INTERNATIONAL MONETARY FUND

International Stablecoin Flows, Capital Controls and Currency Crises

Marco Reuter

WP/26/211

IMF Working Papers **describe research in progress by the author(s) and are published to elicit comments and to encourage debate.**
The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

2026
OCT

WORKING PAPER

© 2026 International Monetary Fund WP/26/211

IMF Working Paper
Research Department

International Stablecoin Flows, Capital Controls and Currency Crises

Prepared by Marco Reuter*

Authorized for distribution by Maria Soledad Martinez Peria
October 2026

***IMF Working Papers* describe research in progress by the author(s) and are published to elicit comments and to encourage debate.** The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: This paper constructs a country-level panel of estimated international stablecoin flows, covering 188 countries, totaling $21.4 trillion in gross flows between 2018 and 2025. Activity is disproportionately concentrated in emerging markets and developing economies, which hold three quarters of stablecoins. When countries restrict capital outflows, stablecoin activity rises by roughly 30 percent over the following six months, then plateaus, and does not fall back when restrictions are removed. When currency crises hit, stablecoin holdings increase by roughly 75 percent within a year, and the new holdings persist. Both findings suggest that dollarization may now also operate through stablecoins.

JEL Classification Numbers:	C58, F31, G15
Keywords:	stablecoins, capital flows, capital flow management measures (CFMs), currency crises, dollarization, currency substitution, crypto assets
Author's E-Mail Address:	mreuter@imf.org

* This paper supersedes a previous version, "Decrypting Crypto: How to Estimate International Stablecoin Flows", published as IMF WP 2025/141, which developed the region-level methodology this version builds on. I thank Cage Englander for excellent research assistance. I am grateful to Itai Agur, Alexander Copestake, Mai Dao, Alessia De Stefani, Andrés Fernández, Ulf Lewrick (discussant), Gordon Liao (discussant) and Maria Soledad Martinez Peria, and to audiences at seminars at the IMF, the IMF's 7th Annual Macro-Financial Conference, the SNB-CIF Conference on Cryptoassets and Financial Innovation, and the 4th Bonn/Mannheim/Frankfurt Workshop on Digital Money.

WORKING PAPERS

International Stablecoin Flows, Capital Controls and Currency Crises

Prepared by Marco Reuter

1 Introduction

Stablecoins—crypto assets whose value is pegged to a fiat currency, most often the US dollar—have grown from a niche instrument for crypto traders into an asset with potentially significant macroeconomic implications (Adrian et al., 2025). We estimate that stablecoin transfers exceeded $7 trillion in 2025, up from $1.2 billion in 2018, and that roughly 90 percent of this volume crossed borders, while the outstanding stock at year-end was roughly $256 billion.[1] Flows of this size have not gone unnoticed, and policymakers have called for better monitoring of crypto transactions and international crypto asset flows (BIS (2023), EU (2023), G7 (2023) , FATF (2023), FSB (2023), IMF (2023), US Treasury (2023)). Further, they have raised concerns, particularly in Emerging Markets and Developing Economies (EMDEs), that stablecoins could facilitate the circumvention of capital flow management measures (CFMs) and dollarization, acknowledging that they can also make cross-border payments and remittances cheaper and faster (Adrian et al., 2025; Basu et al., 2026). Yet, official statistics currently do not comprehensively measure stablecoin flows, resulting in a blind spot at the center of the policy debate.[2]

In this paper, we narrow this blind spot and make three contributions. First, we develop a methodology that converts pseudonymous stablecoin transactions into country-level capital flow estimates, validated out of sample and against Brazil's official statistics[3] and internal figures from crypto exchanges. Second, we construct a country-level panel of international stablecoin flows, covering 188 countries, totaling $21.4 trillion of gross stablecoin transfers between 2018 and 2025, which is publicly available.[4] Third, to the best of our knowledge, we provide the first event-study evidence that stablecoin activity increases when countries impose CFMs, and that stablecoin holdings increase when countries enter currency crises. In both settings the response is persistent: stablecoin activity does not decline when CFMs are removed, and crisis-driven holdings settle at their higher level rather than reverting. This evidence is consistent with the concerns about stablecoins that have been raised by policy-

[1] We focus on USDC and USDT, the two most popular stablecoins, accounting for over 90% of outstanding stablecoins according to CoinMarketCap.

[2] Efforts to reduce this blind spot are underway. The seventh edition of the IMF's balance of payments manual now specifies how crypto assets, including stablecoins, enter external sector statistics (IMF, 2025a), and the G20 Data Gaps Initiative seeks to develop a common reporting framework for digital money, including stablecoins (IMF, 2025b). The Central Bank of Brazil publishes crypto asset transactions between residents and nonresidents alongside its balance of payments statistics (Banco Central do Brasil, 2024), and the BIS's Project Atlas prototypes cross-border crypto flow statistics from on-chain data (BIS Innovation Hub, 2023).

[3] Brazil's central bank is, to the best of our knowledge, the only central bank that systematically records such activity and publishes a series that can be used for external validation.

[4] Our estimates cover simple stablecoin transfers involving centralized exchanges and sufficiently active self-custodial wallets. They exclude decentralized-finance activity, automated trading, and operational movements inside exchanges. See Schär et al. (2026) for a detailed taxonomy.

makers and resembles the hysteresis that has been long documented for dollar-denominated bank deposits (see, e.g., Guidotti and Rodríguez, 1992; Valev, 2010; Levy Yeyati, 2021).

The reason such a dataset can be built at all is a common misconception about crypto assets. Contrary to popular belief, the vast majority of crypto assets, including stablecoins, do not provide anonymity. Every on-chain transaction is publicly recorded on a freely accessible ledger known as a blockchain. The perception of anonymity arises because blockchain data is pseudonymized: rather than personal information such as names or residences, blockchains log only the wallet addresses of senders and receivers (long hexadecimal strings such as `'0xdFDEe1155E1dd7c01774560C6E98C41B7da945dB'`), which do not directly reveal anything about the user. The measurement problem is therefore not to observe stablecoin flows, which anyone can, but to attach geography to the wallet addresses that send and receive them.

We attach geography in two steps. First, we build training data from wallets whose region can be inferred directly. Some carry human-readable domain names, purchased through services such as the Ethereum Name Service (ENS) to replace the hexadecimal address string for convenience. A large language model reads the linguistic and cultural markers of almost 6 million such names, of which only a small share identify a region. The owner of 'pijiu.eth' (*pijiu* is Chinese for "beer") is assigned to Asia and the Pacific. Others transact almost exclusively with centralized crypto exchanges that serve a single regional market, and are assigned to that region.[5] Together, the two approaches label 437,775 wallets across five world regions. Second, we train a machine learning model on the on-chain behavior of the labeled wallets—their time-of-day activity patterns, their (non)adherence to daylight saving time, and the exchanges, tokens, and smart contracts they use—and apply it to estimate a regional probability distribution for any sufficiently active self-custodial wallet. Regional estimates are then split into country-level estimates using, within each region, the web traffic of the exchanges each wallet interacts with.[6] The key identifying assumption is that, conditional on the behavioral features the model is trained on, labeled and unlabeled wallets follow the same data generating process.

The methodology is validated at each step. First, converted to region-local time, activity among labeled wallets dips at night, and wallets labeled North American, for example, shift by one hour when daylight saving time begins, while wallets labeled Asian do not. Second, the trained model classifies 74 percent of held-out wallets into the correct region, against a

[5] A centralized exchange (CEX), such as Binance or Coinbase, converts fiat currency to crypto assets and custodies them for its users; a self-custodial wallet, in contrast, is controlled directly by its user.

[6] Flows involving China, where web traffic is distorted because exchange websites are blocked, are not attributed on the basis of web traffic, but instead handled by a separate, dedicated ML model. Most websites of centralized crypto exchanges cannot be accessed from China without a VPN. A VPN (Virtual Private Network) routes internet traffic through a remote server, often in a different country, and conceals the user's true location.

benchmark of 20 percent for random guessing among five regions. Third, we compare our estimates to two sources of external validation that do not share the methodology's inputs. Brazil's central bank records residents' purchases of stablecoins from abroad from the foreign exchange transactions that pay for them. We estimate the same purchases from blockchain data alone, and the two series agree within 15 percent. Internal statistics from crypto exchanges, disclosed in enforcement proceedings and in reporting, sit within a few percentage points of the shares we assign to the United States and China at those exchanges. Finally, the associations of inflows with macroeconomic instability and with the cost of traditional remittances established in the literature (Auer et al., 2025b) also appear in our panel.

Turning to our estimated panel dataset, a first set of findings establishes that stablecoin flows are predominantly cross-border, retail-sized, and concentrated in EMDEs. The estimated cross-border share is steady at roughly 90 percent throughout the sample, and a median transfer of $300 against an average of $10,853 is consistent with retail usage. Activity concentrates in EMDEs, which we estimate receive 56 percent of gross flows. Flows between EMDEs are the largest corridor, accounting for roughly a third of cross-border volume, while flows between advanced economies (AEs) account for only 14 percent. In contrast, in traditional payments the proportions are reversed, as 85 percent of the volume processed by SWIFT flows between AEs (Basu et al., 2026). We estimate that, in absolute terms, the largest recipient countries in 2025 were the United States and China[7]. However, relative to the size of the economy, nineteen of the twenty largest recipients are EMDEs.[8] Holdings concentrate in EMDEs even more strongly than flows. Our panel covers roughly $200 billion in stablecoin holdings at the end of 2025 (relative to an outstanding stock of $256 billion), of which we estimate EMDEs hold about three quarters. Further, all of the twenty largest estimated holdings relative to GDP are in EMDEs, reaching 5.6 percent of GDP in Venezuela, 4.1 percent in Ukraine, and 1.4 percent in Nigeria, against 0.04 percent in the United States.

A second set of findings documents the economic drivers of these flows. Policymakers worry that stablecoins facilitate the circumvention of CFMs and dollarization, and we provide the first direct evidence consistent with both concerns. Using a stacked event-study design covering 217 activations of CFMs, we document that the restricting country's cross-border stablecoin activity rises by about 4 percent in the activation month and roughly 20

[7]For methodological reasons, estimates for China in this paper combine mainland China, Hong Kong SAR, and Macao SAR, whose on-chain behavior is too similar to distinguish. Private virtual currencies, including stablecoins, are prohibited in mainland China, while Hong Kong SAR and Macao SAR apply different regulatory regimes. Estimated activity in this combined unit is therefore consistent with usage in Hong Kong SAR and Macao SAR, with circumvention of the mainland prohibition, or both.

[8]The only AE in the top 20 is Andorra, likely an outlier due to its small GDP.

percent within six months, relative to countries that activate no restriction of their own in the surrounding months. The response is largest and most persistent after tightenings of *outflow* restrictions, building to roughly 30 percent within six months and plateauing there through the twelfth month. Its economic magnitude amounts to roughly $200 million per month of additional cross-border stablecoin activity for the median treated country. We also provide evidence that the response ratchets: removals of restrictions are followed by no corresponding decline, and we formally reject the hypothesis that removals fully undo the tightening response.

Further evidence that stablecoins meet demand for dollars comes from currency crises. Using the same stacked event-study design, now around 21 currency crises between 2021 and 2025, identified with the standard criterion of Laeven and Valencia (2020), we estimate that stablecoin holdings increase by roughly 75 percent over the year following crisis entry, relative to similar countries that do not enter a crisis. In total, the increase amounts to roughly $7.9 billion. For the median crisis economy, our estimates imply roughly $59 million of additional net accumulation over that year, and between roughly $16 million and $215 million at the quartiles of economy size. These magnitudes reflect a young ecosystem, as most crises in our sample occurred in 2021 and 2022, when the outstanding stock of stablecoins was roughly half its end-of-2025 level. The accumulation is timed to the crisis, as holdings show no movement ahead of the dated onset and build steadily over the year that follows. The response ratchets here as well: holdings plateau after the first year rather than reverting.

Related Literature. Methodologically, the closest paper is Athey et al. (2016), the proof of concept for inferring the geography of crypto wallets from their on-chain behavior: they train a random forest on 2,858 Bitcoin wallet entities identified from addresses posted on an online forum. We expand their design in three ways. Wallet domain names, which had yet to be invented when their data end in 2015, yield a training set two orders of magnitude larger (437,775 wallets); our model predicts the region of *any* sufficiently active self-custodial wallet rather than of a pre-identified subset; and we use the result to map bilateral flows at the country level.[9]

The empirical literature on international crypto flows relies on the country-level estimates of the commercial provider Chainalysis (Cardozo et al., 2024; Cerutti et al., 2025a; Auer et al., 2025b; Friedrich and Zhao, 2026; Hofmann et al., 2026), which allocates transfers between centralized exchanges in proportion to the web traffic of each exchange's website (Chainalysis,

[9]Heuristic clustering links addresses to common owners (Meiklejohn et al., 2013; Makarov and Schoar, 2021), and Béres et al. (2021) use domain names and time-of-day activity, the raw ingredients of our methodology, to profile individual users. None predicts where wallets are located, the step that turns blockchain records into capital flow statistics.

2025).[10] We contribute the first comprehensive country-level panel that does not rely on this allocation, built instead from the wallet-level attribution described above. Because it attributes wallets rather than exchange websites, our panel systematically captures the activity of self-custodial wallets, and because it records the issuance and redemption of stablecoins, it measures holdings as well as flows, which no dataset built on transfers between exchanges can. Papers in this literature rely on those Chainalysis estimates. On gross flows in the cross section, the object they mainly use, we recover—in a different dataset—the finding that flows are larger under macroeconomic instability and when traditional remittances are costly.[11] The two datasets differ on levels, high-frequency net flows, holdings, and remaining time-series variation, and that is the variation the event studies below use.

We provide the first event-study evidence on whether crypto assets undermine restrictions on capital flows, a concern policymakers articulated early (He et al., 2022). The capital-controls literature evaluates whether restrictions manage the traditional flows they target (Fernández et al., 2016; Forbes et al., 2015; Alfaro et al., 2017; Erten et al., 2021); whether activity migrates onto rails outside their perimeter is the complementary margin, and the existing evidence on it pulls in two directions. von Luckner et al. (2023) document, in the two oldest peer-to-peer Bitcoin markets, that Bitcoin is increasingly used to receive remittances and evade capital controls in emerging markets, von Luckner et al. (2024) show that crypto price premia in restricted economies track excess dollar demand and control stringency, and Hu et al. (2021) attribute over a quarter of Chinese Bitcoin exchange volume to capital flight. Crypto activity is also higher where the capital account is more restricted (Alnasaa et al., 2022; Friedrich and Zhao, 2026). Yet the two papers that look directly at stablecoin flows find no response. Auer et al. (2025b) find small coefficients of opposing signs that are significant in one of four specifications, and Hofmann et al. (2026) find stablecoin flows largely unaffected by broad or stablecoin-specific restrictions, in contrast to foreign-currency deposits, which restrictions do suppress. Tracking 217 individual CFM activations, we find that stablecoin activity in the restricting country rises promptly, builds over the following six months, and does not decline when restrictions are removed.

Finally, we present the first evidence that stablecoin holdings rise when a country enters

[10] Aldasoro et al. (2026) and Auer et al. (2025a) measure related objects, currency-level stablecoin demand from order books of centralized exchanges, and country-level retail adoption based on app usage, but not bilateral country flows. Cong et al. (2026) infer wallet regions from transaction timing and find that web-traffic allocations overstate the cross-regional share of wallet-to-wallet volume.

[11] Auer et al. (2025b) find stablecoin flows larger under high inflation, Friedrich and Zhao (2026) find inflation positive and highly significant in cross-country flows, Cong et al. (2024) document the inflation channel at the account level for an Indian exchange, and Aramonte et al. (2022) describe "cryptoisation" under macroeconomic stress. On remittances, our corridor-level elasticity is in line with those Auer et al. (2025b) report for corridors from AEs to EMDEs, and Du et al. (2026) and Reslow ⓡ al. (2026) document that transfers over stablecoin rails can be substantially cheaper than traditional providers.

a currency crisis, extending the classic literature on crisis-driven dollarization to stablecoins. This literature established for bank deposits that the currency substitution a crisis induces persists long after the crisis ends (Guidotti and Rodríguez, 1992; Kamin and Ericsson, 2003; Reinhart et al., 2014; Valev, 2010). Recent work documents that deposit dollarization remains widespread (Levy Yeyati, 2021) and rationalizes it as insurance, with dollar holdings paying off precisely in depreciation states (Bocola and Lorenzoni, 2020; Dalgic, 2024; Christiano et al., 2021; Ferrante and Gornemann, 2022).[12] The deposit literature measures dollarization as the share of foreign-currency deposits in total deposits, observed annually for broad sets of countries (Honohan, 2007) and monthly only for narrow country groups (Neanidis and Savva, 2009). Our stablecoin holdings are the analogue of that stock, measured monthly for a wide panel of countries. Because they combine frequency with coverage, the path around crisis onset can be estimated against controls. Using our data, we show that stablecoin holdings begin accumulating only at crisis onset, build over the year that follows, and level off rather than revert. Azar et al. (2026) develop a model in which household stablecoin adoption makes capital mobility endogenous, and report higher stablecoin flows to ENS-named wallets in the week of nine stress episodes. Hofmann et al. (2026) treat historical crisis exposure as a country attribute, describing which countries are predisposed to stablecoin dollarization.

The rest of the paper is structured as follows. Section 2 explains which legs of an international crypto flow are observable on-chain. Section 3 describes the data, attaches geography to wallets, and reports the internal and external validation of the resulting estimates. Section 4 maps international stablecoin flows and stablecoin holdings, section 5 analyzes their economic drivers, and section 6 concludes.

2 Background

The next section builds the dataset at the core of this paper, a transaction-level record of stablecoin transfers in which each transfer is allocated across ordered country pairs using estimated probability distributions over countries for its sender and receiver. This section establishes the prerequisite for any such exercise. It follows an international stablecoin payment from beginning to end and clarifies which of its steps leave a record on the blockchain's public ledger ("on-chain") and are therefore observable, and which do not.

The typical international stablecoin flow begins with the exchange of fiat currency for a stablecoin. The dominant venue for this is a centralized exchange (CEX) such as Binance

[12]Beyond deposits, Drenik and Perez (2021) measure the dollarization of the prices posted on online marketplaces across Latin America.

or Coinbase, a firm that converts fiat currency into crypto assets and holds them for its customers. A stablecoin bought this way sits in the exchange's custody rather than the user's. Users transfer fiat currency to a CEX, which then facilitates the conversion into a stablecoin. These initial transfers of money to the CEX occur within the traditional fiat banking system, and therefore we do not observe them in our data. Figure 1 provides a schematic overview of the possible subsequent legs and their observability.

Once a user acquires a stablecoin, observability depends on where the stablecoin moves next. Four of the possible next steps are recorded on-chain. Not every such transfer survives into the panel (Section 3.4). The user can withdraw the stablecoin to a deposit address at another CEX, from where it can be exchanged for local currency and withdrawn into a bank account. The user can instead withdraw it to a self-custodial wallet, which is a wallet the user controls directly through a private key rather than through an intermediary. From that wallet, the stablecoin can be sent on to another self-custodial wallet, or deposited at a CEX. Each of these legs is a transfer between two addresses on a public ledger, and the measurement problem the rest of the paper solves is that the ledger does not record where those addresses are.

One step on the crypto side is not observable. If the sender and the receiver both hold accounts at the same CEX, the exchange can settle between them on its own books, and no transaction reaches the blockchain. We do not see these transfers and cannot measure how large they are. They are also not the object we set out to measure. No stablecoin changes hands on the ledger, and the exchange only reallocates a claim from one of its customers to another, just as a bank does when it moves money between two of its own account holders. What crosses the border is a claim on the exchange rather than a stablecoin.

Once the CEX has received the user's fiat and credited the stablecoins, that fiat pays whoever supplied the stablecoins. In the secondary market the supplier is another customer or the CEX itself. In the primary market the supplier is the issuer, so the CEX typically converts the fiat into dollars and sends it through the traditional financial system to pay Tether or Circle for a newly minted stablecoin. None of these fiat legs is observable in our data. Our estimates therefore measure the stablecoin leg of an international transaction and not the fiat leg that funds it.

Finally, the stablecoins themselves enter and leave circulation through their issuers, Tether for USDT and Circle for USDC. Tokens are created when a customer pays the issuer dollars and the issuer sends newly created tokens in return, a step called minting. Redemption is the reverse step, in which a customer returns tokens to the issuer, which pays out dollars from the reserves and destroys the tokens. Issuers transact with vetted wholesale counterparties, which are often CEXs, in this primary market, so a user who wants a stable-

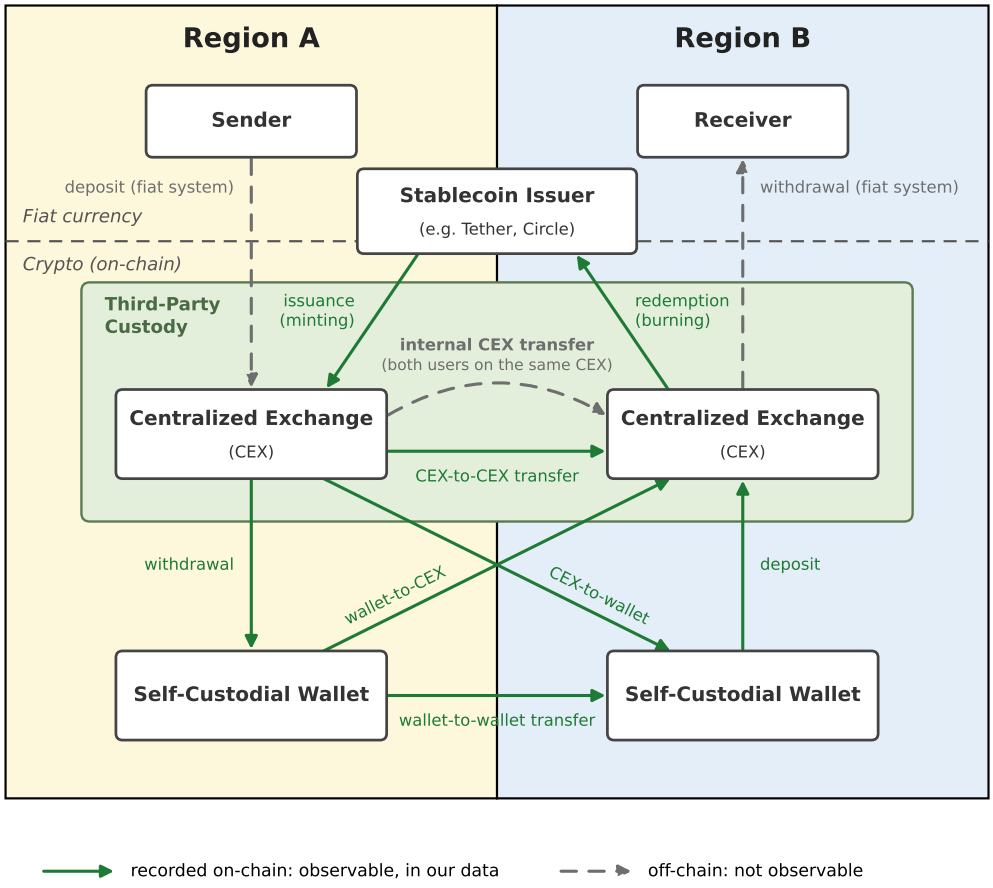


Figure 1: Schematic Overview of International Stablecoin Flows

Notes: The figure follows a stablecoin payment from a sender in one country to a receiver in another. The issuer is drawn outside both countries because it is the source and the sink of the tokens rather than a party in either country, and across the fiat and crypto layers because it receives dollars through the banking system and creates the corresponding tokens on-chain.

coin buys one that already exists, typically on a CEX (Lyons and Viswanath-Natraj, 2023; Ma et al., 2025).[13] The dollar legs of minting and redemption run through the banking system and are not in our data. The token legs are recorded on-chain. Combined with the country attribution in the next section, those issuer legs are what let us estimate how many of them a country holds (section 3.4).

[13] Ma et al. (2025) document that only six counterparties redeem USDT in an average month, against 521 for the more accessible USDC.

3 Data and Measurement

This section builds the country-level panel. It describes the raw data, attaches a geography to the wallets that send and receive stablecoins, states the identifying assumption, and reports the internal and external checks on which the later results rest. Appendix C reports the held-out evaluation of the classification model. Appendix D details the country allocation. Online Appendices A–H record the implementation of the model, the remaining DST figures, the deposit-wallet procedure, the push–pull and remittance analyses, robustness of the event studies, and the detailed Chainalysis comparison.

3.1 Data

The methodology uses four kinds of raw material. Full copies of the Ethereum, Tron, Binance Smart Chain, Optimism, Arbitrum, Base, and Linea blockchains, extracted with the Ethereum ETL package (Medvedev and the D5 team, 2018), supply every stablecoin transfer we allocate and the behavioral information we use to locate wallets.[14] The data run from each chain's genesis block through December 31st 2025 and exceed 35 billion transactions. Wallet addresses carry over across these EVM-compatible chains, so a wallet's activity can be pooled no matter which chain a transaction took place on.[15] Table 3 in Appendix B reports the sample by chain. Online Appendix A records how the extracts were obtained.

Domain name data supply the linguistic labels for the training set. Users purchase human-readable names, most often through the Ethereum Name Service, to replace a 42-character hexadecimal address. We obtain almost 6 million such names from Dune Analytics (Table 4). Web traffic data for centralized-exchange websites, from Semrush at daily frequency by originating country, divide regional estimates across the countries of a region. Supplementary Dune data identify exchange wallets, automated trading programs, smart contracts, and the USDT and USDC contract addresses on each chain. Online Appendix C describes how we identify exchange deposit wallets.

3.2 From Blockchain Data to a Country-level Panel

Figure 2 summarizes the methodology. We first estimate a probability distribution over five regions for every sufficiently active self-custodial wallet. We then split those regional probabilities across countries, using the web traffic of the exchanges each wallet interacts

[14] For Tron, we use a self-adapted version of the Ethereum ETL package.

[15] The most prominent blockchain absent from the sample is Solana. It is not EVM-compatible, so its addresses do not carry over. The omission is small in the aggregate. The blockchains in our sample carry roughly 92 percent of the USDT and USDC outstanding at the end of 2025.

with, except for China, which a dedicated model separates from the rest of Asia and the Pacific. Centralized exchanges receive regional and country shares from their traffic with the wallets we have already located. Each stablecoin transfer is then allocated across ordered country pairs by the product of the sender's and the receiver's country probabilities.

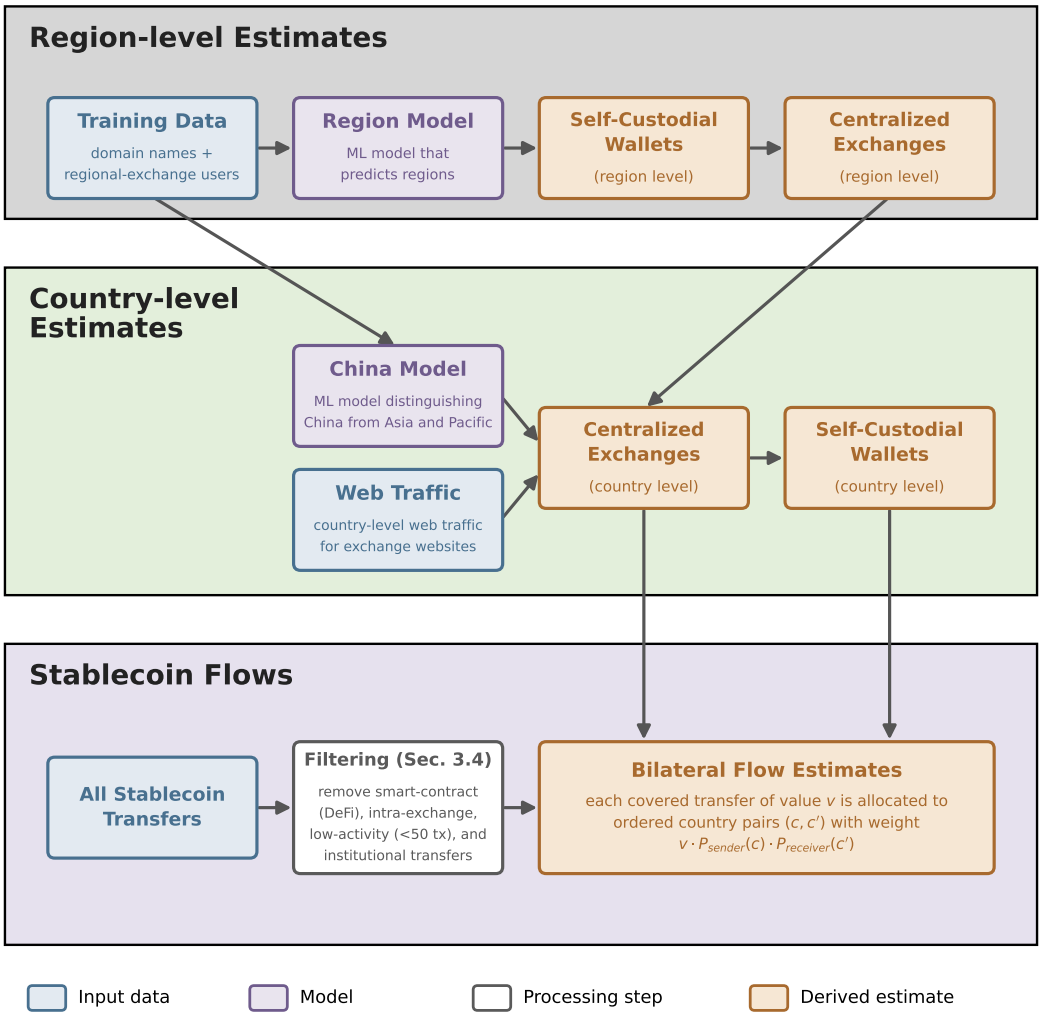


Figure 2: High-level Overview of the Methodology.

A hypothetical example illustrates the methodology from start to finish. Consider a wallet that owns the domain name "nusantara.eth". The language model recognizes that "nusantara" is an Indonesian term meaning "islands" and the name of the planned future capital of Indonesia, labels the wallet Asia and Pacific, and the wallet enters the training data. A second wallet owns no domain name, but the trained model observes that it is active mainly during Asian daytime hours, shows no daylight-savings shift in activity, and interacts mostly with exchanges popular in the region, and might assign it probability 0.8 for Asia

and Pacific. That regional probability is then divided among countries according to the web traffic of the exchanges this wallet interacts with. If Indonesia accounts for 0.75 of that web traffic, the wallet's probability for Indonesia is $0.8 \times 0.75 = 0.6$. When this wallet sends $100 in stablecoins to a wallet that is estimated to be in the United States with probability 0.5, our dataset records, among other pairs, an estimated flow of $30 from Indonesia to the United States.

Training labels. We divide the world into five regions: Africa and the Middle East, Asia and the Pacific, Europe, North America, and Latin America and the Caribbean. Labels come from two sources. First, a large language model reads each domain name for linguistic and cultural markers and assigns a region, or returns "unclassified" when the name carries no such marker. Wallets with more than one classified name take the majority region. Online Appendix D.1 reproduces the prompt. Users may reside in one region and choose a name associated with another, as when migrants pick names that evoke their country of origin. Section 3.3 validates the labels in the aggregate, and Appendix C shows that the classification model partly corrects such heritage-driven mislabeling.

Second, some centralized exchanges serve a single regional market. We classify a self-custodial wallet as belonging to that region if more than 90 percent of its transactions with centralized exchanges are with exchanges focused on that region. Appendix A lists the exchanges. A wallet can qualify for both rules. This occurs for only 126 wallets, which we assign to the regional-exchange group. Together the two approaches label 437,775 wallets, 271,658 through domain names and 166,117 through regional exchanges, proportions that need not be representative of wallets as a whole.

Classification model. We take these labels as training data and train a gradient boosted decision tree (Hastie et al., 2017) on the on-chain behavior of the labeled wallets. The features capture time-of-day activity, (non)adherence to daylight saving time, interactions with particular centralized exchanges, and engagement with popular tokens and smart contracts. We restrict the sample to wallets that have initiated at least 50 transactions, so that these features are measured with enough observations. The classification proceeds in two stages that first separate time zones and then split regions that share a time zone. Online Appendix D describes the features, the two-stage structure, and the training setup, including the balanced sample weights that give each region equal prior weight.

The trained model assigns a probability distribution over the five regions to every wallet with at least 50 transactions, including the wallets in the training data, whose labels it replaces. Our key identifying assumption is that, conditional on the features used to train the

model, labeled and unlabeled wallets of a particular region follow the same data generating process. The assumption fails if wallets that buy domain names, or that patronize regional exchanges, use those features differently from other wallets in the same region. The validation below tests the labels and the model's out-of-sample bite. It cannot directly test the assumption for wallets that never appear in either labeled group.

Country split. Regional probabilities are then divided across countries. For every country except China we use the web traffic of centralized-exchange websites. China is assigned by a separate model, trained to distinguish China-labeled Asia and Pacific wallets from the rest of that region, because exchange websites are blocked in mainland China and traffic is distorted by VPN use.[16] The China model's output replaces web traffic for that country. Only the residual Asia and Pacific share is split by traffic. An exchange's country shares are then propagated back to self-custodial wallets as a transaction-volume-weighted average of the exchanges each wallet has interacted with. Wallet probabilities do not vary over time or by blockchain, because ownership of a self-custodial wallet is tied to an unchangeable private key.

This construction requires that, within a region, web traffic is a reasonable proxy for relative exchange usage across countries, and that average transaction sizes do not differ systematically across those countries. The equal-size assumption is necessary to extrapolate from web traffic, a per-user measure, to volumes. It is a maintained assumption that we cannot test directly. VPN use cannot move volume across regions through this step, because the regional totals are pinned down by the on-chain model, and the dominant known case of large-scale VPN use is excluded from the traffic split. Relative to estimates that allocate the world by web traffic alone, the combination therefore confines the traffic assumptions to a within-region split. Appendix D states the allocation formally.

Centralized exchanges that we have designated as regional are assigned fully to their region. All other exchanges receive time-varying regional shares from their stablecoin traffic with self-custodial wallets, at hourly frequency and separately by blockchain, so that the geography of an exchange can follow the time of day and the chain in use.

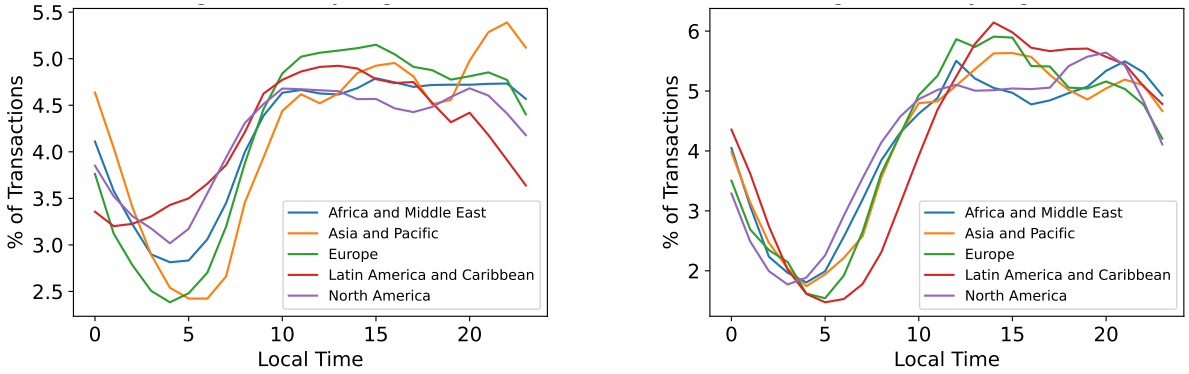
3.3 Validation of the Labels and the Model

The methodology rests on two links that can be checked directly. The first is that the region labels in the training data are correct on average. The second is that the model trained on

[16] The estimate for China should be understood as one estimate containing mainland China, Hong Kong SAR, and Macao SAR, as their on-chain behavior is unlikely to be sufficiently different to be distinguished by a machine learning model.

them has predictive power for wallets it has never seen.

For each labeled wallet we compute an activity profile, the share of its transactions that occur in each hour of the day, and convert the timestamps from UTC to a representative local time for the assigned region.[17] If the labels are correct on average, activity in region-local time should dip at night and peak during waking hours. Figure 3 shows that it does, for wallets labeled through domain names and for wallets labeled through regional exchanges.



