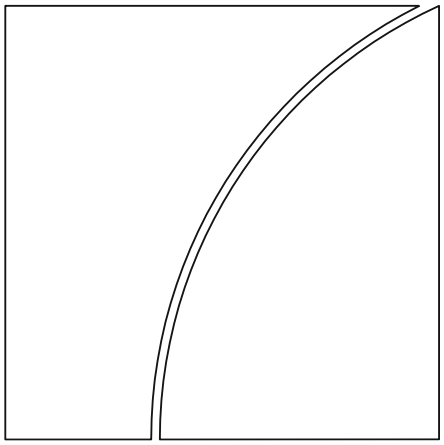

BIS Working Papers

No 1377

Hidden by complexity? Measuring stablecoin, crypto and decentralised finance ecosystems

by Timothy Aerts, Ronald Heijmans, Jan Paulick and Violeta Vuletic

Monetary and Economic Department

September 2026

JEL classification: C81, E42, G15, G23

Keywords: cryptoassets, stablecoins, blockchain, decentralised finance, smart contracts, Bitcoin, Ethereum, Tron

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in this publication are those of the authors and do not necessarily reflect the views of the BIS or its member central banks.

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Hidden by complexity? Measuring stablecoin, crypto and decentralised finance ecosystems*

Timothy Aerts†, Ronald Heijmans‡, Jan Paulick§, Violeta Vuletić¶

August 2026

Abstract

Decentralised finance data presents a distinctive paradox: while every data point is publicly recorded and accessible, deriving meaningful insights is obscured by the scale, fragmentation and complexity of the ecosystem. Key metrics illustrate that the rapidly evolving DeFi ecosystem introduces unique challenges for economic and financial research in accurately capturing financial activity in DeFi. These challenges stem from protocol architecture and the technical execution of transactions that complicate deriving economic meaning. Leveraging blockchain data for Bitcoin, Ethereum and Tron, the paper illustrates three structural sources of measurement divergences: (1) the economically meaningful aggregation of Bitcoin transaction values, (2) the challenge of programmability and proliferation of spurious smart contracts and (3) comparability of use cases across different chains, exemplified by stablecoins. While the examples relate to specific chains and layers of technical execution, the measurement challenges generalise to all blockchains and DeFi ecosystems relying on a similar technical underpinning. We propose measurement approaches based on granular, data-bounded estimates that incorporate explicit assumptions, technical classification and disaggregation to align technical execution with economic meaning. The paper demonstrates that, despite the transparency of public blockchains, widely used indicators of cryptoasset and DeFi activity are highly dependent on methodological choices and underlying assumptions that warrant careful interpretation. The findings imply that on-chain indicators should be treated as noisy approximations rather than direct measures of economic activity.

JEL Codes: C81, E42, G15, G23

Keywords: cryptoassets; stablecoins; blockchain; decentralised finance; decentralised exchanges; smart contracts; measurement; Bitcoin; Ethereum; Tron.

*We thank Italo Borssatto, Miguel Diaz, Martin Diehl, Jon Frost, Ulf Lewrick, Ilhyock Shim and Leanne Zhang as well as the participants of research seminars at the Bank for International Settlements for insightful comments. We thank Maximilian Schrader for excellent research assistance and data preparation. The views expressed in this publication are those of the authors and do not necessarily reflect the views of DNB, the BIS or its member central banks. The authors acknowledge the use of generative AI for the following purposes: proofreading, code testing and debugging. The authors bear full responsibility for the factual accuracy of this paper. References to individual firms and projects are for illustrative purposes and should not be construed as any statement on the soundness or the legal or regulatory status of any firm or project.

