resources*, 2015, 50 (2), 317–372.
- 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**, “The Chainalysis 2021 Crypto Crime Report,” Technical Report, Chainalysis 2021.
- Chalendard, Cyril, Ana M. Fernandes, Gaël Raballand, and Bob Rijkers**, “Corruption in Customs,” *Quarterly Journal of Economics*, 2023, 138 (1), 575–636.
- Chen, Jiafeng and Jonathan Roth**, “Logs with Zeros? Some Problems and Solutions,” *Quarterly Journal of Economics*, 2024, 139 (2), 891–936.
- 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.
- Denizer, Cevdet, Daniel Kaufmann, and Aart Kraay**, “Good Countries or Good Projects? Macro and Micro Correlates of World Bank Project Performance,” *Journal of Development Economics*, 2013, 105, 288–302.
- Dutkowska-Żuk, Agnieszka, Austin Hounsel, Amy Morrill, Andre Xiong, Marshini Chetty, and Nick Feamster**, “How and Why People Use Virtual Private Networks,” in “31st USENIX Security Symposium (USENIX Security 22)” USENIX Association Boston, MA 2022, pp. 3451–3465.
- Egger, Dennis, Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker**, “General Equilibrium Effects of Cash Transfers: Experimental Evidence from Kenya,” *Econometrica*, 2022, 90 (6), 2603–2643.
- 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.
- Graf von Luckner, Clemens, Carmen M. Reinhart, and Kenneth S. Rogoff**, “Decrypting New Age International Capital Flows,” *Journal of Monetary Economics*, 2023, 138, 104–122.
- , **Robin Koepke, and Silvia Sgherri**, “Crypto as a Marketplace for Capital Flight,” Technical Report WP/24/133, International Monetary Fund 2024.
- Griffin, John M. and Amin Shams**, “Is Bitcoin Really Un-Tethered?,” *Journal of Finance*, 2020, 75 (4), 1913–1964.
- **and Kevin Mei**, “How Do Crypto Flows Finance Slavery? The Economics of Pig Butchering. Available at <https://dx.doi.org/10.2139/ssrn.4742235>,” 2024.
- , —, **and Zirui Wang**, “Are Crypto Anti-Money Laundering Policies Effective?,” September 2025. Available at SSRN.
- Guyton, John, Patrick Langetieg, Daniel Reck, Max Risch, and Gabriel Zucman**, “Tax Evasion at the Top of the Income Distribution: Theory and Evidence,” Working Paper 28542, National Bureau of Economic Research 2021.
- Huberman, Gur, Jacob D. Leshno, and Ciamac C. Moallemi**, “Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System,” *Review of Economic Studies*, 2021, 88 (6), 3011–3040.

- International Consortium of Investigative Journalists**, “The Panama Papers,” 2016. Leak exposing offshore holdings of elites. Available at <https://www.icij.org/investigations/panama-papers/>.
- International Telecommunication Union**, “Facts and Figures 2023: Internet Use,” Technical Report, International Telecommunication Union 2023.
- Johannesen, Niels and Gabriel Zucman**, “The End of Bank Secrecy? An Evaluation of the G20 Tax Haven Crackdown,” *American Economic Journal: Economic Policy*, 2014, 6 (1), 65–91.
- , **Patrick Langetieg, Daniel Reck, Max Risch, and Joel Slemrod**, “Taxing Hidden Wealth: The Consequences of US Enforcement Initiatives on Evasive Foreign Accounts,” *American Economic Journal: Economic Policy*, 2020, 12 (3), 312–346.
- Khan, Mohammad Taha, Joe DeBlasio, Geoffrey M. Voelker, Alex C. Snoeren, Chris Kanich, and Narseo Vallina-Rodriguez**, “An Empirical Analysis of the Commercial VPN Ecosystem,” in “Proceedings of the Internet Measurement Conference 2018 (IMC ’18)” ACM Boston, MA, USA 2018, pp. 443–456.
- Kilby, Christopher**, “The Political Economy of Conditionality: An Empirical Analysis of World Bank Loan Disbursements,” *Journal of Development Economics*, 2009, 89 (1), 51–61.
- Kraay, Aart**, “How Large Is the Government Spending Multiplier? Evidence from World Bank Lending,” *Quarterly Journal of Economics*, 2012, 127 (2), 829–887.
- Londoño-Vélez, Juliana and Javier Ávila-Mahecha**, “Enforcing Wealth Taxes in the Developing World: Quasi-Experimental Evidence from Colombia,” *American Economic Review: Insights*, 2021, 3 (2), 131–148.
- Makarov, Igor and Antoinette Schoar**, “Trading and Arbitrage in Cryptocurrency Markets,” *Journal of Financial Economics*, 2020, 135 (2), 293–319.
- Meiklejohn, Sarah, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage**, “A fistful of bitcoins: characterizing payments among men with no names,” in “Proceedings of the 2013 conference on Internet measurement conference” 2013, pp. 127–140.
- Muralidharan, Karthik, Paul Niehaus, and Sandip Sukhtankar**, “Building State Capacity: Evidence from Biometric Smartcards in India,” *American Economic Review*, 2016, 106 (10), 2895–2929.
- Olken, Benjamin A.**, “Monitoring Corruption: Evidence from a Field Experiment in Indonesia,” *Journal of Political Economy*, 2007, 115 (2), 200–249.
- Phua, Kenny, Bo Sang, Chishen Wei, and Gloria Yang Yu**, “The Economics of Financial Scams: Evidence from Initial Coin Offerings,” 2025. Working Paper.
- Ramesh, Reethika, Anjali Vyas, and Roya Ensafi**, ““All of Them Claim to Be the Best”: Multi-perspective Study of VPN Users and VPN Providers,” in “32nd USENIX Security Symposium (USENIX Security 23)” USENIX Association Anaheim, CA 2023.

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.

Silva, J. M. C. Santos and Silvana Tenreyro, “The Log of Gravity,” *Review of Economics and Statistics*, 2006, 88 (4), 641–658.

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

Tørsløv, Thomas, Ludvig Wier, and Gabriel Zucman, “The Missing Profits of Nations,” *Review of Economic Studies*, 2023, 90 (3), 1499–1534.


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.

World Bank, “Investment Project Financing,” 2024.


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. Address-Level Transaction Network Around the Five Largest Aid Disbursements. This figure plots the address-level Bitcoin transaction network around the five largest World Bank disbursements in the sample, which are listed in Table A.3. Panel A plots the network one month before disbursement, Panel B one month after, and Panel C six months after. Node size is proportional to centrality and edge shading to transaction volume. Node color denotes service category: non-tax-haven exchanges (purple), tax-haven exchanges (teal), payment platforms (green), mixers (pink), and gambling and other services (orange).




Panel A. Network One Month Before Aid Disbursements




Panel C. Network Six Months After Aid Disbursements

Figure 2. Clustered Network Analysis Around the Five Largest Aid Disbursements. Larger nodes represent higher centrality, and darker edges indicate greater volume. The panels 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. The two panels are separate difference-network visualizations, so node size and edge shading should be read within each panel rather than as a formal magnitude comparison across panels. Node labels identify the service categories; “Non Haven Exchange” denotes non-tax-haven exchanges.




(a) Network Difference Around Aid Disbursements: One Week Before vs. One Week After




(b) Network Difference Around Aid Disbursements: One Week Before vs. Six Months After


Figure 3. Stacked Event Study of Foreign Aid and Geo-Adjusted Bitcoin Activity, by Wallet Type. This figure plots the dynamic response of geo-adjusted Bitcoin transaction volume and new account creation to foreign-aid disbursements, separately for anonymous and identified wallets and for tax-haven and non-tax-haven exchanges. 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. The left column reports anonymous wallets and the right column identified wallets. Plotted coefficients are the event-time dummies from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes; 18,499 observations per platform), estimated over the window $k = -8, \dots, +5$ in months relative to the tranche-arrival month ($k = 0$) with reference period $k \in [-12, -9]$, on IHS-transformed outcomes by linear TWFE with a no-overlap control group. All specifications include country-by-platform-by-episode (Ctry. $\times$ Plat. $\times$ Ep.) and month fixed effects. Standard errors are clustered at the recipient-country level. Shaded bands denote 95% confidence intervals.




Panel A. Transaction volume: anonymous wallets



Panel B. Transaction volume: identified wallets




Panel C. New accounts: anonymous wallets




Panel D. New accounts: identified wallets


Figure 4. Stacked Event Study of Foreign Aid and Geo-Adjusted Bitcoin Activity before and after the Pandora Papers. This figure plots the dynamic response of geo-adjusted Bitcoin transaction volume and new account creation to foreign-aid disbursements, separately for anonymous and identified wallets, comparing the period before the Pandora Papers leak with the period from October 2021 onward. 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. The left column reports anonymous wallets and the right column identified wallets. Each panel overlays two level event-study paths, whose difference is the triple difference Δ : the pre-Pandora path is the coefficient on the event dummies D_k , and the post-Pandora path adds the $D_k \times \text{Pandora}_{\text{post}}$ interactions, with its confidence band computed from the full variance-covariance matrix by the delta method. Estimates are computed on the flow-based tranche-arrival panel (328 episodes; 18,499 observations per platform) for tax-haven exchanges with no-overlap developing-country controls, over the window $k = -8, \dots, +5$ in months relative to the tranche-arrival month ($k = 0$) with reference period $k \in [-12, -9]$, on IHS-transformed outcomes by linear TWFE. All specifications include country-by-platform-by-episode (Ctry. $\times$ Plat. $\times$ Ep.) and month fixed effects. Standard errors are clustered at the recipient-country level. Shaded bands denote 95% confidence intervals.