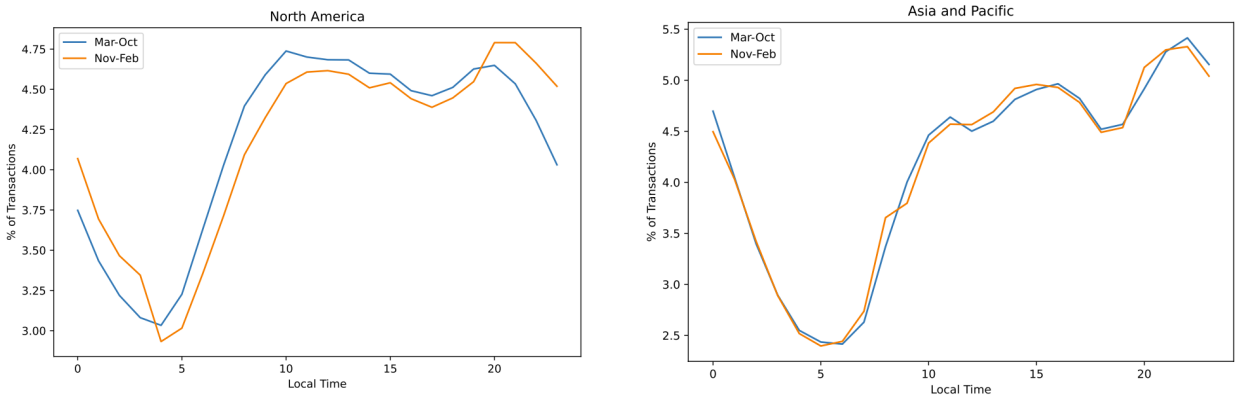
(a) Activity Profile for Wallets Identified by Domain Names

(b) Activity Profile for Wallets Identified through Regional Exchanges

Figure 3: Activity Profiles for Wallets in the Training Data

A sharper test uses daylight saving time, which shifts clocks in some regions and not in others. Figure 4 plots the same activity profiles, still against a fixed UTC offset, separately for the months in which the United States observes daylight saving time and the months in which it does not. Wallets labeled North America and Europe shift by roughly one hour, as the clock change would predict. Wallets labeled Asia and Pacific, and wallets labeled Africa and Middle East, show no such shift, as expected in regions where countries overwhelmingly do not observe daylight saving time. Systematically wrong labels would not produce this pattern. The figures for all five regions are available in Online Appendix B.

[17] We use UTC+2 for Africa and the Middle East, UTC+8 for Asia and Pacific, UTC+2 for Europe, UTC-4 for Latin America and Caribbean, and UTC-6 for North America. The intention is not to take a stance on the most appropriate time zone for regions that span several of them, but to convert UTC to a reasonable proxy of local time.



(a) North America (b) Asia and Pacific

Figure 4: Comparison of the Impact of Daylight Savings Time Across Regions

We then evaluate the model on the 10 percent of labeled wallets held out before training, 43,768 wallets. It assigns 74.1 percent of them to the correct region, against 20 percent for uniform random guessing among five regions. That 74.1 percent weights each region by its share of the testing data. The unweighted average of the five per-region accuracies is 69.3 percent, because Asia and Pacific, the best predicted region, accounts for about half of the testing data. Accuracy ranges from 78.4 percent for Asia and Pacific to 64.2 percent for Latin America and the Caribbean. The methodology propagates each wallet's full probability distribution rather than a hard classification. The corresponding probabilistic measure is the share of a wallet's probability mass assigned to its true region, which averages 62 percent across held-out wallets. In the flow allocation, 62 cents of every dollar sent by the average held-out wallet are assigned to its true region.

These pooled numbers average over the two labeling sources. Regional-exchange wallets are classified almost perfectly (99.9 percent), as expected once the model has learned which exchanges serve which regions, since those labels and the features both reflect the identity of a wallet's exchanges. Domain-name wallets, whose labels have no mechanical link to the features, are the demanding benchmark. Held-out accuracy against those labels is 58.3 percent, against 47.5 percent for always guessing the largest region. Retraining on domain-name wallets alone and applying that model to the 166,117 regional-exchange wallets, whose labels it has never seen, assigns the correct region to 91.5 percent of them. That comparison is not an accuracy estimate for arbitrary unlabeled wallets, which are harder to classify than wallets that concentrate their exchange activity in one regional market. Nor can it be produced by migrants who choose a heritage name and still use a home-region exchange. The activity profiles of the regional-exchange group place those wallets in their labeled regions. The two sources therefore point at the same geography. Appendix C reports the confusion

matrix, the decomposition by labeling source, and the asymmetric errors that are consistent with the model correcting heritage names rather than repeating them.

The China model is evaluated on the 21,997 held-out Asia and Pacific wallets, of which 20.0 percent are labeled China. Always guessing "not China" would already score 80.0 percent. The model identifies 85.2 percent of the wallets labeled China and 81.0 percent of the rest, for a class-balanced accuracy of 83.1 percent and an AUC of 0.92.

3.4 Covered Transfers and Holdings

The dataset begins with all USDT and USDC transfers on the chains in the sample. We then drop five classes of transfer. (i) Either side is a smart contract, which removes decentralized-finance activity.[18] (ii) Both sides are the same centralized exchange, which are operational movements between that exchange's own wallets. (iii) A self-custodial wallet with fewer than 50 initiated transactions, for which the model is not trained. (iv) An automated trading program, institutional trading firm, liquidity provider, or OTC desk. (v) A self-custodial wallet that has never interacted with a centralized exchange, for which the country split is undefined. What remains is simple stablecoin transfers involving centralized exchanges and sufficiently active self-custodial wallets. The estimates describe these covered flows rather than all stablecoin activity. Section 4 reports the share of outstanding stablecoins that the implied holdings account for.

Each remaining transfer of value v is allocated across ordered country pairs by the product of the origin's probability of being in country c and the destination's probability of being in country c'. Summing over pairs returns v, since each endpoint's country probabilities sum to one. Appendix D writes the allocation for wallet-to-wallet, wallet-to-exchange, and exchange-to-exchange transfers.

Holdings. Because the dataset records transfers between countries and the issuers, as well as transfers between countries, cumulating a country's net flows yields an estimate of its stablecoin holdings. Minting and redemption are the source and the sink of the tokens. A dataset of country-to-country transfers alone cannot measure a stock. Its net flows sum to zero across countries in every period. The issuer legs are what make the stock well defined, and they are the reason the crisis analysis in section 5.3 can treat holdings as the analogue of the deposit dollarization stock.

[18] Analyzing decentralized finance requires a different set of tools and is outside the scope of this paper.

3.5 External Checks

The checks above use the same raw material the methodology uses. Two sources do not. Brazil's central bank measures the fiat leg of residents' stablecoin purchases. Internal figures from two large exchanges, disclosed in enforcement proceedings and in reporting, record the country shares of those exchanges' own customers. A separate comparison is with the commercial Chainalysis dataset, which estimates the same object from a different allocation and, like us, uses web traffic. It is the series the existing literature uses, not an independent benchmark.

Official statistics: Brazil. Brazil is, to our knowledge, the only country whose official statistics measure stablecoin transactions with the rest of the world. The Central Bank of Brazil (BCB) compiles them from foreign exchange contracts and publishes them in its balance of payments, separating crypto assets with an issuer, in practice stablecoins, from crypto assets without one (Banco Central do Brasil, 2024).[19] A transaction enters the statistics when its fiat leg touches Brazil's foreign exchange market, and stays outside them when it does not. The transactions most likely to generate a contract are wholesale imports, in which a Brazilian exchange or payment firm converts reais into dollars to buy stablecoins from a nonresident. The BCB thus measures the fiat leg of the flows whose crypto leg our methodology estimates, from a data source that shares no inputs with our methodology.

A Brazilian exchange buying for inventory could obtain newly minted stablecoins from Tether or Circle, or existing stablecoins from other nonresidents. In our data, a purchase from the issuers appears as a mint from Tether's or Circle's treasury to a wallet attributed to Brazil. Over 2021 to 2025, the BCB records $52.3 billion of purchases against $45.5 billion of mints in our estimates, and from 2022 onward the two series track each other year by year (Figure 5).[20] That the mint series tracks the BCB series, and lies only somewhat below it, indicates that wholesale imports are mainly purchases from the issuers. Secondary-market purchases from non-issuers, and measurement error in the primary and secondary market must account for the residual difference, but are quantitatively smaller. The exception is 2021, the first year of the comparison, where our estimate is the larger one.

Cumulated over the same window, the BCB's net purchases are $50.4 billion, while our estimate of Brazil's net stablecoin accumulation is $12.6 billion (Table 1). The gap is explained by the redemption leg. Our data attribute $34.9 billion of redemptions to

[19] The monthly series are published in the memorandum items of the balance of payments tables in the BCB's special tables. We use the July 2026 vintage. Stablecoin purchases account for 95 percent of Brazil's total crypto purchases from nonresidents in 2025.

[20] Using the BCB's data, Fernández et al. (forthcoming) study the drivers of Brazilian stablecoin flows and show that Brazil's stablecoin purchases react to changes in its tax rate on FX transactions.

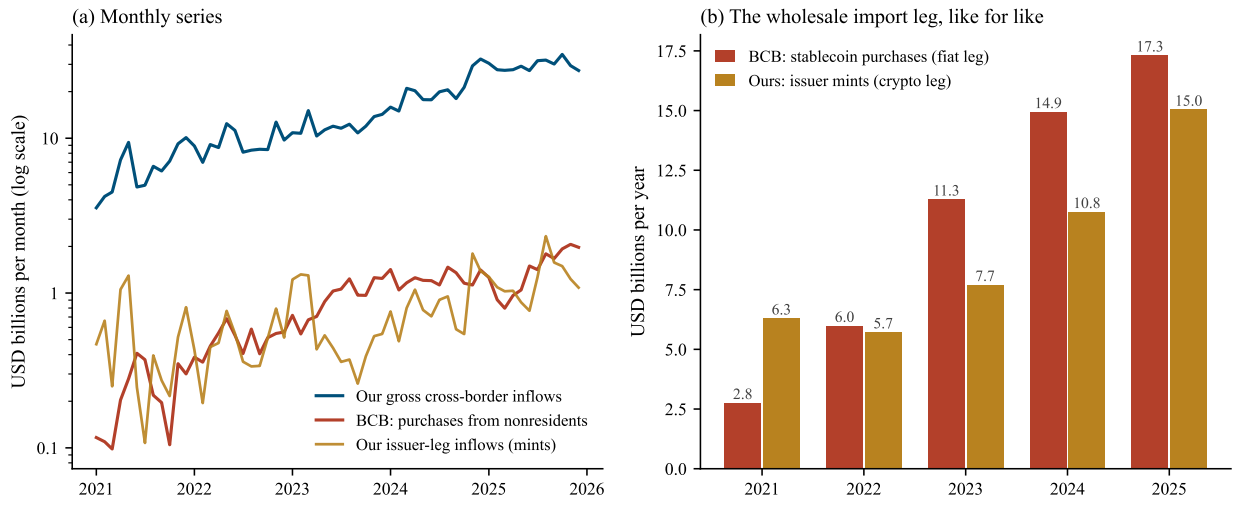


Figure 5: Brazil in our estimates and in the official statistics

Notes: The BCB series is the balance of payments memorandum item on crypto assets with a corresponding liability (stablecoins), compiled from foreign exchange contracts and reported as purchases by residents from nonresidents. Panel (a) plots monthly series: our estimated gross cross-border stablecoin inflows into Brazil, the BCB purchase series, and our estimated issuer-leg inflows (mints from Tether's and Circle's treasuries to wallets attributed to Brazil). Panel (b) compares, by year, the BCB purchase series with our issuer-leg inflows, the two measurements of the wholesale import leg.

Brazil, while the BCB records only $1.9 billion of stablecoin sales by residents. Little of this leg should appear in the foreign exchange records. Where the redeeming exchange is a global platform, the transacting entity is not a Brazilian resident, so no foreign exchange contract arises, even though our attribution assigns to Brazil the share of the crypto leg that corresponds to its customers. Where the exchange is Brazilian, the proceeds are dollars that need not be converted into reais, and a net-buying customer base means payouts to selling customers are covered by other customers' incoming deposits, so nothing forces the conversion into reais that the statistics would record. Subtracting our redemption estimate from the BCB's net purchases leaves $15.5 billion, within $3 billion of our accumulation estimate. Country-to-country flows are not the residual. Brazil's bilateral net positions sum to $2.0 billion across all partners.

The comparison checks issuer-leg mints attributed to Brazil against the BCB's wholesale imports, at the one official margin that exists. The two sources also complement each other. The foreign exchange records include purchases of existing stablecoins that the issuer legs do not capture, and our data include the redemption leg that never reaches the foreign exchange records.

Internal figures from crypto exchanges. For a few exchanges, internal figures describing the same geography have become public. A complaint by the CFTC against Binance

Table 1: Brazil: the official statistics and our estimates, 2021–2025

	BCB *(fx contracts)*	Our estimates *(on-chain transfers)*
Wholesale imports of stablecoins	52.3	45.5
Wholesale exports of stablecoins	1.9	34.9
Net flows with other countries	–	2.0
Net	50.4	12.6

Notes: Cumulative January 2021 to December 2025, in USD billions. In the BCB column, imports are purchases of stablecoins by residents from nonresidents and exports are sales by residents, from the balance of payments memo series on crypto assets with a corresponding liability (stablecoins), compiled from foreign exchange contracts. In our column, imports are issuer mints into Brazil, exports are redemptions from Brazil to the issuers, and net flows with other countries are cumulative bilateral net inflows summed over all partner countries. Net is imports less exports plus net flows with other countries.

quotes the exchange's internal reports from 2020, showing that customers in the United States were between 16 and 19.9 percent of the total. Our assigned US shares over the same months range from 16.8 to 24.8 percent. The internal figures count customers while ours measure volume, and the complaint describes Binance's US customers as disproportionately large traders, which is consistent with our volume shares sitting above customer counts. Internal Binance documents reviewed by the Wall Street Journal put China at roughly 20 percent of the exchange's trading volume in May 2023, excluding its largest traders. We assign 19.3 percent of Binance's volume in that month to China. The CFTC complaint against KuCoin states that between 20 and 50 percent of its customers were based in the United States from July 2019 through June 2023. Over the same window we assign 18.2 percent of KuCoin's volume to the United States, just below the lower edge of that range.

Commercial estimates: Chainalysis. The existing empirical literature on international crypto flows is built on the commercial Chainalysis dataset (Cerutti et al., 2025a; Cardozo et al., 2024; Auer et al., 2025b; Friedrich and Zhao, 2026; Hofmann et al., 2026), which allocates transfers between centralized exchanges in proportion to the web traffic of each exchange's website. We harmonize the two datasets to monthly bilateral flows for the 164 countries present in both over August 2019 through February 2025, restricted to USDT and USDC. Online Appendix H details the harmonization.

The two datasets agree on where stablecoin activity takes place, with a log-log correlation of 0.93 in whole-window gross flows across the 164 common countries. However, the agreement is not corroboration and does not extend beyond the cross-section. Both series scale

with country size and both methodologies use web traffic, and removing country and month fixed effects drops the correlation to 0.26. That residual variation is the margin the CFM event study of Section 5.2 uses. We repeat that design on the Chainalysis series and do not detect a significant response of stablecoin activity to CFM activations. Our estimates are larger in every year, by a factor of 2.6 over the window, because they capture self-custodial activity that a universe built on transfers between exchanges does not. The gap widens with EMDE involvement, from 1.69 for AE–AE flows to 3.19 for EMDE–EMDE flows, and it reaches 14 for China. That wedge is unremarkable before the 2021 enforcement measures and ban, and thereafter lies outside the central 95 percent range of the 40 largest comparison countries (Online Appendix H). The Chainalysis data comprise country-to-country flows only and cannot measure holdings. Treating their cumulative net flows as the holdings analogue, we repeat the crisis event study of Section 5.3 and do not recover a significant accumulation of stablecoins in response to currency crises. Our dataset contains the issuer legs. Cumulative worldwide net flows sum to $146 billion over the comparison window, in line with the holdings accounting of section 4.

4 International Stablecoin Flows

This section presents the basic facts the dataset contributes: the first comprehensive, country-level map of where stablecoins flow and where they are held. Five facts organize the section. First, stablecoin flows have reached macroeconomic scale and are still growing fast. We estimate a gross transfer volume of over $7 trillion in 2025, up from $1.2 billion in 2018, and roughly 90 percent of this volume crosses borders, a share that has been strikingly stable throughout the sample. Second, the typical transfer is retail-sized, with a median of $300 against a mean of $10,853. Third, allocated by the location of the recipient, gross flows split roughly 56/44 between EMDEs and advanced economies (AEs), but holdings concentrate far more sharply: about three quarters of the roughly $200 billion in stablecoin holdings our panel covers at the end of 2025 are held in EMDEs. Fourth, the largest recipient countries in dollar terms are the world's largest economies, while relative to GDP nineteen of the twenty largest recipients are EMDEs. Fifth, sorting cross-border flows into four corridor types by whether sender and receiver are AEs or EMDEs, flows between EMDEs are the largest corridor type, accounting for roughly a third of cross-border volume, while flows between advanced economies account for only 14 percent. All quantities in this section are estimates from the measurement methodology of section 3, subject to the assumptions and coverage summarized there.

4.1 Scale, Growth, and Composition

Merging the country-level probability estimates onto the set of covered stablecoin transfers (section 3.4) yields a bilateral panel of roughly 2 billion transactions between January 1, 2018 and December 31, 2025, totaling $21.4 trillion. These are gross flows, and a stablecoin that enters a country and later leaves it contributes to the total twice. Gross flows therefore measure how intensively stablecoins are used rather than how much capital they move on net. Worldwide remittance flows are smaller and largely one-way: the World Bank estimates $905 billion in 2024 (Ratha et al., 2024). The like-for-like object is gross cross-border payments. Cerutti et al. (2025b) put those at roughly one quadrillion dollars, of which our 2025 volume is less than one percent. Holdings, the cumulation of net flows, are documented at the end of this section. In terms of composition, USDT accounts for approximately 84 percent of total volume in our sample, while USDC accounts for 16 percent.

Figure 6 plots estimated monthly gross flows. Annual volume has grown from just $1.2 billion in 2018 to over $7 trillion in 2025, a compound average growth rate of 244% per year that has moderated to year-over-year growth of 48% in 2024 and 43% in 2025. Throughout this growth, the estimated share of volume in which sender and receiver are located in different countries has remained roughly constant at about 90% (Figure 6, dashed line). That 90 percent includes minting and redemption. Country-to-country corridors exclude those issuer legs and account for 83 percent of gross volume.

The size distribution of transfers is heavily right-skewed. The mean transfer in the sample is $10,853, but the median is $300. The typical stablecoin transfer is therefore retail-sized, closer in scale to a remittance than to a wholesale financial transaction, even though large transfers account for much of total volume.

4.2 Geography and Corridors

In dollar terms, we estimate that the largest recipients of gross stablecoin flows in 2025 were the largest economies: the United States and China.[21] However, scaling by GDP highlights the usage of stablecoins in EMDEs. Nineteen of the twenty largest recipients relative to GDP are EMDEs, and the magnitudes are macroeconomically large for some countries: estimated gross flows in 2025 amount to 89 percent of GDP in Venezuela, 58 percent in Ukraine, and 38 percent in Bolivia.[22]

The corridor composition of these flows differs sharply from that of traditional payments. Of the volume processed by SWIFT, 85 percent moves between AEs (Basu et al., 2026). In

[21] Estimates for China combine mainland China, Hong Kong SAR, and Macao SAR.

[22] All GDP numbers are taken from the IMF WEO database. The only AE in the top 20 is Andorra, which is likely an outlier due to its small GDP of around $4.5 billion.

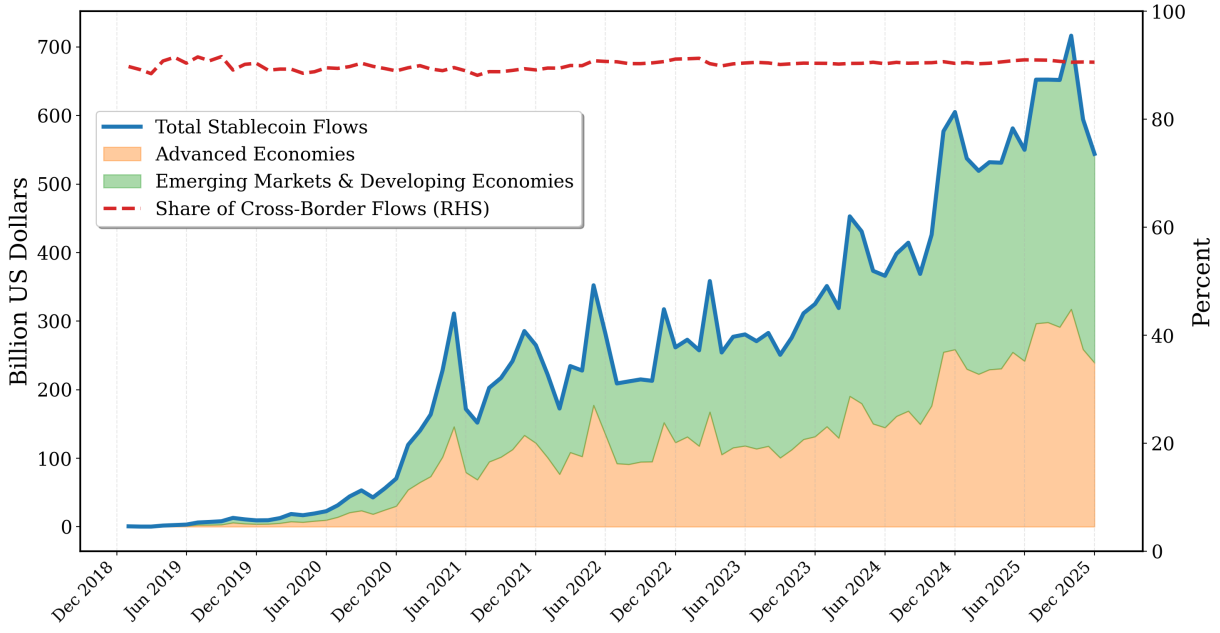


Figure 6: Estimates of Total Stablecoin Flows

Notes: The figure plots estimated gross stablecoin flows at monthly frequency, in billions of US dollars (left axis). Total flows are split into flows received by Advanced Economies (AEs) and by Emerging Markets and Developing Economies (EMDEs) according to the estimated location of the recipient. The dashed line (right axis) plots the estimated share of monthly volume in which sender and receiver are located in different countries. All series are estimates from the classification methodology of section 3. The horizontal axis begins in December 2018; estimated flows before then ($1.2 billion in all of 2018) would not be visible at the scale of the figure.

cross-border stablecoin flows, the proportions are reversed. Flows between AEs account for only 14 percent of cross-border volume, while flows involving at least one EMDE make up 86 percent (Table 2). Flows between EMDEs are the largest of the four corridor types, accounting for roughly a third of cross-border volume, and the two mixed corridor types are nearly symmetric at roughly 26 percent each. Like the cross-border share itself, this composition has been quite stable. In every full year of the sample, the EMDE-to-EMDE share remains between 31 and 35 percent and the AE-to-AE share between 13 and 16 percent.

4.3 Stablecoin Holdings

Because the dataset records not only flows between countries but also flows between countries and issuers, cumulating a country's net flows yields an estimate of its stablecoin holdings (section 3.4 develops why issuance flows are essential to this construction). We estimate total holdings of roughly $200 billion at the end of 2025, accounting for approximately 78 percent

Table 2: Corridor composition of cross-border stablecoin flows

	AE→AE	AE→EMDE	EMDE→AE	EMDE→EMDE
Volume, full sample ($ trillion)	2.56	4.69	4.65	5.98
Share of volume, full sample (%)	14.3	26.2	26.0	33.4
Share of volume, 2025 (%)	14.7	26.5	26.2	32.6

Notes: The table decomposes estimated cross-border stablecoin flows between countries by the corridor type defined by the IMF WEO classification of sender and receiver into advanced economies (AEs) and emerging market and developing economies (EMDEs). Cross-border flows between countries total $17.9 trillion over the full sample (January 2018 through December 2025), or 83% of the $21.4 trillion in gross flows; the remainder is domestic volume, with sender and receiver in the same country (10%), and flows to and from the issuer treasuries (7%), neither of which defines a corridor between two countries. The denominator therefore differs from that of the 56/44 EMDE/AE split of gross flows reported in the text, which allocates all gross volume by recipient. Shares sum to 100 up to rounding. All quantities are estimates from the classification methodology of section 3.

of the $256 billion in stablecoins outstanding at that time.[23] As Figure 7 shows, estimated holdings grow from essentially zero at the start of 2019 to roughly $200 billion at the end of 2025, with a pronounced drawdown in 2022–23 and rapid accumulation since. In 2023, in the wake of the collapse of FTX and the failure of Silicon Valley Bank, which briefly caused USDC to trade at 90 cents on the dollar, estimated holdings in AEs fell sharply, approaching zero, while holdings in EMDEs only marginally decreased (Figure 7). This relative drawdown is consistent with stablecoins having been used more as a vehicle for trading crypto assets in AEs and more as a store of value in EMDEs. Since 2023, holdings in AEs have recovered, and total holdings have risen to their end-of-sample peak.

While gross flows split roughly 56/44 between EMDEs and AEs, holdings are substantially more pronounced in EMDEs. We estimate that around three quarters of the stablecoin holdings our panel covers at the end of 2025 were held in EMDEs, and that the EMDE share has trended upward since 2020 (Figure 7, dashed line). This concentration is consistent with stablecoins serving as a store of value in EMDEs (Adrian et al., 2025).

At the country level, holdings in the largest economies are large in dollars and small relative to GDP: 0.04 percent of GDP in the United States and 0.09 percent in China. Further highlighting the popularity of stablecoins in EMDEs, we estimate that fifteen of the twenty largest holdings in dollar terms, and all twenty of the largest relative to GDP, are in EMDEs. For example, estimated holdings reach $13.3 billion in Türkiye (0.83 percent of

[23]The 22% of outstanding stablecoins that our estimates do not cover break down into 8% outstanding on blockchains outside our dataset and 14% held by entities our dataset does not cover, such as liquidity providers and OTC desks. That residual varies over the sample. It widens in 2022–23, when estimated holdings fall by more than the outstanding stock, consistent with tokens moving to those uncovered entities rather than being redeemed.

GDP), $8.6 billion in Ukraine (4.05), $6.5 billion in Vietnam (1.31), $5.8 billion in Pakistan (1.42), $5.6 billion in Venezuela (5.61), $5.1 billion in Argentina (0.75), and $4.1 billion in Nigeria (1.42). Venezuela and Ukraine hold the most relative to GDP of any country in the panel.

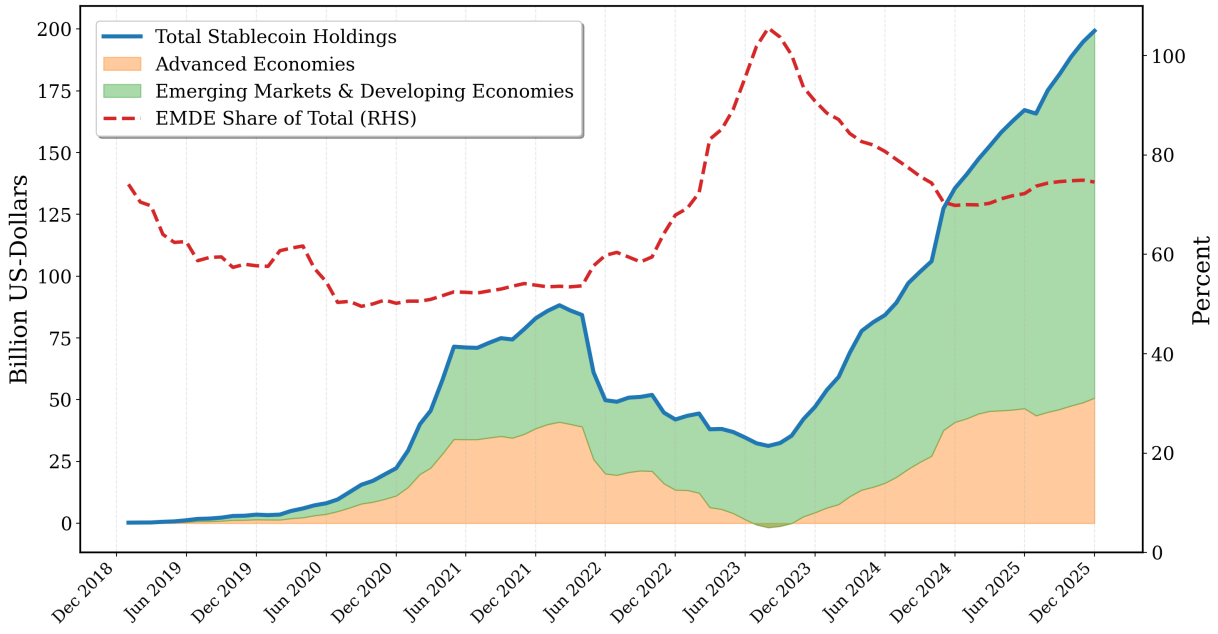


Figure 7: Estimates of Stablecoin Holdings

Notes: The figure plots estimated stablecoin holdings at monthly frequency, in billions of US dollars (left axis), constructed as cumulative net flows since the start of the sample. Total holdings are split into holdings in Advanced Economies (AEs) and in Emerging Markets and Developing Economies (EMDEs) according to the estimated location of the holder. The dashed line (right axis) plots the EMDE share of total holdings; it briefly exceeds 100 percent in late 2023 because estimated AE holdings turn slightly negative. Holdings estimates are not constrained to be non-negative, and small negative values reflect coverage gaps and measurement error rather than negative positions (Appendix G).

5 Economic Drivers of Stablecoin Flows

Having mapped the scale and geography of international stablecoin flows in Section 4, we now turn to what drives them. A single idea organizes the analyses in this section. Stablecoins function as an outside option to the domestic financial system, providing access to a dollar-denominated instrument that offers users an alternative if the default option is costly, restricted, or unstable. The three subsections that follow examine these margins in turn.

First, Section 5.1 documents conditional associations that corroborate, in independently constructed data, results established in the existing literature. Macroeconomic instability

is associated with larger stablecoin inflows in EMDEs but not in AEs, and corridors in which traditional remittances are more expensive see systematically larger stablecoin flows. Because our dataset is constructed independently of the commercial estimates on which this literature builds, these confirmations serve two purposes. They establish that known drivers operate in our data, and they provide external corroboration of our measurement that the existing literature cannot supply for itself (Section 3.5 provides the direct dataset comparison).

Second, we present new evidence on the reaction of stablecoins to CFMs (Section 5.2) and currency crises (Section 5.3), two drivers at the core of international macroeconomics. Both analyses use a stacked event-study design. When a country tightens capital flow restrictions, we find that its cross-border stablecoin activity rises promptly and persistently relative to countries that activate no restriction of their own in the surrounding months, by roughly 20 percent six months after activation and roughly 30 percent after tightenings of outflow restrictions specifically. When a country enters a currency crisis, we estimate that its residents accumulate stablecoin holdings of roughly 0.07 percent of pre-crisis GDP over the year after onset, and the crisis countries' holdings rise by roughly 75 percent. In both designs the response persists. Stablecoin activity does not detectably recede when CFMs are removed, and a formal test rejects that CFM removals fully undo the tightening response. The holdings a crisis builds settle onto a higher plateau, showing no signs of reverting.

One caveat applies throughout. Our flow measure records where stablecoins move, not why. It aggregates gross transfers across all use cases (trading, payments, remittances, the accumulation of a dollar-denominated store of value) and does not decompose any given flow into its motives. The estimates below therefore capture the response of total flows to each driver, and each subsection calibrates its language accordingly. All regressions in this section are estimated on samples beginning in January 2021, when stablecoin flows had become economically meaningful (see the gross flow series in Figure 6).

5.1 Known Drivers: Push and Pull Factors and Remittances

This subsection summarizes the two drivers of stablecoin flows that the literature has already established, the macroeconomic push and pull factors of aggregate flows and the cost of traditional remittances. Online Appendix E reports both analyses in full, including the specifications, the estimates, and their robustness; here we outline what we do and what we find.

Push and pull factors. Following the capital-flow literature (Calvo et al., 1996; Fratzscher, 2012; Koepke, 2019), adopted for crypto flows by Auer et al. (2025b) and Cerutti et al.

(2025a), we decompose the drivers of stablecoin flows into *push* factors, global conditions that move flows everywhere, and *pull* factors, country-specific conditions that determine where flows go. We regress the log of monthly cross-border stablecoin inflows on both sets of factors in a panel of 188 countries from January 2021 through December 2025, with country fixed effects throughout and, in our preferred specification, year–month fixed effects that absorb all global shocks; the pull slopes are allowed to differ between AEs and EMDEs. None of this variation is randomly assigned, so we read all estimates as conditional associations.

On the pull side, where identification is sharper, given that the panel is relatively short but wide, the two country groups part ways. The association with inflation diverges sharply, robustly negative in AEs but positive at the point estimate in EMDEs. A natural hypothesis is that because stablecoins typically pay no interest they lose appeal wherever interest-bearing local-currency instruments are available that reliably outpace inflation. Exchange rate volatility is the strongest pull association in the exercise, and it is detectable only in EMDEs: a one-standard-deviation increase in EMDE volatility is associated with 9.7 percent larger inflows (95% CI: 4.3 to 15.4). A natural reading is the well-documented EMDE pattern. When the local currency is volatile, households move into dollars. In total, on the pull side, our independently constructed flows thus reproduce the instability–use association documented in the Chainalysis-based literature and locate it where the hedging motive is plausible, in EMDEs, where recent theory rationalizes household dollar assets precisely as insurance against depreciation states (Bocola and Lorenzoni, 2020; Dalgic, 2024; Christiano et al., 2021; Ferrante and Gornemann, 2022).

On the push side, identified from time-series variation alone, flows co-move essentially one-for-one with the outstanding stock of stablecoins, with an elasticity of 1.03 (s.e. = 0.10) with respect to the combined market capitalization of USDT and USDC. The remaining, cyclical push coefficients read more like stablecoins co-moving with risk appetite than flight to safety, with larger flows when crypto sentiment improves and smaller flows when the federal funds rate rises, the reverse of the reading in Auer et al. (2025b), but in line with Aldasoro et al. (2025).[24] However, the cyclical associations in our panel and in Auer et al. (2025b) are both estimated from a few years of aggregate time series, and we therefore refrain from overinterpreting either.

Remittances. A frequently cited use case for stablecoins is the sending of remittances, which remain costly through traditional intermediaries: according to the World Bank's Remittance Prices Worldwide database, the global average cost of sending remittances stood

[24] Aldasoro et al. (2025) find on identified US monetary policy shocks that a tightening lowers stablecoin market capitalization while it raises the assets of prime money market funds.

at 6.4% of the amount sent in the third quarter of 2025, more than twice the 3% target by the G20. We merge corridor-level remittance costs from this database with quarterly stablecoin flows in the same direction as each World Bank corridor, for 348 corridors in which at least one endpoint is an EMDE, the corridors in which formal remittance costs are high and stablecoins are a plausible alternative. Corridor and year–quarter fixed effects absorb all time-invariant bilateral characteristics and global developments, and because the remaining within-corridor variation in costs is not randomly assigned, we read the estimate as a conditional association. Within a corridor, quarters with more expensive traditional remittances see systematically larger stablecoin flows: the estimated elasticity is 0.138 (s.e. = 0.046), so a 10 percent increase in the cost of traditional remittances is associated with 1.4 percent larger stablecoin flows in the same corridor and quarter. The elasticity is statistically indistinguishable between AE-to-EMDE and EMDE-to-EMDE corridors, corroborating the corridor-level evidence of Auer et al. (2025b) and extending it to corridors between EMDEs, where traditional remittances are most expensive.

5.2 Stablecoins and Capital Flow Management Measures

In this subsection, we provide novel evidence that stablecoin activity responds to CFMs. Using an event-study design around CFM activations, we document three facts. First, a country's cross-border stablecoin activity rises promptly and persistently after it tightens capital flow restrictions, relative to countries that activate no restriction of their own in the surrounding months: about 4 percent higher in the activation month and roughly 20 percent six months later. Second, the outflow-direction path is the largest and the only one that is individually significant, building to roughly 30 percent after six months and still around that level a year later. Within those tightenings, the response is concentrated in the capital account category, the subcategory that most directly constrains residents' ability to move wealth abroad. Third, the response ratchets: we detect no corresponding decline when restrictions are removed, and we formally reject the hypothesis that removals fully undo the tightening response.

We read these facts as evidence consistent with circumvention, i.e., that residents facing binding restrictions shift cross-border transactions onto stablecoin rails, which lie largely beyond the regulatory perimeter. Alternative explanations are taken up in the identification discussion below. The analysis proceeds in five steps. We first describe the data and the event-study design. We then present the pooled response. Next, we decompose the response by the direction and category of the restriction tightened, and gauge its economic magnitude in dollars. We then confront the central threat to identification, the endogenous timing of

controls, and show that the timing of the response, its geography across corridors, and its composition are difficult to reconcile with confounding factors. We close with the mirror-image experiment, CFM loosenings, with particular attention to removals that follow an earlier tightening by the same country and show that they do not lead to a significant reduction in stablecoin activity.