†Bank for International Settlements (BIS). timothy.aerts@bis.org

‡De Nederlandsche Bank (DNB), ronald.heijmans@dnb.nl

§Bank for International Settlements (BIS). jan.paulick@bis.org

¶Bank for International Settlements (BIS), violeta.vuletic@bis.org

1 Introduction

Over the past decade, the financial landscape has undergone a profound transformation, shaped by rapid technological innovation, evolving regulatory approaches and shifting user demands. A prominent manifestation of this transformation is the growth of stablecoins, cryptoassets and decentralised finance (DeFi) into a sizeable segment of global financial markets. According to industry estimates, the total market capitalisation of cryptoassets reportedly reached $3 trillion by the end of 2025, after peaking at $4.2 trillion earlier that year. The stablecoin market alone grew to over $300 billion, representing a 50% annual increase.

For policymakers and central banks, these developments are potentially consequential. Relative to conventional payment systems, transaction volumes in crypto markets are still small. However, cryptoassets, including stablecoins, play an increasingly prominent role in trading, payments and on-chain financial intermediation. DeFi protocols facilitate lending, trading and leveraged positions without centralised intermediaries, while stablecoins act as a bridge between traditional finance and DeFi ecosystems. As privately issued, blockchain-based proxies for fiat currencies, stablecoins could have significant macro-financial implications if their use for payments and as a store of value grows (Aldasoro et al., 2026; BIS, 2026). The evolution of DeFi could create channels for risk transmission through emerging actors, potentially resulting in liquidity spirals, run dynamics and greater interconnectedness with traditional financial systems. Therefore, accurate measures of economic activity are essential for evaluating the structure, dynamics and potential implications of stablecoins and DeFi for safeguarding financial stability and supporting resilient payment infrastructures.

Analysing DeFi dynamics depends on reliable and interpretable data. Every transaction on a public blockchain is permanently recorded and publicly accessible. This transparency is a foundational design principle: public ledgers are tamper-resistant because their contents are verifiable by anyone. This apparent abundance of data conceals deeper measurement challenges. There is a gap between transparent transaction records and deriving economic meaning, which is obscured by scale, fragmentation and complexity.

Headline indicators—such as market capitalisation, transfer volumes and activity metrics—are widely used to assess the performance and dynamics of markets. Market capitalisation reflects the total value of outstanding assets, while transfer volumes measure the aggregate value or number of transactions over a given period. Activity metrics, such as total value locked (TVL), capture the level of engagement or participation with DeFi applications. For example, TVL represents the total value of assets locked in a smart contract, often used as a proxy for user trust and the intensity of usage. These indicators inherently rely on assumptions about ownership, valuation and protocol behaviour, which can cause ostensibly similar measures to produce significantly different interpretations of underlying conditions.

The paper demonstrates that, despite the transparency of public blockchains, widely used indicators of cryptoasset and DeFi activity are highly dependent on methodological choices and underlying assumptions. We showcase three separate measurement challenges across different technical layers, namely the blockchain protocol, the smart contract layer and individual assets issued across different chains. These challenges impact different measures to varying degrees and magnitudes over time, dependent on user behaviour and technical underpinnings. While data providers offer a plethora of measures, the methodology and underlying assumptions are not always clear and structural measurement challenges remain under-identified. In addition, many indicators are geared towards trading rather than gauging economic activity. The paper identifies three structural sources of measurement divergences that arise in DeFi and demonstrates their empirical significance across three chains—Bitcoin, Ethereum and Tron—that together account for the majority of global cryptoasset and stablecoin activity. While the use cases described here relate to specific chains and layers of technical execution, the measurement challenges generalise to all blockchains and DeFi ecosystems relying on a similar technical underpinning, namely blockchains employing a UTXO architecture and those enabling smart contract programmability.

First, the aggregation challenge arises in Bitcoin-type blockchains, where transfer values and change outputs cannot be easily separated. Value is stored in discrete unspent transaction outputs (UTXO) rather than account balances. The design complicates the identification of change outputs, a technical artefact that artificially inflates transfer values. This has a cascade of measurement consequences. Trans-

actions must consume entire UTXOs and therefore frequently return change. Thus, on-chain transfer volumes inherently conflate economic transfers with technology-implied change outputs. We show that Bitcoin transfer values vary by up to a factor of six, depending on the methodology applied. Market capitalisation is similarly sensitive: accounting for lost coins or valuing supply at realised prices rather than current prices, markedly reduces headline figures in most time periods. In periods of price surges, simple measures of market capitalisation have risen to as much as four times the realised capitalisation.