Panel A. Transaction volume: anonymous wallets



Panel B. Transaction volume: identified wallets



Panel C. New accounts: anonymous wallets



Panel D. New accounts: identified wallets

Figure 5. Turnover-Adjusted Net Leakage. This figure plots net leakage $\lambda_{\text{net}}(m) = \lambda_{\text{gross}}/m$ against the turnover multiplier m , the number of times a diverted dollar is observed as it passes through successive wallets and off-ramps. The two curves correspond to the two conventions for the gross ratio. The solid line uses the cumulative post-disbursement response of anonymous tax-haven transaction volume over $k = 0, \dots, 4$ ($\lambda_{\text{gross}} = 11.2\%$), the convention reported throughout the paper, and the dashed line uses the arrival-month impulse at $k = 0$ ($\lambda_{\text{gross}} = 13.3\%$). Both divide the implied level change by the mean tranche flow of \$238 million across treated episodes, and the shaded region between them spans the range implied by the choice of convention. Both curves begin at $m = 1$, the case in which no dollar is observed twice and gross leakage equals net leakage. The horizontal dashed line marks the 7.5% offshore-banking leakage estimate of Andersen et al. (2022), and the annotated point $m^* = 1.49$ is the multiplier at which the cumulative curve crosses that line.






Figure 6. Sectoral Capture Intensity and Continuing World Bank Funding. This figure plots measured sectoral capture against the World Bank’s continuing funding of the same sectors. Capture is defined as each sector’s arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) on anonymous tax-haven transaction volume, reported in Table A.13. Panel A plots the sectoral capture ranking with 95% confidence intervals; filled bars denote sectors whose impulse is significant at the 10% level, and the number of episodes in each sector is reported in parentheses. Panel B plots the same impulse on the horizontal axis against newly approved World Bank commitments from April 2024 to June 2026, after the close of the study window, on the vertical axis. Filled markers denote sectors whose impulse is significant at the 10% level, marker area is proportional to post-window disbursement, and dotted lines mark the sample medians of the two axes.



Panel A. Which sectors capture more?



Panel B. Supervisory exposure: capture intensity vs. continuing funding

Tables

Table 1. 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 arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes), estimated over the extended window $k = -8, \dots, +5$ with reference period $k \in [-12, -9]$ on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Joint tests of the clean-window pre-trend coefficients ($k \in [-8, -5]$) fail to reject in every column (Panel A: $p \in [0.79, 0.98]$). Anonymous transactions account for 91 percent of tax-haven and 90 percent of non-tax-haven transaction volume over the estimation panel. All specifications include country-by-platform-by-episode ($\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$) and month fixed effects. Standard errors in parentheses, clustered at the recipient-country 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 $\times$ Post	0.866*** (0.234)	0.864*** (0.232)	0.301* (0.168)	1.396*** (0.273)	1.387*** (0.270)	0.622*** (0.189)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel B: Inflows						
Treat $\times$ Post	0.775*** (0.215)	0.773*** (0.211)	0.220 (0.161)	1.267*** (0.250)	1.256*** (0.246)	0.554*** (0.179)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel C: Outflows						
Treat $\times$ Post	0.778*** (0.215)	0.776*** (0.214)	0.245* (0.136)	1.272*** (0.251)	1.265*** (0.249)	0.482*** (0.155)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499

Table 2. 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 arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes), estimated over the extended window $k = -8, \dots, +5$ with reference period $k \in [-12, -9]$ on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Joint tests of the clean-window pre-trend coefficients ($k \in [-8, -5]$) fail to reject in every column (Panel A: $p \in [0.22, 0.91]$). Anonymous accounts constitute 99.9997 percent of new-account mass on tax-haven exchanges and 99.9991 percent on non-tax-haven exchanges over the estimation panel, so the All and Anonymous columns coincide to three decimals. All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country 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 $\times$ Post	1.775*** (0.445)	1.775*** (0.445)	0.044 (0.108)	2.832*** (0.550)	2.832*** (0.550)	0.260** (0.114)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel B: Inflows						
Treat $\times$ Post	1.685*** (0.417)	1.685*** (0.417)	0.005 (0.083)	2.739*** (0.532)	2.739*** (0.532)	0.198** (0.092)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel C: Outflows						
Treat $\times$ Post	1.636*** (0.427)	1.636*** (0.427)	0.033 (0.088)	2.643*** (0.515)	2.643*** (0.515)	0.073 (0.107)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499

Table 3. Impact of Foreign Aid on Off-Chain Transactions: Platform Heterogeneity and Country Characteristics.

This table reports the effect of foreign-aid disbursements on off-chain crypto activity, separately by exchange classification and by recipient-country institutional quality. Panel A reports the effect of lagged foreign aid on two outcomes. Columns 1–3 use the one-quarter change in log exchange volume, $\log(EX_t) - \log(EX_{t-1})$. Columns 4–6 use the one-quarter change in the log ratio of exchange volume to benchmark volume, $\log(EX_t/BM_t) - \log(EX_{t-1}/BM_{t-1})$, where the benchmark BM is transaction volume on the major regulated exchanges 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, and country and time fixed effects. Standard errors in parentheses, clustered at the recipient-country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 _{lagged}	2.39*** (0.55)	1.51 (1.16)	0.88*** (0.31)	0.39*** (0.11)	−0.03 (0.14)	0.42*** (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 _{lagged}	2.38*** (0.64)	1.97*** (0.55)	0.41*** (0.05)	0.37** (0.16)	0.03*** (0.01)	0.34*** (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 _{lagged}	2.41*** (0.64)	1.47** (0.58)	0.94*** (0.03)	0.39** (0.16)	−0.00 (0.00)	0.39*** (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	

Table 3 (continued): Impact of Foreign Aid on Off-Chain Transactions: Platform Heterogeneity and Country Characteristics.

Panel B. Heterogeneity by Country Characteristics								
	CPIA		Domestic Credit		Control for Corruption		Disclosure	
	<i>High</i>	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>	<i>Low</i>	<i>High</i>	<i>Low</i>
Foreign Aid _{lagged}	0.03 (0.02)	0.03*** (0.01)	0.02 (0.01)	0.03*** (0.01)	0.02** (0.01)	0.03*** (0.01)	0.01 (0.04)	0.12*** (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 4. 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, benchmarked against tax-haven countries. May 2021 marks Huobi's relocation of its headquarters from Seychelles to Singapore. Panel A reports results for the one-quarter change in log exchange volume, $\log(EX_t) - \log(EX_{t-1})$. Panel B reports results for the one-quarter change in the log ratio of exchange volume to benchmark volume, $\log(EX_t/BM_t) - \log(EX_{t-1}/BM_{t-1})$, where the benchmark BM is transaction volume on the major regulated exchanges Coinbase, Kraken, and Gemini. Coefficients are multiplied by 100. All specifications control for the two-quarter lagged foreign-aid amount and one-quarter lagged BTC volume, and include country and time fixed effects. Standard errors in parentheses, clustered at the recipient-country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Hong Kong	Singapore	Seychelles	Cayman Isl.	Malta	Huobi Effect
(coefficients multiplied by 100)						
Panel A: Lead - lag						
Foreign Aid _{lagged}	3.35*** (0.704)	1.71*** (0.62)	2.80*** (0.57)	3.87*** (0.85)	1.90*** (0.52)	0.04 (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*** (0.06)	0.01 (0.03)	0.16*** (0.04)	0.04** (0.02)	0.18** (0.08)	0.02** (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 5. Crypto Laundering after the Pandora Papers. This table tests whether the aid-crypto relationship intensified after the Pandora Papers leak (October 2021). We estimate a pooled stacked specification $y \sim \text{DisbPost} \times \text{PandoraPost} \mid \alpha_i + \gamma_t$, where DisbPost equals one in post-disbursement months for treated episodes and PandoraPost indicates calendar months on or after the leak. $\hat{\beta}_{\text{pre}}$ is the coefficient on DisbPost; $\hat{\beta}_{\text{post}}$ is the sum of that coefficient and the DisbPost $\times$ PandoraPost interaction; Δ is the interaction coefficient. The Pandora indicator equals one for months on or after October 2021. Outcomes are IHS-transformed on-chain aggregates, estimated by linear TWFE with a no-overlap control group. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C); 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. $\times$ Plat. $\times$ Ep. and month fixed effects. Standard errors in parentheses, clustered at the recipient-country level; the standard error on Δ is that of the interaction term, and the standard error on $\hat{\beta}_{\text{post}}$ is computed by the delta method from the same variance-covariance matrix. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Tax Haven			Non-Tax Haven		
	All	Anonymous	Identified	All	Anonymous	Identified
Panel A: Transaction Volume						
Treat $\times$ Post, pre-Pandora	0.984*** (0.303)	0.970*** (0.299)	0.717*** (0.233)	1.289*** (0.328)	1.267*** (0.323)	0.974*** (0.257)
Treat $\times$ Post, post-Pandora	0.874*** (0.236)	0.875*** (0.235)	0.185 (0.113)	1.543*** (0.333)	1.542*** (0.333)	0.504*** (0.165)
Δ (post – pre)	–0.110 (0.379)	–0.095 (0.375)	–0.531** (0.245)	0.253 (0.481)	0.275 (0.477)	–0.469 (0.311)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel B: Transaction Frequency						
Treat $\times$ Post, pre-Pandora	1.002*** (0.307)	0.999*** (0.306)	0.546*** (0.191)	1.287*** (0.327)	1.278*** (0.325)	0.820*** (0.226)
Treat $\times$ Post, post-Pandora	0.839*** (0.222)	0.838*** (0.222)	0.289*** (0.102)	1.390*** (0.305)	1.387*** (0.304)	0.588*** (0.160)
Δ (post – pre)	–0.163 (0.377)	–0.162 (0.376)	–0.258 (0.207)	0.104 (0.458)	0.108 (0.456)	–0.232 (0.283)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel C: New Accounts						
Treat $\times$ Post, pre-Pandora	1.924*** (0.577)	1.924*** (0.577)	0.288** (0.126)	2.709*** (0.650)	2.709*** (0.650)	0.437*** (0.160)
Treat $\times$ Post, post-Pandora	1.622*** (0.408)	1.622*** (0.408)	0.012 (0.044)	3.089*** (0.630)	3.089*** (0.630)	0.097* (0.054)
Δ (post – pre)	–0.302 (0.703)	–0.302 (0.703)	–0.275** (0.130)	0.380 (0.914)	0.380 (0.914)	–0.340** (0.155)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499