Data and design. Our design leverages the granular dataset on CFMs by Bergant et al. (2026) in a stacked event-study difference-in-differences design, comparing stablecoin flows between countries when CFMs are tightened. Combining our dataset with theirs, we create a monthly frequency panel of 182 countries spanning from 2021 to 2024.[25] Bergant et al. (2026) record, for each country and month, a count of binding cross-border capital flow restrictions, and distinguish restrictions targeted at inflows, at outflows, or that are "neutral".[26] The count is an integer tally across the eight subcategories of their taxonomy (listed in Appendix E.2), rising by one for each tightening action and falling by one for each loosening action. Those actions include both the establishment or removal of a restriction and the tightening or loosening of an existing one; a month-over-month increase thus means that restrictions tightened on net, whether by a new measure or by a recalibration. We define an event as a country–month in which the total count (irrespective of which flow direction is restricted) has increased month-over-month. Because CFM activations are staggered across countries and can recur within the same country, we estimate a stacked event-study design in the spirit of Cengiz et al. (2019): each event e (a country–month in which the total restriction count rises) is assigned its own dataset over an event window of six months on either side of the activation, comprising the treated country and its clean controls, the countries of the panel with no CFM activation of their own within that event's window, and these event-specific datasets are then stacked. This yields 217 events, 87 of which are increases in outflow CFMs, 43 are increases in inflow CFMs, 60 are increases in neutral CFMs, and 27 are increases in multiple categories at the same time.

Our outcome variable is cross-border stablecoin inflows. While circumvention of a CFM can register both an inflow, when the resident first receives stablecoins, and then an outflow, when they send them abroad, we note that a country's cross-border inflows and outflows are nearly perfectly correlated (see Section 4.1). Therefore, using inflows, outflows, or the sum of in- and outflows, carries the same variation as the baseline design we present. Accordingly, we interpret the results of this section as the impact of CFMs on stablecoin activity more

[25] The sample ends in 2024, as data on CFMs in 2025 by Bergant et al. (2026) is not yet available.

[26] They define a CFM as neutral if it affects inflows and outflows equally, or if its focus cannot be determined.

broadly, not as the impact on stablecoin inflows alone.[27]

In the baseline, we pool these events and estimate the following specification

$$\ln\left(SC_{c,t}\right) = \sum_{\substack{k=-6 \\ k\neq -1}}^{6} \beta_k \, 1\{t - t_e = k\} \times Treat_{c,e} + \alpha_{e,k} + \delta_{e,c} + \varepsilon_{c,t,e} \tag{1}$$

where $SC_{c,t}$ denotes stablecoin inflows into country c in month t; $Treat_{c,e}$ is an indicator equal to one if country c is the treated country in event e; and $1\{t - t_e = k\}$ indicates event time k relative to the activation month t_e, with the month immediately preceding the activation ($k = -1$) serving as the omitted reference period. The specification includes event-by-event-time fixed effects $\alpha_{e,k}$, which absorb common shocks within each event window, and event-by-country fixed effects $\delta_{e,c}$, which absorb time-invariant differences across countries within each stacked event. Standard errors are clustered at the country level, allowing for arbitrary correlation across the multiple events and event windows in which a given country appears.

We handle overlapping events by truncating each event's window whenever the same country activates another CFM nearby: the window is cut off after the most recent preceding activation and before the next subsequent activation, so that no observation within an event's window is contaminated by a different activation of the treated country.[28] Activations by other countries are handled by the clean-control requirement instead: a country that activates inside an event's window drops out of that event's control pool, but remains a control for events whose windows it does not contaminate.

Baseline: stablecoin activity rises after CFM activations. Figure 8 (panel a) plots the pooled event-study path; the full event-time coefficient estimates underlying both panels are reported in Table 11 in Appendix E. Two features stand out. First, the estimated coefficients in the pre-activation months are small and statistically indistinguishable from zero, supporting the identifying assumption of parallel trends. Second, following activation, stablecoin inflows, measured relative to control countries in the same calendar months, rise immediately and build monotonically over the event window: inflows are about 4 percent

[27] Cross-border transactions switching from traditional channels due to the CFM can raise both legs as well. Appendix E.1 describes the logic behind both stablecoin inflows and outflows increasing due to CFMs in more detail, including the stablecoin premium presented in von Luckner et al. (2024) that explains the rationale for agents to facilitate the inflow leg of stablecoins into the country.

[28] In practice, 88 events have their pre-activation window truncated on this account, and an equal number their post-activation window. Twenty-nine events retain no pre-activation months at all. These events contribute their control observations and the shape of the response across post-activation horizons, while the level of the event-time path relative to the omitted reference month is identified from the events with a surviving pre-window.

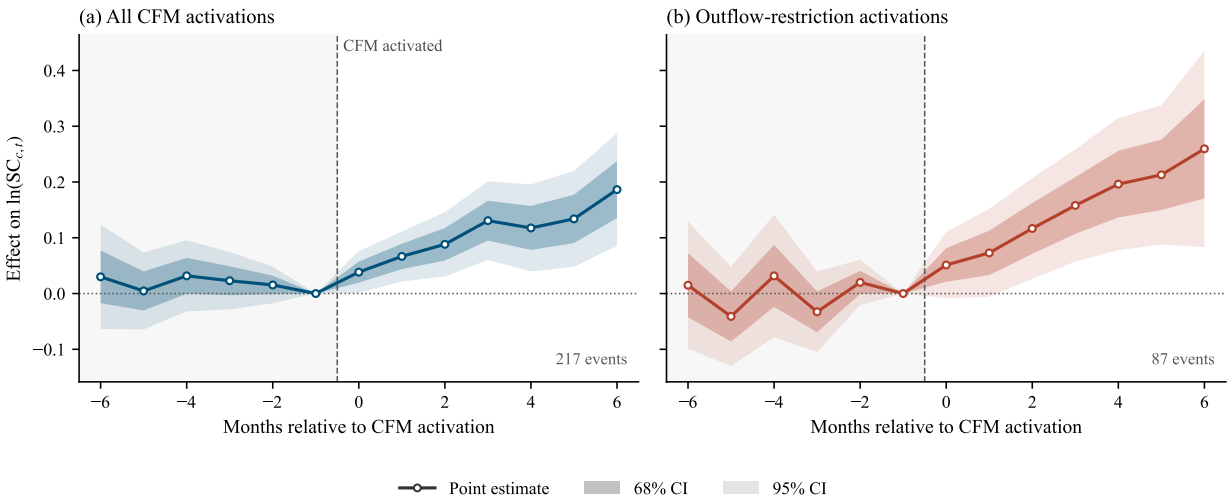


Figure 8: Stablecoin inflows around CFM activations

Notes: The figure plots stacked event-study estimates of the effect of new capital flow management measures (CFMs) on stablecoin inflows. Panel (a) pools all CFM activations. Panel (b) plots the outflow-direction path from the joint directional regression, in which each restriction direction has its own event-time path; the paths for the other directions are reported in columns (3)–(5) of Table 11 in Appendix E. Coefficients are relative to the month immediately preceding the activation ($k = -1$, omitted, dashed vertical line), with event time binned at ±6 months. Shaded areas show 68 and 95 percent confidence intervals based on standard errors clustered at the country level. Both specifications include event-by-event-time and event-by-country fixed effects and are estimated on 283,294 country–month observations across the 217 stacked events.

higher in the activation month, 9 percent higher after two months, and roughly 20 percent higher six months after activation (s.e. = 0.051).[29]

Which restrictions drive the response? We then ask which type of CFM drives this response. To this end, we attribute each event to the direction of the restriction that was tightened (inflow, outflow, or directionally neutral, with events in which multiple directions tighten simultaneously assigned to a separate category) and re-estimate the specification with the treated-side event-time dummies interacted with the event's attributed direction. Formally, let $d(e) \in \mathcal{D} = \{\text{out}, \text{in}, \text{neutral}, \text{multiple}\}$ denote the direction attributed to event

[29]The path is robust to the design's discretionary choices and to the estimator. Using not-yet-treated countries as controls, retained up to the month of their own first activation, or restricting the control pool to countries in the treated country's group (AEs or EMDEs), leaves the estimates essentially unchanged. So does a local-projections difference-in-differences estimator in the spirit of Dube et al. (2025), and so does excluding from the stacked design the 29 events without a surviving pre-activation window, which the LP-DiD sample drops by construction. These checks are in Online Appendix F.

e. We estimate

$$\ln(SC_{c,t}) = \sum_{d \in \mathcal{D}} \sum_{\substack{k=-6 \\ k \neq -1}}^{6} \beta_k^d \, 1\{t - t_e = k\} \times Treat_{c,e} \times 1\{d(e) = d\} + \alpha_{e,k} + \delta_{e,c} + \varepsilon_{c,t,e} \qquad (2)$$

where all variables and fixed effects are defined as in the baseline specification, and β_k^d traces out a separate event-time path for each restriction direction. Because the direction indicators $1\{d(e) = d\}$ partition the set of events, this is a single joint regression: all directions share the same fixed effects and clean-control pool, and each β_k^d is identified from the comparison of the treated country in events of direction d against the common set of controls.

Figure 8 (panel b) plots the resulting outflow-direction path; the full set of directional estimates is reported in columns (2)–(5) of Table 11 in Appendix E. The outflow path is the largest in magnitude and the only one that is individually significant and monotonically building throughout the post-event window: stablecoin inflows are approximately 12 percent higher two months after an outflow tightening, approximately 22 percent higher after four months, and approximately 30 percent higher after six months (s.e. = 0.089).[30] Directionally neutral restrictions display at most marginally significant increases at isolated horizons, and inflow restrictions elicit a more modest increase that never attains significance at the 5 percent level. Wald tests do not reject equality of the directional paths, whether horizon by horizon or averaging the post-activation coefficients within each direction (Table 17 in Online Appendix F.1).[31] A response to inflow or neutral tightenings would not cut against circumvention either, since those restrictions also impede cross-border transactions for which stablecoins are a substitute. The composition evidence below provides the sharpest discrimination the data allow.

Economic magnitude. To gauge whether these percentages amount to an economically meaningful volume of flows, we translate them into dollars for the treated countries. Over the six months preceding activation, the median treated country receives stablecoin inflows of approximately $375 million per month (interquartile range $114 million to $1.8 billion), and sends almost exactly as much. Applying the outflow-path estimates to both legs, the

[30] Nor does the response revert at longer horizons: re-estimating the directional design with a twelve-month post-activation window leaves stablecoin inflows approximately 26 to 33 percent above controls throughout months seven to twelve after an outflow tightening, with 95 percent confidence intervals excluding zero at every one of these horizons (Online Appendix F.1).

[31] The average post-activation effect over $k = 0, \ldots, 6$ is 15.2 percent (s.e. = 4.8) for outflow restrictions, against 10.3 (s.e. = 5.7) for inflow and 8.2 (s.e. = 4.8) for neutral restrictions ($p = 0.50$ and 0.27 for the averaged differences). With 43 inflow events across 26 treated countries, against 87 outflow events across 48, these comparisons have limited power.

implied increment in cross-border activity for the median outflow tightening is approximately $205 million per month (interquartile range $64 million to $840 million).

Identification: endogenous timing and coincident shocks. A natural concern with this design is that CFM activations are not randomly timed: outflow restrictions in particular tend to be imposed in response to capital-flight or depreciation pressure (Cardoso and Goldfajn, 1998; Bhargava et al., 2023), and such pressure could independently raise demand for dollar-denominated stablecoins even absent any interaction with the restriction itself. Note first that the event-by-country fixed effects absorb the level of any country-specific conditions prevailing over the event window, so crisis circumstances per se do not confound the estimates; a confounding effect requires that stress deteriorates within the window, with timing that tracks the activation itself. The timing of the response is difficult to reconcile with this. Adjustments to controls are infrequent, slow-moving, and historically deployed well into stress episodes that themselves extend over several quarters (Fernández et al., 2015; Eichengreen and Rose, 2014; Acosta-Henao et al., 2025; Bhargava et al., 2023; Kaplan and Rodrik, 2002; Forbes and Warnock, 2012), so if stablecoin inflows tracked such pressure directly, they should already be drifting upward in the months preceding activation. The pre-activation coefficients show no such drift. This is not because stablecoin demand is insensitive to stress (Section 5.3 documents that it is not), but because around these activations, inflows were not rising in the way contemporaneous confounding would require. And since the pressure motivating controls typically predates activation yet elicits no pre-activation response, its persistence after activation cannot readily account for the steadily growing post-activation effect either.

The corridor dimension of the dataset sharpens this argument in two steps, each holding one side of a corridor fixed. The first step holds the origin fixed. For each activation, we compare flows from each origin country into the activating country with the same origin's flows, in the same month, into countries without any activation in the event window, so that global conditions and origin-specific shocks difference out exactly (Appendix E.3). The estimates are close to the baseline, and the within-origin pre-activation coefficients are flat: before the restriction is imposed, inflows into the activating country were not rising relative to its peer destinations either.

This first comparison still measures the activating country against other destination countries, so a shock specific to the activating country in the activation month remains a possible confound. The second step therefore stays within the activating country. We measure each corridor's exposure as the share of the activating country's residents' bank deposits abroad that is held in the partner country, taken from the BIS locational banking

statistics in the quarter before the activation, and re-estimate the event study with fixed effects for the activating country in each event month. The comparison then runs across a single country's corridors at a single date, and any shock to the country as a whole is absorbed exactly, whatever its source or timing. After an outflow tightening, stablecoin inflows rise most from the partner countries where the activating country's residents hold their deposits. There is no such gradient before the activation, none after inflow or neutral tightenings, and no detectable gradient when exposure is measured by bank lending into the country instead of by residents' deposits abroad (Appendix E.4). A shock that merely coincides with the activation raises a country's activity on all of its corridors at once. It produces this gradient only if the shock itself is stronger on the corridors where residents hold deposits. We return to the gradient and its magnitude below.

The gradient test has limited precision, so we also address the coincident-shock concern episode by episode, for the cases in which CFMs and other factors, such as sanctions and currency stress, hit within weeks of one another. We find that the outflow response is not driven by crisis or sanctions episodes: it survives leave-one-country-out estimation essentially unchanged, remains qualitatively similar when sanctioned countries are excluded, and is, if anything, larger when all countries experiencing a currency crisis over the sample period are excluded before events are defined, a cut that removes 45 of the 87 outflow events (Online Appendix F.1).[32] Conditioning on month-over-month inflation and trailing twelve-month exchange-rate volatility leaves the path essentially unchanged (Online Appendix F.1). A remaining alternative explanation is that the activation itself may signal the severity of underlying conditions and thereby prompt precautionary demand for dollar assets, rather than circumvention of the restriction. Even under this informational reading, the restriction causes the response. The removals studied below also cut against a purely informational account. A removal signals that conditions have improved, so precautionary dollar demand should fall when restrictions are lifted, yet removals that follow a tightening leave inflows entirely flat.

The composition of the response is consistent with circumvention. We decompose the 87 outflow tightenings into the eight subcategories of the Bergant et al. (2026) taxonomy (Appendix E.2). The response is concentrated in tightenings of capital account restrictions, the subcategory that most directly constrains residents' ability to move wealth abroad. Capital account tightenings are followed by stablecoin inflows approximately 40

[32] Currency crises are dated with the Laeven and Valencia (2020) criterion of Section 5.3. Further, a dose-response specification points in the same direction: the response, if anything, declines with the intensity of the tightening, whereas crisis-driven policy episodes are precisely the ones imposing the most sweeping restrictions. For details, see Online Appendix F.1.

percent higher after six months. No subcategory is identified from more than 15 countries, so we base the decomposition's inference on randomization tests, which are exact under the sharp null regardless of the cluster count. Capital account restrictions are the only subcategory that survives, at a randomization p of 0.06 for the average effect over the six months following activation, and dropping any one of its eleven treated countries leaves that average essentially unchanged. We read the remaining subcategories as descriptive, even though several carry sizeable point estimates suggestive of natural circumvention margins, such as restricted transfer payments and export proceeds held offshore in stablecoins.

The geography of the response is consistent with circumvention as well. After an outflow tightening, inflows rise most from the partner countries in which the activating country's residents already hold deposits abroad, the gradient documented in the identification discussion. Stablecoin activity thus grows most along the financial connections residents had already established. The gradient is modest in magnitude, however. A corridor one standard deviation higher in its deposit share shows inflows approximately 2 percent higher six months after the tightening, against the roughly 30 percent average response, so most of the response is common to all of a country's corridors, consistent with cross-border payments of all kinds shifting onto stablecoin rails, not just wealth transfers alone.

Loosenings and the ratchet. The mirror-image experiment is a month-over-month *decrease* in the number of restrictions. A reversal of the tightening response would be expected in the removals that follow an in-sample tightening by the same country. There are 187 such matched removals, with a median gap of only six months, so the typical one occurs while the post-tightening elevation documented above is near its peak. A matched removal is a net loosening by a country that tightened earlier in the sample, but the data do not identify whether it lifts the very measure that tightening introduced. Three results, reported in full in Appendix E.5 (Table 10), establish the ratchet. First, restricted to matched removals, the post-removal path is flat throughout (average post-removal effect 0.011, s.e. = 0.050, for the same-direction match). Second, measured against the *pre-tightening* baseline, treated months after a removal show no shift off the elevated post-tightening path, so the twelve-month persistence is not carried by already-removed events. Third, estimating the tightening and matched-loosening designs jointly rejects the hypothesis that removals fully undo the tightening response ($p = 0.02$). Inflows do not recede after a matched removal at the point estimate, and a give-back of more than about half the tightening effect lies outside the 95 percent confidence interval.

Loosenings in general point the same way. Pooled across all 351 loosening events, the effect on stablecoin inflows is essentially nil, and in particular we find no evidence that

inflows decline when outflow restrictions are removed (Online Appendix F.2).[33] We refrain from a causal reading of these pooled paths, because outflow loosenings are preceded by significant pre-event dynamics. But what they cannot be is a reversal of the tightening response, since the late-horizon negative drift is concentrated entirely in removals that no in-sample tightening precedes (Appendix E.5). What the twelve-month window and the matched-removal tests establish is that the rise in cross-border stablecoin activity after a tightening persists for at least a year and does not detectably recede when restrictions are later removed, consistent with circumvention.

5.3 Stablecoins and Currency Crises

In this subsection, we provide event-study evidence that the onset of a currency crisis triggers increased dollarization through stablecoins. The outcome is the stock of stablecoins a country's residents hold, the object that corresponds to how dollarization is measured in the traditional literature, and we document three facts about it. First, following the onset of a currency crisis, holdings rise by roughly 0.07 percentage points of pre-crisis GDP over the year after onset, relative to matched EMDEs that do not enter a crisis. That accumulation is about $7.9 billion, or roughly 75 percent of the crisis countries' combined pre-crisis stock. Second, the accumulation is timed to the crisis itself: the pre-onset path is flat, and the response builds steadily over the months following onset. Third, the accumulation is concentrated in the year after onset and then persists. Extending the window to twenty-four months, the buildup slows and the stock settles onto its higher plateau, with no evidence of reversion at any horizon. The analysis proceeds in four steps. We describe how crises are dated and which episodes enter the sample, define the holdings outcome and its scaling, set out the event-study design and its inference, and present the response together with the evidence addressing the principal threats to interpreting it as crisis-driven dollarization.

Our outcome is the stablecoin analogue of the deposit dollarization stock that this literature measures, and it has one advantage for a crisis design. Holdings are measured in dollar units and cumulated from net flows (Section 4.3), so a change in holdings is net new accumulation by construction. This matters because the treatment is the depreciation itself. A deposit ratio rises mechanically when the depreciation revalues the existing dollar stock, which is why studies of deposit dollarization compute constant-exchange-rate adjustments (Levy Yeyati, 2006), while our outcome needs no such adjustment.

[33] The sample period 2021–2024 is, on net, a liberalizing one: it contains 351 loosenings against 217 tightenings. The design thus measures responses to tightenings imposed against a backdrop of gradual liberalization, which is worth bearing in mind when extrapolating to episodes of broad-based retrenchment.

Crisis dating and sample. We date currency crises with the criterion of Laeven and Valencia (2020): a country-month is in crisis when the twelve-month depreciation against the dollar reaches at least 30 percent and exceeds its own year-earlier rate by at least 10 percentage points. The published crisis dates of Laeven and Valencia (2020) end in 2017, and the database's 2026 update extends only its banking-crisis component (Laeven and Valencia, 2026), so we apply the criterion to the exchange-rate panel to date the crises in our window. The acceleration condition matters in our sample period: it separates a country-specific crisis from a common move against the dollar, and 2021–2025 contains a dollar tightening that moved most currencies at once.[34] The criterion is mechanical. It does not distinguish disorderly depreciations from policy-initiated exchange-rate unifications or floats, and the crisis label refers to the criterion throughout.

We include a country as a crisis event only when it meets that twelve-month criterion. We date the onset as the first month, in the year before the criterion is met, in which three-month depreciation exceeds a 30 percent annualized pace. The three-month window only moves the date; it does not add episodes that never reach the twelve-month threshold. Dating at the criterion month would push part of any response into the pre-period, because by then the slide has already been underway. Entry estimates are robust to detection windows of two to six months instead (Appendix F). Classification uses each country's full exchange-rate history, while the regressions use the sample from January 2021 onward, when stablecoin volumes become economically meaningful (Section 3.1). Each country contributes at most one event, its first *estimable* entry. An entry is estimable if its event window begins inside the regression sample and if, where an earlier episode overlaps the pre-window, at least three clean pre-months remain once the window is truncated at that episode's end (Appendix F). Currency crises recur, but a recurrence cannot enter as an event of its own: its pre-window lies in the aftermath of the previous spell, so it offers no clean pre-period against which to measure a response. The estimates below are therefore the response to the arrival of a country's first crisis of the sample period, not an average over repeat spells; we take up the implications below, for the long-horizon paths and for external validity. This yields 22 crisis entries, listed with their onset dates in Table 12, among them Ukraine and Russia in 2022, Sri Lanka, Ghana, Egypt, and Argentina in 2022, and Nigeria and Türkiye in 2023. One event, Suriname, is excluded from the baseline as an extreme influence outlier, so that no single event drives the pooled estimates, even though including it raises every estimate we report; the influence diagnostics appear with the robustness discussion below. The baseline

[34] The results are robust to dating crises with the Frankel and Rose (1996) variant instead, which applies the same rule at a 25 percent depreciation threshold; the estimates are very similar, and Appendix F reports them.

sample therefore comprises 21 events. Countries that are ever in crisis during the sample are excluded from every control pool, whether or not their episode is estimable, and this applies to Suriname as well.

All designs are anchored on entry only. Crisis *onsets* are sharply identified by the definition; crisis *ends* are not: under mechanical exit rules, episodes are routinely reclassified as over while inflation and depreciation remain far above pre-crisis levels, so an exit-anchored design would measure the effect of a currency settling into a new trend rather than of a recovery (Appendix F). Every estimate below should accordingly be read as "in the year after the onset of a currency crisis."

Outcome and scaling. The outcome is the holdings measure of Section 4.3: cumulative net stablecoin flows to a country since the start of the flow estimates in January 2018, with issuer treasuries excluded, expressed in percent of pre-crisis GDP. Four features of this construction are choices worth defending. First, the outcome is scaled by GDP because raw dollar holdings would not survive pooling: the treated countries' within-window ranges of dollar holdings span a factor of more than thirty thousand (Türkiye $4.13 billion against Burundi $122,000), so a pooled coefficient on dollar holdings would describe the three largest economies and nothing else. Second, holdings enter in levels rather than in logarithms. The share of GDP measures the object of interest, how much of an economy moves into stablecoins. Estimated holdings are also sometimes negative, since they are measured as cumulative net flows and are not constrained to be non-negative (Appendix G), so logarithms could drop exactly the observations where holdings are falling, and transformations defined at negative values, such as the inverse hyperbolic sine, estimate quantities that depend on the units of measurement (Chen and Roth, 2024). Third, the baseline denominator is WEO nominal GDP of the calendar year before the event's anchor year, a country-by-event constant fixed before treatment and measured near the event. The vintage matters because a currency crisis itself moves dollar GDP. Scaling by crisis-year GDP would inflate the outcome mechanically as the denominator collapses. A common pre-sample vintage, each country's 2020 GDP, is also fixed before treatment and yields somewhat larger estimates with the same inference (Online Appendix G), so the choice affects the measured magnitude, not the finding, and the baseline is the conservative one. Fourth, the denominator is GDP rather than total deposits, which would express the outcome in the units of the traditional literature. Deposits are predominantly in domestic currency, so the depreciation that defines the event would shrink that denominator mechanically, the crisis-year GDP problem in sharper form, and repairing it would require exactly the constant-exchange-rate adjustments the outcome was chosen not to need. Deposit stocks are also observed annually, so the ratio could not be

formed at the design's monthly frequency.

Specification and inference. We estimate a stacked event-study design, as in Section 5.2, because crisis entries are staggered across countries. Each entry event e is assigned its own dataset over a window from six months before to twelve months after the anchor, comprising the treated country and its clean controls: emerging market and developing economies that are never in crisis during the sample, drawn from a matched pool described below. We estimate

$$h_{c,t} = \sum_{\substack{k=-6 \\ k\neq -1}}^{12} \beta_k \, 1\{t - t_e = k\} \times Treat_{c,e} + \alpha_{e,k} + \delta_{e,c} + \varepsilon_{c,t,e}, \tag{3}$$

where $h_{c,t}$ is holdings in percent of pre-crisis GDP and the fixed effects are defined as in Section 5.2. The pooled coefficient weights events by their precision, and an equal-weighted average of the event-specific effects yields a nearly identical estimate (Online Appendix G). The stacked implementation accommodates the two features this setting requires, matched control pools and randomization inference.[35]

The control pool is built on three principles. Controls are emerging market and developing economies, as every crisis country is. An advanced economy is not a plausible counterfactual for a country in a currency crisis. Countries whose crisis status could never be evaluated are excluded rather than treated as calm, because a country with no usable price and exchange-rate history cannot be certified crisis-free.[36] Finally, each treated country is compared with the $k = 15$ controls most similar to it within its own region, matched on pre-window holdings level, holdings volatility, and economy size, all measured over the twelve months before the event window opens. The volatility distribution of the unmatched control pool is far more right-skewed than that of the crisis countries it would stand in for, with a mean 1.7 times theirs. With inference limited by the number of treated clusters, adding controls therefore widens the placebo distribution rather than sharpening the test (Appendix F).

Inference is based on randomization tests (see, e.g., Conley and Taber, 2011; MacKinnon and Webb, 2020; Young, 2019): for each of 1,000 draws, one control country per event is assigned placebo-treated status at the event's actual anchor month, and every reported statistic is referred to its own placebo distribution; confidence bands invert these tests horizon

[35] With 21 treated clusters, inference based on asymptotics in the number of treated groups is unreliable (MacKinnon and Webb, 2020); randomization tests are exact under the sharp null regardless of the cluster count.

[36] Venezuela's inflation series is not continuously available in our data sources, so its crisis status could never be evaluated, and a naive never-in-crisis screen would admit it as a clean control.

by horizon (Appendix F details the protocol). Standard errors clustered at the country level accompany the point estimates and, at 21 treated clusters, deliver the same conclusions.

The holdings response to crisis onset. Crisis onset is followed by a steady accumulation of stablecoin holdings, with no evidence of anticipation. Figure 9 plots the estimated path. A natural first concern in any event study around crises is that the outcome begins adjusting before the dated onset, either because agents anticipate the crisis or because the anchor mismeasures its start. The pre-onset coefficients show no such pattern: all six lie between -0.014 and $+0.003$ percentage points, and the joint randomization test of pre-trend flatness is passed comfortably ($p = 0.48$). From onset, holdings build gradually — roughly 0.02 percentage points of pre-crisis GDP by the third month, 0.04 by mid-year, and 0.072 (95 percent randomization interval $[0.023, 0.128]$) at month twelve — averaging 0.040 percentage points over the post-onset year (s.e. 0.014), significant by randomization inference ($p = 0.015$; cluster-robust $p = 0.010$). The average and the month-twelve level describe the same path: for a response that builds steadily from zero, the average over the post-onset year is mechanically about half the end-of-year level, and we use the average as the primary test statistic because a single parameter concentrates power.

These magnitudes are large relative to what the crisis countries held going in, and modest on any macroeconomic yardstick. Summing the month-twelve responses across the 21 events, scaled by each event's own pre-crisis GDP, the crisis countries accumulate roughly $7.9 billion against combined pre-crisis holdings of $10.4 billion, so their aggregate stablecoin holdings rise by roughly 75 percent over the year following onset.[37] The median crisis country entered its episode with a stock of 0.099 percent of pre-crisis GDP. In dollar terms, for the median crisis economy in the sample (pre-crisis GDP of roughly $83 billion, Table 12), the estimate corresponds to roughly $59 million of additional net stablecoin accumulation over the year following onset. This figure is net accumulation rather than gross activity, so it is smaller by construction than the two-way transfer increments of Section 5.2. The estimate is a share of GDP, so its dollar counterpart scales with economy size: at the quartiles of pre-crisis GDP, roughly $22 billion and $300 billion, the same response amounts to roughly $16 million and $215 million. Set against traditional deposit dollarization, the comparison runs in opposite directions on the two margins that matter. In stocks, these are small numbers: an economy counts as meaningfully deposit-dollarized already at 10 percent of broad money, and Reinhart et al. (2014) treat a reversal as lasting only once foreign-currency deposits settle below 20

[37] The aggregate is a translation of the event-specific month-twelve estimates, each stack fit alone as in the equal-weighted aggregation, into dollars; the inference remains that of the pooled path (Figure 9). It is dollar-weighted by construction, though excluding Russia and Ukraine, the two events for which the war in Ukraine is a competing explanation, lowers the rise only to roughly 66 percent.

percent of broad money, so the deposit stocks through which crises have historically been weathered dwarf the stablecoin stocks we measure. In speed, the comparison reverses, and the first-year response can be put side by side. In deposits, the average answer is small. The within-year movement of the dollar share is estimated at a quarter to a half of even the mechanical valuation effect of the depreciation (Honohan, 2007).[38] The most extreme deposit-ratio jumps on record are of the same proportional size: in the 2014–15 depreciations, the foreign-currency deposit share rose from 20 to 35 percent in Russia and from 37 to 65 percent in Kazakhstan between the start of 2014 and the third quarter of 2015, proportional jumps of roughly 75 percent (Krupkina and Ponomarenko, 2016). Our 75 percent is similar in size while containing none of their valuation component: it is net new dollar accumulation in its entirety, and a first-year pace that the deposit literature documents only over multi-year horizons around inflation peaks (Kamin and Ericsson, 2003; Reinhart et al., 2014). Agur (2023) measures that horizon directly: across threshold-based episode definitions in a panel of 173 countries, deposit substitution raises the dollar share by 4 to 12 percentage points per year when it happens, and a typical episode plays out over five years. The finding is therefore not that stablecoins rival bank deposits as a crisis refuge in scale. It is that a channel that barely existed in 2020 delivers first-year dollar accumulation at a pace the deposit margin has not been documented to match.

The response accumulates gradually from onset rather than arriving as a step, as would be expected if conversion of local currency continued as the crisis eroded it rather than occurring once at onset. The flat pre-period, against onsets defined by exchange-rate collapse, indicates that residents did not accumulate ahead of depreciations that, by construction, reach at least 30 percent over the following year: the substitution is a response to the crisis, not a leading indicator of it. What happens after the first year, the twelve-month window cannot say, so we extend it. Re-estimating the identical design with the post-window lengthened to twenty-four months — all but one event contribute through month twenty-four — shows the accumulation decelerating and the stock settling onto a higher plateau: holdings stand at 0.115 percentage points of pre-crisis GDP two years after onset (95 percent randomization interval $[0.009, 0.215]$), and the change between months twelve and twenty-four is +0.044 (randomization $p = 0.23$) — continued accumulation at roughly half the first-year pace, statistically indistinguishable from a flat stock (Appendix F, Figure 14). The data thus show no reversion over the two years we can observe. We do not push the claim further:

[38] Honohan (2007) estimates factors of 0.23 (full sample) to 0.55 (excluding outliers) on the mechanical valuation impact of an exchange-rate change on the deposit dollar share: within the year, depositors on average partially rebalance away rather than pile in. At the 30 percent depreciation that defines our crises and typical EMDE dollarization shares, these factors imply a within-year rise in the share of roughly 4 to 10 percent of its level.

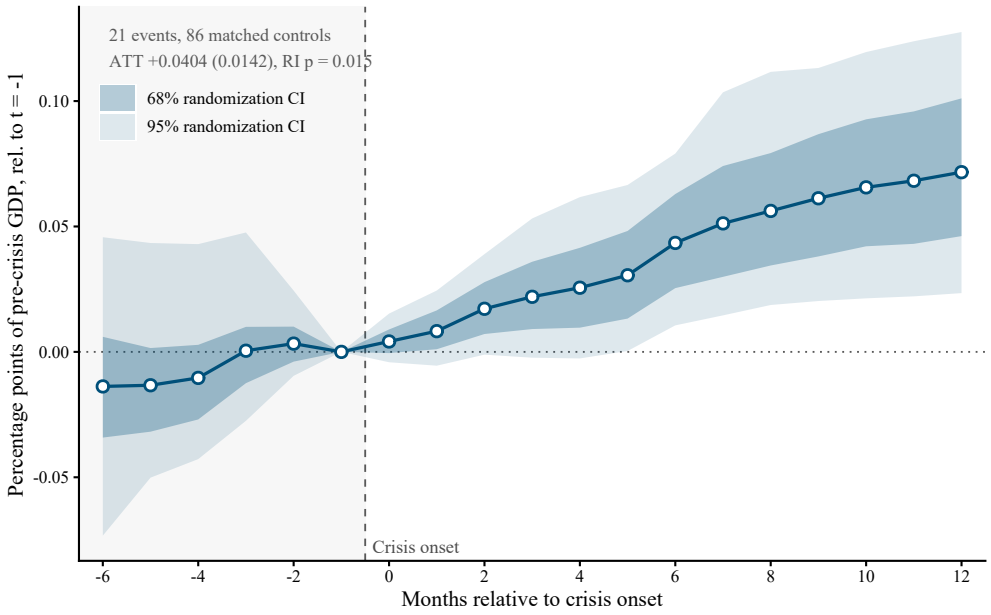


Figure 9: Stablecoin holdings around currency crises

Notes: The figure plots stacked event-study estimates of stablecoin holdings, in percentage points of pre-crisis GDP, around entry into a currency crisis under the Laeven and Valencia (2020) definition (21 events, 86 matched control countries, 5,909 observations; Suriname excluded as described in the text). Coefficients are relative to $k = -1$ (omitted, dashed vertical line). Shaded areas are 68 and 95 percent randomization confidence intervals, constructed by test inversion against the placebo distribution of 1,000 matched draws (one control country per event); the design includes event-by-event-time and event-by-country fixed effects, with matched EMDE control pools as described in the text. The corresponding results under the Frankel and Rose (1996) threshold are very similar and are reported in Appendix F.

distinguishing a permanently higher stock from slow eventual reversal would require longer horizons and defensible crisis end-dates, and neither is available.[39] Longer horizons would also change what is measured: currency crises recur, and because each country enters the design once, a second spell falls inside the first event's extended post-window, so distant horizons increasingly reflect recurrence rather than the aftermath of the original episode (Appendix F).

Threats and robustness. We address five concerns that could affect the estimates. Online Appendix G collects the corresponding estimates, and none alters the conclusions.

The first is influence of outliers. With economies spanning three orders of magnitude, a pooled coefficient could in principle describe a handful of large events, or a single outlier, rather than the phenomenon. Re-estimating the design on the full 22-event sample excluding one event at a time moves the average effect within $[0.064, 0.074]$ for every event but Suri-

[39] Reversal of dollarization, however, is rare in the deposit analogue. In the 1980–2001 sample of Reinhart et al. (2014), only four of eighty-five countries achieved a large and lasting reversal of deposit dollarization.

name, whose exclusion defines the baseline. Suriname is the smallest economy in the sample, and its estimated response of 0.65 is four times the next largest and concentrated in a single month. Movements this large in this small an economy are first order in a pooled estimate of a level outcome scaled by GDP. Including it raises the average effect to 0.068 (s.e. 0.030) and adds an impact-month response of 0.030, against 0.004 in the baseline, that is specific to that one event. No other estimate in this section moves materially under any single exclusion (Online Appendix G). Single exclusions cannot speak to Russia and Ukraine, whose 2022 crises arrived together with the war in Ukraine. The war moved stablecoin use for reasons of its own, through demand possibly connected to the sanctions environment in Russia and wartime transfers in Ukraine, so we exclude the pair jointly. This lowers the average effect to 0.032 (randomization $p = 0.051$), so the war episodes contribute to the response but do not carry it. Size does not drive the estimate either. Fitting each event's stack separately and averaging the event-specific effects with equal weights yields 0.041 against the pooled 0.040, so the pooled estimate is not an artifact of its implicit precision weights. The GDP-vintage variant speaks to the denominator side of the same concern. Scaling every country by its 2020 GDP, a common pre-sample vintage, yields somewhat larger estimates with the same inference (0.047, randomization $p = 0.010$), because pre-crisis-year GDP exceeds 2020 GDP for most crisis countries. The baseline's event-specific pre-crisis vintage is therefore the conservative choice, and it is the one measured near the event.