Second, the programmability challenge stems from the smart contract layer, which enables complex, automated and programmable transactions. With Ethereum's Turing-complete design, which allows for highly flexible programming, any logic can be encoded in a smart contract. This freedom in programmability means that classifying and identifying relevant economic activities hinges on careful technical dissection at scale. While 54 million smart contracts remain outside the scope of technical classification, we successfully classify 13 million active contracts based on common technical standards. We find token proliferation to be vast and highly concentrated. There are 1.4 million smart contracts issuing tokens, while transfer and trading activity on decentralised exchanges is concentrated among a handful of them. In addition, we observe widespread symbol reuse: USDT appears across roughly 7,000 token contracts, mimicking the issuance contract of the largest stablecoin to date. Increased stablecoin activity appears as a driver of spurious contract activity from 2020 onwards.

Third, the comparability challenge arises from the same asset being issued and transferred across different chains. Comparing aggregate activity at the asset level can hinder analysis of use cases, if analysts do not consider the different use cases implied by different infrastructures. A narrow focus on the asset level often fails to account for the distinct economic functions of these assets across blockchains. One example is the multi-chain issuance of the same stablecoins across heterogeneous blockchain architectures. We show that USDT held by smart contracts on Ethereum has reached over 20% of total holdings, while the equivalent figure on Tron has remained markedly low at just around 1%. This is a consequence of architectural and fee structure differences affecting use cases across chains. USDT activity on Tron appears more related to smaller value transactions and hints at factors such as store-of-value use and crypto exchange holdings, while USDT activity on Ethereum seems more heavily linked to DeFi services such as liquidity provision and trading. Consequently, aggregating USDT activity across chains conflates different economic activities. This highlights the need for a nuanced approach considering the asset itself as well as the specific characteristics of the chains on which it operates for a clearer understanding of the roles these assets play within the broader ecosystem.

The three measurement challenges stem from three sources: (i) architectural ambiguity, due to uncertainty inherent in the protocol's execution semantics and, hence, data model; (ii) behavioural factors, introduced by how participants interact with protocols; and (iii) compositional noise, uncertainty arising from the high degree of freedom in application design that fragments measurement.

The paradox of DeFi measurement lies in the fact that, while ledger data is highly transparent, the architectural flexibility of programmable chains creates a volume and complexity of records that complicate economic interpretation. We propose using bounded estimates and a technical classification framework, combined with methodological clarity and an analysis of its impact based on granular data, to address some of the key measurement challenges in DeFi and provide a better understanding of these markets. We argue that sound metrics rest as much on technical knowledge of DeFi infrastructures as on careful analysis of economic behaviour that governs interactions with these applications.

Our analysis draws on blockchain data from Mercurius, a research initiative operated by De Nederlandsche Bank (DNB), which was developed together with the Bank for International Settlements (BIS) Innovation Hub and Deutsche Bundesbank. The platform ingests and structures full node data from Bitcoin, Ethereum and Tron, comprising a total of 100 billion records.

This paper is organised as follows. Section 2 provides an overview of the literature, focusing on the implications and measurement of stablecoins, cryptoassets and DeFi. Section 3 describes the technical features, data structure and measurement challenges of Bitcoin, Ethereum and Tron as illustrative examples for DeFi analysis. Section 4 presents key metrics on native cryptoasset transfers and smart contracts activity with a focus on stablecoin activity derived from granular blockchain data. Finally, Section 5 concludes and provides policy recommendations for central bankers.