Table 6. Placebo Test of Natural Disasters Without World Bank Disbursement. This table reports a placebo test in which the treated cohort is defined by crisis exposure rather than by aid inflows. Treated countries are those that experienced a natural disaster recorded in EM-DAT (2018–2024) but received no World Bank disbursement; event time is centered on the disaster month, and the control group consists of countries with no World Bank aid in the window. The reported coefficient is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on a disaster-centered episode panel built with the same stacking machinery as the baseline tranche-arrival panel, estimated on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C); columns distinguish all, anonymous, and identified transactions by platform type. The panel comprises 120 disaster episodes and is larger than the 18,499-cell baseline stack because controls are restricted only by the no-disbursement rule, and not additionally to developing economies as in the baseline, so many more clean control countries enter each episode. All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country 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 $\times$ Post	0.088 (0.116)	0.087 (0.114)	0.060 (0.093)	0.156 (0.159)	0.153 (0.155)	0.110 (0.125)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	129,381	129,381	129,381	129,381	129,381	129,381
Panel B: Transaction Frequency						
Treat $\times$ Post	0.093 (0.115)	0.093 (0.115)	0.044 (0.073)	0.154 (0.156)	0.153 (0.154)	0.087 (0.106)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	129,381	129,381	129,381	129,381	129,381	129,381
Panel C: New Accounts						
Treat $\times$ Post	0.084 (0.201)	0.084 (0.201)	0.022 (0.082)	0.346 (0.329)	0.346 (0.329)	0.002 (0.101)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	129,381	129,381	129,381	129,381	129,381	129,381

Table 7. Impact of Foreign Aid on Crypto Activity: VPN Exclusion Robustness. This table re-estimates the baseline arrival-month impulse after excluding high-VPN recipient countries. The reported coefficient is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from the stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes), re-estimated after dropping high-VPN recipients from the full sample under three rules: (i) a manual list (Indonesia, Turkey, Pakistan), flagged in the policy and industry literature as the primary VPN-penetration concerns in the sample; (ii) the top decile of recipient countries by Google Trends search intensity for the term “VPN”; and (iii) the top decile by GWI VPN adoption (the share of internet users aged 16–64 using a VPN; *Digital 2026 Global Overview Report* by DataReportal and We Are Social, Q2 2025 wave). Outcomes are IHS-transformed; each coefficient uses its rule’s sample, and the three Obs rows report the corresponding counts. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C). All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country 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						
(i) Manual list (Indonesia, Turkey, Pakistan)	0.948*** (0.250)	0.944*** (0.247)	0.380** (0.175)	1.514*** (0.290)	1.504*** (0.287)	0.713*** (0.198)
(ii) Top 10% Google Trends (VPN)	0.878*** (0.242)	0.875*** (0.239)	0.319* (0.173)	1.413*** (0.281)	1.404*** (0.278)	0.637*** (0.195)
(iii) Top 10% GWI adoption	0.898*** (0.247)	0.895*** (0.244)	0.329* (0.177)	1.450*** (0.286)	1.442*** (0.283)	0.656*** (0.199)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs. under rule (i)	16,879	16,879	16,879	16,879	16,879	16,879
Obs. under rule (ii)	18,366	18,366	18,366	18,366	18,366	18,366
Obs. under rule (iii)	17,016	17,016	17,016	17,016	17,016	17,016
Panel B: Transaction Frequency						
(i) Manual list (Indonesia, Turkey, Pakistan)	0.963*** (0.246)	0.962*** (0.245)	0.376*** (0.140)	1.424*** (0.281)	1.419*** (0.280)	0.699*** (0.175)
(ii) Top 10% Google Trends (VPN)	0.895*** (0.239)	0.894*** (0.238)	0.333** (0.137)	1.326*** (0.273)	1.321*** (0.271)	0.635*** (0.171)
(iii) Top 10% GWI adoption	0.913*** (0.244)	0.912*** (0.243)	0.341** (0.140)	1.361*** (0.278)	1.357*** (0.276)	0.652*** (0.175)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs. under rule (i)	16,879	16,879	16,879	16,879	16,879	16,879
Obs. under rule (ii)	18,366	18,366	18,366	18,366	18,366	18,366
Obs. under rule (iii)	17,016	17,016	17,016	17,016	17,016	17,016
Panel C: New Accounts						
(i) Manual list (Indonesia, Turkey, Pakistan)	1.906*** (0.475)	1.906*** (0.475)	0.080 (0.112)	3.062*** (0.584)	3.063*** (0.584)	0.316*** (0.118)
(ii) Top 10% Google Trends (VPN)	1.782*** (0.460)	1.782*** (0.460)	0.040 (0.112)	2.865*** (0.566)	2.865*** (0.566)	0.265** (0.118)
(iii) Top 10% GWI adoption	1.812*** (0.469)	1.812*** (0.469)	0.039 (0.115)	2.936*** (0.576)	2.936*** (0.576)	0.272** (0.121)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs. under rule (i)	16,879	16,879	16,879	16,879	16,879	16,879
Obs. under rule (ii)	18,366	18,366	18,366	18,366	18,366	18,366
Obs. under rule (iii)	17,016	17,016	17,016	17,016	17,016	17,016

Table 8. Sensitivity to the Tranche-Spike Thresholds. This table re-estimates the baseline stacked event study under alternative definitions of a treated episode. Each row applies a different spike rule, requiring a monthly World Bank disbursement flow of at least \$ a M and at least $m \times$ the country's trailing twelve-month median, with no qualifying spike in the prior four months, and re-estimates the specification on the resulting panel. The reported coefficient is the anonymous arrival-month impulse $\hat{\beta}_0$ (the coefficient on the $k=0$ dummy) for transaction volume and new accounts, on tax-haven (TH) and non-tax-haven (NTH) exchanges. The baseline rule used throughout the paper, (\$50M, $3 \times$), is reported in bold. Impulses for identified wallets and joint tests of the clean pre-period ($k \in \{-8, \dots, -5\}$) are estimated for every rule but are not reported. All specifications include Ctry. $\times$ Plat. $\times$ Ep. and month fixed effects. Standard errors in parentheses, clustered at the recipient-country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Abs. (\$M)	Rel. ($\times$)	Episodes	Volume ($\hat{\beta}_0$, ano)		New accounts ($\hat{\beta}_0$, ano)	
			TH	NTH	TH	NTH
25	2	461	0.636*** (0.168)	1.037*** (0.200)	1.380*** (0.324)	2.237*** (0.416)
25	3	425	0.572*** (0.187)	1.107*** (0.218)	1.273*** (0.361)	2.341*** (0.450)
25	5	346	0.882*** (0.209)	1.406*** (0.240)	1.751*** (0.395)	2.928*** (0.490)
50	2	355	0.883*** (0.207)	1.256*** (0.234)	1.762*** (0.394)	2.587*** (0.476)
50	3	328	0.864*** (0.231)	1.387*** (0.259)	1.775*** (0.439)	2.832*** (0.526)
50	5	281	1.054*** (0.247)	1.630*** (0.275)	2.097*** (0.466)	3.364*** (0.557)
100	2	254	0.981*** (0.257)	1.714*** (0.296)	1.924*** (0.490)	3.630*** (0.596)
100	3	242	1.033*** (0.271)	1.910*** (0.307)	2.071*** (0.517)	4.054*** (0.614)
100	5	207	1.149*** (0.299)	2.149*** (0.332)	2.248*** (0.565)	4.552*** (0.664)

Table 9. Robustness Check: Alternative Geographic Identifier (Visit-Duration Ratio). This table re-estimates the baseline specification with the visits-ratio geographic identifier replaced by a visit-duration ratio, defined as each country-month's share of global crypto-exchange visit duration. The reported coefficient is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel, estimated on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C); columns distinguish all, anonymous, and identified transactions by platform type. All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country 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 $\times$ Post	1.609*** (0.343)	1.606*** (0.340)	0.754*** (0.248)	1.921*** (0.382)	1.913*** (0.379)	1.112*** (0.281)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel B: Transaction Frequency						
Treat $\times$ Post	1.630*** (0.337)	1.629*** (0.336)	0.813*** (0.205)	1.833*** (0.372)	1.829*** (0.370)	1.146*** (0.257)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel C: New Accounts						
Treat $\times$ Post	2.909*** (0.597)	2.909*** (0.597)	0.202 (0.146)	3.338*** (0.672)	3.338*** (0.672)	0.540*** (0.164)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499