The second is dating error. Varying the real-time detection window from two to six months moves the average effect between 0.034 and 0.042, with randomization $p \leq 0.043$ in every case. The joint pre-trend test rejects in one of the five designs ($p = 0.04$). The direction of any residual dating error is attenuation. Where the pre-onset coefficients deviate from zero they lie below the reference month and rise toward it, consistent with accumulation already underway before the dated onset, so mistimed anchors move response months into the pre-period rather than the reverse.

The third is implementation dependence. The Callaway and Sant'Anna (2021) estimator replaces our regression-based aggregation with nonparametric group-time comparisons, and the Sun and Abraham (2021) estimator is robust to heterogeneous treatment effects contaminating individual event-study coefficients. Both are estimated without the matched control pools, and both reproduce the stacked estimate to within one standard error, so the estimate depends on neither the fixed-effects machinery nor the matching. Assessed with the randomization protocol used throughout, both reject the null (randomization $p = 0.004$ for Callaway–Sant'Anna and $p = 0.007$ for Sun–Abraham). Online Appendix G reports the estimated paths, the common GDP scale the comparison requires, and the implementation

details.[40] The comparison also shows what the matched pools contribute. Estimated on the unmatched panel, both estimators display a drift in holdings before onset, and the matched pools remove it.

The fourth is a policy confound. Crises rarely arrive alone. Governments often restrict capital flows just as the currency collapses, and Section 5.2 shows that such restrictions move stablecoin flows on their own. We control for an indicator that a capital flow measure was activated in the country during the current or preceding six months, which is the treatment of Section 5.2 and the horizon over which its effect builds there. This leaves the estimate essentially unchanged (0.039, against 0.040 without the control), so the holdings response is not a relabeled restriction response. The activation indicator is the only control we add. Contemporaneous macroeconomic variables such as inflation, depreciation, and interest rates are the crisis itself, the treatment and its direct outcomes, so conditioning on them would absorb part of the response the design is built to measure. That said, some of the measured accumulation may be dollar demand that existed anyway and was pushed onto stablecoin rails by crisis-time restrictions on official channels, e.g., through the channel of Section 5.2. Our estimates cannot separate newly created from rerouted dollar demand, and both readings constitute dollarization through stablecoins.

The fifth is concentration. The events are not spread evenly across the cycle. Fourteen of the 21 onsets fall between the second half of 2021 and the end of 2022, the period of the global dollar tightening, and twelve of the 21 events are in Africa and the Middle East. The design addresses the common component directly. Event-by-event-time fixed effects absorb calendar-time shocks, the acceleration condition in the crisis definition screens out countries moved only by the common dollar cycle, and matched controls are drawn from the same region and period, so each crisis country is compared with neighbors facing the same global environment. The design cannot, however, supply a second global cycle. The evidence is therefore closer to "many countries, few episodes of global stress" than the event count alone suggests, and we see this as the main limit on external validity. The estimand is limited in a second way. It is the response to a country's first crisis of the sample period, and the design cannot estimate the response to recurrence, given the lack of clean pre-periods. Under the adoption-cost reading of the persistence evidence, a repeat spell finds residents already holding wallets, accounts, and stablecoin balances, so the recurrence response could plausibly be larger than the first-crisis response. Most onsets, moreover, fall in 2021 and 2022, so these are the responses of a young stablecoin ecosystem. The outstanding stock then was roughly half its end-of-2025 level, so the dollar magnitudes in this section describe a smaller and less mature market than the one that exists now.

[40] The results are also robust to using synthetic control methods. Details are available upon request.

6 Conclusion

We estimate that stablecoin transfers have grown to trillions of dollars a year. However, official statistics do not yet comprehensively measure these flows. This paper narrows that blind spot. We estimate the geography of the wallets that send and receive stablecoins, and build a panel that allocates $21.4 trillion in transfers between 2018 and 2025. With it, we first establish descriptive facts about stablecoins as a channel of international capital flows. Our estimates show that stablecoins operate mainly as a dollar rail for emerging market and developing economies. We estimate that roughly 90 percent of the volume crosses borders, and that about three quarters of the stablecoins the panel covers are held there.

In our event-study designs, we show that a country's cross-border stablecoin activity rises by roughly 30 percent in the six months after it tightens restrictions on capital outflows, relative to countries that activate no restriction of their own. When a country enters a currency crisis, residents accumulate stablecoins, raising the crisis countries' combined holdings by roughly 75 percent of the pre-crisis stock over the year following onset, relative to similar countries that do not enter a crisis. We also find that both responses persist. Formal tests reject that removing a restriction fully undoes the rise it triggered, and the stablecoin holdings accumulated in a currency crisis settle at their higher level rather than reverting. A restriction after which activity rises may not be followed by a decline when it is lifted.

Our estimates agree with Brazil's official statistics on wholesale imports. They also match the United States and China shares recorded in Binance's internal figures, and the United States share recorded in KuCoin's.

However, caveats remain. Our estimates are probabilistic, and they attach a geography to each flow but not a motive. Decomposing flows by use, into remittances, trading, and the accumulation of dollar savings, would sharpen both the measurement and the designs built on it. The same methodology could attach a geography to transfers of other crypto assets, though that is not what this panel measures. The panel itself is freely available, for research beyond the questions posed here. In the long run, the statistical initiatives now underway may bring stablecoin flows comprehensively into official statistics. Until then, our estimates offer a measurement that is available today.

References

Acosta-Henao, Miguel, Laura Alfaro, and Andrés Fernández (2025) "Sticky Capital Controls," *Journal of International Economics*, 157.

Adrian, Tobias, Parma Bains, Marianne Bechara et al. (2025) "Understanding Stablecoins," *Departmental Papers*, 2025 (009), 1, 10.5089/9798229024075.087.

Agur, Itai (2023) "Measuring the Speed of Currency Substitution," *Applied Economics Letters*, 30 (10), 1385–1389, 10.1080/13504851.2022.2056122.

Aldasoro, Iñaki, Paula Beltrán, and Federico Grinberg (2026) "Stablecoin Flows and Spillovers to FX Markets," BIS Working Papers 1340, Bank for International Settlements.

Aldasoro, Iñaki, Giulio Cornelli, Massimo Ferrari Minesso, Leonardo Gambacorta, and Maurizio Michael Habib (2025) "Stablecoins, Money Market Funds and Monetary Policy," *Economics Letters*, 247, 112203, 10.1016/j.econlet.2025.112203.

Alfaro, Laura, Anusha Chari, and Fabio Kanczuk (2017) "The Real Effects of Capital Controls: Firm-Level Evidence from a Policy Experiment," *Journal of International Economics*, 108, 191–210.

Alnasaa, Marwa, Nikolay Gueorguiev, Jiro Honda, Eslem Imamoglu, Paolo Mauro, Keyra Primus, and Dmitriy Rozhkov (2022) "Crypto-assets, corruption, and capital controls: Cross-country correlations," *Economics Letters*, 215, 110492, https://doi.org/10.1016/j.econlet.2022.110492.

Aramonte, Sirio, Wenqian Huang, and Andreas Schrimpf (2022) "Tracing the Footprint of Cryptoisation in Emerging Market Economies," *BIS Quarterly Review, Basel, March*.

Athey, Susan, Ivo Parashkevov, Vishnu Sarukkai, and Jing Xia (2016) "Bitcoin Pricing, Adoption, and Usage: Theory and Evidence," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2826674, Stanford University Graduate School of Business Research Paper No. 16-42.

Auer, Raphael, Giulio Cornelli, Sebastian Doerr, Jon Frost, and Leonardo Gambacorta (2025a) "Crypto Trading and Bitcoin Prices: Evidence from a New Database of Retail Adoption," *IMF Economic Review*, 10.1057/s41308-025-00275-0, Published online 23 April 2025.

Auer, Raphael, Ulf Lewrick, and Jan Paulick (2025b) "DeFiying Gravity? An Empirical Analysis of Cross-Border Bitcoin, Ether and Stablecoin Flows," BIS Working Papers 1265, Bank for International Settlements, https://www.bis.org/publ/work1265.htm.

Azar, Pablo D., Maryam Farboodi, and Nish D. Sinha (2026) "Stablecoins Meet the Mundell–Fleming Trilemma," Staff Reports 1202, Federal Reserve Bank of New York, 10.59576/sr.1202.

Banco Central do Brasil (2024) "External Sector Statistics: Press Release, July 2024," https://www.bcb.gov.br/content/statistics/externalsectorstatistics_prev/202407_External_sector_statistics_text.pdf.

Bank for International Settlements (2023) "BIS Annual Economic Report 2023."

Basu, Suman, Marianne Bechara, Eugenio Cerutti et al. (2026) "Foreign Currency Stablecoins: Initial Considerations on Policies to Enhance Macrofinancial Stability," *IMF Departmental Papers*.

Béres, Ferenc, István András Seres, András A. Benczúr, and Mikerah Quintyne-Collins (2021) "Blockchain is Watching You: Profiling and Deanonymizing Ethereum Users," in *2021 IEEE International Conference on Decentralized Applications and Infrastructures (DAPPS)*, 69–78: IEEE, 10.1109/DAPPS52256.2021.00013.

Bergant, Katharina, Andrés Fernández, Ken Teoh, and Martín Uribe (2026) "Expanding the Landscape of Cross-Border Flow Restrictions: Modern Tools and Historical Perspectives,"Technical report, National Bureau of Economic Research, Inc.

Bhargava, Apoorv, Romain Bouis, Annamaria Kokenyne, Manuel Perez Archila, Umang Rawat, and Ratna Sahay (2023) "Capital Controls in Times of Crisis—Do They Work?" in *Shocks and Capital Flows: Policy Responses in a Volatile World*, Washington, DC: International Monetary Fund.

BIS Innovation Hub (2023) "Project Atlas: Mapping the World of Decentralised Finance,"Technical report, Bank for International Settlements, https://www.bis.org/publ/othp76.pdf.

Bocola, Luigi and Guido Lorenzoni (2020) "Financial Crises, Dollarization, and Lending of Last Resort in Open Economies," *American Economic Review*, 110 (8), 2524–2557, 10.1257/aer.20180830.

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

Calvo, Guillermo A., Leonardo Leiderman, and Carmen M. Reinhart (1996) "Inflows of Capital to Developing Countries in the 1990s," *Journal of Economic Perspectives*, 10 (2), 123–139.

Cardoso, Eliana and Ilan Goldfajn (1998) "Capital Flows to Brazil: The Endogeneity of Capital Controls," *IMF Staff Papers*, 45 (1), 161–202.

Cardozo, Pamela, Andrés Fernández, Jerzy Jiang, and Felipe Rojas (2024) "On Cross-Border Crypto Flows: Measurement, Drivers, and Policy Implications," *IMF Working Papers*, https://www.imf.org/en/Publications/WP/Issues/2024/12/20/On-Cross-Border-Crypto-Flows-Measurement-Drivers-and-Policy-Implications-559166.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019) "The effect of minimum wages on low-wage jobs," *The Quarterly Journal of Economics*, 134 (3), 1405–1454.

Cerutti, Eugenio, Jiaqian Chen, and Martina Hengge (2025a) "A Primer on Bitcoin Cross-Border Flows: Measurement and Drivers," *Journal of International Money and Finance*, 159, 103424,

10.1016/j.jimonfin.2025.103424.

Cerutti, Eugenio, Melih Firat, and Martina Hengge (2025b) "Global Cross-Border Payments: A $1 Quadrillion Evolving Market?" IMF Working Paper WP/25/120, International Monetary Fund.

Chainalysis (2025) "The 2025 Geography of Cryptocurrency Report,"Technical report, Chainalysis, https://www.chainalysis.com/reports/2025-geo-crypto-report/.

Chen, Jiafeng and Jonathan Roth (2024) "Logs with Zeros? Some Problems and Solutions," *Quarterly Journal of Economics*, 139 (2), 891–936.

Christiano, Lawrence J., Husnu C. Dalgic, and Armen Nurbekyan (2021) "Financial Dollarization: Efficient Intranational Risk Sharing or Prescription for Disaster?" Working Paper 29034, National Bureau of Economic Research.

Cong, Lin William, Pulak Ghosh, Jiasun Li, and Qihong Ruan (2024) "Inflation Expectation and Cryptocurrency Investment," https://www.nber.org/papers/w32945, NBER Working Paper No. 32945.

Cong, Lin William, Gordon Y. Liao, Eswar Prasad, and Tony Zhang (2026) "Stablecoins as Tokenized Money," Working Paper.

Conley, Timothy G. and Christopher R. Taber (2011) "Inference with "Difference in Differences" with a Small Number of Policy Changes," *Review of Economics and Statistics*, 93 (1), 113–125.

Dalgic, Husnu C. (2024) "Financial Dollarization in Emerging Markets: An Insurance Arrangement," *International Economic Review*, 65 (3), 1189–1219, 10.1111/iere.12686.

Drenik, Andrés and Diego J. Perez (2021) "Domestic Price Dollarization in Emerging Economies," *Journal of Monetary Economics*, 122, 38–55.

Du, Wenxin, Catherine Huang, and David Scharfstein (2026) "Competing rails for cross-border payments: banks, fintechs, and stablecoins," *Harvard Business School, working paper*.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor (2025) "A Local Projections Approach to Difference-in-Differences," *Journal of Applied Econometrics*, 40, 741–758, 10.1002/jae.70000.

Eichengreen, Barry and Andrew K. Rose (2014) "Capital Controls in the 21st Century," *Journal of International Money and Finance*, 48, 1–16.

Erten, Bilge, Anton Korinek, and José Antonio Ocampo (2021) "Capital Controls: Theory and Evidence," *Journal of Economic Literature*, 59 (1), 45–89.

European Union (2023) "Regulation (EU) 2023/1113 on Information Accompanying Transfers of Funds and Certain Crypto-assets."

Fernández, Andrés, Michael Gottschalk, and Tatsuya Sugiura (forthcoming) "On The Drivers of Cross-Border Stablecoin Flows: A Case Study," IMF Working Paper (forthcoming), International Monetary Fund.

Fernández, Andrés, Michael W. Klein, Alessandro Rebucci, Martin Schindler, and Martín Uribe (2016) "Capital Control Measures: A New Dataset," *IMF Economic Review*, 64 (3), 548–574.

Fernández, Andrés, Alessandro Rebucci, and Martín Uribe (2015) "Are Capital Controls Countercyclical?" *Journal of Monetary Economics*, 76, 1–14.

Ferrante, Francesco and Nils Gornemann (2022) "Devaluations, Deposit Dollarization, and Household Heterogeneity," International Finance Discussion Papers 1336, Board of Governors of the Federal Reserve System, 10.17016/IFDP.2022.1336.

Financial Action Task Force (2023) "Virtual Assets: Targeted Update on Implementation of the FATF Standards on VAs and VASPs."

Financial Stability Board (2023) "High-level Recommendations for the Regulation, Supervision and Oversight of Crypto-Asset Activities and Markets."

Fisher, Ronald A. (1935) *The Design of Experiments*, Edinburgh: Oliver and Boyd.

Forbes, Kristin, Marcel Fratzscher, and Roland Straub (2015) "Capital-Flow Management Measures: What Are They Good For?" *Journal of International Economics*, 96 (S1), S76–S97.

Forbes, Kristin J. and Francis E. Warnock (2012) "Capital Flow Waves: Surges, Stops, Flight, and Retrenchment," *Journal of International Economics*, 88 (2), 235–251.

Frankel, Jeffrey A. and Andrew K. Rose (1996) "Currency Crashes in Emerging Markets: An Empirical Treatment," *Journal of International Economics*, 41 (3–4), 351–366.

Fratzscher, Marcel (2012) "Capital Flows, Push versus Pull Factors and the Global Financial Crisis," *Journal of International Economics*, 88 (2), 341–356.

Friedrich, Christian and Laura Zhao (2026) "Patterns and Determinants of Global Cryptocurrency Flows," Staff Working Paper 2026-15, Bank of Canada.

G7 (2023) "G7 Finance Ministers and Central Bank Governors' Communiqué," https://g7.utoronto.ca/finance/230513-finance.html, Niigata, Japan, May 13, 2023; crypto-asset language in the "Financial Digitalization" section.

Guidotti, Pablo E. and Carlos A. Rodríguez (1992) "Dollarization in Latin America: Gresham's Law in Reverse?" *IMF Staff Papers*, 39 (3), 518–544.

Guillame-Bert, Mathieu, Sebastian Bruch, Richard Stotz, and Jan Pfeifer (2023) "Yggdrasil Decision Forests: A Fast and Extensible Decision Forests Library," in *Proceedings of the 29th ACM*

SIGKDD Conference on Knowledge Discovery and Data Mining, KDD 2023, Long Beach, CA, USA, August 6-10, 2023, 4068–4077, 10.1145/3580305.3599933.

Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2017) *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*: Springer.

He, Dong, Annamaria Kokenyne, Xavier Lavayssière, Inutu Lukonga, Nadine Schwarz, Nobuyasu Sugimoto, and Jeanne Verrier (2022) "Capital Flow Management Measures in the Digital Age: Challenges of Crypto Assets," https://www.imf.org/en/Publications/fintech-notes/Issues/2022/05/09/Capital-Flow-Management-Measures-in-the-Digital-Age-516671.

He, Haibo and Edwardo A Garcia (2009) "Learning from imbalanced data," *IEEE Transactions on knowledge and data engineering*, 21 (9), 1263–1284.

Hofmann, Boris, Aaron Mehrotra, and Jan Paulick (2026) "Dollarisation and Monetary Control: What Lessons for the Rise of Stablecoins?" BIS Working Papers 1370, Bank for International Settlements, Basel, https://www.bis.org/publ/work1370.pdf.

Honohan, Patrick (2007) "Dollarization and Exchange Rate Fluctuations," Policy Research Working Paper 4172, The World Bank.

Hu, Maggie R., Adrian D. Lee, and Tālis J. Putniņš (2021) "Evading Capital Controls via Cryptocurrencies: Evidence from the Blockchain," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3956933, SSRN Working Paper No. 3956933, last revised September 2023.

Imbens, Guido W. and Donald B. Rubin (2015) *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, New York: Cambridge University Press.

IMF (2023) "Elements of Effective Policies for Crypto Assets."

——— (2025a) *Balance of Payments and International Investment Position Manual, Seventh Edition (BPM7)*, Washington, DC: International Monetary Fund.

——— (2025b) "Third Phase of the G20 Data Gaps Initiative (DGI-3): Third Progress Report,"Technical report, International Monetary Fund, Published on behalf of the Inter-Agency Group on Economic and Financial Statistics and the Financial Stability Board.

Kamin, Steven B. and Neil R. Ericsson (2003) "Dollarization in Post-Hyperinflationary Argentina," *Journal of International Money and Finance*, 22 (2), 185–211.

Kaplan, Ethan and Dani Rodrik (2002) "Did the Malaysian Capital Controls Work?" *NBER Working Paper No. 8142; in Preventing Currency Crises in Emerging Markets, S. Edwards and J. Frankel (eds.), University of Chicago Press.*

King, Gary and Langche Zeng (2001) "Logistic regression in rare events data," *Political analysis*,

9 (2), 137–163.

Koepke, Robin (2019) "What Drives Capital Flows to Emerging Markets? A Survey of the Empirical Literature," *Journal of Economic Surveys*, 33 (2), 516–540.

Krupkina, Anna and Alexey Ponomarenko (2016) "Deposit Dollarization and National Currency Depreciation in Russia and Kazakhstan," Research and Forecasting Department Analytical Note 1, Bank of Russia.

Laeven, Luc and Fabián Valencia (2020) "Systemic Banking Crises Database II," *IMF Economic Review*, 68 (2), 307–361.

Laeven, Luc and Fabian Valencia (2026) "Systemic Banking Crises Database: 1970-2025," IMF Working Paper 2026/94, International Monetary Fund.

Levy Yeyati, Eduardo (2006) "Financial Dollarization: Evaluating the Consequences," *Economic Policy*, 21 (45), 62–118, 10.1111/j.1468-0327.2006.00154.x.

——— (2021) "Financial Dollarization and De-dollarization in the New Millennium," working paper, Latin American Reserve Fund (FLAR).

von Luckner, Clemens Graf, Robin Koepke, and Silvia Sgherri (2024) "Crypto as a Marketplace for Capital Flight," IMF Working Paper WP/24/133, International Monetary Fund.

von Luckner, Clemens Graf, Carmen M Reinhart, and Kenneth Rogoff (2023) "Decrypting New Age International Capital Flows," *Journal of Monetary Economics*, 138, 104–122.

Lyons, Richard K. and Ganesh Viswanath-Natraj (2023) "What keeps stablecoins stable?" *Journal of International Money and Finance*, 131, 102777, https://doi.org/10.1016/j.jimonfin.2022.102777.

Ma, Yiming, Yao Zeng, and Anthony Lee Zhang (2025) "Stablecoin Runs and the Centralization of Arbitrage," NBER Working Paper 33882, National Bureau of Economic Research, https://www.nber.org/papers/w33882.

MacKinnon, James G. and Matthew D. Webb (2020) "Randomization Inference for Difference-in-Differences with Few Treated Clusters," *Journal of Econometrics*, 218 (2), 435–450.

Makarov, Igor and Antoinette Schoar (2021) "Blockchain Analysis of the Bitcoin Market," https://www.nber.org/papers/w29396, NBER Working Paper No. 29396.

Medvedev, Evgeny and the D5 team (2018) "Ethereum ETL," https://github.com/blockchain-etl/ethereum-etl.

Meiklejohn, Sarah, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage (2013) "A fistful of bitcoins: characterizing payments among men

with no names," in *Proceedings of the 2013 conference on Internet measurement conference*, 127–140.

Neanidis, Kyriakos C. and Christos S. Savva (2009) "Financial Dollarization: Short-Run Determinants in Transition Economies," *Journal of Banking & Finance*, 33 (10), 1860–1873.

Ratha, Dilip, Sonia Plaza, and Eung Ju Kim (2024) "In 2024, Remittance Flows to Low- and Middle-Income Countries Are Expected to Reach $685 Billion, Larger than FDI and ODA Combined,"Technical report, World Bank, https://blogs.worldbank.org/en/peoplemove/in-2024--remittance-flows-to-low--and-middle-income-countries-ar, People Move blog, December 18, 2024.

Reinhart, Carmen M., Kenneth S. Rogoff, and Miguel A. Savastano (2014) "Addicted to Dollars," *Annals of Economics and Finance*, 15 (1), 1–51.

Reslow, André ⓡ Edona Reshidi ⓡ Andrew Usher ⓡ Jeanne Verrier ⓡ Marcello Miccoli (2026) "Deconstructing the Stablecoin Sandwich: The Cost of Cross-Border Stablecoin Transfers," IMF Working Paper forthcoming, International Monetary Fund, Forthcoming.

Schär, Fabian, Anneke Kosse, Tara Rice, Takeshi Shirakami, and Jirapat Siridhasanakul (2026) "The Anatomy of Stablecoin Transactions," BIS Working Papers 1359, Bank for International Settlements, https://www.bis.org/publ/work1359.htm.

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

United States Department of the Treasury (2023) "Illicit Finance Risk Assessment of Decentralized Finance," https://home.treasury.gov/system/files/136/DeFi-Risk-Full-Review.pdf.

United States Department of the Treasury, Office of Foreign Assets Control (2026) "Sanctions Programs and Country Information," https://ofac.treasury.gov/sanctions-programs-and-country-information.

Valev, Neven T. (2010) "The Hysteresis of Currency Substitution: Currency Risk vs. Network Externalities," *Journal of International Money and Finance*, 29 (2), 224–235.

Young, Alwyn (2019) "Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results," *Quarterly Journal of Economics*, 134 (2), 557–598.

A Regional Centralized Exchanges

This appendix lists all centralized exchanges that we classify as particularly regional. As described in section 3, an exchange is designated as regional by manually inspecting the services it offers and

the markets it targets, and the designation is cross-validated through the geographic composition of the exchange's web traffic.

Africa and Middle East: Altcointrader, Artis Turba Exchange, Arzpaya.com, Bit2c, Bitoasis, Nobitex, Roqqu, Valr.
Asia and Pacific: Bitbank, Bitbns, Bitflyer, Bithumb, Bitkub, Bitopro, Btc Markets, Coincheck, Coindcx, Coinhako, Coinjar, Coinone, Coins.ph, Coinspot, Coinswitch, Digital Surge, Gdac, Gmo Coin, Gopax, Indodax, Korbit, Maicoin, Max Exchange, Reku, Tokocrypto, Upbit, Wazirx.
Europe: Anycoin Direct, Bitci, Bitcoin Meester, Bitcoin Suisse, Bitpanda, Bitvavo, Btcturk, Coinmetro, Exmo, Firi, Kanga Exchange, Kuna, Litebit, Norwegian Block Exchange, Paribu, Revolut, Swissborg, Yobit, Zonda.
Latin America and Caribbean: Bitso, Brasil Bitcoin, C-patex, Lemon Cash, Mercado Bitcoin, Orionx, Ripio, Southxchange.
North America: Binance US, Bitbuy, Coinsquare, FTX US, Ndax, Netcoins, Shakepay, Wealthsimple.

B Data Overview

This appendix collects overview statistics of the raw data (section 3.1) and of the training data (section 3.3): the blockchain sample by chain (Table 3), the domain name sample by name service (Table 4), and the training data by region and identification approach (Table 5).

Blockchain	**Genesis Block**	**# Transactions**
Ethereum	2015, Jul 30	3,180,948,859
Tron	2018, Jun 25	12,535,209,471
Binance Smart Chain	2020, Aug 29	10,658,700,651
Optimism	2021, Jan 14	895,021,749
Arbitrum	2021, May 28	2,220,931,313
Base	2023, Jun 15	5,169,516,168
Linea	2023, Jul 06	292,030,980
Total		**34,952,359,191**

Table 3: Overview of Blockchain Data

C Classification Model: Held-out Evaluation

This appendix reports the held-out evaluation of the regional classification model. The prompt, the features, the two-stage training setup, and the ROC curves are in Online Appendix D.

The model is evaluated on the 10 percent of labeled wallets held out before training, 43,768 wallets. Figure 10 displays the confusion matrix, with true regions in the rows and predicted regions

Name Service	# of Registered Domains
Ethereum Name Service	3,413,426
Uxlink	875,564
Spaceid	474,021
Basenames	437,329
Linea ENS Subdomains	429,956
Arbid	303,632
Total	**5,933,958**

Table 4: Overview of Domain Name Services

Region	Domain	Regional CEX	Total	% of Total
Africa and Middle East	19,731	21,123	40,854	9.3%
Asia and Pacific	129,516	90,472	219,988	50.3%
Europe	94,937	40,466	135,403	30.9%
Latin America and Caribbean	3,905	5,621	9,526	2.2%
North America	23,569	8,435	32,004	7.3%
Total	**271,658**	**166,117**	**437,775**	**100.0%**

Table 5: Overview of Training Data

in the columns, normalized so that each row sums to 100 percent. A perfectly accurate model would display 100 percent along the main diagonal. Uniform random guessing would place 20 percent in every cell. Overall accuracy is 74.1 percent. The 74.1 percent figure weights each region by its share of the testing data. The unweighted average of the five per-region accuracies on the diagonal is 69.3 percent. The weighted figure is higher because Asia and Pacific, the best predicted region, accounts for about half of the testing data.

Table 6 summarizes held-out performance of each component model, of the combined five-region classification, and of the cross-source validation exercise.

Performance by labeling source. The last two columns of Table 6 split held-out accuracy by the source of the wallet's label. Regional-exchange wallets are classified almost perfectly (99.9 percent), as expected once the model has learned which exchanges serve which regions. Domain-name wallets, whose labels have no mechanical link to the features, are the demanding benchmark. Combined five-region accuracy is 58.3 percent (95 percent CI: 57.7–58.9). On the probabilistic allocation measure, the average probability mass assigned to the true region is 41 percent for domain-name-labeled wallets and 98 percent for regional-exchange-labeled wallets.

Cross-source validation. The two labeling sources are generated by entirely different mechanisms. Retraining the full architecture on domain-name-labeled wallets alone and applying it to all 166,117 regional-exchange wallets, whose labels it has never seen, assigns the correct region to 91.5 percent of them (95 percent CI: 91.4–91.7; class-balanced accuracy 88.2 percent). Because the exchange-to-region associations this model relies on were learned entirely from domain-name-labeled wallets, the two sources point at the same geography.

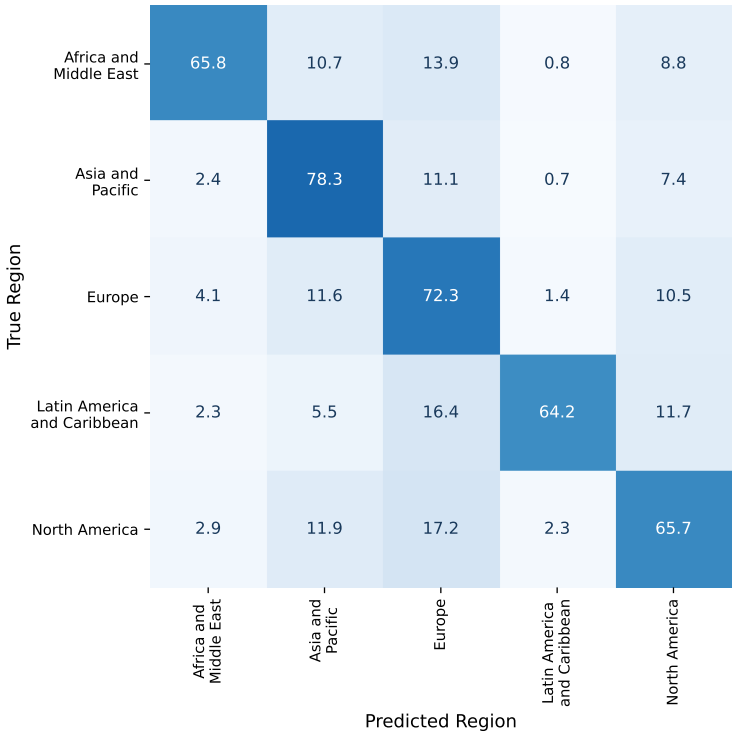


Figure 10: Confusion Matrix of Classification Model

Table 6: Classification Performance by Model and Labeling Source

Model	N	Accuracy [95% CI]	Macro	Majority	Domain	Reg. CEX
Stage 1 (three groups)	43,768	77.1 [76.7, 77.5]	75.1	50.3	63.1	99.93
Stage 2: Africa and Middle East vs. Europe	17,622	86.7 [86.2, 87.2]	82.7	76.8	79.6	99.90
Stage 2: Latin America vs. North America	4,148	89.1 [88.1, 90.0]	82.8	77.0	83.8	100.00
China vs. rest of Asia and Pacific	21,997	81.9 [81.4, 82.4]	83.1	80.0	69.1	100.00
Combined five-region classification	43,768	74.1 [73.7, 74.5]	69.3	50.3	58.3	99.90
Cross-source test (trained on domain names only)	166,117	91.5 [91.4, 91.7]	88.2	54.5	–	–

Notes: All values in percent. Accuracy is the share of correctly classified held-out wallets, with Wilson 95% confidence intervals. Macro is the unweighted average of per-class recalls; any strategy that always guesses the same class scores $1/K$ on this measure, where K is the number of classes. Majority is the accuracy of always guessing the largest class; while this strategy can score high plain accuracy in an imbalanced sample, it has no discriminating power, as it never correctly classifies a wallet outside the largest class. The last two columns split held-out accuracy by the source of the wallet's label (domain names vs. regional exchanges). The final row evaluates a model trained exclusively on domain-name-labeled wallets on all 166,117 regional-exchange-labeled wallets, none of which entered its training.

The pattern of off-diagonal errors in Figure 10 is consistent with the model correcting heritage names rather than repeating them. Wallets labeled Africa and Middle East are classified as North America 8.8 percent of the time, against only 2.9 percent in the reverse direction. Among domain-name-labeled wallets alone the same asymmetry is 18.5 percent against 3.9 percent. The same signature appears for Latin America vis-à-vis North America, and for Africa and Latin America vis-à-vis Europe. Off-diagonal errors also pull toward Asia and Pacific and Europe, the two largest

labeled groups, a majority-class bias that the balanced sample weights mitigate but do not eliminate. Errors involving Asia and Pacific are roughly symmetric, consistent with emigration from that region offsetting the majority-class pull. A model trained on domain-name wallets alone loses about one percentage point of accuracy on held-out domain-name wallets relative to the deployed model, which is why training uses both sources.

D Country Allocation of Transfers

This appendix writes the country allocation of section 3. The filters that define the covered universe are listed there. Each remaining transfer of value v is assigned across ordered country pairs as follows.

D.1 Regional shares of centralized exchanges

Let $V_{x,r,t,b}$ be the estimated stablecoin transaction volume between centralized exchange x and region r through self-custodial wallets, at hour t on blockchain b. Each transfer between exchange x and a self-custodial wallet contributes its value multiplied by that wallet's estimated probability of belonging to region r. The proportion of exchange x assigned to region r is

$$P_x(r \mid t, b) = \frac{V_{x,r,t,b}}{\sum_{r'} V_{x,r',t,b}}. \tag{4}$$

Allowing the probability to vary with the hour of the day and the blockchain used captures heterogeneity that stems from the popularity of blockchains in certain regions, and from increased or decreased activity of regions during certain times of the day. The construction assumes that, within a given hour and blockchain, an exchange's stablecoin traffic with self-custodial wallets is representative of the geography of its stablecoin traffic overall. Centralized exchanges that have been determined to be particularly regional are not split in this manner and are assigned fully to their respective regions.

D.2 Country split

Denote by $W_{c,x,t}$ the web traffic of the website for centralized exchange x that has originated from country c on day t. The regional share $P_x(r \mid t, b)$ is subdivided across the countries c in region r as

$$P_x(c \mid t, b) = \frac{W_{c,x,t}}{\sum_{c' \in r} W_{c',x,t}} \cdot P_x(r \mid t, b). \tag{5}$$

Because web traffic is observed at a daily frequency while $P_x(r \mid t, b)$ varies hourly, each day's traffic shares are applied to all hours of that day. For China the first factor is replaced by the

China model's probability, as described in section 3.2, and only the residual Asia and Pacific share enters the traffic split.

For any self-custodial wallet, the country-level breakdowns of centralized exchanges are propagated backwards as a transaction-volume-weighted average of the exchanges the wallet has interacted with:

$$P_w(c) = \frac{\sum_x \left(\sum_{t,b} V_{w,x,t,b} \cdot P_x(c \mid t, b) \right)}{\sum_x \left(\sum_{t,b} V_{w,x,t,b} \right)}. \tag{6}$$

Wallet probabilities do not vary over time or by blockchain. For a wallet that has never interacted with a centralized exchange, this country-level split is undefined. Such wallets are removed from the sample, and only a negligible fraction of wallets with at least 50 transactions falls into this category.

The country split uses web traffic under two assumptions. First, web traffic is a reasonable proxy for the relative usage of centralized crypto exchanges across countries within a region. A requirement for this is that the proportions are not systematically distorted through the usage of VPNs, which is the reason we do not rely on web traffic data for China. Remaining VPN-driven distortion is second-order for our construction. Because the region-level totals are pinned down by the on-chain estimates, web traffic only splits an exchange's regional share within a region, so VPN usage cannot move volume across regions through this step. Second, within each region, the average transaction size of each country is the same. This assumption is necessary to extrapolate from web traffic, which is a per-user measure, to transaction volumes. It is a maintained assumption that we cannot test directly.

Relative to purely web-traffic-based estimates of stablecoin flows, the combination of the machine learning models with web traffic relaxes the traffic-related assumptions. Purely traffic-based measures have to assume that web traffic is a reasonable proxy for relative usage across countries in the whole world, that estimates for China are not significantly tainted by VPN usage, and that average transaction size is the same across countries in the whole world.

D.3 Allocation of transfers

For a transfer of value v from self-custodial wallet w to wallet w', the amount assigned to ordered country pair (c, c') is

$$v \cdot P_w(c) \cdot P_{w'}(c'), \tag{7}$$

which is independent of t and b because wallet probabilities do not vary over time or by blockchain. For a transfer from wallet w to centralized exchange x at time t on blockchain b it is

$$v \cdot P_w(c) \cdot P_x(c' \mid t, b), \tag{8}$$

with transfers from an exchange to a wallet handled analogously. For a transfer between centralized exchanges x and x' it is

$$v \cdot P_x(c \mid t, b) \cdot P_{x'}(c' \mid t, b). \tag{9}$$

Summing over all country pairs returns the original value v in every case, since each endpoint's country probabilities sum to one.

E CFM Event Studies

This appendix collects the analyses that interpret and identify the capital-flow-measure results of Section 5.2. Appendix E.1 sets out the accounting by which circumvention registers in gross activity rather than in net flows. Appendix E.2 decomposes the outflow response into subcategories of outflow restrictions. Appendices E.3 and E.4 report the corridor-level and bank-deposit-exposure designs. Appendix E.5 reports the matched-removal tests of the tightening–loosening ratchet. Robustness of the directional paths, pooled loosenings, and an alternative estimator are in Online Appendix F.