2 Literature on DeFi

The emergence of cryptoassets and blockchain technology has spurred a growing body of research across economics, finance and computer science. However, few studies investigate how underlying data structures across blockchains affect economic analysis. Our paper bridges the burgeoning literature on technology of DeFi and empirical analysis leveraging blockchain data. We contribute to the literature on cryptoassets and the DeFi ecosystem by examining the foundational data and transaction mechanisms that affect empirical analysis. We show how the architectural setup of different chains affects analytical scope and the calibration of indicators.

The ecosystem has evolved considerably since the publication of the Bitcoin white paper by Nakamoto (2008) and the subsequent launch of the Bitcoin network in early 2009, which laid the foundation for cryptoassets by pioneering blockchain technology. The Ethereum white paper (Buterin, 2014) introduced a Turing-complete gas-metered execution environment for smart contracts, providing a programmable architecture that supports smart contracts and decentralised applications. Stablecoins are gaining importance in global finance, drawing attention to their issuance on platforms like Tron. The Tron blockchain (TRON-Foundation, 2017) employs a similar architecture to Ethereum while emphasising high transfer throughput.

Blockchain data, by design, is publicly available, providing a rich set of information for empirical analysis. Given its pioneering role, a number of studies, for example Ron and Shamir (2013) and Makarov and Schoar (2021) leverage Bitcoin's transaction history for quantitative analysis. Meiklejohn et al. (2016) use idiosyncratic technical features for address clustering in the Bitcoin network. Clustering is employed by Jahanshahloo et al. (2023) to identify activity and user type holdings. The paper also connects to the analysis of large amounts of blockchain data, for example Liu et al. (2023) who identify user types in the Bitcoin network based on transaction patterns and Guo et al. (2019) who describe the statistical properties of the Ethereum network. In this vein, Mao et al. (2024) highlight data extraction of Tron data and the network's role in supporting stablecoin transactions. The paper is also closely connected to the analysis of transaction-level data for stablecoins which we investigate in detail for the case of the largest stablecoin USDT on Ethereum and Tron. Kosse et al. (2026) offer a comprehensive overview and detailed examination of stablecoin transactions on the Ethereum blockchain.

Early studies on Bitcoin, such as Böhme et al. (2015), provide a comprehensive overview of Bitcoin's economic implications and technological features. Halaburda et al. (2022) outline the ecosystem and emerging literature on cryptoassets. Schilling and Uhlig (2019) discuss the economics of Bitcoin. Due to its speculative nature and wide price swings, research on Bitcoin often investigates price formation and volatility. Among others, Prat and Walter (2021) and Biais et al. (2023) develop structural models showing how miners' incentives, entry dynamics, and investor beliefs jointly determine Bitcoin's price. Auer et al. (2023a) find strong links between new user inflows and short-term Bitcoin price dynamics. Gandal et al. (2018) investigate price manipulation on crypto exchanges via suspicious trading activity. Bewaji et al. (2024) identify triangular arbitrage trading opportunities of major cryptoassets and stablecoins, highlighting market inefficiencies. Employing crypto exchange data, Cong et al. (2023) document that trading volumes are often vulnerable to manipulations such as wash trading.

Increasingly, macroeconomic factors are identified as factors influencing crypto markets. Di Casola et al. (2023) investigate cross-border Bitcoin trading and find that macro-financial conditions, and exchange-rate volatility significantly shape Bitcoin trading volumes. Auer et al. (2025) study a comprehensive dataset on cryptoasset and stablecoin cross-border flows and find they respond strongly to global funding conditions, while stablecoins also serve transactional and remittance demand. The evolving role of Bitcoin in relation to monetary policy decisions is documented by Karau (2023) and Pietrzak (2023) who show that Bitcoin systematically reacts to monetary policy shocks. Hofmann et al. (2026) show that stablecoins have implications for monetary control in the context of dollarisation and that traditional capital flow restrictions do not affect stablecoin flows. Aquilina et al. (2025) discuss financial stability implications posed by DeFi markets. ECB Crypto-Assets Task Force (2019) and Aldasoro et al. (2026) study the implications of cryptoassets for financial stability, monetary policy and the international currency hierarchy.