Table 10. Sectoral Capture Intensity and Continuing World Bank Funding. This table relates each project category’s capture intensity to its subsequent World Bank funding. Capture intensity is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) on anonymous tax-haven transaction volume, estimated separately on each category’s subset of the 328 spike episodes and reported in full in Table A.13, with its standard error in parentheses. “Disbursed” is gross disbursement on the legacy loan book over April 2024 to June 2026, measured from the live IBRD/IDA Statement of Loans and Credits net of the March 2024 baseline. “Committed” and “Approvals” are new board-approved commitments and their count over the same window, from the World Bank Projects API. Annualised commitment rates divide in-sample commitments by 6 years and post-window commitments by 2.25 years, and “Accel.” is the ratio of the two. The title-based sector map covers 98% of post-window disbursed dollars, 95% of newly committed dollars, and 92% of new approvals; unclassified loans are excluded from the category totals. Environment, Climate, and Disaster contains no qualifying spike episode and is omitted. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Project category	Capture		Post-window funding			Commitment rate (US\$ bn/yr)		
	$k=0$ impulse (IHS)	Episodes (n)	Disbursed (US\$ bn)	Committed (US\$ bn)	Approvals (n)	In-sample	Post	Accel. ($\times$)
Transport	2.676*** (0.897)	18	54.4	35.8	117	4.98	15.93	3.20
Water & San.	1.528** (0.773)	15	36.2	9.2	73	4.55	4.09	0.90
Soc. Protection	1.518*** (0.565)	58	45.8	13.2	106	9.96	5.87	0.59
Finance	0.913 (0.708)	39	50.7	25.5	103	7.16	11.32	1.58
Governance	0.743** (0.373)	101	123.2	49.8	250	25.69	22.12	0.86
Energy	0.582 (0.628)	27	57.9	28.9	138	8.95	12.85	1.44
Agriculture	0.521 (1.554)	11	70.3	14.6	121	6.06	6.50	1.07
Education	0.405 (0.622)	13	49.4	14.1	105	4.94	6.26	1.27
Digital	0.029 (0.048)	3	7.4	3.7	32	1.48	1.65	1.12
Health	-0.255 (0.507)	43	42.2	29.3	91	8.95	13.02	1.45
All 10 categories		328	537.6	224.1	1,136	82.72	99.60	1.20

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 in the baseline analysis is a monthly World Bank tranche-arrival flow, constructed from first differences of loan-level cumulative disbursement records. Treated episodes are high-flow arrival months that meet the \$50 million and three-times-trailing-median spike rule used in the stacked event-study design. Quarterly lagged aid measures enter only the auxiliary off-chain exchange specifications. 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	Monthly World Bank disbursement flow, constructed by first-differencing loan-level cumulative disbursement records; the baseline treatment is a high-flow tranche-arrival month. Quarterly lagged aid is used only in auxiliary off-chain specifications.
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.

B Overview of Aid Recipients and Crypto Transaction Platforms

This section reports the descriptive overview of the aid-recipient sample and the cryptocurrency transaction platforms underlying the on-chain data, discussed in Section 3.

Table A.2. Overview of Aid Recipients and Crypto Transaction Platforms. This table summarizes key statistics on aid distribution and crypto exchange activity, organized in five panels. Panel A reports the top 15 recipient countries by total disbursed aid (USD millions) and their share of the estimation sample. Disbursement is measured as the sum of positive monthly disbursement flows, constructed by first-differencing loan-level end-of-period disbursed stocks, over April 2018 through March 2024, across the 94 recipient countries that contribute at least one treated tranche-arrival episode. The final line reports the subtotal arriving in the 328 treated tranche-arrival months, which is the denominator of the leakage ratio in Section 4.5. 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 E reports the anonymous share of measured on-chain activity by platform type for transaction volume, transaction frequency, and new-account creation, computed from the attention-adjusted series over April 2018 through March 2024.

Panel A. Top 15 Aid Recipient Countries by Disbursed Amount		
Country	Disbursed Amount <i>(USD millions)</i>	Fraction
India	20,153	0.08
Indonesia	10,739	0.05
Nigeria	9,826	0.04
Bangladesh	9,737	0.04
Ethiopia	9,728	0.04
Philippines	8,827	0.04
Pakistan	8,500	0.04
Ukraine	8,330	0.04
Kenya	7,166	0.03
Colombia	6,901	0.03
Argentina	6,178	0.03
Morocco	5,941	0.02
Egypt, Arab Republic of	5,618	0.02
Peru	5,381	0.02
Mexico	5,168	0.02
Total (top 15)	128,192	0.54
All 94 sample recipients	237,701	1.00
<i>of which, in treated tranche arrivals</i>	<i>78,166</i>	<i>0.33</i>

Panel B. Regulatory-Type Distribution of Crypto Exchanges			
	Tax Haven	Regulated	Crypto Friendly
Proportion	0.45	0.18	0.70

Table A.2 (continued): Overview of Aid Recipients and Crypto Transaction Platforms.

Panel C. Platform Distribution by New Accounts and 24h Volume				
Platform	# Accounts <i>(thousands)</i>	% Accounts	Volume 24h <i>(thousands)</i>	% Volume
CoinJoinMess	31,404	0.77	1,579	0.02
Binance	2,136	0.05	28,254	0.34
Kraken	2,051	0.05	20,808	0.25
CoinPayments	1,856	0.05	539	0.01
Luno	1,634	0.04	1,164	0.01
Huobi	1,230	0.03	19,029	0.23
Bittrex	496	0.01	3,482	0.04
YABTCL	17	0.00	2	0.00
999Dice	14	0.00	4	0.00
Bitstamp	13	0.00	5,639	0.07
Xapo	13	0.00	670	0.01
BetMoose	1	0.00	1	0.00
HitBTC	< 1	0.00	3	0.00
Cex	< 1	0.00	7	0.00
BitZino	< 1	0.00	< 1	0.00
Total	40,864	1.00	83,301	1.00

Panel D. Top Countries by Off-chain Crypto Exchange Volume

Country	Volume 24h <i>(USD thousands)</i>	Fraction
Malta	71,954	0.33
Singapore	60,677	0.27
Seychelles	26,542	0.12
South Korea	12,888	0.06
Hong Kong	11,673	0.05
UAE	10,126	0.05
USA	6,096	0.03
Australia	5,758	0.03
Estonia	3,794	0.02
Cayman Islands	2,793	0.01
Total (top 10)	212,301	0.97

Panel E. Anonymous Share of On-Chain Activity by Platform Type

Measure	Tax-haven	Non-tax-haven
	<i>anonymous share</i>	
Transaction volume	91.2%	89.8%
Transaction frequency	98.0%	95.7%
New accounts	99.9997%	99.9991%

We examine the aggregate network dynamics of on-chain Bitcoin transactions around the largest foreign-aid disbursement events. We identify the five largest single country-month World Bank disbursement tranches (IBRD + IDA) between April 2018 and March 2024, where each flow is the month-over-month change in loan-level disbursed amounts and each event is the arrival month of the tranche in the recipient country, and track how network-level activity, including transaction volume, the number of active addresses, and the flow of funds across exchanges, evolves in the months surrounding each event. Table A.3 lists these five disbursement tranches, the corresponding amounts, and the receiving countries. The events span a range of recipient economies, from Peru and Ukraine to India, Indonesia, and Kenya, allowing us to assess whether the network response varies across large aid flows.

Table A.3. Five Largest Disbursement Tranches. The five largest single country-month World Bank disbursement flows (IBRD + IDA) between April 2018 and March 2024. Flows are month-over-month changes in loan-level disbursed amounts; each event is the arrival month of the tranche in the recipient country.

Arrival Month	Disbursed Flow (in millions)	Receiving Country
August 2020	1,835	Peru
December 2023	1,696	India
July 2023	1,510	Ukraine
December 2022	1,214	Indonesia
June 2023	1,183	Kenya

C Geographic Attribution of Crypto Activity

This appendix collects the material underlying our web-traffic-based geographic attribution of crypto activity: the web-traffic response to disbursements, the external validation of the geo-adjusted measures against an independent index, and a robustness analysis addressing VPN-masked traffic.

C.1 Web Traffic

Table A.4 reports the $k = 0$ arrival-month impulse of foreign-aid disbursements on crypto-exchange web traffic from tax-haven platforms across three control-group definitions. Panel A shows that foreign-aid receipt is associated with a large and highly significant increase in total visits to tax-haven crypto-exchanges, with a coefficient of approximately 2.21 in the baseline and no-neighbors samples and a somewhat smaller 1.69 when the control pool is restricted to developing economies, in every case significant at the 1 percent level. Given the IHS specification, the estimate implies that visits to tax-haven exchange websites from treated countries rise by an order of magnitude relative to controls during the post-disbursement window, a response far too large to be attributed to organic growth in crypto adoption or to confounding macroeconomic shocks common to aid recipients. Panel B documents a similarly strong response in total visit duration, with a coefficient of approximately 2.44 (and 1.90 in the developing-economy control sample), suggesting that the visit increase reflects sustained engagement rather than brief or incidental page loads. Taken together, the two panels indicate that treated-country users are not only navigating to these platforms in much greater numbers but also spending substantially more time on them, a pattern consistent with active account setup, transaction execution, or prolonged browsing of the exchange interface rather than passive clicks driven by, for instance, search-engine referrals. Both estimates are broadly stable across the no-overlap, no-overlap+developing, and no-overlap+no-neighbors samples, with only a modest attenuation when the control pool is restricted to developing economies, indicating that the results are not sensitive to the composition of the control group, and in particular are not driven by spillover from geographically or economically adjacent non-recipient countries.