E.1 Why circumvention registers in activity rather than in net flows

Circumvention of capital flow restrictions, including capital flight, registers in gross cross-border stablecoin activity rather than in net flows or holdings. This appendix sets out the accounting behind the outcome choice of Section 5.2.

Follow a transfer of wealth abroad through its full round trip. The resident first converts local currency into dollar-denominated tokens, receiving a cross-border stablecoin transfer that our data record as an inflow into the restricting country. Sending the tokens abroad is recorded as an outflow of the same size, and selling them there for foreign currency leaves no further trace on-chain. The wealth has left the country, yet net stablecoin flows are zero. The value moves through the two legs no on-chain measure records, the local-currency payment at home and the dollar sale abroad. A completed transfer therefore raises two-way stablecoin activity, with the inflow and outflow legs moving together, and builds neither net flows nor holdings.

The round trip begins with tokens flowing into the restricting country, and the restrictions themselves generate the incentive to supply them. von Luckner et al. (2024) document that crypto assets in restricted economies trade at premia over the official exchange rate, and that the premia track excess dollar demand and the stringency of controls. At such a premium, a counterparty converting dollars into the restricted currency obtains more local currency by delivering tokens than by transacting through official channels.

The same signature arises without any relocation of wealth. When restrictions make transfers through traditional channels more costly, cross-border transactions of all kinds shift onto stable-

coin rails, and because the migrating activity runs in both directions, the two legs rise together. Rising two-way activity is thus consistent with residents moving wealth abroad and with a broader migration of cross-border payments, and the corridor evidence of Section 5.2 speaks to their relative weight. After an outflow tightening, inflows rise most from the partner countries where the activating country's residents hold deposits abroad, while most of the response is common to all of a country's corridors. Holdings would rise only if residents kept the tokens rather than passing them on, the store-of-value margin that the currency-crisis design of Section 5.3 measures. The near-perfect co-movement of a country's inflows and outflows, and the matched response of the two legs to outflow tightenings reported in Section 5.2, are accordingly the pattern this accounting implies, not evidence that the restrictions left flows unchanged.

E.2 Which outflow restrictions drive the response?

This appendix decomposes the outflow result of Section 5.2 into the eight subcategories of outflow restrictions distinguished by Bergant et al. (2026): (i) restrictions in foreign exchange markets; (ii) restrictions on payments and receipts; (iii) restrictions on resident and nonresident accounts; (iv) restrictions on import payments; (v) export proceeds requirements, i.e., repatriation and surrender requirements on exporters' foreign exchange earnings; (vi) restrictions on invisibles, i.e., on payments for services and on cross-border income and current transfers such as dividend, interest, and profit remittances, travel allowances, and personal transfers; (vii) capital account restrictions, i.e., controls on capital and money market instruments, credit operations, direct investment, and real estate transactions; and (viii) restrictions on the financial sector, i.e., provisions specific to commercial banks and institutional investors.

The events for this exercise are exactly the 87 outflow tightenings underlying column (2) of Table 11: country–months in which the outflow restriction count rises, and neither the inflow nor the neutral count rises. Each event is attributed to the single subcategory of outflow restriction whose count rose in that month; 62 events are attributed to a single subcategory, and the remaining 25 events, in which two or more subcategories were tightened simultaneously, are assigned to a separate "multiple" category. All other features of the design are identical to the baseline: a six-month event window, truncation of each event's window around any nearby CFM activation, of any direction, by the same country (38 events have their pre-window and 32 their post-window truncated on this account), a clean-control pool consisting of countries without any CFM activation within the event window, and event-by-event-time and event-by-country fixed effects. Inference is the one feature that differs, for the reason set out below. The subcategory paths are estimated in a single joint regression analogous to equation (2), with the treated-side event-time dummies interacted with the event's attributed subcategory; pooling the treated-side dummies across subcategories instead reproduces the outflow-direction estimates of column (2) in Table 11.

Inference for this table rests on randomization tests rather than on country-clustered standard errors. The clustering unit is the country, not the event. The pooled outflow design of Section 5.2

is identified from 48 treated countries, where country-clustered asymptotics are unproblematic, but splitting its 87 events across eight subcategories leaves between three and 15 treated countries per column. At those counts country-clustered standard errors understate the dispersion of the estimator against the randomization distribution described next, by a median factor of 2.3 across the nine columns and by as much as 5.2, and the significance they imply is spurious.

The protocol follows the currency-crisis design of Section 5.3. In each of 1,000 draws, the treated country's observations are dropped from every stack and one country from that stack's clean-control pool is assigned placebo-treated status at the event's actual anchor month, keeping the event's attributed subcategory. The joint regression is refit on each draw, and every coefficient is referred to its own placebo distribution. A country with several events keeps a single placebo country across all of its stacks, and placebo countries are drawn without replacement within a draw, so each draw reproduces the design's distinct-treated-country counts. Randomization tests are exact under the sharp null whatever those counts are (MacKinnon and Webb, 2020).

A small randomization p does not establish that a column rests on more than one country, which at these counts is a live possibility. Table 7 therefore reports each subcategory's post-activation average alongside the value that average takes under the most adverse single-country drop.

Capital account restrictions are the only subcategory whose response survives both tests. Capital account tightenings are followed by an increase in stablecoin inflows of approximately 25 percent already in the activation month, building to approximately 37 to 40 percent five to six months after activation. The average effect over the six months following activation is 0.241 log points, with a randomization p of 0.057, and no one of the eleven treated countries carries it, since the most adverse drop still leaves 0.181. Coefficients at the earliest pre-event horizons are mildly negative, but the randomization distribution places them well inside ordinary variation, at $p = 0.27$ and $p = 0.34$ for $k = -6$ and $k = -5$, so the path gives no evidence of anticipation.

Financial sector tightenings and simultaneous tightenings of several subcategories do not survive. Both build positive point estimates over the post-activation window, reaching approximately 15 and 23 percent at six months, and country-clustered standard errors would make several of their coefficients significant at the 1 percent level. Referred to their own placebo distributions, none is: the average effect over the six months following activation carries a randomization p of 0.379 for financial sector restrictions and 0.221 for multiple tightenings. The financial sector column is identified from 15 countries, more than any other single subcategory, and its country-clustered standard errors are still too small by a factor of 2.5.

The remaining six subcategories rest on three to eight activations each, and the invisibles column shows why their point estimates cannot be read on their own. Tightenings of restrictions on invisibles carry the largest coefficients in the table and a randomization p of 0.019 for the average effect over the six months following activation. The category contains three activations in three countries, and dropping one of them takes that average from 1.079 to 0.154 log points. That country is subject to sanctions, which Section 5.2 treats as a distinct confound, so we do not read

the column as evidence on circumvention.

None of the other five columns reaches significance at any horizon, and we report their point estimates as suggestive of natural circumvention margins. Export proceeds requirements (eight activations) carry sizeable positive point estimates on the order of 30 to 45 percent at horizons of three to six months, a natural margin for stablecoin-based circumvention, since export proceeds subject to repatriation or surrender requirements could be received or held offshore in stablecoins. Restrictions on resident and nonresident accounts (five activations) show large positive point estimates at the five- and six-month horizons. Restrictions on import payments are followed by small positive point estimates, and restrictions in foreign exchange markets and on payments and receipts show no systematic pattern.

E.3 A Corridor-Level Comparison

The baseline design compares a country's total stablecoin inflows with inflows into countries without any CFM activation in the event window. The corridor dimension of the dataset supports a sharper comparison. For each of the 217 activations, we construct, for every origin country, the difference between the log of flows from that origin into the activating country and the average log of the same origin's flows into clean destinations in the same month. Clean destinations are countries covered by the CFM data with no activation of any direction anywhere in the event window, excluding the activating country and the issuer treasuries; the median event has 142 of them, observed from 187 origins. Events, event windows, the truncation of overlapping own-windows, and the directional attribution are those of the baseline design. We regress the within-origin difference on event-time dummies with event-by-origin fixed effects, and cluster standard errors on the activating country.[41] Because both terms of the difference share the same origin and month, any shock common to the origin or to the global crypto cycle drops out of the outcome exactly, rather than on average through the control group. The identifying comparison is the activating country against its peer destinations, within each origin and month.

Figure 11 plots the corridor-level event-time paths and Table 8 reports the coefficients. The estimates are close to the baseline. The average post-activation effect is 0.090 (s.e. = 0.024) across all activations, and 0.138 (s.e. = 0.040) for outflow tightenings, against 0.152 (s.e. = 0.048) in the baseline design.[42] The pre-activation coefficients for outflow tightenings average −0.002, and none is distinguishable from zero. The flat pre-trends of the baseline design are therefore not an artifact

[41] Every event's control set draws on the same clean destinations, and event windows overlap in calendar time, so residuals may correlate across event-country clusters in a way country-level clustering does not capture. A randomization placebo that reassigns each country's activation dates at random within the sample window, holding the design fixed, produces no statistic as large as the actual response in 500 draws, for either the pooled or the outflow path ($p < 0.002$).

[42] The within-origin difference weights origins equally. Estimating instead on the corridor-level micro panel, with event-by-origin-by-event-time and event-by-corridor fixed effects and fifteen control destinations per event, gives post-activation averages of 0.085 (s.e. = 0.025) across all activations and 0.127 (s.e. = 0.040) for outflow tightenings.

Table 7: Stablecoin Inflows around Outflow CFM Tightenings, by Subcategory of Restriction

	(1) FX markets	(2) Paym. & receipts	(3) Resid. accounts	(4) Import payments	(5) Export proceeds	(6) Invisi-bles	(7) Capital account	(8) Financial sector	(9) Multiple
$k = -6$	−0.016	0.182	—	0.252	0.108	−0.182	−0.183	−0.023	0.199
	[0.976]	[0.385]		[0.551]	[0.645]	[0.488]	[0.269]	[0.869]	[0.283]
$k = -5$	−0.149	0.059	−0.087	0.286	0.041	−0.108	−0.147	−0.048	−0.008
	[0.641]	[0.722]	[0.800]	[0.457]	[0.854]	[0.661]	[0.340]	[0.752]	[0.962]
$k = -4$	−0.145	0.175	0.245	0.066	0.224	−0.045	−0.092	−0.028	0.072
	[0.670]	[0.340]	[0.449]	[0.873]	[0.348]	[0.828]	[0.469]	[0.834]	[0.603]
$k = -3$	−0.183	0.081	0.110	−0.065	−0.154	−0.067	−0.092	0.016	0.036
	[0.537]	[0.582]	[0.664]	[0.806]	[0.408]	[0.717]	[0.392]	[0.909]	[0.700]
$k = -2$	0.328	0.035	0.097	0.043	−0.087	0.026	0.017	−0.036	0.011
	[0.168]	[0.765]	[0.619]	[0.730]	[0.530]	[0.863]	[0.845]	[0.728]	[0.890]
$k = 0$	−0.147	−0.134	0.088	0.052	0.025	0.407*	0.226*	0.033	0.010
	[0.413]	[0.388]	[0.587]	[0.639]	[0.825]	[0.064]	[0.050]	[0.720]	[0.878]
$k = 1$	−0.055	−0.215	−0.019	0.047	0.220	0.618**	0.211*	0.062	0.024
	[0.786]	[0.277]	[0.934]	[0.706]	[0.238]	[0.037]	[0.088]	[0.595]	[0.800]
$k = 2$	0.226	−0.170	0.078	−0.011	0.146	0.838**	0.141	0.093	0.112
	[0.340]	[0.414]	[0.755]	[0.949]	[0.411]	[0.019]	[0.290]	[0.458]	[0.299]
$k = 3$	0.116	−0.041	0.080	−0.006	0.279	0.813	0.214	0.144	0.145
	[0.689]	[0.878]	[0.816]	[0.966]	[0.207]	[0.155]	[0.153]	[0.318]	[0.213]
$k = 4$	0.323	−0.385	0.278	0.040	0.372	0.687	0.244	0.178	0.159
	[0.301]	[0.279]	[0.535]	[0.880]	[0.102]	[0.311]	[0.104]	[0.284]	[0.196]
$k = 5$	0.442	−0.192	0.343	−0.007	0.332	1.398**	0.317*	0.118	0.116
	[0.227]	[0.543]	[0.416]	[0.961]	[0.163]	[0.037]	[0.050]	[0.492]	[0.361]
$k = 6$	0.035	−0.210	0.502	0.110	0.316	2.789***	0.335*	0.143	0.210
	[0.904]	[0.539]	[0.273]	[0.712]	[0.225]	[0.000]	[0.059]	[0.449]	[0.167]
Post-activation average	0.134	−0.193	0.193	0.032	0.242	1.079**	0.241*	0.110	0.111
	[0.574]	[0.409]	[0.484]	[0.840]	[0.171]	[0.019]	[0.057]	[0.379]	[0.221]
dropping most influential country	−0.052	−0.058	0.021	−0.005	0.074	0.154	0.181	0.081	0.090
Events per category	6	4	5	7	8	3	12	17	25
Treated countries	6	4	4	4	7	3	11	15	14
Event × event-time FE					Yes				
Event × country FE					Yes				
Events					87				
Observations					113,537				
R^2					0.991				

Notes: The dependent variable is the log of stablecoin inflows. Events are the 87 outflow tightenings underlying column (2) of Table 11, i.e., country–months in which the outflow restriction count rises and neither the inflow nor the neutral count rises; each event is attributed to the single subcategory of outflow restriction whose count rose in that month (column (9) collects events in which several subcategories tightened simultaneously). All columns come from a single joint regression in which the treated-side event-time dummies are interacted with the event's attributed subcategory, so that all subcategories share the same fixed effects and the same pool of clean controls. Coefficients are event-time effects relative to the month immediately preceding the activation ($k = -1$, omitted), with event time binned at $k = \pm 6$. The $k = -6$ coefficient for column (3) is not identified (no treated observations remain at that horizon after window truncation). The post-activation average is the mean coefficient over $k = 0$ to $k = 6$. Each subcategory is identified from between three and 15 treated countries, too few for country-clustered standard errors to be reliable, so inference is by randomization test. In each of 1,000 draws the treated country's observations are dropped from every stack and one clean control is assigned placebo-treated status at that event's anchor month, keeping the event's attributed subcategory; a country with several events keeps one placebo country across all of them, so each draw reproduces the treated-country counts reported above. Randomization p-values, two-sided, in brackets, and stars are computed from them. A small randomization p does not establish that a column rests on more than one country, so the row beneath the post-activation average reports the most adverse single-country drop: each treated country's events are removed from its subcategory in turn and the joint regression refit, and the reported figure is the resulting post-activation average that is smallest in absolute size. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

of the composition of the control group: before the restriction is imposed, inflows into the activating country were not drifting upward relative to its peer destinations from the same origins.

Three limits of this comparison are worth recording. First, it absorbs origin-side and global conditions, not conditions specific to the activating country in the activation month itself; those are addressed by the bank-deposit exposure design of the next subsection and the exclusion exercises in Online Appendix F.1. Second, the directional contrast does not sharpen: the averaged outflow–inflow difference is 0.058 (s.e. = 0.062, $p = 0.35$), in line with the equality tests in Online Appendix F.1, so the additional granularity leaves the power limits of the 87-against-43-event comparison in place. Third, corridor-level flows inherit the attribution error of the estimation methodology at a finer granularity than the country aggregates. The corridor comparison accordingly serves as a check on the baseline design rather than as a replacement for it.

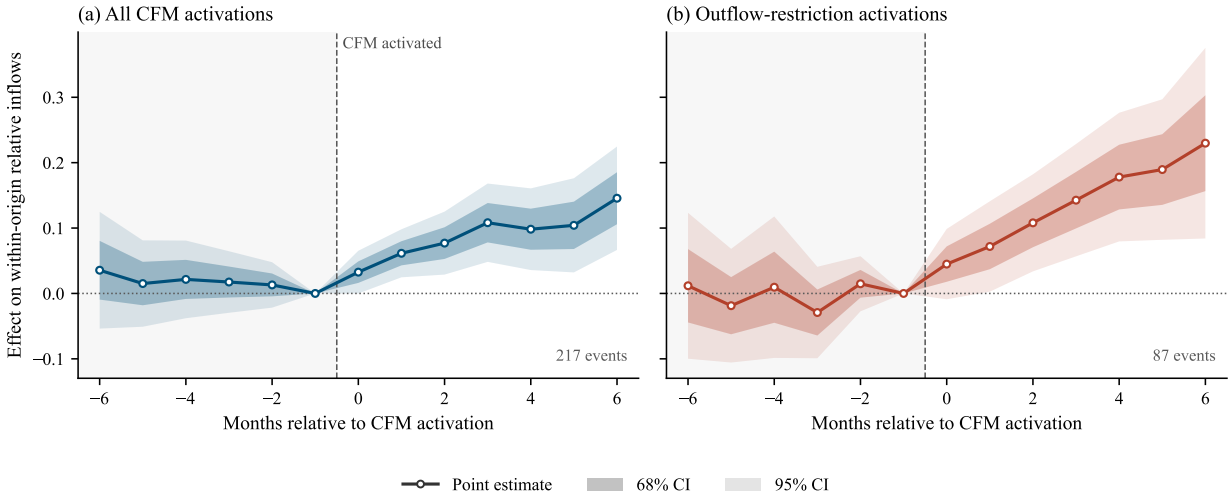


Figure 11: Stablecoin Inflows around CFM Activations: Corridor-Level Event Studies

Notes: The outcome is the difference, within each origin country and month, between the log of stablecoin flows from that origin into the activating country and the average log of the same origin's flows into clean destinations, defined as countries covered by the CFM data with no activation of any direction within the event window, excluding the activating country and the issuer treasuries. Panel (a) pools all 217 activations; panel (b) traces the outflow-restriction path from the joint directional regression. Coefficients are relative to the month immediately preceding the activation ($k = -1$, dashed vertical line), from a stacked specification with event-by-origin fixed effects. Shaded areas show 68 and 95 percent confidence intervals based on standard errors clustered on the activating country.

E.4 Bank-Deposit Exposure

The corridor comparison above holds the origin fixed and still measures the activating country against other destinations, so a shock specific to the activating country in the activation month remains outside its reach. The design of this subsection holds the activating country fixed instead. It asks whether, within the activating country and month, the rise in stablecoin inflows concentrates

Table 8: Stablecoin Inflows around CFM Activations: Corridor-Level Estimates

	(1) All CFMs	(2) Outflow	(3) Inflow	(4) Neutral	(5) Multiple
$k = -6$	0.036	0.012	0.114	0.039	−0.057
	(0.045)	(0.056)	(0.148)	(0.062)	(0.083)
$k = -5$	0.015	−0.019	0.028	0.078	−0.108
	(0.033)	(0.044)	(0.066)	(0.074)	(0.104)
$k = -4$	0.021	0.010	0.010	0.068	−0.066
	(0.030)	(0.054)	(0.042)	(0.055)	(0.104)
$k = -3$	0.017	−0.029	−0.002	0.069	0.075
	(0.024)	(0.035)	(0.036)	(0.054)	(0.065)
$k = -2$	0.013	0.015	0.026	0.035	−0.073*
	(0.017)	(0.021)	(0.035)	(0.038)	(0.043)
$k = 0$	0.033*	0.045	0.062*	0.027	−0.042
	(0.016)	(0.027)	(0.035)	(0.027)	(0.047)
$k = 1$	0.061***	0.072**	0.069	0.057*	0.030
	(0.018)	(0.035)	(0.045)	(0.033)	(0.035)
$k = 2$	0.077***	0.108***	0.076*	0.052	0.034
	(0.024)	(0.037)	(0.041)	(0.047)	(0.102)
$k = 3$	0.108***	0.143***	0.083	0.106	0.035
	(0.030)	(0.043)	(0.052)	(0.064)	(0.091)
$k = 4$	0.098***	0.178***	0.091	0.044	−0.037
	(0.031)	(0.049)	(0.056)	(0.058)	(0.077)
$k = 5$	0.104***	0.189***	0.067	0.062	−0.071
	(0.036)	(0.054)	(0.080)	(0.061)	(0.070)
$k = 6$	0.146***	0.230***	0.110	0.107*	−0.042
	(0.040)	(0.073)	(0.089)	(0.064)	(0.060)
Post-activation avg. ($k = 0 \dots 6$)	0.090***	0.138***	0.080	0.065	−0.013
	(0.024)	(0.040)	(0.048)	(0.044)	(0.056)
Event × origin FE			Yes		
Events			217		
Observations			386,040		
R^2			0.983		

on the corridors that carry the country's residents' offshore bank deposits.

We measure a corridor's exposure as the share of the activating country's residents' bank deposits abroad that is held in the partner country. From the BIS locational banking statistics, we take the cross-border liabilities of banks in the partner country to the activating country's non-bank sector, read them in the last quarter before the activation month, and divide by the sum across all reporting partners. Exposure is standardized within each event, so a coefficient reads as the effect of a one-standard-deviation higher deposit share among that country's own partners. Events, event windows, the truncation of overlapping own-windows, and the directional attribution are those of the baseline design. A corridor enters if it is observed in every month of the window. An event enters if at least five of its corridors carry measurable exposure. These screens keep 154 of the 217 events, among them 62 of the 87 outflow tightenings, across 82 countries. The median event has 21 measurable partners, and the measurable corridors carry roughly 37 percent of the median event country's cross-border volume. We regress log corridor inflows on the interactions of exposure with event time, and we cluster standard errors on the activating country. The regression also includes interactions of distance, contiguity, and common language with event time. Exposure is correlated with proximity within an event, so without them any tendency of nearby corridors to move differently around an activation would load onto the exposure coefficient. There are two sets of fixed effects. The first is an effect for each corridor, an event-by-partner pair, which absorbs the corridor's average level over the window. The second is an effect for each event, event month, and within-event corridor-size quintile, so the exposure coefficient compares corridors of the same country, in the same month, in the same size class. Absorbing corridor size this way matters because size works against the gradient: large corridors mean-revert within the window, and exposure correlates positively with size, so a specification without the size interactions biases the exposure coefficient toward zero. The post-activation average for outflow tightenings is 0.004 with event-by-event-time fixed effects alone and 0.012 once the size quintiles are added, which is the specification we report.

Figure 12 plots the interaction paths and Table 9 reports the estimates. After an outflow tightening, the exposure gradient builds steadily to 0.025 six months after activation. The post-activation average is 0.012 (s.e. = 0.007, country-clustered $p = 0.08$), and the pre-activation average is -0.003 (s.e. = 0.005). A permutation test supports this inference. Reassigning the exposure shares at random across each country's own partners, holding the design fixed, produces a post-activation average as large as the actual one in 1.7 percent of 1,000 draws. The gradient is specific in three ways. It appears only after outflow tightenings. After inflow and neutral tightenings the post-activation averages are -0.002 and -0.011, neither distinguishable from zero.[43] It appears only for residents' deposits abroad. Measuring exposure instead by bank lending into the country, the cross-border claims of partner-country banks on the activating country's non-banks, yields a

[43]The inflow column's pre-activation coefficients are themselves negative and significant at $k = -5$ and $k = -3$, so its post-activation zero should be read as uninformative rather than as a precisely estimated null. The outflow column's own pre-activation path is flat, which is the comparison the gradient rests on.

post-activation average of 0.007 (s.e. = 0.006), half the deposit-side gradient and not distinguishable from zero. And it does not depend on exposure measured close to the activation. Freezing every event's exposure at the end of 2020, before the sample's activations, gives a post-activation average of 0.014 (country-clustered $p = 0.03$), so the gradient is not driven by deposits repositioned in response to earlier restrictions.

A shock that coincides with the activation, a currency crisis, a sanction, or the alarm the announcement itself causes, raises a country's stablecoin activity on all of its corridors at once. It has no mechanism for concentrating the increase on the corridors carrying residents' offshore deposits, timed to the activation month, only after outflow tightenings, and only for the deposit direction of the banking relationship. The gradient is therefore evidence that the response to the restriction runs along residents' pre-existing financial relationships, which no country-level confound can mimic.

Three limits are worth recording. First, the gradient is estimated with limited precision. It is significant at the 5 percent level by the permutation test and at the 10 percent level by the country-clustered t-test. Second, it is a statement about composition, not magnitude. One standard deviation of exposure moves the six-month response by approximately 2.5 percent against the roughly 30 percent average, so most of the response is common to a country's corridors, and the gradient identifies the channel rather than accounting for the bulk of the effect. Third, the gradient appears on the inflow leg, where residents acquire stablecoins, and not on the sending leg. The post-activation average for flows from the activating country toward its partners is -0.002 (s.e. = 0.005). The pattern indicates that acquisition runs through the jurisdictions holding residents' deposits; where the funds travel afterwards is not something this design measures.

E.5 The Tightening–Loosening Ratchet

This appendix collects the tests behind the ratchet result of the main text: the matched-removal design, the through-the-cycle shift, and the formal asymmetry tests. Table 10 summarizes.

Loosenings matched to prior tightenings. The loosening events in Online Appendix F.2 are, as a whole, a different population of country–months from the tightenings that drive the main result, so a flat loosening path could in principle coexist with a tightening effect that decays. The subsample where a reversal would actually be expected is the set of loosenings that *follow* an in-sample tightening by the same country. We therefore re-estimate the loosening design restricted to matched loosenings: events in which the loosening country had at least one tightening event earlier in the sample. The design is otherwise unchanged: all loosenings by the same country continue to truncate windows, and the clean-control pool continues to exclude any country with a loosening inside the window; only the set of stacked (treated) events is restricted. Note that for matched events the reference month $k = -1$ sits on the post-tightening level of inflows, so a transitory tightening effect would appear here as *negative* post-removal coefficients.

Table 9: Bank-Deposit Exposure and the Corridor Composition of the CFM Response

Exposure Tightening	(1) Deposits Outflow	(2) Deposits Inflow	(3) Deposits Neutral	(4) Lending Outflow
$k = -6$	−0.001	−0.015	−0.012*	0.006
	(0.008)	(0.010)	(0.006)	(0.008)
$k = -5$	−0.003	−0.022**	−0.006	0.007
	(0.006)	(0.009)	(0.008)	(0.008)
$k = -4$	−0.008	−0.014*	−0.009	0.014*
	(0.007)	(0.008)	(0.008)	(0.007)
$k = -3$	−0.001	−0.018***	−0.004	0.008
	(0.006)	(0.006)	(0.006)	(0.006)
$k = -2$	−0.005	0.003	−0.007	−0.000
	(0.006)	(0.006)	(0.006)	(0.004)
$k = 0$	0.004	−0.003	−0.008	0.000
	(0.004)	(0.003)	(0.007)	(0.005)
$k = 1$	0.008	−0.000	−0.011	0.008
	(0.006)	(0.004)	(0.010)	(0.005)
$k = 2$	0.015	−0.001	−0.008	0.005
	(0.010)	(0.006)	(0.010)	(0.006)
$k = 3$	0.006	0.000	−0.008	0.008
	(0.005)	(0.008)	(0.010)	(0.007)
$k = 4$	0.010	−0.001	−0.014	0.004
	(0.006)	(0.008)	(0.010)	(0.007)
$k = 5$	0.019*	−0.007	−0.016	0.011
	(0.010)	(0.008)	(0.013)	(0.008)
$k = 6$	0.025	−0.004	−0.013	0.010
	(0.016)	(0.006)	(0.013)	(0.008)
Post-activation avg. ($k = 0 \dots 6$)	0.012*	−0.002	−0.011	0.007
	(0.007)	(0.005)	(0.010)	(0.006)
Pre-activation avg. ($k = -6 \cdots -2$)	−0.003	−0.013**	−0.008	0.007
	(0.005)	(0.006)	(0.005)	(0.006)
Events	62	31	47	60
Observations	13,352	7,196	11,021	10,098
Event countries	46	23	36	44

Notes: Each column is a separate stacked event-study regression of log corridor inflows into the activating country. The reported coefficients are the interactions of event time with exposure, the corridor's share of the activating country's residents' bank deposits abroad (columns (1)–(3)) or of partner-country bank lending into the activating country (column (4)), from the BIS locational banking statistics in the last quarter before the activation, standardized within event. All specifications include event-by-event-time-by-size-quintile and event-by-partner fixed effects and distance, contiguity, and common-language interactions with event time. Standard errors, clustered on the activating country, in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% level by the country-clustered t-test; the permutation p-values quoted in the text are computed separately.

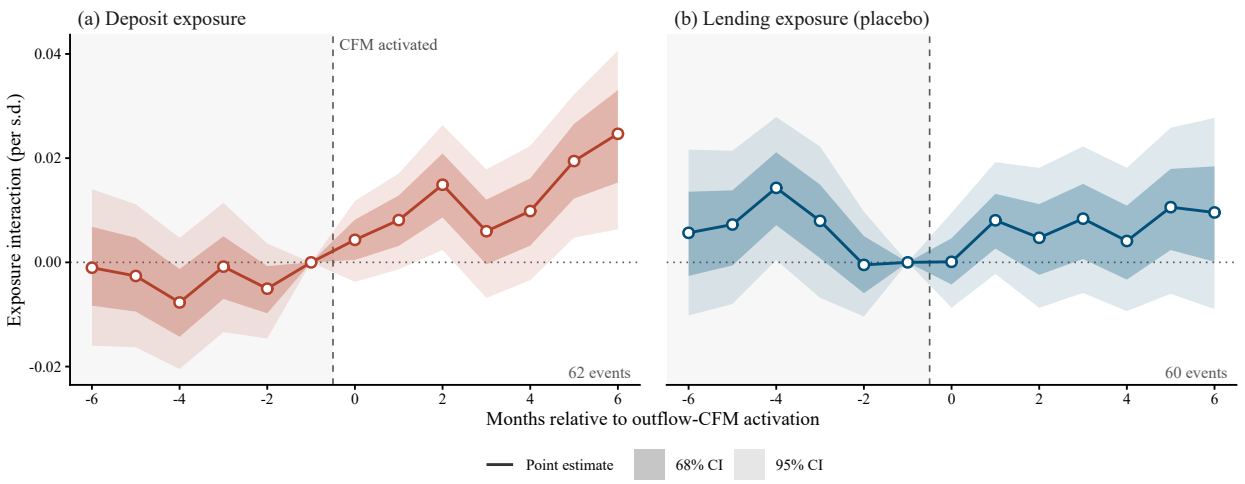


Figure 12: Bank-Deposit Exposure and the Corridor Composition of the CFM Response

Notes: The figure plots the exposure-by-event-time interaction coefficients around outflow-CFM activations. The outcome is log stablecoin inflows at the corridor level. Exposure is the corridor's share of the activating country's residents' bank deposits abroad, from the BIS locational banking statistics in the last quarter before the activation, standardized within event. Panel (a) measures exposure by deposits held in the partner country; panel (b) measures it instead by partner-country bank lending into the activating country, which should carry no gradient under the circumvention reading. Coefficients are relative to the month immediately preceding the activation ($k = -1$, dashed vertical line), from stacked specifications with event-by-event-time-by-size-quintile and event-by-partner fixed effects and gravity-by-event-time controls. Shaded areas show 68 and 95 percent randomization confidence intervals, constructed by test inversion against 1,000 permutations of exposure across each country's own partners, with each panel permuted separately. We use randomization inference here rather than the country-clustered standard errors reported in Table 9 because this design has only 46 clusters and estimates 48 interaction coefficients, so cluster-robust asymptotics are unreliable at this scale.

Of the 351 loosenings, 187 follow at least one in-sample tightening by the same country (81 outflow, 33 inflow, 30 neutral, and 43 multi-direction loosenings), and 110 follow a prior tightening of the *same* direction, including 75 of the 152 outflow loosenings. The median matched loosening occurs six months after the most recent prior tightening (interquartile range 2 to 11 months), so the typical matched removal falls squarely within the horizon over which the tightening response is at or near its peak. The post-removal path for outflow loosenings is flat in both matched subsamples (Figure 13), with coefficients between -0.02 and $+0.04$ at every post-removal horizon and average post-removal effects of 0.005 (s.e. = 0.046) for the any-direction match and 0.011 (s.e. = 0.050) for the same-direction match. The negative drift at the five- and six-month horizons in the full sample of outflow loosenings (-0.141 and -0.150, marginally significant) is, in point estimates, concentrated entirely in the unmatched removals: estimating the matched and unmatched outflow-loosening paths jointly yields average post-removal effects of $+0.005$ (s.e. = 0.046) and -0.097 (s.e. = 0.077), respectively, though the difference between the two is imprecisely estimated ($p = 0.25$). Whatever the unmatched drift reflects, however, it cannot be a reversal of the tightening response: unmatched removals are liberalizations of a stock of controls predating our sample, with

no documented stablecoin build-up to reverse. The matched pre-event coefficients remain mildly negative at the earliest horizons (e.g., -0.167, s.e. $= 0.093$, at $k = -6$ for the same-direction match), i.e., inflows were still rising in the months leading up to the removal, which for matched events is consistent with the post-tightening build-up still being under way, so the caution about the reference month applies here as well; the reversal question, however, concerns the post-removal path, which is flat.

The ratchet through the cycle. The matched-removal design measures changes relative to the removal month. The stronger ratchet statement concerns the *level* relative to the pre-tightening baseline: once a restriction has been imposed and then removed, do inflows remain on the elevated post-tightening path, or do they fall back? Because loosenings do not truncate the tightening windows (only repeat tightenings do), the twelve-month tightening design in Online Appendix F.1 already contains the post-removal months for the events whose restriction is removed within the window. We therefore add to that specification an indicator for treated months on or after the treated country's first subsequent same-direction loosening. Thirty of the outflow tightening events see such a removal within their truncation-surviving window, with the median event's first post-removal month at horizon $k = 4$. The estimated post-removal shift is $+0.048$ (s.e. $= 0.084$): inflows show no movement off the elevated post-tightening path once the restriction is gone. The event-time path with this shift absorbed is essentially identical to the baseline twelve-month path (e.g., 0.225, s.e. $= 0.099$, at $k = 12$, against 0.246 in the baseline), so the twelve-month persistence is not an artifact of already-removed events carrying the late horizons.

Formal asymmetry tests. Because the tightening and matched-loosening designs are estimated on separate stacked datasets, comparing their averages requires the covariance between them. We therefore stack both designs into a single joint regression, in which each event retains its own event-by-event-time and event-by-country fixed effects and clustering at the country level carries the covariance between the two sets of estimates, and test linear restrictions across them, in the spirit of the entry–exit asymmetry tests of Section 5.3. The average post-tightening effect for the outflow direction is 0.152 (s.e. $= 0.048$, $p = 0.002$); the average post-removal effect for same-direction matched removals is 0.011 (s.e. $= 0.050$, $p = 0.83$). Under full reversal, the two would sum to zero; the data reject this hypothesis ($t = 2.30$, $p = 0.022$). Expressed as the share of the tightening effect reversed after removal, $-\bar{\beta}^L/\bar{\beta}^T$, the point estimate is -0.07 with a delta-method 95 percent confidence interval of $[-0.71, 0.57]$; the estimates are centered on a pure ratchet and reject full reversal, though a partial give-back of up to roughly half cannot be excluded.

Table 10: Ratchet Tests on the Tightening–Loosening Path

	Estimate	s.e.	p-value
Avg. post-tightening effect, $k = 0, \ldots, 6$	0.152	(0.048)	0.002
Avg. post-removal effect, matched removals	0.011	(0.050)	0.832
Through-the-cycle shift after removal	0.048	(0.084)	0.567
Full-reversal test, H_0: $\bar{\beta}^T + \bar{\beta}^L = 0$	0.163	(0.071)	0.022
Share of tightening effect reversed	−0.07	(0.33)	
95% confidence interval		[−0.71, 0.57]	

Notes: All statistics concern the outflow direction. The first two rows report the average post-event treatment effect over horizons $k = 0, \ldots, 6$ from a single joint regression stacking the tightening design of Section 5.2 and the loosening design restricted to removals preceded by an in-sample tightening of the same direction by the same country (75 events); country-level clustering in the joint regression carries the covariance between the two sets of estimates. The through-the-cycle shift is estimated within the twelve-month tightening design of Online Appendix F.1: it is the coefficient on an indicator for treated months after the treated country's first subsequent same-direction loosening (30 events), so that a negative value would indicate inflows falling back toward the pre-tightening baseline once the restriction is removed. The full-reversal test refers the sum of the two average effects ($\bar{\beta}^T + \bar{\beta}^L$, zero under full reversal) to a $t(G-1)$ distribution, where G is the number of country clusters. The share reversed is $-\bar{\beta}^L/\bar{\beta}^T$, with a delta-method standard error based on the joint country-clustered covariance matrix; a value of 0 corresponds to a pure ratchet and 1 to full reversal. Standard errors clustered at the country level.