Table A.4. Impact of Foreign Aid on Crypto Exchange Web Traffic. This table reports the effect of foreign-aid disbursements on web traffic to tax-haven crypto-exchange sites, the only country-varying on-chain outcome, re-timed onto the flow-based tranche-arrival panel. The $\text{Treat} \times \text{Post}$ row is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from the extended-window stacked event-study difference-in-differences on IHS-transformed outcomes, reported across three control-group definitions. Panels distinguish total visits (Panel A) and total visit duration in seconds (Panel B). All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country level. $^* p < 0.10$, $^{**} p < 0.05$, $^{***} p < 0.01$.

	Baseline	+ Developing	+ No neighbors
Panel A: Total visits			
$\text{Treat} \times \text{Post}$	2.212 ^{***} (0.343)	1.686 ^{***} (0.328)	2.212 ^{***} (0.343)
$\text{Ctry.} \times \text{Plat.} \times \text{Ep. FE}$	✓	✓	✓
Time FE	✓	✓	✓
Obs.	119,691	13,643	118,682
Panel B: Total visit duration (seconds)			
$\text{Treat} \times \text{Post}$	2.441 ^{***} (0.364)	1.904 ^{***} (0.362)	2.441 ^{***} (0.364)
$\text{Ctry.} \times \text{Plat.} \times \text{Ep. FE}$	✓	✓	✓
Time FE	✓	✓	✓
Obs.	119,691	13,643	118,682

C.2 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

¹See <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 artifact.

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.

C.3 VPN Robustness: Country Ranking and Interaction Tests

This section supports the VPN robustness analysis in Section 5.3. Table A.5 reports the top fifteen recipient countries ranked by their Google Trends VPN search intensity (on a 0–100 scale, averaged over the sample period), drawn from the Google Trends API for the query “VPN” at country resolution. This ranking forms the descriptive basis for the top-decile exclusion in rule (ii) of Table 7. The manual exclusion in rule (i) uses VPN penetration scores for Indonesia (0.550), Turkey (0.320), and Pakistan (0.190); we report these values here rather than as a table column.

Table A.5. Top 15 Recipient Countries by Google Trends VPN Intensity. This table reports the country-level VPN intensity measure used in the robustness checks of Section 5.3. Rows are recipient countries only, ordered from highest to lowest Google Trends VPN search intensity on the 0–100 scale, averaged over the sample period. “Trends (0–100)” is the Google Trends VPN intensity score. The final three columns are computed on the current 328-episode tranche-spike panel: “Sample share (%)” is the country’s share of panel observations, “Treated rows” is the number of treated country-month observations, and “Episodes” is the number of treated tranche-spike episodes, not episode-stack appearances. The earlier version of the Episodes column counted episode-stack appearances, including control appearances, which overstated the number of episodes contributed by several countries. The manual VPN-concern scores (rule i) and GWI adoption rates (rule iii) are reported in the surrounding text.

Country	Trends (0–100)	Sample share (%)	Treated rows	Episodes
South Sudan	71.0	0.2	12	2
Iran, Islamic Republic of	41.0	1.6	12	2
Uganda	29.0	0.3	36	6
Afghanistan	25.0	1.1	24	4
Benin	25.0	0.0	12	2
Ethiopia	23.0	0.1	30	5
Sudan	22.0	2.3	6	1
Ghana	19.0	0.0	18	3
Somalia	19.0	0.2	18	3
Pakistan	18.0	0.1	48	8
Bangladesh	16.0	0.1	30	5
Tanzania	16.0	0.1	42	7
Togo	14.0	0.0	18	3
Yemen, Republic of	14.0	0.1	36	6
Madagascar	13.0	0.1	42	7

For rule (iii), the GWI-covered recipients and their Q2 2025 VPN-adoption rates are Turkey (35.3), Indonesia (31.4), Vietnam (28.8), Kenya (24.1), Nigeria (24.1), Philippines (22.1), Egypt (21.6), South Africa (21.0), China (19.6), Colombia (19.5), Brazil (18.8), Romania (18.4), Mexico (18.1), Croatia (17.7), Chile (16.9), Serbia (16.7), Argentina (14.2), Ghana (12.1), and Morocco (5.7), where each value is the percentage

of internet users aged 16–64 who report using a VPN or proxy server. Because GWI surveys a panel of larger economies, it overlaps with only one of the fifteen Trends-ranked countries in Table A.5 (Indonesia); the Google Trends ranking and the GWI survey therefore emphasize different sets of countries, with Trends elevating small, high-search recipients and GWI covering larger economies. The top-decile cutoff excludes Indonesia and Turkey.

Table A.6 reports the interaction specification in which the arrival-month impulse is interacted with VPN intensity directly, both as a discrete top-half indicator (HighVPN, above the median of Google Trends VPN interest) and as the continuous standardized Google Trends score (VPNScore). Across all eighteen cells, neither interaction is statistically distinguishable from zero: the HighVPN interaction is uniformly negative and the VPNScore interaction uniformly positive, but every coefficient carries a standard error that swamps its point estimate. The interactions are identified off the treated recipients alone, which leaves them underpowered.

Under VPN-mediated mismeasurement that would threaten the baseline result (non-recipient users reverse-routing through exit nodes in the treated country), higher recipient-country VPN penetration should produce a positive and significant interaction, scaling the measured response with the intensity of the contaminating traffic. We find no such gradient: the interaction coefficients are statistically indistinguishable from zero and economically small relative to the arrival-month impulse, so the data do not exhibit the systematic VPN-intensity dependence that the biasing channel requires. Google Trends search intensity is an imperfect proxy for actual VPN usage, so we read this interaction test as corroborating rather than dispositive, alongside the exclusion-based tests in Table 7.

Table A.6. Impact of Foreign Aid on Crypto Activity by VPN Intensity. Arrival-month impulse ($k = 0$) and its interaction with recipient VPN intensity on the flow-based tranche-arrival panel (328 episodes), entered both as a discrete top-half indicator ($\text{Treat} \times \text{Post} \times \text{HighVPN}$, where HighVPN flags recipients above the median of Google Trends VPN interest) and as the continuous standardized Google Trends VPN score ($\text{Treat} \times \text{Post} \times \text{VPNScore}$, mean zero and unit standard deviation). The interactions are identified off treated recipients; the VPN main effects are absorbed by the fixed effects. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C); within each panel columns are split by Tax Haven and Non-Tax Haven platform type and further into all, anonymous, and identified transactions. Outcomes are IHS-transformed. All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects; per-column within-adjusted R^2 is reported. Standard errors in parentheses, clustered at the recipient-country 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: Transaction Volume						
Treat $\times$ Post	0.880*** (0.248)	0.878*** (0.245)	0.307* (0.179)	1.432*** (0.286)	1.424*** (0.283)	0.638*** (0.199)
Treat $\times$ Post $\times$ HighVPN	-0.310 (0.673)	-0.309 (0.668)	-0.137 (0.412)	-0.806 (0.680)	-0.800 (0.676)	-0.362 (0.453)
Treat $\times$ Post $\times$ VPNScore	0.096 (0.517)	0.093 (0.513)	0.136 (0.299)	0.288 (0.511)	0.281 (0.508)	0.314 (0.331)
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01
Panel B: Transaction Frequency						
Treat $\times$ Post	0.902*** (0.244)	0.901*** (0.243)	0.330** (0.142)	1.343*** (0.277)	1.338*** (0.276)	0.641*** (0.175)
Treat $\times$ Post $\times$ HighVPN	-0.273 (0.659)	-0.271 (0.657)	-0.070 (0.391)	-0.725 (0.655)	-0.720 (0.653)	-0.295 (0.444)
Treat $\times$ Post $\times$ VPNScore	0.160 (0.505)	0.161 (0.503)	0.067 (0.294)	0.329 (0.490)	0.328 (0.489)	0.240 (0.333)
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01
Panel C: New Accounts						
Treat $\times$ Post	1.804*** (0.467)	1.804*** (0.467)	0.046 (0.109)	2.915*** (0.572)	2.916*** (0.572)	0.257** (0.119)
Treat $\times$ Post $\times$ HighVPN	-0.642 (1.085)	-0.642 (1.085)	-0.061 (0.115)	-1.848* (1.120)	-1.848* (1.120)	0.050 (0.211)
Treat $\times$ Post $\times$ VPNScore	0.613 (0.814)	0.612 (0.815)	0.493*** (0.050)	0.830 (0.805)	0.830 (0.805)	0.467*** (0.139)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01

D Construction of the Anonymous Tax-Haven Flow Measure

This appendix makes precise what “anonymous transaction volume on tax-haven exchanges,” the outcome in Table 1, is at the level of an individual Bitcoin transaction, and how the three conventions a reader familiar with these data will ask about, counterparty attribution, edge direction, and the treatment of change outputs, are resolved. Section 3.2 states the clustering and labeling steps; here we walk a single deposit and a single withdrawal through them.

The unit is the wallet cluster, a set of addresses inferred to share a controller by the common-input-ownership heuristic (addresses spent together as inputs to one transaction) together with the standard change-address heuristic. Each cluster carries at most one label from WalletExplorer and allied forensic sources. A cluster that resolves to a named service is identified; a cluster with no such label is anonymous. This is a residual definition, and the resulting misclassification is one-sided (Section 3.2): an unlabeled service cluster is counted as anonymous, never the reverse. A tax-haven exchange cluster is an identified cluster whose service is an exchange domiciled in a tax-haven jurisdiction (for example the Seychelles or Malta; see Table A.2, Panel D).

Each on-chain transfer is then recorded on the ledger of the platform cluster it touches as a directed edge between the tax-haven exchange cluster and its counterparty cluster, weighted by the BTC value moved. The direction follows the coin: a transfer into the exchange cluster (the counterparty is the sender) is an inflow, or deposit, while a transfer out of the exchange cluster (the counterparty is the recipient) is an outflow, or withdrawal. “Anonymous transaction volume on tax-haven exchanges” then sums the value of edges for which (i) one endpoint is a tax-haven exchange cluster and (ii) the counterparty is an anonymous cluster. Panel B of Table 1 restricts this sum to inflow edges, Panel C to outflow edges, and Panel A takes their union; identified volume is the identical construction with an identified counterparty.

Change outputs are resolved at the clustering stage, before any flow is measured. A Bitcoin spend consumes whole unspent outputs and returns the unspent remainder to a fresh change address controlled by the sender. The change-address heuristic assigns that fresh address to the sending cluster, so the change is absorbed back into the sender rather than becoming a transfer to a separate counterparty; by the time

flows are summed, the ledger already records only transfers between distinct clusters, with no self-edge left to count. The flow measure therefore sums the value actually delivered to a counterparty, not the gross input value of a transaction. Together with the counterparty rule above, this keeps change out of both the anonymous and the tax-haven totals: a transfer back to the platform's own cluster is not a distinct counterparty at all, and a transfer to any named service is recorded as identified rather than anonymous volume.

A single worked example ties these conventions together. Suppose an anonymous cluster spends a single 3.00 BTC output in one transaction, sending 1.20 BTC to a deposit address of a Seychelles-domiciled exchange and returning 1.79 BTC as change, with a 0.01 BTC miner fee. The 1.79 BTC change output clusters back to the sender and is dropped as a self-loop. The remaining edge, anonymous cluster → tax-haven exchange cluster, contributes 1.20 BTC, not the 3.00 BTC gross input, to anonymous inflow volume on tax-haven exchanges, entering Panels A and B. If the exchange later withdraws 0.90 BTC to a different unlabeled cluster, that edge, tax-haven exchange cluster → anonymous cluster, contributes 0.90 BTC to anonymous outflow volume, entering Panels A and C. Had either counterparty instead resolved to a named service, the same value would have been recorded as identified rather than anonymous volume. The other two outcome families are built on these same edges: transaction frequency counts the qualifying edges rather than summing their value, and new-account creation counts the distinct anonymous counterparty clusters transacting with a tax-haven exchange for the first time in the event window.

E Additional Results and Robustness

This appendix collects supplementary results and robustness checks supporting the main analysis. We first report the on-chain transaction-frequency results, then a sequence of robustness exercises: an alternative fixed-effect specification, an alternative control-group composition, a diagnostic on the omitted reference period, and an exclusion of the 2021 DeFi boom. We close with a test distinguishing aid-linked anonymity-seeking from ordinary privacy demand.

E.1 On-Chain Transaction Frequency

Table A.7 reports the arrival-month impulse ($k = 0$) 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., 1.310 versus 0.889 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., 0.888 versus 0.889 on tax-haven exchanges), while identified transactions show substantially smaller responses, with a coefficient of 0.327. 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.7. 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 arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes), estimated over the extended window $k = -8, \dots, +5$ with reference period $k \in [-12, -9]$ on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Joint tests of the clean-window pre-trend coefficients ($k \in [-8, -5]$) fail to reject in every column (Panel A: $p \in [0.89, 0.99]$). All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country 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 $\times$ Post	0.889*** (0.231)	0.888*** (0.230)	0.327** (0.133)	1.310*** (0.265)	1.306*** (0.263)	0.628*** (0.166)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel B: Inflows						
Treat $\times$ Post	0.800*** (0.205)	0.799*** (0.204)	0.266** (0.123)	1.198*** (0.243)	1.193*** (0.241)	0.541*** (0.154)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel C: Outflows						
Treat $\times$ Post	0.774*** (0.216)	0.771*** (0.215)	0.239** (0.105)	1.146*** (0.240)	1.141*** (0.239)	0.492*** (0.136)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499

E.2 Alternative Specification

The baseline specification employs IHS-transformed outcomes estimated by linear TWFE with fully interacted $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Because this combination imposes two potentially consequential modelling choices (the IHS functional form and the saturation of the fixed-effect structure), we assess whether the main findings are sensitive to either.

Table A.8 varies the fixed-effect structure while holding the IHS functional form fixed, replacing the interacted FE with additive country and month fixed effects. Across all three outcome families, the arrival-month ($k = 0$) $\text{Treat} \times \text{Post}$ impulse remains positive and statistically significant for both platform types

and for the all- and anonymous-transaction groups (at the 1 percent level throughout); the identified-transaction coefficients are markedly smaller but stay positive and significant at conventional levels. The qualitative ordering established in the baseline is preserved: non-tax-haven effects exceed tax-haven effects, and estimates for identified transactions are markedly smaller than those for anonymous transactions. The stability of the sign, significance, and relative ordering indicates that the main results are properties of the data rather than artifacts of the saturation imposed by the interacted FE structure.

The second modelling choice is the IHS functional form itself, which we examine separately below in Table A.9.

Table A.8. Impact of Foreign Aid on On-Chain Crypto Activity (Additive Fixed Effects). This table re-estimates the baseline arrival-month impulse under additive country and month fixed effects, replacing the interacted $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ baseline. A separate platform fixed effect is not identified on the spike panel, whose two platform types define the table columns. The reported coefficient is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes), estimated on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C); columns distinguish all, anonymous, and identified transactions by platform type. Per-column within-adjusted R^2 is reported. Standard errors in parentheses, clustered at the recipient-country 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: Transaction Volume						
Treat $\times$ Post	0.810*** (0.218)	0.802*** (0.215)	0.339** (0.164)	1.489*** (0.258)	1.474*** (0.255)	0.716*** (0.184)
Country FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01
Panel B: Transaction Frequency						
Treat $\times$ Post	0.806*** (0.216)	0.804*** (0.216)	0.322** (0.129)	1.391*** (0.252)	1.383*** (0.251)	0.677*** (0.160)
Country FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01
Panel C: New Accounts						
Treat $\times$ Post	1.600*** (0.417)	1.600*** (0.417)	0.161* (0.090)	3.151*** (0.533)	3.151*** (0.533)	0.294** (0.116)
Country FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01

The second modelling choice is the IHS functional form. Because the on-chain outcome is zero in the large majority of treated pre-disbursement cells, the IHS coefficient is not scale-invariant: its magnitude depends on the units in which Bitcoin activity is measured, and the implied average effect is not point-identified once the extensive margin moves (Chen and Roth, 2024). We therefore re-estimate the identical stacked Treat $\times$ Post (the same Ctry. $\times$ Plat. $\times$ Ep. and month fixed effects, the same inverse-episode stack weights, the same no-overlap controls, and the same cluster-robust standard errors) under three additional estimators that do not depend on the IHS scale. First, PPML (Poisson pseudo-maximum-

likelihood on levels; Santos Silva and Tenreyro, 2006) is scale-equivariant, so rescaling the outcome shifts only the absorbed constant and whatever verdict $\hat{\beta}$ delivers, significant or not, is a property of the data rather than of the Bitcoin unit in which activity is measured; the coefficient reads as a semi-elasticity with $IRR = \exp(\hat{\beta})$. Implemented natively with an iteratively reweighted least-squares routine that projects out the fully interacted fixed effects, PPML converges on the complete interacted-FE panel, not a restricted subsample, so it delivers a like-for-like comparison with the baseline. Second, the extensive-margin estimator is a linear probability model for $\mathbf{1}\{y > 0\}$, the percentage-point change in the probability of any flow. Third, the intensive-margin estimator is $\log y$ among the positive cells only. Together the extensive and intensive rows decompose what the single IHS coefficient is averaging over.

Table A.9 reports the linear TWFE–IHS baseline alongside the three alternative estimators. On the sparse flow-based tranche-arrival panel the levels-based PPML impulse is imprecise: the arrival-month level change is dominated by a handful of large-tranche cells at nearby event times, so the $k = 0$ PPML coefficient is estimated with wide standard errors and is not statistically significant. PPML is therefore uninformative on this panel and neither confirms nor overturns the baseline. The scale-free confirmation comes instead from the extensive margin, which does not depend on the IHS units: treatment raises the probability of any anonymous flow by about 13 percentage points on tax-haven and 18 percentage points on non-tax-haven exchanges ($p < 0.01$), consistent with aid episodes activating anonymous on-chain activity that was previously absent. The intensive margin, identified off a small, selected subsample of already-active cells, is insignificant throughout; we report it for completeness but do not rely on it. The anonymity gradient persists, with the anonymous column tracking the all-activity column across estimators. Overall, the baseline arrival-month impulse is corroborated on the scale-free extensive margin and is best understood as a predominantly extensive-margin effect.