F Currency Crises: Classification and Design

Crisis classification. A country-month is classified as in currency crisis when cumulative twelve-month log depreciation against the dollar is at least log(1.30) and exceeds the corresponding twelve-month rate one year earlier by at least 10 percentage points. This is the Laeven and Valencia (2020) definition, implemented on rolling twelve-month windows of our monthly exchange-rate panel. The Frankel and Rose (1996) variant applies the identical rule at log(1.25). Both are published as annual rules; a crisis is a period meeting the criterion, and because the twelve-month depreciation is a rolling sum, an episode is persistent by construction without any appended post-crash window. Classification uses each country's full exchange-rate history, before any restriction to the estimation sample, so that rolling statistics are computed on complete windows. The acceleration condition is what makes these definitions usable on 2018–2025 data: rules without it (depreciation-only thresholds of 15 or 25 percent) flag much of the advanced world in 2022, when the dollar's broad appreciation moved most currencies at once, and would leave few clean controls; the acceleration clause separates a country-specific collapse from that common move.

Real-time anchoring. Each event is a month that meets the twelve-month criterion. We date its onset as the first month, in the year before that criterion date, in which cumulative depreciation over a trailing three-month window exceeds $\log(1.30) \times 3/12$. The three-month window only moves the date. A three-month burst that is not followed by the twelve-month criterion is not an event.

Table 11: Stablecoin Inflows around CFM Activations: Stacked Event-Study Estimates

	(1) All CFMs	(2) Outflow	(3) Inflow	(4) Neutral	(5) Multiple
$k = -6$	0.030	0.015	0.122	0.016	−0.062
	(0.047)	(0.058)	(0.163)	(0.061)	(0.072)
$k = -5$	0.005	−0.041	0.013	0.091	−0.134
	(0.035)	(0.045)	(0.065)	(0.077)	(0.106)
$k = -4$	0.032	0.032	0.004	0.085	−0.087
	(0.032)	(0.056)	(0.043)	(0.060)	(0.106)
$k = -3$	0.023	−0.033	−0.012	0.095	0.069
	(0.026)	(0.037)	(0.035)	(0.062)	(0.067)
$k = -2$	0.015	0.020	0.020	0.043	−0.082*
	(0.017)	(0.021)	(0.032)	(0.037)	(0.050)
$k = 0$	0.038**	0.051*	0.066*	0.041	−0.054
	(0.019)	(0.030)	(0.036)	(0.031)	(0.063)
$k = 1$	0.067***	0.073*	0.073	0.069*	0.041
	(0.023)	(0.040)	(0.049)	(0.037)	(0.051)
$k = 2$	0.088***	0.117**	0.082*	0.063	0.070
	(0.029)	(0.046)	(0.044)	(0.053)	(0.128)
$k = 3$	0.131***	0.158***	0.117*	0.121*	0.100
	(0.036)	(0.051)	(0.064)	(0.070)	(0.119)
$k = 4$	0.118***	0.196***	0.117*	0.051	0.023
	(0.040)	(0.060)	(0.070)	(0.064)	(0.108)
$k = 5$	0.134***	0.213***	0.117	0.075	0.015
	(0.043)	(0.063)	(0.091)	(0.061)	(0.095)
$k = 6$	0.186***	0.260***	0.150	0.152*	0.030
	(0.051)	(0.089)	(0.096)	(0.080)	(0.086)
Event × event-time FE			Yes		
Event × country FE			Yes		
Events			217		
Observations			283,294		
R^2			0.990		

Notes: The dependent variable is the log of stablecoin inflows. Coefficients are event-time effects relative to the month immediately preceding the CFM activation ($k = -1$, omitted), with event time binned at $k = \pm 6$. Column (1) reports the aggregate stacked event-study estimates pooling all CFM activations. Columns (2)–(5) come from a single joint regression in which the treated-side event-time dummies are interacted with the direction of the restriction tightened at each event: outflow, inflow, directionally neutral, or multiple directions simultaneously. The outflow, inflow, neutral, and multiple categories comprise 87, 43, 60, and 27 events across 48, 26, 41, and 16 distinct treated countries, respectively; Wald tests of the equality of the directional paths are reported in Table 17. All specifications include event-by-event-time and event-by-country fixed effects. Standard errors, clustered at the country level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

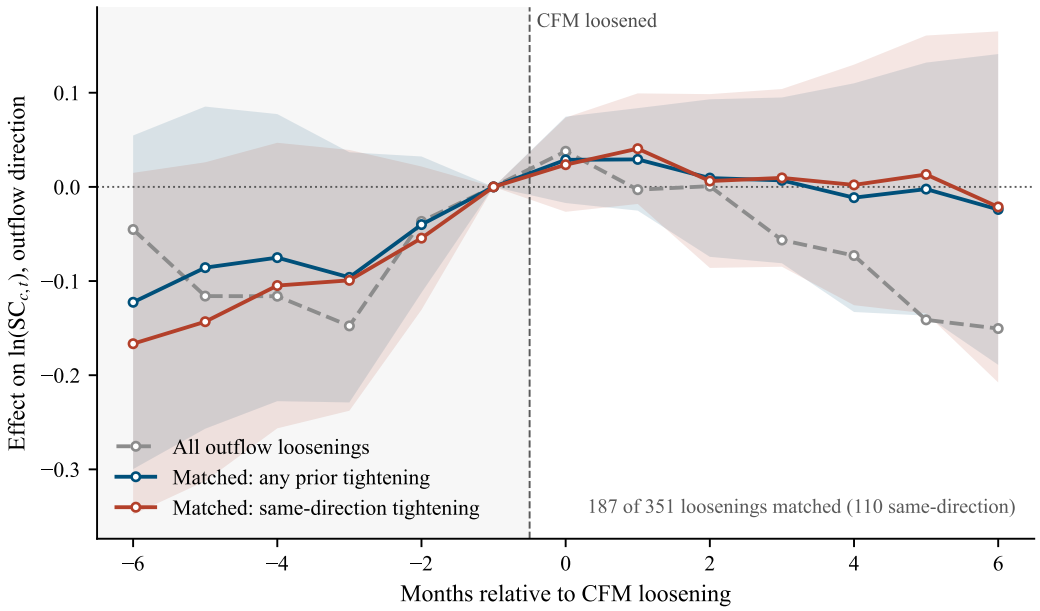


Figure 13: Loosenings matched to prior in-sample tightenings

Notes: The figure plots the outflow-direction event-time path from the loosening design estimated on the full sample of loosenings, on the subsample of loosenings preceded by an in-sample tightening by the same country, and on the subsample preceded by an in-sample tightening of the same direction. All loosenings by the same country truncate windows and contaminate the clean-control pool in every specification; only the set of stacked events differs. Coefficients are relative to the month immediately preceding the loosening ($k = -1$, dashed vertical line). Shaded areas show 95 percent confidence intervals based on standard errors clustered at the country level.

The twelve-month rule lags the start of the slide, and anchoring at the month it fires would push part of any response into the pre-period. The anchor may never reach back into a previous episode: it is floored at the month after the country's prior crisis spell ends, so that a recurring crisis cannot be dated inside its predecessor. Online Appendix G varies the detection window from two to six months.

Entry rule and estimable events. Every transition from out-of-crisis to in-crisis is a candidate event, and a country's first *estimable* entry is used: the event window must begin inside the regression sample (January 2021 onward, Section 3.1). Where a country's earlier episode overlaps the six-month pre-window, the window is truncated at the earlier episode's end, as in the CFM design's treatment of overlapping windows, and the event is retained only if at least three pre-months survive; the truncation applies to the whole stack, controls included, so each stack remains balanced in calendar time. Taking the first estimable entry rather than the first entry ever matters for recurring crises: Türkiye's June 2023 entry, the event in Table 12, follows a 2020–21 spell whose own entry predates the estimable sample, and a first-entry-ever rule would discard the country entirely. The mirror image of the rule is that recurrences drop out: a country contributes at most one event, and any later spell either falls inside the first event's (extended) post-window or fails the clean-pre

screen, so every estimate in the paper is a first-estimable-entry response, not a first-crisis-ever response. Countries in continuous crisis from before the estimable window (their onsets predate it) cannot appear as events under any rule; they remain excluded from every control pool through the ever-in-crisis screen. Table 12 lists the resulting 22 Laeven–Valencia and 25 Frankel–Rose events; the estimation baseline excludes Suriname, as described in the main text and in the influence block below. The two sets differ in dating as well as membership: the looser threshold fires earlier for eight countries, by one to three months for most, and by fifteen months for Malawi, whose 2022 slide meets only the 25 percent criterion while its 2023 acceleration meets both.

Why entry only. No exit-anchored design is estimated, for a structural reason: exit rules identify the end of *dislocation*, not of distress. Rules that release an episode when the exchange rate stabilizes read a currency that has collapsed and is now falling smoothly as recovered; in our sample period, such mechanical reclassifications routinely declare episodes over while twelve-month inflation remains between 25 and 32 percent, in one case while the currency is still depreciating by more than 40 percent year on year. An event study anchored on such dates would estimate the effect of a currency settling into a new trend, not of a recovery. Onset dates have no analogous problem, because the crisis definition itself supplies them. All designs are therefore anchored on entry, and no estimate in the paper conditions on when a crisis ends.

Control pool and matching. Controls are countries never in crisis during the sample, restricted to emerging market and developing economies. Countries whose crisis status could never be evaluated for lack of a usable price history are excluded outright rather than treated as calm. Each treated country is then matched to the $k = 15$ most similar control candidates within its own region, by Euclidean distance over three standardized characteristics measured over the twelve months ending one month before the event window opens: the level of holdings in percent of GDP, the volatility of its monthly change, and log GDP. Standardization is over the stack's own candidate pool, and candidates missing any characteristic are dropped, since a country with no measurable pre-window behavior cannot be shown to be a good match. Matching within region additionally absorbs regional common shocks, since every comparison is between contemporaneous neighbors. The motivation is the composition of the unmatched pool: its volatility distribution in monthly holdings changes is far more right-skewed than the crisis countries', with a mean 1.7 times theirs, and a placebo distribution built from it is correspondingly too dispersed, making randomization inference conservative rather than sharp. More controls are not better here; what binds is the number of treated clusters.

Randomization inference. For each of 1,000 permutations, one clean control country is drawn uniformly at random within each event's matched pool, assigned placebo treated status at the event's actual anchor month, and the true treated country is dropped; every reported statistic — the average post-onset effect, the month-twelve level, and the joint pre-trend Wald statistic — is recomputed on each draw and referred to its own placebo distribution. Two-sided p-values are the share of permuted statistics at least as large in absolute value as the observed one. Confidence

Table 12: Currency-Crisis Events according to the Laeven–Valencia and Frankel–Rose Criteria

Country	Region	Onset		Pre-crisis GDP
		Laeven–Valencia	Frankel–Rose	($bn, year)
Chile	Latin America and Caribbean	—	2021-07	254.1 (2020)
Pakistan	Asia and Pacific	2021-08	2021-08	300.4 (2020)
Zimbabwe	Africa and Middle East	2021-11	2021-10	40.0 (2020)
Hungary	Europe	2021-11	2021-11	158.5 (2020)
Lao P.D.R.	Asia and Pacific	2021-11	2021-11	18.5 (2020)
Sierra Leone	Africa and Middle East	2021-12	2021-11	6.7 (2020)
Colombia	Latin America and Caribbean	—	2022-01	318.5 (2021)
Russia	Europe	2022-03	2022-01	1,828.9 (2021)
Ukraine	Europe	2022-02	2022-01	199.8 (2021)
Argentina	Latin America and Caribbean	2022-03	2022-02	486.0 (2021)
Ghana	Africa and Middle East	2022-02	2022-02	79.5 (2021)
Egypt	Africa and Middle East	2022-04	2022-03	423.3 (2021)
Sri Lanka	Asia and Pacific	2022-03	2022-03	88.6 (2021)
Malawi	Africa and Middle East	2023-09	2022-06	12.5 (2021)
Botswana	Africa and Middle East	2022-07	2022-07	18.8 (2021)
Suriname†	Latin America and Caribbean	2022-07	2022-07	3.1 (2021)
Angola	Africa and Middle East	2022-11	2022-11	83.0 (2021)
Zambia	Africa and Middle East	2022-12	2022-12	22.1 (2021)
Belarus	Europe	2023-02	2023-01	73.7 (2022)
Lebanon	Africa and Middle East	2023-02	2023-02	25.7 (2022)
Kenya	Africa and Middle East	—	2023-04	114.7 (2022)
Burundi	Africa and Middle East	2023-05	2023-05	3.9 (2022)
Nigeria	Africa and Middle East	2023-06	2023-06	645.7 (2022)
Türkiye	Europe	2023-06	2023-06	924.8 (2022)
Ethiopia	Africa and Middle East	2024-08	2024-08	159.8 (2023)

Notes: Events are country-months meeting the mechanical depreciation criteria of Laeven and Valencia (2020) and Frankel and Rose (1996); the "crisis" label refers to these criteria throughout. The table lists the entries surviving every sample screen (estimable onset, matched control pool, minimum pre-window). Onsets are real-time anchor dates. A dash in the Laeven–Valencia column marks the three events only the looser Frankel–Rose threshold selects. †Suriname is excluded from the estimation baseline as an extreme influence outlier (see the text); the baseline comprises the remaining 21 (Laeven–Valencia) and 24 (Frankel–Rose) events. Pre-crisis GDP is IMF WEO nominal GDP of the calendar year before the anchor year – the denominator of the holdings outcome.

intervals invert these tests horizon by horizon: the interval for β_k is $[\hat{\beta}_k - q_{hi}(P_k),\ \hat{\beta}_k - q_{lo}(P_k)]$, where P_k is the placebo distribution at horizon k; because that distribution need not be centered at zero, the intervals are not symmetric around the point estimates, and at 1,000 draws a 95 percent interval rests on 25 draws in each tail. The equal-weighted aggregation in Online Appendix G applies the same one-draw-per-event protocol to the equal-weighted statistic, with the placebo event-level effects computed stack by stack. Cluster-robust standard errors accompany all point estimates; at 21–24 treated clusters they agree with the randomization tests, and randomization inference remains primary because its validity does not rest on asymptotics in the number of treated clusters (Fisher, 1935; Imbens and Rubin, 2015; MacKinnon and Webb, 2020).

Extended horizon. The main text's window ends twelve months after onset. Figure 14 re-estimates the Laeven–Valencia design with the post-window extended to twenty-four months, leaving anchors, truncation, matching, and the randomization protocol unchanged. The control pool does not depend on the horizon, because controls are countries never in crisis during the sample. The extended design therefore draws the same 86 matched controls as the baseline and reproduces its months-0–12 average of 0.040 exactly. All but one event contribute through month twenty-four; Ethiopia's window is cut by the sample end after month sixteen, and the events-per-horizon count is annotated in the figure. The path decelerates and settles: the month-twenty-four level is 0.115 (95 percent randomization interval $[0.009, 0.215]$), and the change between months twelve and twenty-four is +0.044 with a randomization p of 0.23, indistinguishable both from a flat stock and from continued accumulation at a diminished pace. Extending further to thirty-six months, the point estimates continue to drift upward with no reversion at any horizon (0.200 at month thirty-six, 95 percent randomization interval $[-0.010, 0.375]$; month-thirty-six minus month-twelve +0.128, randomization $p = 0.15$), though we do not lean on the long-horizon slope, for two reasons. Only 14 of the 21 events remain estimable that far out, as the later crises' windows are cut by the sample end, so the slope partly reflects composition. And several of the crisis countries enter a *second* crisis spell within thirty-six months of the first onset, so at long horizons the path partially measures crisis recurrence rather than the aftermath of the original episode alone. The crisis spells themselves are mostly over well before these horizons, with the median spell ending about a year after onset, which is what makes the non-declining path informative: a flat stock after the spell ends means zero net flows, and a decline would require residents to actively sell their accumulated holdings back into the local currency. No horizon in either extension shows that.

Influence and the Suriname exclusion. Fitting each event's stack alone shows a right-skewed distribution of event-level effects, and one extreme outlier. Suriname's event effect of 0.65 is four times the next largest (0.17), and it is concentrated in a single month of net inflows. Suriname is also the smallest economy in the sample, with pre-crisis GDP of $3.1 billion, so in an outcome measured as a level in percent of GDP its movements are first order in any pooled average. Re-estimating the pooled design excluding one event at a time confirms that no other event matters comparably: the leave-one-out range of the average effect on the full 22-event sample is $[0.040, 0.074]$, with the

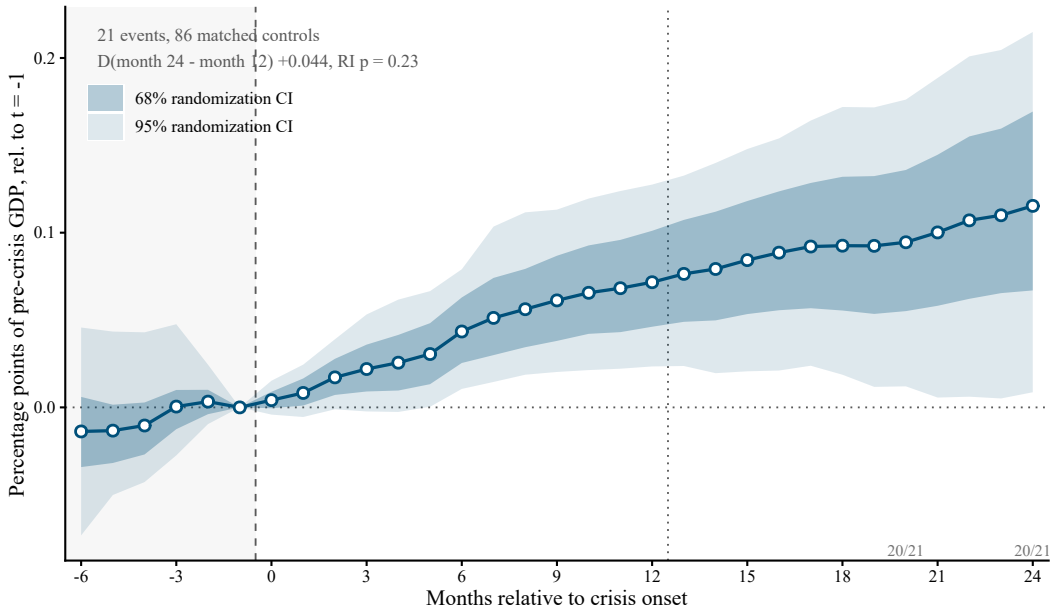


Figure 14: Holdings around currency crises: extended horizon

Notes: Stacked event-study estimates of stablecoin holdings in percentage points of pre-crisis GDP, as in Figure 9, with the post-onset window extended to twenty-four months. Shaded areas are 68 and 95 percent randomization confidence intervals over 1,000 matched draws. The dotted vertical line marks month twelve, the end of the main-text window; counts below the axis report the number of events contributing at horizons where the panel is no longer complete.

minimum at Suriname's exclusion and every other exclusion inside $[0.064, 0.074]$. The baseline therefore excludes Suriname, even though including it raises every estimate reported in the paper: including Suriname, the Laeven–Valencia average effect is 0.068 (s.e. 0.030; randomization $p < 0.001$), the month-twelve level is 0.106, and the path acquires an impact-month response of 0.030, against 0.004 in the baseline, that is specific to this one event. The exclusion also sharpens the estimates. The baseline average effect of 0.040 carries a cluster-robust t-statistic of 2.9, against 2.3 when Suriname is included.

G Negative Holdings Estimates

Because holdings are estimated as cumulative net flows, nothing constrains the estimates to be non-negative country by country, and in our full dataset 25.9 percent of country–month estimates fall below zero. The negative estimates are overwhelmingly small: 88 percent lie within $100 million of zero, only five countries ever fall below −$1 billion (none below −$5 billion), and at the end of the sample negative estimates sum to $6.0 billion against $206 billion of positive holdings. Negative estimates arise from coverage gaps, where residents acquire stablecoins through channels outside our universe (Section 3.4) and dispose of them through channels within it, and from measurement error in the probabilistic country assignment. They should be read not as debt or short positions,

which are impossible, but as holdings near zero measured with error.

Online Appendix

International Stablecoin Flows, Capital Controls and Currency Crises

Marco Reuter

(For Online Publication Only)

A Raw Extracts and Supplementary Lists

Section 3.1 lists the four inputs. This appendix records how they were obtained.

The blockchain extracts are full node copies, synchronized with each network and decoded with the Ethereum ETL package (Medvedev and the D5 team, 2018). Tron uses an adapted version of that package. Each chain enters at its genesis block, Ethereum on 30 July 2015, and runs through 31 December 2025. The raw extract exceeds 30 terabytes. A transaction in these extracts is any signed state change, not only an asset transfer.

Addresses are shared across the EVM-compatible chains in the sample, so a wallet's activity can be pooled. Solana is omitted because it is not EVM-compatible and because a full copy is in the petabyte range.

Domain names come from Dune Analytics. Users pay from single- to triple-digit dollar amounts for a name that replaces the 42-character hex address. That personal choice is what makes the names informative.

The Dune supplementary lists do four jobs. Exchange wallets, including the deposit wallets identified in Online Appendix C, mark centralized-exchange endpoints. Automated-trading wallets are dropped because their on-chain behavior is unlike the labeled wallets the model is trained on. Smart-contract addresses supply features and also identify the DeFi transfers that filter (i) removes. The USDT and USDC contract addresses on each chain identify the transfers the panel allocates.

B Daylight Savings Time Graphs

Figure 15 repeats the daylight savings time comparison of section 3.3 for all five regions. Each panel plots the region's activity profile separately for the daylight savings time months and the standard-time months.

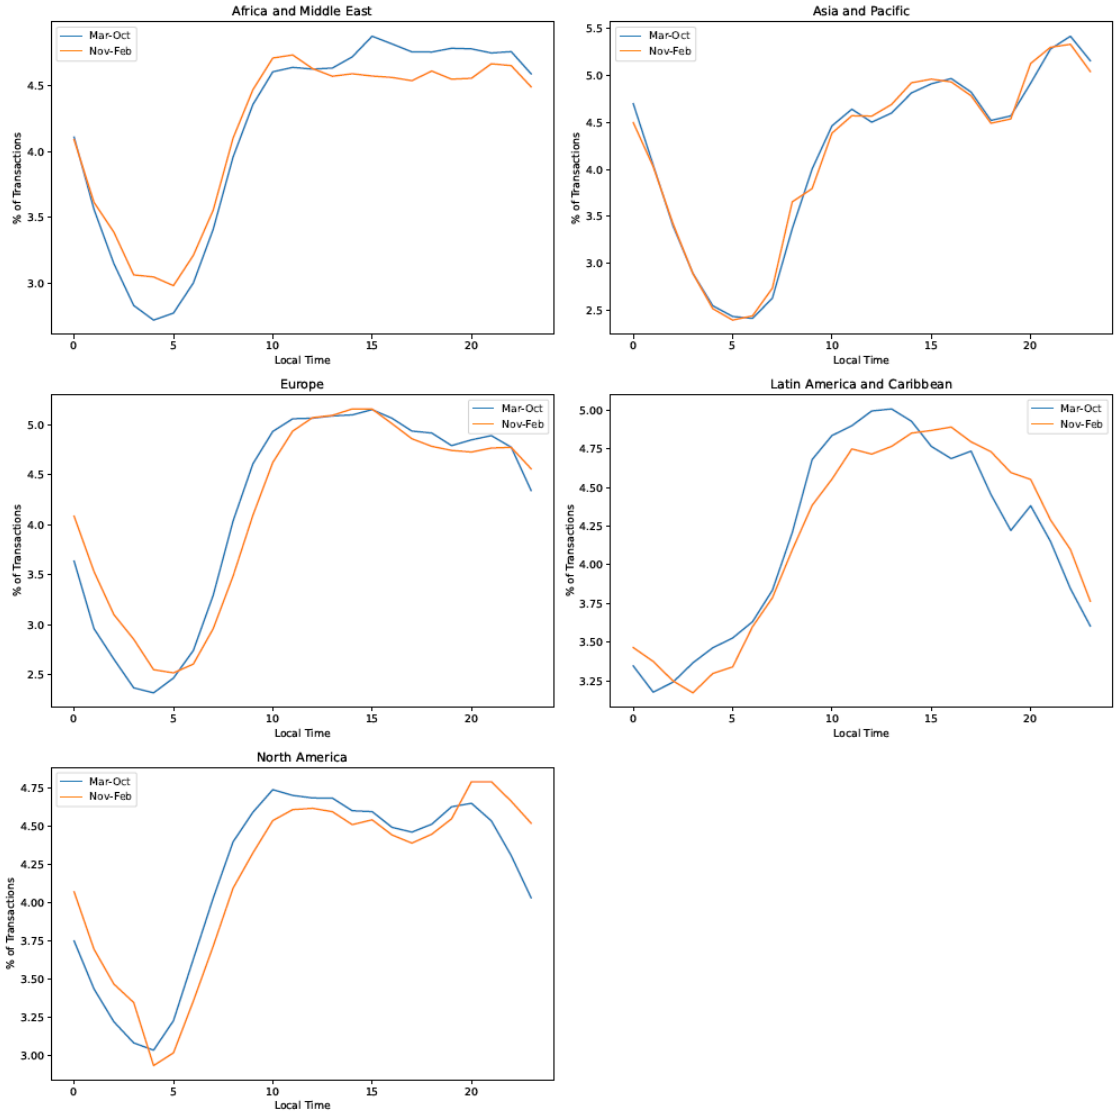


Figure 15: Activity Profiles During DST Months and Standard-Time Months, by Region

C Identifying Additional Centralized Exchange Wallets

We briefly describe how centralized exchanges structure their wallets. There are four types: hot wallets, cold wallets, deposit wallets and gas supplier wallets. Hot wallets hold pooled customer funds and send funds from the exchange to other wallets in on-chain transactions. Cold wallets are used for the long-term storage of funds and are not relevant for our analysis. Gas supplier wallets provide gas to other exchange wallets when needed.[44] Wallets of these three types are present in

[44]Gas measures the computational effort needed to process on-chain transactions. Users must pay for gas used, typically with the native protocol crypto asset, which for the blockchains in our sample is Ether, except

the dataset that we obtain from Dune.

Deposit wallets are created for individual customers when they want to deposit funds into the exchange, and they are absent from the Dune dataset. Typically, these addresses immediately forward deposits into the exchange's hot wallets and do not engage in any other activity. We identify a wallet as a deposit wallet belonging to a particular exchange if it (1) only forwards funds to hot wallets of that exchange or (2) has received gas from a gas supplier wallet of that exchange.

D Classification Model: Features, Training, and Evaluation

This appendix records the implementation that Section 3.2 defers: the prompt, the features, the two-stage tree, the training setup, the regional-exchange labeling rule, and the China model. Held-out scores are in Appendix C. The models are trained with the Yggdrasil Decision Forests library (Guillame-Bert et al., 2023).

D.1 LLM Prompt for Domain Name Classification

As described in section 3, we query the LLM for every domain name in our sample and instruct it to analyze the domain name and guess the country, language, and region the user is from and to provide a short reason for the classification. The exact prompt reads:

> "You are trained to classify ENS domain names by country, language, region, and provide a short reason for the classification. Consider references to culture, language, localities, memes. Be creative. Be mindful of the language commonly used in crypto and web3; and of languages that are spoken in many parts of the world, such as English, French and Spanish. If you cannot classify any of these, output 'unclassified' for that particular attribute. Classify the domain: {domain name}."

The prompt includes some further instructions for the LLM to restrict its answers for the region to the set {Africa and the Middle East, Asia and Pacific, Europe, Latin America and Caribbean, North America, unclassified}. We restrict the maximum number of tokens of the reply to 150 (roughly 110-120 English words) and set the temperature and "top_p" to 0.4 each, which makes the LLM's responses more deterministic and focused by reducing randomness. Table 13 presents one example classification, with the LLM's stated reason, for each possible region output.

The language model is used because the names include Arabic, Chinese, Korean, and Japanese scripts and even emojis, which a rule-based classifier would miss. Some wallets have more than one domain name. We assign the majority region among classified names, excluding "unclassified"

on Binance Smart Chain, which uses Binance Coin, and Tron, which uses TRX.

Name	Region	Country	Language
vucoworld	unclassified	unclassified	unclassified
Reason: The term 'vucoworld' does not clearly reference a specific country, language, or region. It appears to be a coined term possibly related to virtual or digital worlds, which are universal concepts in the crypto and web3 space.			
الكيماويات	Africa and Middle East	unclassified	Arabic
Reason: The domain name is in Arabic, which is widely spoken in the Middle East and North Africa, and it translates to 'chemicals', indicating a potential relevance to industries in those regions.			
pijiu	Asia and Pacific	China	Chinese
Reason: 'Pijiu' means 'beer' in Chinese, reflecting a popular beverage in Chinese culture.			
philippzinner	Europe	Germany	German
Reason: The name 'philippzinner' suggests a German origin, likely a personal name, which is common in Germany and associated with the German language.			
laplazart	Latin America and Caribbean	Mexico	Spanish
Reason: The name 'laplazart' suggests a connection to 'La Plaza', which is Spanish for 'The Square', a common term in many Latin American countries, particularly in Mexico, where plazas are central to community life and culture.			
lakings	North America	United States	English
Reason: The term 'lakings' likely refers to the Los Angeles Kings, a professional ice hockey team based in Los Angeles, California, which is a significant part of American sports culture.			

Table 13: Example of domain name classifications.

replies. A wallet with ten names, six unclassified, three Europe, and one North America, is assigned to Europe.

D.2 Regional-exchange labels

We designate an exchange as regional from its market and from its web traffic. Indodax, whose site is in Indonesian and which markets itself as an Indonesian exchange, is the type of case. The 126 wallets that also carry a classified domain name are assigned to the regional-exchange group, because concentration at a regional exchange is the stronger signal.

D.3 Features

We use the following features to train the classification model:

Time of Day Features. For every wallet, we calculate the percentage of transactions that are conducted within every hour of the day. To condense the data and avoid potential overfitting issues, we then estimate a third degree polynomial describing the activity profile and use the coefficients as features.[45]

[45]We use a third degree polynomial to be able to accommodate both a minimum and a maximum in the activity profile.

Daylight Savings Time Features. For every wallet, we calculate its activity profile during the months of daylight savings time and during the remainder of the year. This feature exploits regional heterogeneity in the use of daylight savings time. While DST is common in Europe and North America, it is not common in many other regions and non-existent in Asia.

Top 5 Centralized Exchanges. For every wallet, we count the number of transactions with a centralized exchange (i.e., withdrawals and deposits). We then sort descending by exchange name and create a categorical variable for the most used centralized exchange, second most used centralized exchange, and so on.

Top 10 ERC-20 Tokens and Smart Contract Namespaces. These are two similar features. For the top 10 ERC-20 tokens, we count, for every wallet, the number of transactions using ERC-20 contracts, sort them descending by count and create categorical variables for the top 10 most used ERC-20 contracts. For smart contracts in general, we count, for every wallet, the number of transactions with smart contracts of the same "namespace"[46], sort them in descending order by count and create categorical variables for the top 10 most interacted namespaces. While the first feature offers some more granularity exploiting heterogeneity on the ERC-20 level, the second feature groups smart contracts based on functionalities and exploits heterogeneity along this dimension.

D.4 Training and Evaluation Setup

To train and evaluate the model, we split the data into a training data set (90% of the data) and a testing data set (10% of the data). We then train 3 models that we combine to estimate a probability distribution for the regions from the set {Africa and Middle East, Asia and Pacific, Europe, Latin America and Caribbean, North America} that is assigned to the wallet. The process is illustrated by the decision tree in Figure 16.

That is, we first train a model that assigns probabilities for the following three classifications to an address: {North America, Latin America and Caribbean} or {Africa and Middle East, Europe} or Asia and Pacific. Instead of immediately estimating a probability for North America, or Latin America and Caribbean, the model assigns a probability that the wallet could be in either region. This approach maximizes the usefulness of the time of day based features, as the grouped regions largely share the same time zones. We then proceed to train two more models. One that splits {North America, Latin America and Caribbean} into its respective regions and one that splits {Africa and Middle East, Europe} into its respective regions. These further models then largely exploit the other features to distinguish between the respective regions. For each of the three models trained, the *YDF* package automatically splits off some of the data for validation to avoid overfitting any given model. To address imbalance in the training data, we calculate balanced sample weights (that is, each observation is weighted inversely proportional to its class frequency), ensuring that underrepresented classes contribute equally to the estimation process (see e.g., King

[46]The "namespace" groups smart contracts that offer similar functionalities as provided by Dune Analytics.

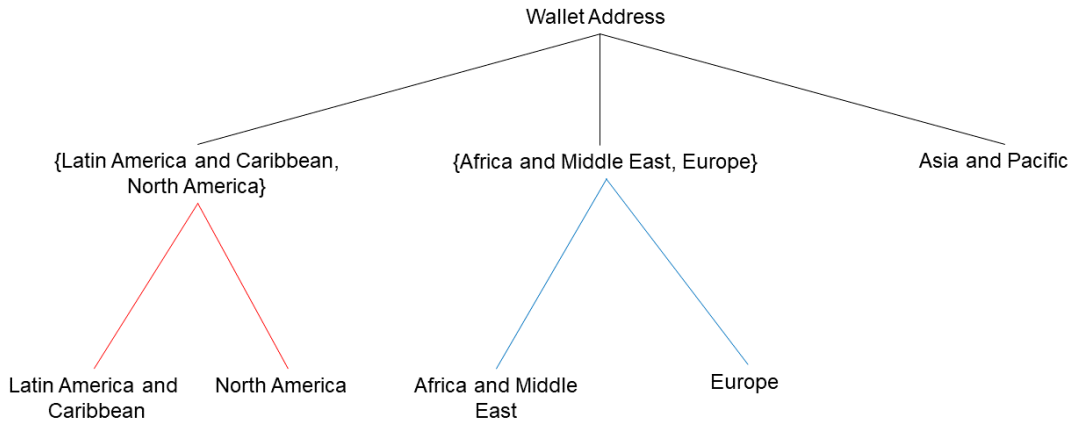


Figure 16: Decision Tree Structure of Classification Models

and Zeng (2001), He and Garcia (2009)).

D.5 China model

The China model uses the same features. It is trained on every labeled Asia and Pacific wallet, split into wallets whose domain names the language model assigned to China and all remaining labeled wallets in that region. The scores in Section 3.3 come from an identically specified model trained only on the 90 percent training split. The deployed model is trained on the full labeled Asia and Pacific set. Its output replaces web traffic for China. Only the residual Asia and Pacific share is then split by traffic, for wallets and for exchanges.

D.6 Component-model ROC curves

Figure 17 reports one-versus-rest receiver operating characteristic curves for the three component models of the regional classification and for the China model, on held-out wallets. For each class, the curve traces the share of that class's wallets the model correctly identifies against the share of other wallets it incorrectly assigns to the class, as the probability threshold is varied. A model with no discriminating power lies on the 45-degree line. The area under the curve ranges from 0.5 (random ranking) to 1 (perfect separation) and does not depend on class shares or on the threshold. Held-out accuracy, the source split, and the cross-source test are in Appendix C.

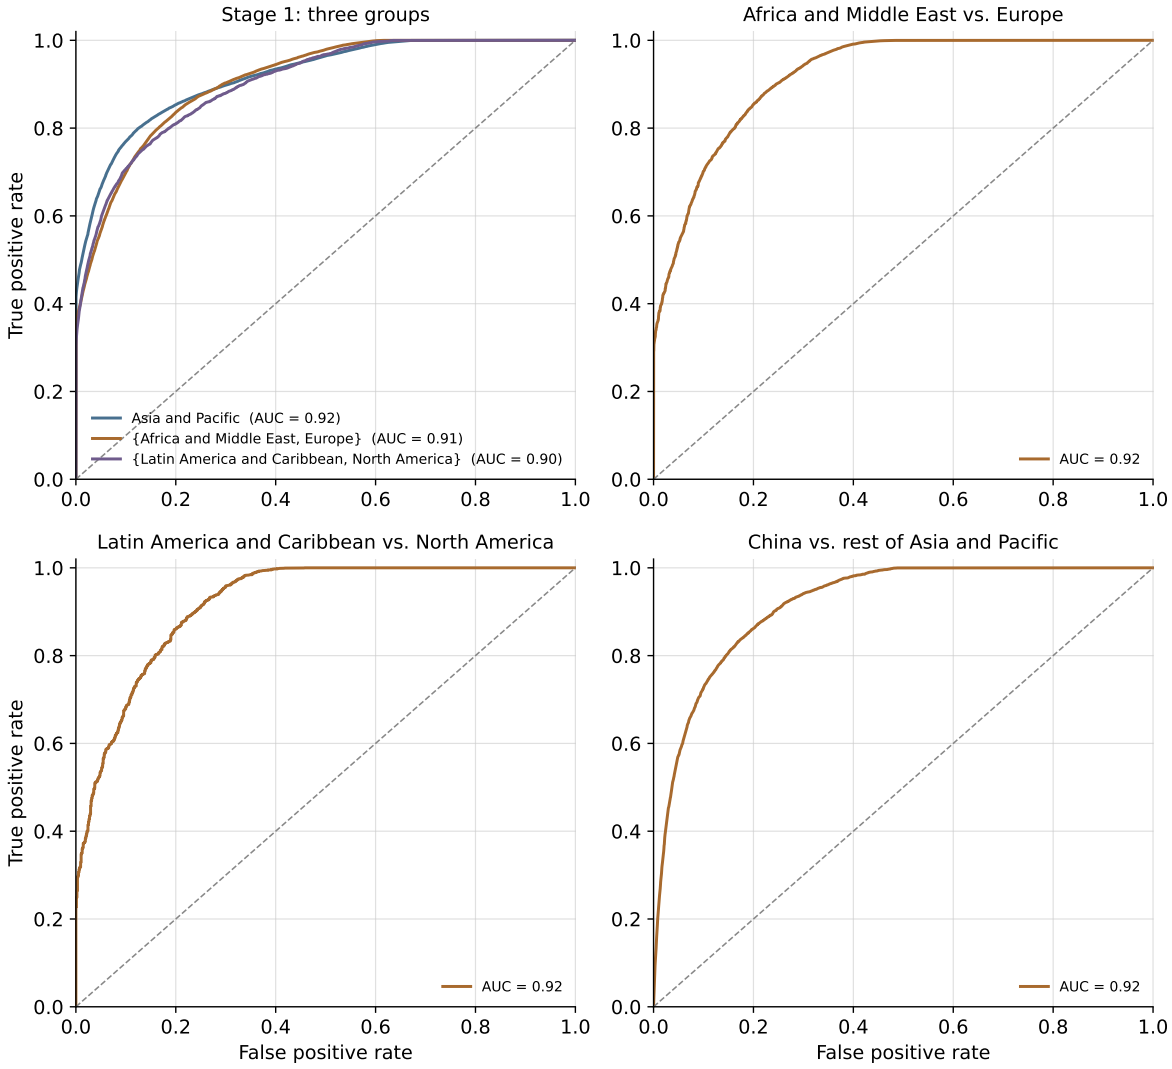


Figure 17: One-vs-rest ROC curves for the three component models of the regional classification and for the China model, evaluated on held-out wallets.