Table A.9. Identification Robustness: Alternative Estimators of the Stacked Treat×Post. This table re-estimates the baseline average treatment effect with estimators that do not rely on the inverse-hyperbolic-sine (IHS) scale, addressing the concern that IHS coefficients are not scale-invariant when the outcome contains many zeros (Chen and Roth, 2024). The three alternative estimators are benchmarked against the linear TWFE–IHS baseline (top row of each block; the arrival-month impulse reported in Tables 1–A.7). PPML is Poisson pseudo-maximum-likelihood on levels (Santos Silva and Tenreiro, 2006), which is scale-equivariant, so its verdict cannot depend on the Bitcoin unit ($IRR = \exp(\hat{\beta})$); Extensive is a linear probability model for $\mathbf{1}\{y > 0\}$; Intensive is $\log y$ among positive cells. Each estimator is a full stacked event study on the flow-based tranche-arrival panel (window $k = -8, \dots, +5$, reference $k \in [-12, -9]$); the reported coefficient is the arrival-month ($k = 0$) dummy, estimated with a no-overlap control group; columns distinguish all, anonymous, and identified activity by platform type. On the sparse spike panel the PPML impulse is estimated imprecisely, because levels are dominated by a few large-tranche cells at adjacent event times. All specifications include Ctry.×Plat.×Ep. and month fixed effects. Standard errors in parentheses. Standard errors for the linear TWFE–IHS baseline row are clustered at the recipient-country level, matching Tables 1–A.7; standard errors for the PPML, extensive, and intensive rows are 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: Transaction Volume						
Linear TWFE–IHS (baseline)	0.866*** (0.234)	0.864*** (0.232)	0.301* (0.168)	1.396*** (0.273)	1.387*** (0.270)	0.622*** (0.189)
PPML (levels)	−0.369 (0.663)	−0.390 (0.668)	−0.166 (0.694)	0.927 (1.148)	0.861 (1.100)	1.420 (1.564)
Extensive $\mathbf{1}\{y > 0\}$	0.128*** (0.030)	0.128*** (0.030)	0.128*** (0.030)	0.179*** (0.035)	0.179*** (0.035)	0.179*** (0.035)
Intensive $\log y \mid y > 0$	−0.226 (0.638)	−0.226 (0.638)	−0.278 (0.640)	−0.289 (0.613)	−0.286 (0.614)	−0.348 (0.606)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel B: New Accounts						
Linear TWFE–IHS (baseline)	1.775*** (0.445)	1.775*** (0.445)	0.044 (0.108)	2.832*** (0.550)	2.832*** (0.550)	0.260** (0.114)
PPML (levels)	−2.434 (1.661)	−2.432 (1.661)	−0.022 (1.122)	−0.561 (0.883)	−0.554 (0.882)	0.847 (1.936)
Extensive $\mathbf{1}\{y > 0\}$	0.127*** (0.030)	0.127*** (0.030)	0.011 (0.023)	0.181*** (0.035)	0.181*** (0.035)	0.097*** (0.024)
Intensive $\log y \mid y > 0$	−0.588 (0.778)	−0.588 (0.778)	1.107 (1.446)	−0.342 (0.677)	−0.341 (0.677)	−2.931** (1.151)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499
Panel C: Transaction Frequency						
Linear TWFE–IHS (baseline)	0.889*** (0.231)	0.888*** (0.230)	0.327** (0.133)	1.310*** (0.265)	1.306*** (0.263)	0.628*** (0.166)
PPML (levels)	−1.353 (0.920)	−1.370 (0.927)	−0.483 (0.696)	0.822 (1.037)	0.792 (1.019)	1.457 (1.477)
Extensive $\mathbf{1}\{y > 0\}$	0.128*** (0.030)	0.128*** (0.030)	0.128*** (0.030)	0.179*** (0.035)	0.179*** (0.035)	0.179*** (0.035)
Intensive $\log y \mid y > 0$	−0.253 (0.635)	−0.252 (0.635)	−0.351 (0.636)	−0.352 (0.612)	−0.351 (0.612)	−0.393 (0.609)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓ 20	✓	✓	✓	✓
Obs.	18,499	18,499	18,499	18,499	18,499	18,499

E.3 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.10 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 arrival-month ($k = 0$) impulse is closely aligned with the baseline across both specifications. Anonymous transaction volume on tax-haven exchanges is 0.864 and 0.863 under the two restrictions, essentially identical to the headline arrival-month impulse. Transaction frequency and new anonymous account creation show the same stability, and the anonymous-to-identified gap is preserved throughout. The observation counts shift only modestly (from 18,499 to 18,274 under the no-neighbor filter), 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.10. Impact of Foreign Aid on Crypto Activity across Alternative Control Compositions. This table re-estimates the baseline arrival-month impulse under two alternative control-group restrictions: Dev. Ctrl limits controls to developing economies with no disbursement in the window (the headline sample), and No Nbr additionally drops countries sharing a land border with a treated country. The reported coefficient is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes), estimated on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C); columns distinguish all, anonymous, and identified transactions by platform type. All specifications include $\text{Ctry.} \times \text{Plat.} \times \text{Ep.}$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country 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	0.866*** (0.234)	0.865*** (0.234)	0.864*** (0.232)	0.863*** (0.231)	0.301* (0.168)	0.300* (0.168)	1.396*** (0.273)	1.396*** (0.273)	1.387*** (0.270)	1.387*** (0.270)	0.622*** (0.189)	0.622*** (0.190)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274
Adj. R ²	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Panel B: Transaction Frequency												
Treat×Post	0.889*** (0.231)	0.888*** (0.231)	0.888*** (0.230)	0.887*** (0.230)	0.327** (0.133)	0.326** (0.133)	1.310*** (0.265)	1.310*** (0.265)	1.306*** (0.263)	1.306*** (0.264)	0.628*** (0.166)	0.628*** (0.166)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274
Adj. R ²	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Panel C: New Accounts												
Treat×Post	1.775*** (0.445)	1.771*** (0.444)	1.775*** (0.445)	1.771*** (0.444)	0.044 (0.108)	0.044 (0.108)	2.832*** (0.550)	2.833*** (0.550)	2.832*** (0.550)	2.833*** (0.550)	0.260** (0.114)	0.260** (0.114)
Ctry.×Plat.×Ep. FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs.	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274	18,499	18,274
Adj. R ²	0.01	0.01	0.01	0.01	0.00	0.00	0.01	0.01	0.01	0.01	0.01	0.01

E.4 Reference-Period Diagnostic

Figure A.1 plots the anonymous tax-haven transaction-volume event study under the two candidate omitted references discussed in Section 4.1. The two coefficient paths differ only by a vertical shift of under 0.1 IHS units: the arrival-month impulse is 0.85 under the far-pre reference $k \in [-12, -9]$ (the paper’s choice) and 0.77 under $k = -1$, both significant at the one percent level, and the clean pre-period $k \in \{-8, \dots, -5\}$ is flat and jointly insignificant under either normalization. The only pre-arrival movement is a shallow, localized dip at $k \in \{-4, -3\}$ (shaded), which the four-month cleanliness screen mechanically enforces; normalizing to $k = -1$ places the reference at the top of the recovery from this dip and thereby turns it into an apparent anticipatory “ramp,” whereas the far-pre reference reveals it as a confined dip against an otherwise flat pre-period.

Figure A.1. Sensitivity of the Event Study to the Omitted Reference Period. This figure plots stacked event-study coefficients for geo-adjusted anonymous transaction volume on tax-haven exchanges under two normalizations of the omitted reference period: $k = -1$, the convention standard in the literature, and the far-pre bin $k \in [-12, -9]$ used throughout the paper. Estimates are computed on the flow-based tranche-arrival panel with developing-country controls, over the panel window $k = -12, \dots, +5$ in months relative to the tranche-arrival month ($k = 0$), on IHS-transformed outcomes by linear TWFE with a no-overlap control group. The clean pre-period is $k \in \{-8, \dots, -5\}$, and the window $k \in \{-4, \dots, -2\}$ is mechanically quiet by construction of the four-month screen that defines a treated episode. All specifications include country-by-platform-by-episode (Ctry. $\times$ Plat. $\times$ Ep.) and month fixed effects. Standard errors are clustered at the recipient-country level. Shaded bands denote 95% confidence intervals.



E.5 Mixer Crowd-Out and Ordinary Privacy Demand

This subsection reports the mixer service-comparison test discussed in Section 5.4. The table compares CoinJoinMess mixer transaction counts with gambling, payment, CEX, and DEX services around aid disbursement, using no-overlap stacked TWFE with standard errors clustered at calendar month.

Table A.11. Mixer Activity Relative to Other Crypto Services around Aid Disbursement. This table reports the effect of foreign-aid disbursements on mixer activity measured relative to other on-chain service categories. The outcome is the IHS transaction-count difference between CoinJoinMess mixers and the comparison service category, constructed from global monthly time series, so that a positive coefficient indicates that mixer transaction counts rise more than those of the comparison service after disbursement. Estimates are stacked TWFE with a no-overlap control group. Columns $k = 0$ and $k = 1$ report immediate post-disbursement event-study coefficients, and “Full post” averages over $k = 0, \dots, 4$ (omitting $k = -1$), matching the post window used in the main text. All specifications include country $\times$ episode fixed effects; calendar-month fixed effects are omitted because the global service outcomes are constant within calendar month. Standard errors in parentheses, clustered at the calendar-month level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)
Contrast	$k = 0$	$k = 1$	Full post ($k = 0, \dots, 4$)
Mixer – Gambling	0.006 (0.047)	0.018 (0.077)	0.097 (0.140)
Mixer – Payment	0.139*** (0.038)	0.279*** (0.054)	0.642*** (0.103)
Mixer – CEX	0.085*** (0.033)	0.170*** (0.044)	0.410*** (0.083)
Mixer – DEX	0.092 (0.075)	0.196** (0.092)	0.622*** (0.203)
Country $\times$ episode FE	✓	✓	✓
Obs.	75,423	75,423	75,423
Adj. R ²	0.00	0.00	0.00

E.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 to September 2021 window from both the treatment and control samples.

Table A.12 reports the results. The arrival-month ($k = 0$) impulse remains 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 0.931, and anonymous transaction frequency is 0.963; new anonymous account creation is 1.943 on tax-haven and 3.031 on non-tax-haven exchanges, all closely tracking the headline arrival-month magnitudes. The anonymous-to-identified gap is preserved throughout, and the modest reduction in observation counts (from 18,499 to 17,736) confirms that few episodes fall within the excluded window. These results confirm that the 2021 DeFi boom does not drive our findings.

Table A.12. Robustness Check: DeFi-Boom Exclusion. This table re-estimates the baseline arrival-month impulse after dropping the May to September 2021 DeFi window. The reported coefficient is the arrival-month impulse (the coefficient on $\text{Treat} \times \text{Post}$ at event time $k = 0$) from a stacked event-study difference-in-differences on the flow-based tranche-arrival panel (328 episodes), estimated on IHS-transformed outcomes by linear TWFE with a no-overlap control group. Panels distinguish transaction volume (Panel A), transaction frequency (Panel B), and new accounts (Panel C); columns distinguish all, anonymous, and identified transactions by platform type. Per-column within-adjusted R^2 is reported. Standard errors in parentheses, clustered at the recipient-country 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: Transaction Volume						
Treat $\times$ Post	0.936 ^{***} (0.233)	0.931 ^{***} (0.230)	0.345 ^{**} (0.167)	1.474 ^{***} (0.284)	1.462 ^{***} (0.281)	0.679 ^{***} (0.196)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	17,736	17,736	17,736	17,736	17,736	17,736
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01
Panel B: Transaction Frequency						
Treat $\times$ Post	0.965 ^{***} (0.231)	0.963 ^{***} (0.230)	0.365 ^{***} (0.133)	1.378 ^{***} (0.277)	1.372 ^{***} (0.275)	0.672 ^{***} (0.172)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	17,736	17,736	17,736	17,736	17,736	17,736
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01
Panel C: New Accounts						
Treat $\times$ Post	1.943 ^{***} (0.441)	1.943 ^{***} (0.441)	0.049 (0.118)	3.031 ^{***} (0.577)	3.031 ^{***} (0.577)	0.272 ^{**} (0.127)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓	✓
Obs.	17,736	17,736	17,736	17,736	17,736	17,736
Adj. R^2	0.01	0.01	0.01	0.01	0.01	0.01

F Project-Category Heterogeneity

This appendix reports the project-category heterogeneity that underlies the capture intensities used in Section 6. Table A.13 re-estimates the arrival-month impulse separately within each project category, on that category’s own subset of the 328 spike episodes. The specification, the panel construction, and the fixed effects are identical to the baseline; only the estimation sample changes, so each column is a within-category replication of the headline design rather than an interaction model.

Episodes are assigned to categories as follows. We classify the universe of World Bank projects approved in fiscal years 2018–2024 into eleven sector categories on the basis of the project title, using case-insensitive regular-expression matching with word-boundary anchors against category-specific keyword sets (Health; Education; Transport; Energy; Water and Sanitation; Agriculture and Rural Development; Governance and Public Sector; Finance and Private Sector; Social Protection and Jobs; Digital and Communications; Environment, Climate, and Disaster). Each spike episode is then tagged by its dominant project sector, defined as the sector receiving the largest IBRD/IDA disbursement flow in the episode’s arrival month. We use the title-based sector map rather than the World Bank’s own aggregate sector codes because the latter pool at a coarser level than the heterogeneity we want to detect.

The table reports two outcome families, transaction volume and new accounts, each split into all, anonymous, and identified activity, and each estimated separately on tax-haven and non-tax-haven exchanges. Two categories carry little information. Environment, Climate, and Disaster contains no qualifying spike episode and is reported as “—” throughout. Digital and Communications rests on only three episodes, so its coefficients, including the significant non-tax-haven entries, should be read with caution. The final row reports the number of country-by-platform-by-episode cells contributing to each column, which varies with the number of episodes assigned to the category.

The row carried into the main text is anonymous transaction volume on tax-haven exchanges, the paper’s headline outcome. Those ten estimates and their standard errors are the capture intensities that enter Table 10 and Figure 6.

Table A.13. Project-Category Heterogeneity in the Arrival-Month Impulse. This table reports the arrival-month impulse estimated separately for each project category. Each of the 328 spike episodes is assigned to the project category accounting for the largest IBRD/IDA disbursement flow in its arrival month, and the specification is re-estimated on that category's episode subset. The reported coefficient is the arrival-month impulse (the coefficient on $Treat \times Post$ at event time $k = 0$) from the baseline stacked event-study difference-in-differences on IHS-transformed outcomes. Rows report transaction volume and new accounts, each split into all, anonymous, and identified activity, for tax-haven and non-tax-haven exchanges. Environment, Climate, and Disaster contains no qualifying spike episode and is reported as "—"; Digital and Communications is estimated on three episodes. Category abbreviations: Gov. = Governance and Public Sector; Env. = Environment, Climate, and Disaster; Agric. = Agriculture and Rural Development; Educ. = Education; Fin. = Finance and Private Sector; Trans. = Transport; SocProt = Social Protection and Jobs. All specifications include $Ctry. \times Plat. \times Ep.$ and month fixed effects. Standard errors in parentheses, clustered at the recipient-country level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	Water	SocProt	Gov.	Health	Env.	Fin.	Trans.	Agric.	Educ.	Digital	Energy
Transaction volume											
<i>Tax Haven</i>											
All	1.510 [*] (0.789)	1.528 ^{***} (0.569)	0.745 ^{**} (0.376)	-0.263 (0.514)	—	0.920 (0.718)	2.713 ^{***} (0.913)	0.503 (1.559)	0.386 (0.635)	0.028 (0.048)	0.579 (0.634)
Anonymous	1.528 ^{**} (0.773)	1.518 ^{***} (0.565)	0.743 ^{**} (0.373)	-0.255 (0.507)	—	0.913 (0.708)	2.676 ^{***} (0.897)	0.521 (1.554)	0.405 (0.622)	0.029 (0.048)	0.582 (0.628)
Identified	0.468 (0.456)	0.820 [*] (0.421)	0.112 (0.246)	-0.323 (0.364)	—	0.364 (0.519)	1.568 ^{**} (0.669)	-0.113 (0.626)	0.118 (0.534)	-0.016 (0.019)	-0.022 (0.377)
<i>Non-Tax Haven</i>											
All	1.438 (1.130)	1.737 ^{***} (0.599)	1.722 ^{***} (0.510)	-0.178 (0.583)	—	1.058 (0.735)	2.588 ^{**} (1.090)	0.586 (1.517)	0.447 (0.695)	2.016 ^{**} (0.949)	2.056 ^{***} (0.793)
Anonymous	1.440 (1.112)	1.731 ^{***} (0.593)	1.710 ^{***} (0.504)	-0.168 (0.577)	—	1.046 (0.722)	2.548 ^{**} (1.071)	0.616 (1.514)	0.454 (0.681)	1.976 ^{**} (0.935)	2.040 ^{***} (0.784)
Identified	0.549 (0.713)	0.794 [*] (0.435)	0.743 ^{**} (0.332)	-0.200 (0.415)	—	0.582 (0.577)	1.690 [*] (0.880)	-0.443 (0.652)	0.192 (0.558)	0.462 (0.350)	0.717 (0.473)
New accounts											
<i>Tax Haven</i>											
All	3.196 ^{**} (1.463)	2.553 ^{**} (1.072)	1.486 ^{**} (0.679)	-0.371 (0.960)	—	2.200 (1.426)	5.491 ^{***} (1.938)	1.609 (2.597)	0.979 (1.465)	0.045 (0.103)	1.347 (1.273)
Anonymous	3.196 ^{**} (1.463)	2.553 ^{**} (1.072)	1.486 ^{**} (0.679)	-0.371 (0.960)	—	2.200 (1.426)	5.491 ^{***} (1.939)	1.608 (2.597)	0.979 (1.465)	0.045 (0.103)	1.347 (1.273)
Identified	-0.818 (0.888)	0.120 (0.219)	-0.057 (0.114)	-0.261 (0.228)	—	0.229 (0.366)	1.259 ^{**} (0.597)	-0.152 (0.411)	-0.071 (0.347)	-0.012 (0.012)	0.002 (0.147)
<i>Non-Tax Haven</i>											
All	3.482 (2.215)	3.240 ^{***} (1.161)	3.600 ^{***} (1.032)	-0.663 (1.160)	—	2.130 (1.492)	5.323 ^{**} (2.202)	1.327 (3.017)	0.787 (1.465)	5.921 ^{**} (2.744)	4.294 ^{**} (1.698)
Anonymous	3.483 (2.215)	3.240 ^{***} (1.161)	3.600 ^{***} (1.032)	-0.662 (1.160)	—	2.130 (1.492)	5.322 ^{**} (2.202)	1.326 (3.018)	0.788 (1.465)	5.921 ^{**} (2.744)	4.294 ^{**} (1.698)
Identified	-0.750 (0.811)	0.380 (0.267)	0.244 (0.148)	0.131 (0.207)	—	0.249 (0.359)	1.469 [*] (0.784)	-0.188 (0.495)	-0.479 [*] (0.269)	-0.066 (0.049)	0.259 (0.243)
Ctry. $\times$ Plat. $\times$ Ep. FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Month FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
N	847	2,891	5,934	2,336	—	2,065	1,078	661	793	228	1,666