E Known Drivers: Full Results and Robustness

This appendix reports in full the two analyses that Section 5.1 summarizes: the push–pull decomposition of monthly stablecoin inflows and the corridor-level analysis of remittance costs. For each, we describe the data and specification, present the estimates, and probe their robustness.

E.1 Push and Pull Factors

We estimate the push and pull associations on a monthly panel of 188 countries from January 2021 through December 2025, for 9,775 country–month observations; the outcome is the log of monthly cross-border stablecoin inflows into country c, excluding domestic transfers, as in the

event studies of Sections 5.2 and 5.3. The push factors are the effective federal funds rate, the VIX, the crypto Fear & Greed index (a sentiment measure ranging from 0, maximally bearish, to 100, maximally bullish), the Federal Reserve's broad dollar index, and the log of the combined market capitalization of USDT and USDC;[47] the pull factors are month-over-month CPI inflation and exchange rate volatility, measured as the standard deviation of monthly log changes in the country's US dollar exchange rate over the trailing twelve months.[48] All country-varying regressors are winsorized at the pooled 1st and 99th percentiles in every specification, a choice we probe below. Our baseline specification is

$$\ln(SC_{c,t}) = \beta' P_t + \gamma'_{g(c)} Z_{c,t} + \alpha_c + \delta_t + \varepsilon_{c,t}, \tag{10}$$

where P_t collects the push factors, $Z_{c,t}$ the pull factors, and $g(c) \in \{\text{AE}, \text{EMDE}\}$ indexes group-specific pull slopes. Country fixed effects α_c absorb time-invariant country characteristics; year–month fixed effects δ_t, where included, absorb all global shocks, including the push factors themselves, which is why the specification ladder in Table 14 introduces them last. None of this variation is randomly assigned, so we read all estimates as conditional associations. Standard errors are two-way clustered by country and year–month.

Push: flows scale with the stock of stablecoins; the cyclical layer is suggestive. Table 14 reports the results. The push factors enter columns (1) and (2), and their estimates barely differ, so we describe column (2). We cannot reject unit elasticity with respect to the stablecoin stock: monthly flows scale proportionally with the market capitalization of USDT and USDC, a scale fact that anchors the level of flows. The remaining push coefficients form the cyclical layer: a one-standard-deviation improvement in crypto sentiment (21 points on the Fear & Greed index) is associated with about 5 percent larger flows, and a one-standard-deviation increase in the federal funds rate (2.1 percentage points) with roughly 9 percent smaller flows, signs that, taken at face value, make stablecoin flows in our data look like risk assets rather than a flight-to-safety vehicle, the reverse of the reading in Auer et al. (2025b).[49] But cyclical push coefficients, theirs and ours, are identified from a few years of aggregate time series, conditional on different sets of global controls and different flow constructions, so we place little weight on any individual one. Our preferred

[47] Inflation and exchange rates are from the IMF; the federal funds rate and the broad dollar index from FRED; the VIX from CBOE; the crypto Fear & Greed index from alternative.me; and the combined market capitalization of USDT and USDC from the data platform Allium.

[48] Auer et al. (2025b) measure volatility as the within-quarter coefficient of variation of the bilateral exchange rate; with monthly exchange rate data, the trailing twelve-month standard deviation of log changes is the less noisy analogue available at the frequency of our panel. The dollar exchange rate is the economically relevant one here: the stablecoins in our data are dollar-pegged, so its volatility measures the instability of the local currency against the asset being acquired.

[49] For both USDT and USDC they report significantly positive coefficients on the VIX and on the federal funds rate (their Table 6), estimated jointly with high-yield corporate spreads, US growth, and crypto risk factors, and on that basis read stablecoins as the safe leg within crypto. In our estimates, by contrast, the VIX association is negative but not significant, and the broad dollar index shows none.

specification in column (3) absorbs all global variation with year–month fixed effects, and none of the pull conclusions below depends on the push estimates.

Pull: the instability associations are detectable only in EMDEs. The pull side is where the groups part ways, and column (1) shows why pooling hides it: the common inflation slope is indistinguishable from zero (-0.017, s.e. $= 0.014$), the average of two group-level slopes that differ sharply. Column (3), our preferred specification, identifies the pull slopes from within-country movements in the pull variables, compared across countries within the same month, net of all global shocks. For inflation, two findings are robust. The AE slope is negative, with a one-percentage-point increase in monthly inflation associated with 7.4 percent smaller inflows, significant at the 1 percent level in every robustness column; and the two slopes diverge, with equality rejected at $p < 0.001$. The EMDE slope itself is positive (0.020, $p = 0.088$), with a point estimate that is stable across robustness cuts while its significance is not; read at the point estimates, the association reverses sign across groups, at -4.7 percent per standard deviation in AEs against $+2.3$ percent in EMDEs, but the claim the table robustly supports is the divergence, with the reversal as the point-estimate reading. The positive EMDE slope is the margin on which our estimates line up with Auer et al. (2025b), who find that both USDT and USDC flows are significantly larger when the sending or receiving country experiences high inflation, with flows roughly 15 percent larger for USDT and 50 percent for USDC.[50]

Exchange rate volatility is the strongest pull association in the exercise, and it is detectable only in EMDEs. A one-standard-deviation increase in EMDE volatility (2.3 percentage points of monthly volatility) is associated with 9.7 percent larger inflows (95% CI: 4.3 to 15.4). The association survives every treatment of the extreme-devaluation tail and strengthens slightly when Lebanon is excluded, while the unwinsorized estimate is positive but attenuated, consistent with the flow response saturating at collapse-level volatility, the territory of the currency-crisis analysis in Section 5.3. The AE slope, by contrast, is uninformative rather than zero: its per-unit point estimate (3.08, s.e. $= 5.23$) is close to the EMDE one, so equality is not rejected ($p = 0.85$), but AE dollar volatility moves so little (a standard deviation of 0.6 percentage points, against 2.3 in EMDEs) that even at the point estimate a one-standard-deviation move implies only about a 2 percent change in flows (95% CI: -4 to $+8$). What the data support is therefore not a contrast in responses but an asymmetry in detection: a robust positive association in EMDEs, and in AEs a volatility margin too quiet to identify anything.[51]

[50] Their inflation variable is a high-inflation indicator rather than a continuous rate, and high-inflation episodes in their sample occur overwhelmingly in EMDEs, so their coefficient corresponds most closely to our EMDE slope. The negative AE slope, and with it the divergence, is a margin their pooled design does not isolate; it is also why our pooled slope in column (1) is zero rather than positive.

[51] In column (2), the AE volatility slope appears positive and significant; it collapses once year–month fixed effects are added, reflecting that AE exchange rate volatility against the dollar spikes precisely when global risk conditions deteriorate. The EMDE slope is essentially unchanged across the two columns. The direction of the EMDE result is consistent with Auer et al. (2025b), who report positive though smaller and

A natural reading of this pattern, which we do however not explicitly test, runs through remuneration. Stablecoins typically pay no interest, so as an inflation hedge they are dominated wherever interest-bearing local-currency instruments keep pace with prices, as they broadly do in AEs, and all the more so when AE inflation raises nominal rates and with them the opportunity cost of holding a zero-yield asset. Where inflation and depreciation could erode the local currency faster than local remuneration compensates, as in some EMDEs, even a non-remunerated dollar claim protects principal. Where identification is sharpest, our independently constructed flows thus reproduce the instability–use association documented in the Chainalysis-based literature and locate it where the hedging motive is plausible, in EMDEs; where identification is weakest, in time-series-identified global coefficients, the datasets disagree.

The robustness cuts referenced above, reported in Table 15, re-estimate the preferred specification with the pull variables winsorized at the 2.5th and 97.5th percentiles instead of the baseline 1st and 99th, and with Lebanon, the largest depreciation episode in the sample, excluded altogether. Neither changes the conclusions.

E.2 Remittances

A frequently cited use case for stablecoins is the sending of remittances. Remittances sent through traditional intermediaries remain costly: according to the World Bank's Remittance Prices Worldwide (RPW) database, the global average cost of sending remittances stood at 6.4% of the amount sent in the third quarter of 2025, more than twice the 3% target by the G20. Against this backdrop, stablecoins, which can be transferred at low cost, around the clock, and without recourse to correspondent banking chains, have been highlighted as a potentially attractive alternative channel, particularly in EMDEs (Adrian et al., 2025); Du et al. (2026) and Reslow ⓡ al. (2026) document that stablecoins can be cheaper than traditional providers, and Auer et al. (2025b) show in a gravity framework that stablecoin flows from AEs to EMDEs are systematically larger in corridors where traditional remittances are more expensive. We document this relationship in our independently constructed data and extend it to corridors between EMDEs, where traditional remittance costs tend to be highest.

Data and specification. We merge corridor-level remittance costs from the RPW database, which reports the total cost of sending $500 (as a share of the amount sent) at a quarterly frequency,[52] with our bilateral flows aggregated to the corridor–quarter level, restricting the sample to corridors in which at least one endpoint is an EMDE, the corridors in which formal remittance costs are high and stablecoins are a plausible alternative: 348 corridors (231 AE-to-EMDE and 117

less uniformly significant coefficients on bilateral exchange rate volatility for stablecoin flows.

[52]The RPW database also reports the cost of sending $200. Results are similar with this measure: the pooled elasticity is 0.168 (s.e. = 0.053), with corridor-type-specific elasticities of 0.154 (s.e. = 0.080) for AE-to-EMDE and 0.180 (s.e. = 0.070) for EMDE-to-EMDE corridors.

Table 14: Push and pull factors of stablecoin flows

Dependent variable: $\ln(SC_{c,t})$	(1)	(2)	(3)
Push factors (global)			
Federal funds rate	−0.040	−0.044*	
	(0.024)	(0.025)	
VIX	−0.010	−0.010	
	(0.008)	(0.008)	
Crypto Fear & Greed index	0.002*	0.002*	
	(0.001)	(0.001)	
Broad dollar index	−0.013	−0.012	
	(0.010)	(0.010)	
ln(stablecoin market capitalization)	1.024***	1.027***	
	(0.099)	(0.097)	
Pull factors (country-specific)			
Inflation	−0.017		
	(0.014)		
FX volatility	4.192***		
	(1.231)		
Inflation × AE		−0.121***	−0.077***
		(0.030)	(0.022)
Inflation × EMDE		−0.005	0.020*
		(0.014)	(0.011)
FX volatility × AE		11.658*	3.077
		(5.862)	(5.233)
FX volatility × EMDE		4.045***	4.102***
		(1.215)	(1.151)
p-value: Inflation, AE = EMDE			< 0.001
p-value: FX volatility, AE = EMDE			0.846
Country FE	Yes	Yes	Yes
Year–month FE	No	No	Yes
Observations	9,775	9,775	9,775
Within R^2	0.279	0.282	0.015

Notes: The dependent variable is the log of monthly stablecoin inflows into country c. Push factors are global variables, identical across countries within a month; they are absorbed by the year–month fixed effects in column (3). Stablecoin market capitalization is the combined market capitalization of USDT and USDC. Inflation is month-over-month CPI inflation in percent; FX volatility is the standard deviation of monthly log changes in the country's US dollar exchange rate over the trailing twelve months. Both pull variables are winsorized at the pooled 1st and 99th percentiles (Online Appendix E reports robustness). In columns (2) and (3), the pull variables are fully interacted with country-group indicators, so each reported coefficient is the total slope for that group; the p-values, reported for the preferred specification in column (3), test equality of the AE and EMDE slopes. Standard errors, two-way clustered by country and year–month, in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

Table 15: Push and pull factors: robustness of the pull estimates

Dependent variable: $\ln(SC_{c,t})$	Baseline (1/99) (1)	Winsorized (2.5/97.5) (2)	Excl. Lebanon (3)	Unwinsorized (none) (4)
Inflation × AE	−0.077***	−0.081***	−0.077***	−0.078***
	(0.022)	(0.022)	(0.022)	(0.022)
Inflation × EMDE	0.020*	0.018	0.017	0.013*
	(0.011)	(0.015)	(0.012)	(0.007)
FX volatility × AE	3.077	3.460	3.229	2.402
	(5.233)	(5.216)	(5.225)	(5.188)
FX volatility × EMDE	4.102***	5.483***	4.683***	1.171
	(1.151)	(1.694)	(1.175)	(0.858)
Country FE	Yes	Yes	Yes	Yes
Year–month FE	Yes	Yes	Yes	Yes
Observations	9,775	9,775	9,715	9,775
Within R^2	0.015	0.014	0.017	0.007

Notes: All columns re-estimate the specification of column (3) of Table 14. Column (1) repeats the baseline, with both pull variables winsorized at the pooled 1st and 99th percentiles. Column (2) winsorizes both pull variables at the 2.5th and 97.5th percentiles instead. Column (3) uses the baseline winsorization but excludes Lebanon, the country with the highest exchange rate volatility in the sample. Column (4) leaves both pull variables unwinsorized. Standard errors, two-way clustered by country and year–month, in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

EMDE-to-EMDE) observed over 18 quarters from 2021Q1 to 2025Q3, for 6,011 corridor–quarter observations.[53] Our baseline specification is

$$\ln\left(SC_{ij,t}\right) \;=\; \beta\,\ln\left(Cost_{ij,t}\right) \;+\; \alpha_{ij} \;+\; \gamma_t \;+\; \varepsilon_{ij,t}, \tag{11}$$

where $SC_{ij,t}$ denotes stablecoin flows from country i to country j in quarter t, $Cost_{ij,t}$ is the average cost of sending \$500 in the same corridor and quarter, and the corridor fixed effects α_{ij} and year–quarter fixed effects γ_t absorb all time-invariant bilateral characteristics (e.g., distance, common language, migrant stocks) and global developments such as aggregate crypto market cycles.[54]

[53] Coverage is dictated by the RPW database, which is restricted to corridors with a material volume of traditional remittance flows. The RPW sample nominally contains two further corridor types, but neither supports meaningful inference: it includes only 3 AE-to-AE corridors (54 corridor–quarters), all running from Western Europe to Central and Eastern European EU member states, and only 3 EMDE-to-AE corridors (12 corridor–quarters), all of which originate in Russia, where the post-2022 sanctions environment plausibly moved both measured remittance costs and stablecoin flows for reasons unrelated to remittances. We therefore exclude both corridor types from the analysis. The same concern does not drive our results within the retained sample: the EMDE-to-EMDE elasticity is essentially unchanged, at 0.135 (s.e. = 0.061), when the ten remaining Russian corridors (all EMDE-to-EMDE, 40 corridor–quarters) are also excluded.

[54] Results are essentially unchanged when we replace the common year–quarter fixed effects with corridor-type-by-quarter fixed effects, which allow each corridor type its own time path: the pooled elasticity is 0.132

Identification thus comes from within-corridor variation in remittance costs over time, which is substantial (the within-corridor standard deviation of log costs is 0.27); because this variation is not randomly assigned, we read the estimates below as conditional associations rather than causal effects. Standard errors are clustered at the corridor level.[55]

Costlier remittances are associated with larger stablecoin flows. Table 16 reports the results. Within a corridor, quarters with more expensive traditional remittances see systematically larger stablecoin flows: the pooled elasticity is 0.138 (s.e. = 0.046), so a 10 percent increase in the cost of traditional remittances is associated with 1.4 percent larger stablecoin flows in the same corridor and quarter. Column 2 allows the elasticity to differ by corridor type: the estimates are 0.157 (s.e. = 0.071) for AE-to-EMDE and 0.125 (s.e. = 0.060) for EMDE-to-EMDE corridors, statistically indistinguishable from one another ($p = 0.74$), so the association is a broad feature of corridors involving EMDEs rather than an artifact of a single corridor type, extending the AE-to-EMDE evidence of Auer et al. (2025b) to the corridors where traditional remittances are most expensive.[56]

F CFM Event Studies: Robustness

This appendix collects robustness exercises for the capital-flow-measure results of the paper: formal tests of equality across restriction directions, sensitivity to individual countries and to crisis and sanctions episodes, a dose-response in the stringency of the tightening, alternative control pools, contemporaneous inflation and exchange-rate volatility, a twelve-month post-activation window, pooled loosenings, and a local-projections difference-in-differences estimator.

F.1 Robustness of the Outflow Response

This appendix collects robustness exercises for the directional results of Section 5.2. We test the equality of the directional event-time paths formally, probe the sensitivity of the outflow response to individual countries and to crisis and sanctions episodes, estimate a dose-response specification in the stringency of the tightening, vary the composition of the clean-control pool, condition on

(s.e. = 0.046), with corridor-type-specific elasticities of 0.150 (AE to EMDE) and 0.119 (EMDE to EMDE).

[55] With corridor and year–quarter fixed effects included, the dominant remaining source of residual correlation is serial correlation within corridors, which corridor-level clustering accounts for. The pooled estimates are robust to two-way clustering by corridor and year–quarter as well as to CRV3 jackknife inference over the 18 year–quarter clusters, which is conservative given the small number of time clusters. The corridor-type-specific estimates remain significant at the 10% level or better under these alternatives.

[56] Auer et al. (2025b) report elasticities with respect to remittance costs lagged by one quarter. Our estimates are similar when we lag the cost measure accordingly: the pooled elasticity is 0.116 (s.e. = 0.046), with corridor-type-specific elasticities of 0.123 (AE to EMDE) and 0.111 (EMDE to EMDE), the former imprecise under corridor clustering ($p = 0.11$). The lagged AE-to-EMDE elasticity is the estimate directly comparable to the quarterly elasticities they report, which amount to 0.127 for USDT and 0.055 for USDC.

Table 16: Remittance costs and stablecoin flows

Dependent variable: $\ln(SC_{ij,t})$	(1)	(2)
$\ln(Cost_{ij,t})$	0.138***	
	(0.046)	
$\ln(Cost_{ij,t})\times$ AE to EMDE		0.157**
		(0.071)
$\ln(Cost_{ij,t})\times$ EMDE to EMDE		0.125**
		(0.060)
Corridor FE	Yes	Yes
Year–quarter FE	Yes	Yes
Observations	6,011	6,011
Within R^2	0.005	0.005

Notes: The dependent variable is the log of quarterly stablecoin flows in corridor ij. The sample comprises corridors in which at least one endpoint is an EMDE (231 AE-to-EMDE and 117 EMDE-to-EMDE corridors); see the text for details on the sample restriction. $Cost_{ij,t}$ is the average total cost of sending \$500 via traditional remittance service providers (World Bank Remittance Prices Worldwide), measured contemporaneously. In column (2), the cost variable is fully interacted with corridor-type indicators, so each reported coefficient is the total elasticity for that corridor type. Standard errors, clustered by corridor, in parentheses. *, **, *** denote significance at the 10%, 5%, and 1% level, respectively.

contemporaneous inflation and exchange-rate volatility, and extend the post-activation window to twelve months.

Equality of the directional paths. Table 11 reports the full event-time coefficient estimates underlying Figure 8 in the main text: column (1) the pooled baseline, columns (2)–(5) the directional decomposition. Because the directional estimates of columns (2)–(5) come from a single joint regression, the equality of the paths can be tested directly from the estimated covariance matrix. Table 17 reports, at each horizon, the difference between the outflow coefficient and the inflow and neutral coefficients, together with joint Wald tests over the post-activation horizons $k = 0, \ldots, 6$. The outflow–inflow difference is positive at every post-activation horizon and grows with the horizon, reaching approximately 11 percent at $k = 6$, but it is statistically indistinguishable from zero throughout (smallest post-activation p-value 0.38; joint $p = 0.89$). The outflow–neutral difference is marginally significant at the four- and five-month horizons ($p \approx 0.07$ and 0.08) but the joint test does not reject ($p = 0.15$). These comparisons reflect limited power as much as substance: the inflow direction comprises 43 events across only 26 treated countries, and the joint tests spread that limited power over seven restrictions. A more powered test of the claim at stake in the main text is the single restriction that the *average* post-activation effect ($k = 0, \ldots, 6$) is equal across directions. The bottom panel of Table 17 reports this test: the average post-activation effect is 0.152 (s.e. = 0.048) for the outflow direction, against 0.103 (s.e. = 0.057) for inflow and 0.082 (s.e.

$= 0.048$) for neutral restrictions, and the averaged differences of 0.049 (s.e. $= 0.072$, $p = 0.50$) for outflow–inflow and 0.071 (s.e. $= 0.063$, $p = 0.27$) for outflow–neutral do not reject equality, referred to a $t(G-1)$ distribution based on the same country-clustered covariance matrix. The corridor-level comparison of Appendix E.3 reproduces both the ordering and the insignificance of these contrasts under a sharper control structure, which supports reading them as a limit of power rather than of design. The directional evidence should accordingly be read as we present it in the main text: the outflow path is the largest, the most persistent, and the only one individually robust across horizons, but the data cannot rule out that other restriction types elicit attenuated versions of the same response, most naturally because stablecoins can be used to circumvent those restrictions as well.

Influential countries and crisis episodes. The timing arguments of Section 5.2 address gradually deteriorating conditions but not discrete shocks arriving in the activation month itself, episodes in which outflow controls, sanctions, and currency stress are imposed within weeks of one another. We probe this concern in three ways, in each case re-estimating the directional specification and tracing the outflow-direction path. First, we estimate leave-one-out paths, dropping each of the 84 treated countries in turn, both its events and its observations as a control. The resulting paths are plotted in Figure 18 (panel (a)) and highlight that the estimates are essentially unchanged. Second, we exclude all countries subject to broad-based sanctions programs of the US Treasury's Office of Foreign Assets Control (OFAC (2026)) during the sample period (Russia, Belarus, Iran, and Venezuela), again dropping them both as treated and as control units before events are defined. (Figure 18, panel (b)). Third, we exclude the 28 countries meeting the Laeven and Valencia (2020) currency-crisis criterion of Section 5.3, a twelve-month depreciation against the dollar of at least 30 percent that exceeds its year-earlier rate by at least 10 percentage points, in at least one month of the sample period, so that the remaining events are not accompanied by acute currency stress (Figure 18, panel (b)). This screen is demanding: it removes the countries behind 45 of the 87 outflow tightenings. The leave-one-out paths are essentially unchanged from the baseline, the outflow path after the sanctions exclusion is qualitatively similar, and the path estimated on the 42 events in countries without a currency crisis is, if anything, larger than the baseline. We conclude that the outflow response is not an artifact of a handful of crisis or sanctions episodes.

Treatment intensity. The baseline treats any month-over-month increase in the restriction count as an event, so that a mild administrative measure and a sweeping prohibition are weighted identically. To assess whether this matters, we interact the outflow event-time dummies with the intensity of the tightening, holding all other event categories and the fixed-effects structure fixed. Bergant et al. (2026) score each measure by its stringency, from 0.1 for a notification requirement to 1.0 for an outright ban, and we measure an event's intensity as the summed score of the measures activated in the event month. Intensity enters in standardized logs, because its distribution is heavily right-skewed: the median event scores 0.75 and the largest 43. Three events record a

Table 17: Wald Tests of the Equality of Directional Event-Time Paths

	(1) Outflow − Inflow	(2) Outflow − Neutral
$k = -6$	−0.107	−0.001
	(0.173)	(0.080)
$k = -5$	−0.054	−0.132
	(0.081)	(0.092)
$k = -4$	0.028	−0.054
	(0.073)	(0.084)
$k = -3$	−0.021	−0.127*
	(0.051)	(0.074)
$k = -2$	−0.000	−0.023
	(0.032)	(0.040)
$k = 0$	−0.015	0.010
	(0.038)	(0.046)
$k = 1$	0.000	0.004
	(0.066)	(0.053)
$k = 2$	0.035	0.053
	(0.058)	(0.069)
$k = 3$	0.041	0.037
	(0.078)	(0.081)
$k = 4$	0.079	0.145*
	(0.089)	(0.079)
$k = 5$	0.095	0.138*
	(0.111)	(0.079)
$k = 6$	0.109	0.108
	(0.129)	(0.112)
Joint test, $k = 0, \ldots, 6$: F-statistic	0.422	1.563
Joint test, p-value	0.888	0.149
Avg. post effect, $k = 0, \ldots, 6$: difference	0.049	0.071
	(0.072)	(0.063)
Averaged test, p-value	0.497	0.267

Notes: Each cell reports the difference between the outflow-direction event-time coefficient and the inflow-direction (column (1)) or neutral-direction (column (2)) coefficient at the same horizon, computed from the joint regression underlying columns (2)–(5) of Table 11, with standard errors of the difference (in parentheses) based on the country-clustered covariance matrix. The joint tests are Wald tests of the equality of the two paths over all post-activation horizons $k = 0, \ldots, 6$, referred to an $F(7, G-1)$ distribution, where G is the number of country clusters. The averaged test refers the difference in the average post-activation effect ($k = 0, \ldots, 6$) between the two directions to a $t(G-1)$ distribution, based on the same country-clustered covariance matrix. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

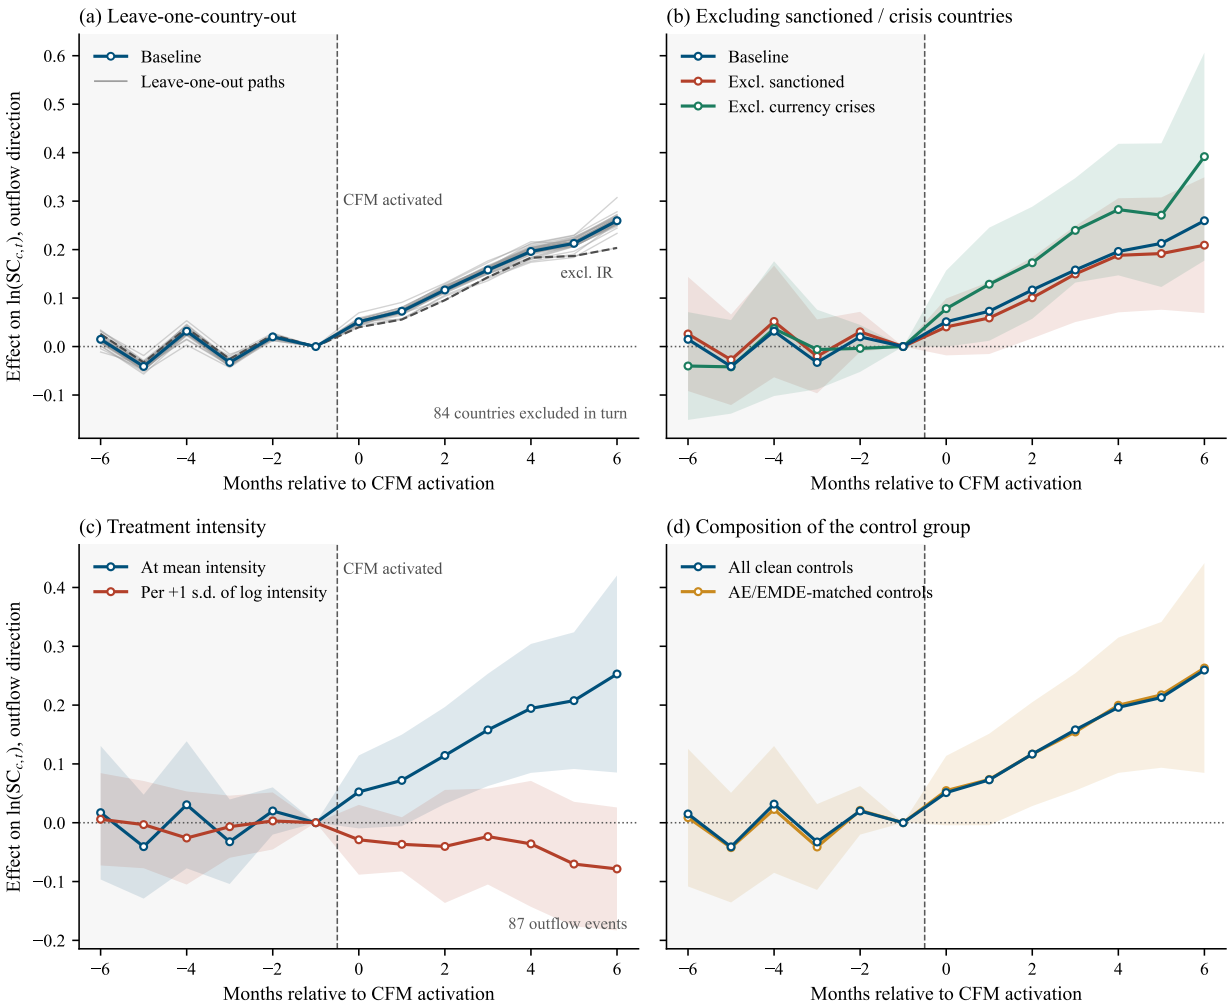


Figure 18: Robustness of the Outflow Response

Notes: All panels plot the outflow-direction event-time path from the directional specification of Section 5.2 under alternative designs. Panel (a) re-estimates the path dropping each treated country in turn (grey lines), both as a treated unit and from all control pools, alongside the baseline path (blue); the dashed grey line marks the single most influential exclusion. Panel (b) plots the baseline path alongside the paths obtained after excluding countries under broad-based OFAC sanctions programs during the sample period and after excluding the 28 countries meeting the Laeven–Valencia currency-crisis criterion of Section 5.3 in at least one month of the sample period; excluded countries are removed before events are defined, so their activations neither generate events nor enter any control pool. Panel (c) interacts the outflow event-time dummies with the standardized log intensity of the tightening in the event month, using the intensity scores of Bergant et al. (2026), and plots the outflow path at mean intensity together with the change in the path per standard deviation of log intensity, estimated jointly with all other event categories. Panel (d) plots the baseline outflow-direction path alongside the path obtained when clean controls are required to belong to the same country group (advanced economy or emerging market and developing economy) as the treated country. Coefficients are relative to the month immediately preceding the activation ($k = -1$, dashed vertical line). Shaded areas show 95 percent confidence intervals based on standard errors clustered at the country level.

nonpositive intensity change despite the count increase; we floor them at 0.1, the smallest score in the rubric. The specification traces two paths (Figure 18, panel (c)): the outflow response at mean intensity, and the change in that response per standard deviation of log intensity. The mean-intensity path is indistinguishable from the baseline, building to 0.25 (s.e. = 0.09) six months after activation. The intensity slope is flat before activation and, if anything, negative after it, at −0.08 (s.e. = 0.05) per standard deviation at the six-month horizon; raw and rank codings of intensity give the same picture. The response therefore does not grow with the stringency of the tightening. The main channel is an extensive-margin effect of restrictions being imposed, not of how much stringency is added.

Composition of the control group. Treated countries are predominantly emerging markets, while the clean-control pool in the baseline draws on all 182 countries. If stablecoin adoption trends differently across country groups, the event-by-event-time fixed effects would not absorb the difference. We therefore re-estimate the directional specification requiring clean controls to belong to the same country group (advanced economy or emerging market and developing economy) as the treated country in each event. The outflow-direction path under the group-matched pool is virtually indistinguishable from the baseline (Figure 18, panel (d)).

The baseline pool is never-treated: a country is excluded from an event's controls if it activates any restriction anywhere in that event's window. Using not-yet-treated countries as controls, retained up to the month of their own first activation, leaves the six-month coefficients unchanged, at 0.187 against 0.186 in the pooled specification and 0.260 in both outflow-direction specifications. Twenty-nine events retain no pre-activation months, 24 because a predecessor falls in the immediately preceding month and 5 because they fall in the sample's first month. Their treated level is absorbed by the event-by-country fixed effect. Dropping them from the treated stack, which the LP-DiD sample does by construction, moves the outflow six-month coefficient from 0.260 to 0.278.

Macroeconomic controls. A remaining concern is that time-varying country conditions, such as inflation and exchange-rate volatility, coincide with CFM activations and independently raise stablecoin inflows. We re-estimate the stacked design on the countries for which both series are observed in every month of the panel, first without those covariates and then with them, using the same month-over-month CPI inflation and twelve-month FX volatility measures as in Section 5.1. The restriction drops 33 of 217 events. On this common sample, adding the two controls leaves both the pooled path and the outflow-direction path essentially unchanged (Figure 19). The six-month pooled coefficient is 0.149 without the controls and 0.152 with them, and the outflow-direction coefficients are 0.195 and 0.193. The small gap relative to the baseline 0.186 and 0.260 is the sample cut, not the covariates.

Persistence: a twelve-month window. The baseline window extends six months after activation, so persistence beyond that horizon is not measured by the baseline design. We therefore

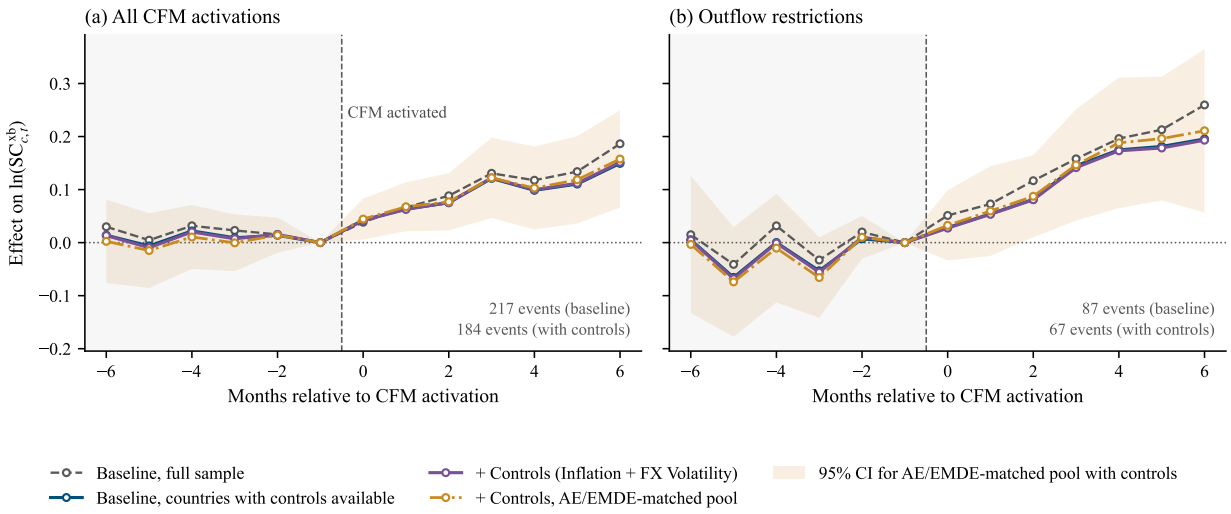


Figure 19: Stablecoin inflows around CFM activations, with inflation and FX volatility

Notes: The figure plots stacked event-study estimates of cross-border stablecoin inflows around CFM activations. Panel (a) pools all activations. Panel (b) plots the outflow-direction path. The dashed line is the baseline on the full sample (217 and 87 events). The solid and dotted lines restrict the panel to the 157 countries whose month-over-month CPI inflation and 12-month FX volatility are observed in every month, first without those covariates and then with them. Coefficients are relative to the month immediately preceding the activation ($k = -1$, dashed vertical line). The shaded area shows 95 percent confidence intervals for the AE/EMDE-matched pool with controls, based on standard errors clustered at the country level.

re-estimate the directional specification with an asymmetric window of six months before and twelve months after each activation, exactly as in the currency-crisis design of Section 5.3; all other features of the design are unchanged: window truncation around repeat activations by the same country, the clean-control pool (defined over the longer window), and the fixed-effects structure. Over the shared horizons $k \leq 6$, the outflow-direction path is essentially identical to the baseline (Figure 20); beyond them, it remains flat at its six-month level rather than reverting, with coefficients of 0.288, 0.255, 0.267, 0.244, 0.233, and 0.246 at horizons seven through twelve (standard errors between 0.087 and 0.105), each with a 95 percent confidence interval excluding zero. The long-horizon sample does thin as repeat activations by the same country truncate the extended post-windows: of the 87 outflow-direction events, 49 retain a treated observation at the six-month horizon and 32 at the twelve-month horizon. The stability of the path across this thinning, and its agreement with the baseline where the two designs overlap, indicate that the persistence is not an artifact of the surviving subsample.

F.2 Loosening of CFMs

We next estimate the mirror image of the baseline experiment: instead of increases, we define an event as any country–month in which the total count of capital flow restrictions *falls*. This yields

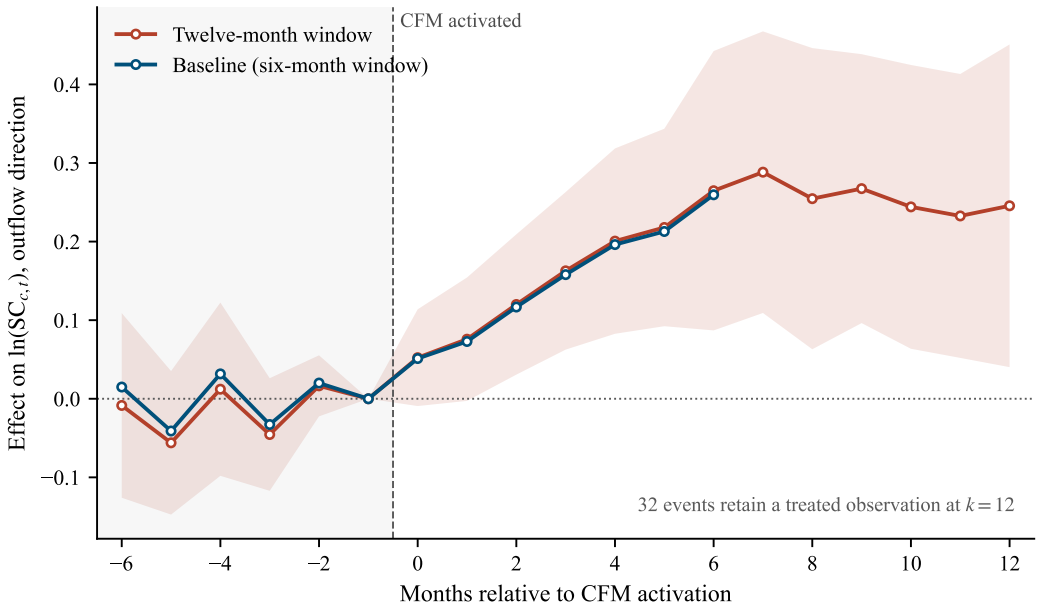


Figure 20: Outflow-direction response over a twelve-month window

Notes: The figure plots the outflow-direction event-time path from the directional specification re-estimated with an asymmetric window of six months before and twelve months after each activation, alongside the baseline six-month-window path. Of the 87 outflow tightenings, 49 retain a treated observation at the six-month horizon and 32 at the twelve-month horizon, as repeat activations truncate the extended post-windows. Coefficients are relative to the month immediately preceding the activation ($k = -1$, dashed vertical line). The shaded area shows 95 percent confidence intervals for the twelve-month path, based on standard errors clustered at the country level.

351 loosening events over 2021–2024, which we attribute, exactly as in the directional decomposition of the main text, to the direction of the restriction that was relaxed: outflow (152 events), inflow (52 events), directionally neutral (72 events), or multiple directions simultaneously (75 events). The design is otherwise identical to the baseline, with the loosening events taking the role of the activations: a six-month window on either side of the event, truncation around repeat loosenings by the same country (179 events have their pre-window and 179 their post-window truncated), a clean-control pool consisting of countries without a loosening of their own within the event window, event-by-event-time and event-by-country fixed effects, and standard errors clustered at the country level.

The results are reported in Table 18. Pooled across all loosenings (column (1)), the effect on stablecoin inflows is essentially nil: apart from a small uptick of approximately 5 percent in the event month itself, every post-event coefficient is close to zero and statistically insignificant, in sharp contrast to the steadily building response to tightenings documented in the main text. The splits in columns (2)–(5) add nuance to this null. Loosenings of *outflow* restrictions (column (2)), the mirror image of the tightenings that drive the main result, are not followed by any significant change in stablecoin inflows at conventional levels; the point estimates drift negative at the longest

horizons (approximately −14 percent at five to six months) but remain only marginally significant. Outflow loosenings are, however, preceded by negative and significant pre-event coefficients, i.e., stablecoin inflows were rising in the months leading up to the removal of the restriction, so the reference month is not a clean counterfactual for this event type and we refrain from a causal reading of these paths in either direction.[57] What the data clearly rule out is a mirror-image response: nothing in the outflow-loosening path comes close, in magnitude or statistical robustness, to the prompt and persistent 30 percent increase that follows outflow tightenings. Columns (3) and (4) split those outflow removals by whether the same country tightened outflow restrictions earlier in the sample, which is the comparison the ratchet turns on. The late-horizon drift is entirely in the removals that no tightening precedes, at −0.207 six months after the event; where a reversal would be expected, the path is flat throughout, running from 0.024 in the event month to −0.021 at six months with no coefficient distinguishable from zero. That subsample carries a pre-event pattern of its own, significantly negative at five and six months before the removal, so inflows are rising into these events. It cuts against a reversal reading rather than for one: a flat path after a rising pre-period is the conservative reading of no give-back. The effect of CFMs on stablecoin adoption thus appears asymmetric, consistent with fixed costs of adoption: users who have set up wallets, exchange accounts, and counterparty relationships in response to a restriction have little reason to abandon them once it is lifted. The pre-event dynamics nonetheless counsel caution in pushing this interpretation too far. The closest thing to a systematic response appears for *inflow* restrictions (column (5)): stablecoin inflows tick up significantly in the month an inflow restriction is removed and drift upward thereafter, reaching approximately 10 percent after six months, though this longer-horizon estimate is significant only at the 10 percent level. While we cannot pin down the mechanism, the pattern would be consistent with stablecoins serving legitimate inbound-transfer use cases, most naturally remittances, that become more attractive, or newly permissible, once restrictions on receiving funds from abroad are relaxed. Directionally neutral removals and simultaneous removals across several directions are not tabulated, since neither shows a response at any horizon.

The stacked design of Section 5.2 may be sensitive to two implementation choices. First, countries treated shortly *before* an event window are admitted into the control pool; given the persistence documented in Appendix F.2, their stablecoin inflows may remain elevated, which would attenuate the estimates. Second, the stacking implicitly weights events by the size of their surviving samples, so that a subset of events could receive disproportionate weight. For robustness, we therefore re-estimate the baseline using the local-projections difference-in-differences (LP-DiD) estimator of Dube et al. (2025), which addresses both. For each horizon h, we estimate the long-

[57] Such pre-event dynamics are unsurprising: authorities plausibly relax controls in response to, or in anticipation of, changing capital flow pressures. Note that this endogeneity concern cuts in the opposite direction for tightenings, where restrictions are typically imposed in response to outflow pressure through *conventional* channels, which would, if anything, bias the estimated stablecoin response in the main text downward.

Table 18: Stablecoin Inflows around CFM Removals: Stacked Event-Study Estimates

	(1)	(2)	(3)	(4)	(5)
			Outflow removals		
	All loosenings	Outflow	Preceded by an outflow tightening	Not preceded	Inflow
$k = -6$	−0.015	−0.045	−0.167*	0.029	−0.036
	(0.041)	(0.069)	(0.093)	(0.093)	(0.058)
$k = -5$	−0.045	−0.116*	−0.143*	−0.104	−0.103**
	(0.037)	(0.061)	(0.086)	(0.082)	(0.045)
$k = -4$	−0.038	−0.116**	−0.105	−0.123	−0.073**
	(0.031)	(0.057)	(0.077)	(0.079)	(0.037)
$k = -3$	−0.048*	−0.148***	−0.099	−0.181**	0.012
	(0.027)	(0.054)	(0.071)	(0.075)	(0.048)
$k = -2$	0.003	−0.037	−0.054	−0.024	0.016
	(0.019)	(0.033)	(0.039)	(0.049)	(0.022)
$k = 0$	0.047***	0.038	0.024	0.052	0.041***
	(0.017)	(0.026)	(0.026)	(0.042)	(0.014)
$k = 1$	0.019	−0.003	0.041	−0.034	0.060*
	(0.023)	(0.039)	(0.030)	(0.064)	(0.031)
$k = 2$	0.009	0.001	0.006	−0.001	0.037
	(0.028)	(0.051)	(0.047)	(0.078)	(0.037)
$k = 3$	−0.012	−0.056	0.010	−0.102	0.022
	(0.033)	(0.055)	(0.048)	(0.085)	(0.043)
$k = 4$	−0.004	−0.073	0.002	−0.125	0.041
	(0.037)	(0.063)	(0.065)	(0.093)	(0.049)
$k = 5$	−0.032	−0.141*	0.013	−0.217**	0.059
	(0.040)	(0.074)	(0.075)	(0.101)	(0.046)
$k = 6$	−0.014	−0.150*	−0.021	−0.207*	0.100*
	(0.050)	(0.086)	(0.095)	(0.114)	(0.056)
Events	351	152	75	77	52
Event × event-time FE			Yes		
Event × country FE			Yes		
Observations			358,591		
R^2			0.989		

Notes: The dependent variable is the log of stablecoin inflows. Coefficients are event-time effects relative to the month immediately preceding the loosening ($k = -1$, omitted), with event time binned at $k = \pm 6$. Column (1) pools all loosenings. Columns (2) and (5) come from a single joint regression in which the treated-side event-time dummies are interacted with the direction of the restriction relaxed. Columns (3) and (4) split the 152 outflow removals by whether an outflow tightening by the same country precedes them, from a separate regression. Column (3) is the same-direction outflow subset of the matched removals in Section 5.2. All specifications include event-by-event-time and event-by-country fixed effects. Standard errors, clustered at the country level, in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

difference regression

$$\ln SC_{c,t+h} - \ln SC_{c,t-1} = \beta_h D_{c,t} + \tau_t + \varepsilon_{c,t}, \tag{12}$$

where $D_{c,t}$ indicates a CFM activation by country c in month t and τ_t are base-month fixed effects, estimated on a horizon-specific sample: treated observations are activations with no other activation by the same country within $[t-6, t+h]$, and control observations are country–months with no activation within the same window. Pre-event placebo horizons $h \in \{-6, \ldots, -2\}$ use the analogous long difference. The cleanliness requirement thus extends to the months preceding the base period, so that previously treated countries do not enter the control pool, and the weighting of events is explicit: plain OLS delivers the variance-weighted LP-DiD estimand, while a reweighted version recovers the equally-weighted average effect across events, so that a comparison of the two reveals whether a subset of events drives the results.

Figure 21 plots the variance-weighted and reweighted LP-DiD estimates alongside the stacked event-study path of column (1) of Table 11. Both LP-DiD paths closely track the stacked estimates throughout the event window: stablecoin inflows rise by approximately 5 percent in the activation month and build to approximately 18 percent at the six-month horizon, against 3.8 and 18.6 in the baseline, and the pre-event coefficients are small and statistically insignificant, mirroring the flat pre-trends of the baseline. The variance-weighted and reweighted estimates are moreover nearly indistinguishable at every horizon. The baseline path is therefore neither attenuated by control-group contamination from previously treated countries nor driven by a heavily-weighted subset of events.

G Currency Crises: Robustness

This appendix reports robustness of the currency-crisis event study: the Frankel–Rose dating variant, alternative detection windows, heterogeneity-robust estimators, and the estimates collected in Table 19.

Threshold sensitivity. The Frankel–Rose variant adds Chile, Colombia, and Kenya and dates eight onsets earlier, for a baseline of 24 events with 85 matched control countries (6,821 observations). Its estimated path is very similar to the Laeven–Valencia path of Figure 9 and slightly steeper: the average effect over months 0–12 is 0.057 (s.e. 0.018; randomization $p = 0.017$, cluster-robust $p = 0.004$) against the baseline 0.040, holdings reach 0.109 percentage points of pre-crisis GDP at month twelve (95 percent randomization interval $[0.046, 0.189]$), and the pre-onset path is flat (joint pre-trend randomization $p = 1.00$). Every conclusion of the main text is unchanged, and the classifier row of Table 19 places the estimate alongside the baseline's robustness battery.

Heterogeneity-robust estimators. The main text reports that the Callaway and Sant'Anna (2021) and Sun and Abraham (2021) estimators reproduce the baseline estimates. Both are estimated on a single monthly panel constructed from the same inputs as the stacked design: emerging

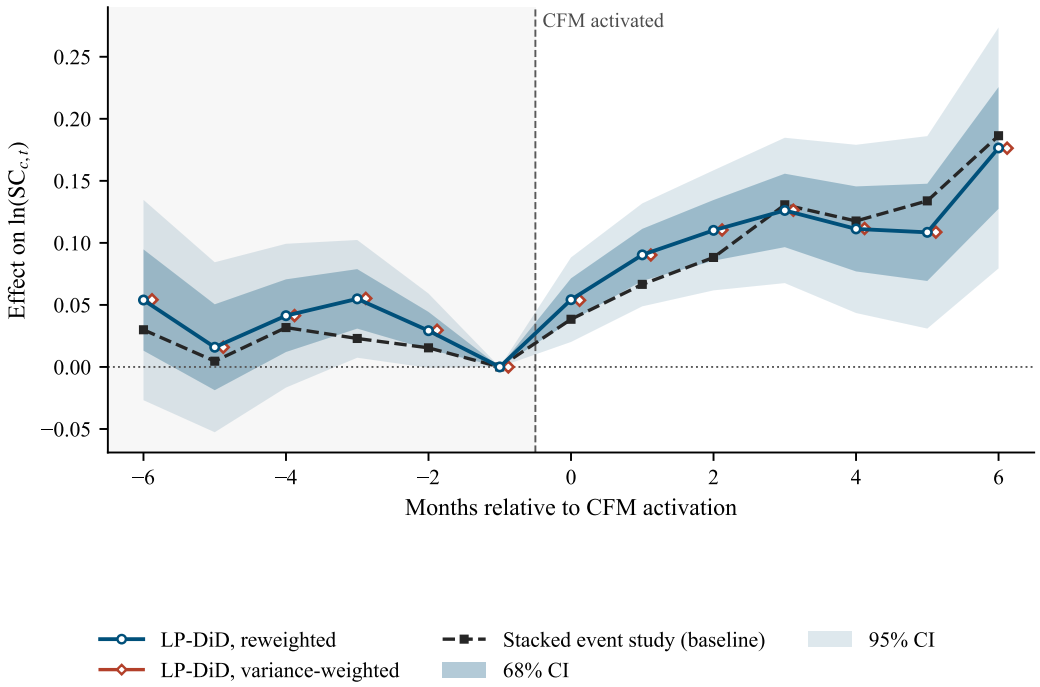


Figure 21: LP-DiD Estimates of Stablecoin Inflows around CFM Activations

Notes: The figure plots local-projections difference-in-differences estimates of the effect of new capital flow measures (CFMs) on stablecoin inflows at each horizon h relative to the activation month, alongside the stacked event-study estimates of column (1) of Table 11. At each horizon, treated observations are activations with no other activation by the same country within $[t-6, t+h]$, and control observations are country–months with no activation within the same window; all regressions include base-month fixed effects. The variance-weighted estimates are obtained by OLS; the reweighted estimates recover the equally-weighted average effect across events. Coefficients are relative to the month immediately preceding the activation ($k=-1$, dashed vertical line). Shaded areas show 68 and 95 percent confidence intervals for the reweighted estimates, based on standard errors clustered at the country level. The variance-weighted estimates are displaced slightly to the right for visibility.

market and developing economies only, countries whose crisis status cannot be evaluated excluded, ever-in-crisis countries excluded unless they are estimable events, and the 21 baseline anchors as treatment cohorts. These estimators take a single outcome column, which cannot carry the event-specific pre-crisis GDP vintages of the stacked design, since the same control country serves different events with different vintages. The outcome is therefore scaled by each country's 2020 WEO GDP, a common pre-sample vintage that is fixed before every event in the sample and coincides with the baseline's pre-crisis vintage for the 2021 cohorts; the comparable stacked estimate is the pre-sample-vintage variant in Table 19, not the pre-crisis-vintage baseline. The Callaway–Sant'Anna estimator is implemented with not-yet-treated controls, outcome regression, and a universal base period; the Sun–Abraham estimator with never-treated cohorts as the control group. The average effects over months zero to twelve are 0.052 and 0.054, against 0.047 for the stacked design under the same scaling. Estimated on the panel including Suriname, the Callaway–Sant'Anna and Sun–Abraham

estimates rise to 0.081, so the influence diagnosis is not an artifact of the stacked implementation either. Two further lessons come from the comparison. First, estimated without matched control pools, both estimators display a cumulative drift in holdings before onset, rising from about −0.05 six months before onset to zero at the reference month (Figure 22). The matched pools of the stacked design remove this drift, which is why the design matches, and its direction implies that any residual trend difference attenuates the post-onset estimates rather than inflating them. Second, the analytic inference these estimators provide is unreliable in this setting by their own diagnostics, with one treated country per cohort and 21 treated clusters, which is why the paper reports their point estimates and relies on randomization inference throughout. Run through the paper's randomization protocol (1,000 matched draws, one placebo control per event drawn from that event's matched pool, with draws distinct within a permutation because a single panel cannot assign one country two cohort dates), the observed estimates are rejected against their placebo distributions at $p = 0.004$ for Callaway–Sant'Anna and $p = 0.007$ for Sun–Abraham.

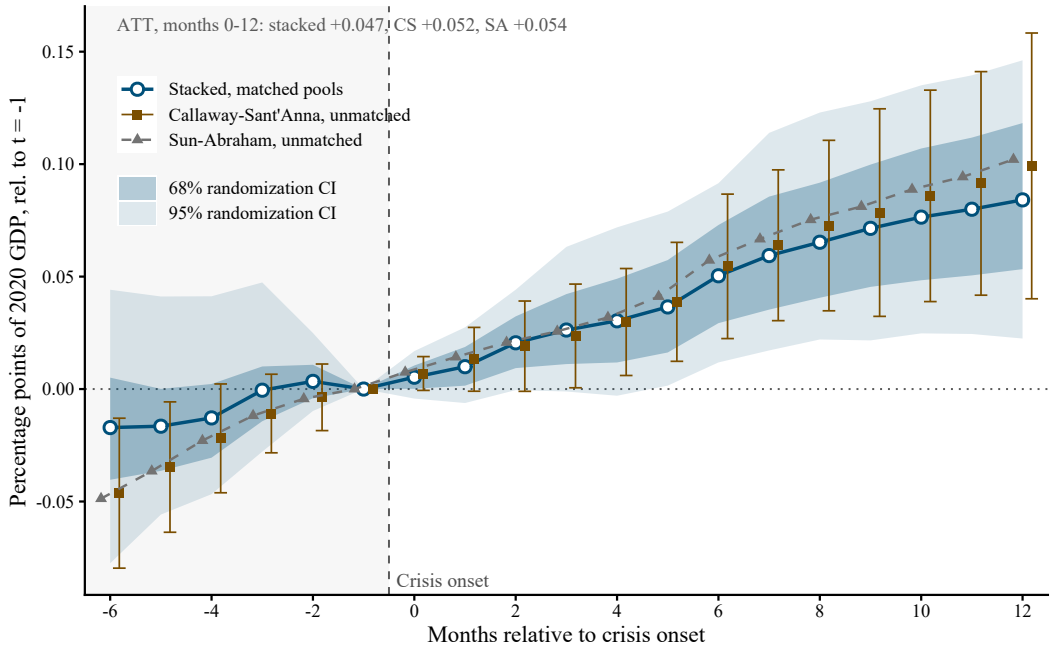


Figure 22: Holdings around currency crises: matched design and unmatched heterogeneity-robust estimators

Notes: Stablecoin holdings in percentage points of 2020 GDP, the common pre-sample vintage described in the text, on the 21-event baseline. The solid line is the stacked design on its matched control pools; shaded areas are 68 and 95 percent randomization confidence intervals over 1,000 matched draws. Squares and triangles are the Callaway and Sant'Anna (2021) and Sun and Abraham (2021) paths estimated on the unmatched panel. Error bars on the Callaway–Sant'Anna path are pointwise ±1.96 analytic standard errors, shown for completeness; the text explains why inference relies on randomization tests instead. Both unmatched paths drift before onset and the matched design removes the drift, while the post-onset paths and the average effects coincide.

Robustness table. Table 19 collects the estimates referenced in the main text, all on the 21-event

baseline apart from the Frankel–Rose row and its 24 events. Three entries warrant detail beyond the table notes.

Table 19: Holdings around Currency Crises: Robustness

	Estimate	(s.e.)	RI p	Pre-trend p	Events
Baseline (pooled, anchor window 3)	0.040	(0.014)	0.015	0.480	21
Equal-weighted event average	0.041	(0.013)	0.015	—	21
Including Suriname	0.068	(0.030)	< 0.001	0.654	22
Leave-one-out range (full sample)	[0.040, 0.074]	—	—	—	22
Excluding Russia and Ukraine	0.032	(0.013)	0.051	0.205	19
Frankel–Rose classifier (25% threshold)	0.057	(0.018)	0.017	0.999	24
Anchor window 2 months	0.041	(0.015)	0.022	0.933	21
Anchor window 4 months	0.036	(0.012)	0.017	0.036	21
Anchor window 5 months	0.042	(0.012)	0.002	0.185	22
Anchor window 6 months	0.034	(0.012)	0.043	0.338	22
CFM activation control	0.039	(0.015)	0.024	0.207	21
Pre-sample (2020) GDP denominator	0.047	(0.016)	0.010	0.368	21

Notes: Each row re-estimates the baseline average months-0–12 holdings response (first row) under one departure from the baseline. The equal-weighted row fits each event's stack separately and averages the event-specific effects with equal weights; its standard error is the standard error of that mean across events, and its randomization p applies the same one-draw-per-event placebo protocol to the equal-weighted statistic. The Including-Suriname row re-estimates the baseline on the full sample; the leave-one-out row reports the range of the pooled estimate across all single-event exclusions from that full sample, with cluster-robust inference per refit, and its minimum is the exclusion of Suriname, which is the baseline. The excluding-Russia-and-Ukraine row drops the two war and sanctions episodes jointly, the analogue for this design of the CFM analysis's sanctions screen; single-event exclusions cannot detect a pattern carried by the pair. The classifier row dates crises with the Frankel–Rose 25 percent threshold instead of the baseline's Laeven–Valencia 30 percent, which adds Chile, Colombia, and Kenya and dates eight onsets earlier; like the baseline, it excludes Suriname. Anchor-window rows vary the real-time detection window used to date onset (baseline three months, threshold scaled proportionally). The CFM row controls for an indicator that a capital-flow-measure activation, defined exactly as in Section 5.2, occurred in the country during the current or preceding six months; its sample ends in 2024:12 with the CFM data. The GDP-vintage row scales holdings by each country's 2020 WEO GDP, a common pre-sample vintage fixed before every event, instead of the event-specific pre-crisis vintage. Randomization p-values use 1,000 matched draws; the pre-trend column is the joint randomization test of the pre-onset coefficients.

Anchor windows. A longer detection window demands the same cumulative depreciation over more months and therefore dates onset later; any early months of the response then fall into the pre-period and the reference month itself absorbs part of the buildup. Across detection windows of two to six months the average effect stays within 0.034–0.042, with randomization $p \leq 0.043$ in every design; the pre-trend test rejects only at the four-month window ($p = 0.036$), one rejection across five tests at the 5 percent level. Reading the pre-onset coefficients directly gives the direction of any residual dating error: where they deviate from zero they are *below* the reference month and rising toward it, accumulation already underway before the dated onset, so residual dating error

attenuates the baseline estimate rather than inflating it.

GDP vintage. The baseline denominator is WEO nominal GDP of the calendar year before the anchor year, fixed before treatment and measured near the event. The vintage row scales every country instead by its 2020 GDP, a common pre-sample vintage that is also fixed before every event in the sample. It yields a somewhat larger estimate (0.047, randomization $p = 0.010$), because pre-crisis-year GDP exceeds 2020 GDP for most crisis countries, so the baseline is the conservative choice of the two. Two denominators are not reported at all. A contemporaneous-year denominator would inflate the estimate mechanically, because a currency crisis collapses dollar GDP over precisely the treated window. A post-sample snapshot would reweight countries by devaluations that postdate their crises.

Capital-flow-measure control. The control is an indicator equal to one if the country records a capital-flow-measure activation, defined as in Section 5.2, in the current month or in any of the six preceding months. The window matches the six-month horizon over which the response of Section 5.2 builds, so a restriction imposed at the onset of a crisis is absorbed over the months in which it moves flows rather than in the month it is announced. Weighting activations by their stringency instead would condition on a margin that Section 5.2 shows carries no response. The measure's coverage ends in December 2024, which truncates one event's window and removes about 2 percent of the rows, so we fit the identical rows with and without the control to separate the sample change from the covariate. The average effect is 0.040 without the control and 0.039 with it, and the controlled and uncontrolled event-time paths are nearly indistinguishable at every horizon. The control's own coefficient is positive, at 0.024. Activations accompany higher holdings conditional on the crisis, which is the sign the channel of Section 5.2 predicts.

H Additional Detail for the Chainalysis Comparison

Section 3.5 reports the headline comparison. This appendix records the harmonization, the year-by-year, fixed-effects, and country-group tables behind those headlines, the China timing, a contradiction inside the Chainalysis categories, and the holdings-feasibility test.

The web-traffic allocation. Chainalysis allocates transfers between centralized exchanges across country pairs in proportion to the product of the exchanges' website traffic shares (Cerutti et al., 2025a; Chainalysis, 2025). If exchange 1 draws 70 percent of its traffic from country X and exchange 2 draws 50 percent from country Z, then 35 percent of their volume is recorded as a flow from X to Z. Self-custodial wallets fall outside that universe. Chainalysis approximates them with an *indirect* category of exchange-to-exchange transfers routed through a self-custodial wallet, alongside a *direct* category of transfers that pass from one exchange to another. How our methodology uses web traffic, and how that confines the traffic assumptions to a within-region split, is in Section 3.2 and Appendix D.

Harmonization and the comparison sample. Both datasets are aggregated to monthly bilateral country flows and joined on country pairs and months. The window is August 2019 through February 2025, the period covered by both.[58] Hong Kong SAR and Macao SAR are merged into China in both series. The comparison uses the 164 countries present in both; countries only we cover carry 0.01 percent of our gross flows, and countries only they cover carry 0.5 percent of theirs.[59] Both series are restricted to USDT and USDC. Auer et al. (2025b) report that the Chainalysis data cover USDC on Ethereum and USDT on Ethereum and Tron.

Table 20 reports gross flows by year. Table 21 reports the country-level correlations after

Table 20: Gross Stablecoin Flows by Year: Our Dataset vs. Chainalysis

Year	Ours	Direct	Indirect (USD bn)	Direct+Indirect	Ratio Ours/(D+I)
2019	39.4	7.3	5.8	13.1	3.00
2020	314.5	55.1	52.0	107.1	2.94
2021	1,992.6	582.8	364.1	946.8	2.10
2022	2,323.9	567.0	505.2	1,072.1	2.17
2023	2,856.5	481.4	450.4	931.8	3.07
2024	4,363.8	758.3	840.5	1,598.8	2.73
2025	899.2	136.7	145.0	281.6	3.19
Total	12,790	2,536	2,313	4,849	2.58

Notes: Common sample of 164 countries covered by both datasets, August 2019 to February 2025, monthly frequency, with Hong Kong SAR and Macao SAR merged into China in both datasets. The 2019 rows cover August–December and the 2025 rows January–February. Direct and Indirect are the two transaction categories in the Chainalysis dataset; flows involving issuer treasury wallets are excluded here (see Section 3.5).

removing period means and after additionally removing country means. Table 22 reports the four AE/EMDE corridor pairs.

The remaining numbers are not in those tables. Cumulative country nets over the window correlate at 0.63 once issuer flows are excluded from our data. The largest net recipients in both rankings are an overlapping set of EMDEs. Yearly net flows into EMDEs, in Table 23, are close in the summed Chainalysis series and far apart in each component.

The China timing. Figure 23 is the timing behind the text's claim that the China gap is a measurement break at the 2021 measures. The log wedge is unremarkable before May 2021, jumps by about 0.6 log points between the May enforcement measures and the September ban, and thereafter lies outside the central 95 percent range of the 40 largest comparison countries.

[58] Our access to the Chainalysis data ends with a partial month in March 2025, which we drop.

[59] The jurisdictions only Chainalysis covers are those missing from our web traffic data (for example Jersey, Guadeloupe, Réunion, Myanmar, Syria, Cuba, Libya).

Table 21: Country-Level Correlation Between the Datasets, With and Without Fixed Effects

Sample	Raw	Month FE	Country+Month FE	N
Panel A: gross flows (logs)				
Monthly, all countries	0.858	0.806	0.261	9,869
Monthly, top 30 countries	0.898	0.671	0.254	2,010
Quarterly, all countries	0.886	0.842	0.323	3,436
Quarterly, top 30 countries	0.908	0.677	0.265	690
Panel B: net flows (signed log)				
Monthly, all countries	0.068	0.081	0.035	10,988
Monthly, top 30 countries	0.109	0.141	0.046	2,010
Quarterly, all countries	0.088	0.099	0.036	3,772
Quarterly, top 30 countries	0.149	0.168	0.024	690

Notes: Pearson correlations between our estimates and the Chainalysis estimates (direct+indirect) on the country–period panel of the common sample. "Month FE" correlates the series after removing period means (cross-sectional agreement, period by period); "Country+Month FE" additionally removes country means (agreement on within-country movements). Gross flows enter in logs (observations with positive flows in both datasets); net flows enter as $\text{sign}(x)\log(1+|x|)$. Top 30 countries by combined gross flows.

Table 22: Gross Flows Between Advanced Economies and EMDEs: Levels and Dynamics

Flow	Ours (USD bn)	Chainalysis (USD bn)	Ratio	Corr. (logs)	Corr. (growth)
AE → AE	1,834	1,088	1.69	0.987	0.847
AE → EMDE	3,351	1,272	2.63	0.988	0.843
EMDE → AE	3,325	1,248	2.66	0.986	0.823
EMDE → EMDE	4,280	1,343	3.19	0.979	0.758

Notes: Common sample, cumulated over August 2019–February 2025. Countries are classified as advanced economies (AE) or emerging market and developing economies (EMDE) following the IMF World Economic Outlook (April 2026). Correlations are between the two datasets' monthly series for each group pair, in log levels and in month-over-month log growth.

The comparison does not run one way. For Iran they report eight times more gross activity than we do ($34 billion against $4 billion). Iran is the one country in our data whose negative implied holdings are large relative to its own activity, 46 percent of its cumulative gross flows against 0.07–0.3 percent for the other countries that ever turn materially negative (Appendix G). That is a coverage gap on our side. The jurisdictions they cover and we do not are those our web traffic inputs omit.

The direct–indirect contradiction. Table 23 is the annual counterpart of the 2024 split. In that year the Chainalysis direct category reports $60.8 billion of net flows into EMDEs and

Table 23: Net Stablecoin Flows into EMDEs by Year: Our Dataset vs. Chainalysis

	Ours	Chainalysis		
Year		Direct	Indirect	Direct+Indirect
		(USD bn)		
2019	0.02	0.11	−0.04	0.07
2020	−1.86	1.91	−0.24	1.67
2021	2.36	5.95	−3.54	2.41
2022	1.62	22.39	−24.33	−1.94
2023	7.88	25.48	−18.99	6.49
2024	13.88	60.83	−47.82	13.01
2025	1.78	12.52	−10.08	2.44

Notes: Net flows into EMDE countries (inflows minus outflows), common sample; ex-issuer flows in our dataset, so that both datasets measure the same redistribution concept. The 2019 row covers August–December and the 2025 row January–February.

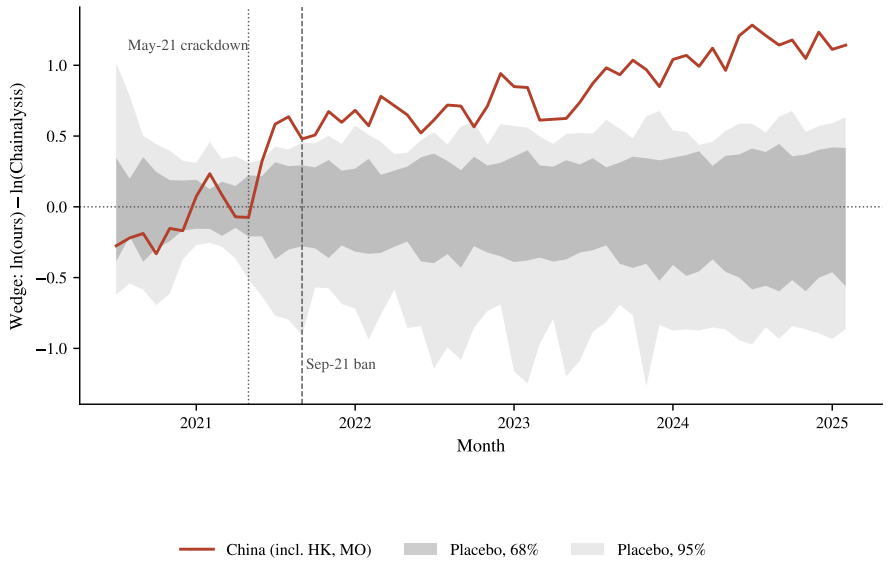


Figure 23: The China measurement wedge around the 2021 enforcement measures and ban

Notes: The figure plots the log wedge between our estimate of China's stablecoin gross flows and the Chainalysis estimate (Hong Kong SAR and Macao SAR merged into China in both), on the volume-matched sample described in the text: countries whose monthly gross flows exceed $100 million in both datasets by mid-2020 (balanced panel, once in always in). Month fixed effects are removed, and each country's series is anchored to zero over its pre-enforcement window (November 2020 to April 2021). The solid red line is China; the grey bands show the central 68 and 95 percent ranges of the same wedge across the 40 size-matched placebo countries with the largest combined gross flows. Vertical lines mark China's May 2021 mining restrictions (dotted) and September 2021 exchange and transaction ban (dashed).

the indirect category −$47.8 billion, each several times the combined answer and opposite in sign. Across countries, cumulative direct and indirect nets have a correlation of −0.98 and agree on sign for only 11 percent of countries. For the United States, the direct category implies cumulative net outflows of $102 billion and the indirect category net inflows of $82 billion. The indirect series also contains at least one large artifact: in April 2024 it records $198 billion of flows, roughly four times adjacent months, which shifts Brazil's implied cumulative position by about $13 billion in a single month. The foreign exchange records of Section 3.5 show $1.3 billion of stablecoin purchases in that month, in line with adjacent months. Only the sum of the two components is consistent with our data.

Holdings and the feasibility test. Cumulative worldwide net flows have to approximately equal the outstanding stock, and holdings approximated by cumulative nets cannot be negative in a closed system. The Chainalysis data comprise country-to-country flows only, so the first requirement cannot be met and negative holdings must exactly mirror positive ones. Cardozo et al. (2024) document the same limitation against official statistics for Brazil. Our dataset contains the issuer flows, $1.1 trillion gross over the window, which we exclude wherever we compare like with like. Nothing in our construction forces cumulative nets to remain positive. The issuer-inclusive series passes that test, with negative implied holdings of $5 billion against $151 billion of positive holdings at the end of the window and a worst case, the United States at −$5.0 billion during 2023, that recovers before the window ends. When those legs are dropped, negative and positive implied holdings each equal $47 billion, as they must in any closed bilateral panel. Negative holdings in our full panel, as opposed to this comparison window, are discussed in Appendix G.

What each dataset can and cannot measure. For the cross-country ranking of gross activity, either dataset serves, though the agreement rests partly on a shared web-traffic input. Findings on that margin transfer. Levels do not. In restricted markets, only entity-level attribution reaches the activity. Net-flow analysis with the Chainalysis data carries a common signal only in multi-year cumulations of the summed categories. The Chainalysis data cannot deliver holdings. Measuring accumulation requires the issuer flows, and measuring high-frequency net flows requires self-custodial wallets.

We repeat the two event studies of Section 5 on the Chainalysis series. The CFM design compares activating countries in event time with countries that have no activation in the window. The pooled post-activation average is +4.3 percent ($p = 0.44$). Treating their cumulative country nets as the holdings analogue, the crisis design's average post-onset effect is 0.029 percentage points of pre-crisis GDP (randomization $p = 0.43$, 200 draws).



PUBLICATIONS

International Stablecoin Flows, Capital Controls and Currency Crises
Working Paper No. WP/2026/211