While a strand of literature has focused on Bitcoin, Liu and Tsyvinski (2021) show that returns on ma-

jor cryptoassets, are predictable using momentum and attention-based factors that drive crypto markets more broadly. Schär (2021) and Auer et al. (2023b) develop frameworks of DeFi architecture's technical backbone and financial applications such as decentralised exchanges, lending protocols and derivatives protocols. Aramonte et al. (2021) show that DeFi markets are highly concentrated. These dynamics also hint at the wider ecosystem developed around DeFi applications and stablecoins in particular.

Within the wider DeFi ecosystem, stablecoins have become of interest as they aim to avoid price swings via different mechanisms. The most relevant are asset-referenced stablecoins as they bridge the traditional financial system and DeFi. In times of stress, stablecoins may still be subject to price fluctuations due to their heterogeneous nature (Alasadi et al., 2026). Stablecoin adoption, studied by Altavilla et al. (2026), can erode deposit bases, increase wholesale funding dependence and weaken monetary-policy transmission. Azzimonti and Quadrini (2025) use a multi-country model showing stablecoins may increase global demand for dollar assets while also partially substituting them in cross-border settlement.

3 DeFi characteristics and data structure

Although blockchain networks differ in technical design and governance, they share foundational architectural principles: transactions are grouped into blocks, which are cryptographically linked in chronological order to form a tamper-resistant chain. At the same time, the transfer mode and complexity of different blockchains have implications for the scope of economic analysis.

This section describes the key architectural features of Bitcoin, Ethereum and Tron. It links those features to three measurement challenges: the aggregation challenge in Bitcoin (Section 3.1), the programmability and the comparability challenge in Ethereum (Section 3.2) and Tron (Section 3.3). We distinguish between the base layer features governing native cryptoasset transfers and the smart contract layer features enabling programmable applications, adopting a simplified version of the frameworks proposed by Schär (2021) and Auer et al. (2023b).

3.1 The challenge of aggregation: the case of Bitcoin

Unlike account-based transaction models, such as those used in traditional banking systems (e.g. wire transfers), the Bitcoin ledger operates on a UTXO model. Rather than storing balances at addresses, the ledger records discrete outputs which are essentially value chunks—in some ways akin to digital coins. Outputs can carry arbitrary values and are divisible by the smallest unit, so-called Satoshis, where one Bitcoin equals 100 million Satoshis.

As a concrete illustration of the UTXO model, Figure 1 presents a sequence of outputs created and spent in a chain of transactions. UTXOs are spent as inputs in a transaction and generate outputs (newly created UTXOs), which can serve as inputs for later transactions. Each output is assigned a unique identifier and is associated with a transaction, a block and receiving address. Balances are not stored directly but can be inferred via aggregating UTXOs across addresses. The structure enables the analysis of value flows across addresses by examining the relationships between transaction inputs and outputs. For instance, when multiple input addresses fund a single transaction, these addresses are arguably controlled by the same entity.[1] This architectural feature enables address clustering heuristics (see for example, Meiklejohn et al., 2016 and Moeser and Narayanan, 2022).

The UTXO structure shapes Bitcoin's underlying data model. At the core are **blocks**, which serve as containers for all transaction data and include metadata such as the sequential block number and a timestamp. Blocks typically contain multiple **transactions** with information on input and output values. **Outputs** are typically associated with **addresses** to indicate ownership.[2] These components and their relationships provide a granular dataset for the analysis of network activity, transaction flows

[1] One notable exception are so-called "Coinjoins", a technique that combines inputs from multiple users into a single transaction. In addition, there are mixers, third parties which pool funds to obfuscate their origins.

[2] Addresses are shortened, hashed versions of public keys. UTXOs are typically associated with public keys (more precisely, public key hashes), meaning that only the holder of the corresponding private keys can authorise its spending by signing a transaction.

Figure 1: UTXO model—illustrative transaction paths

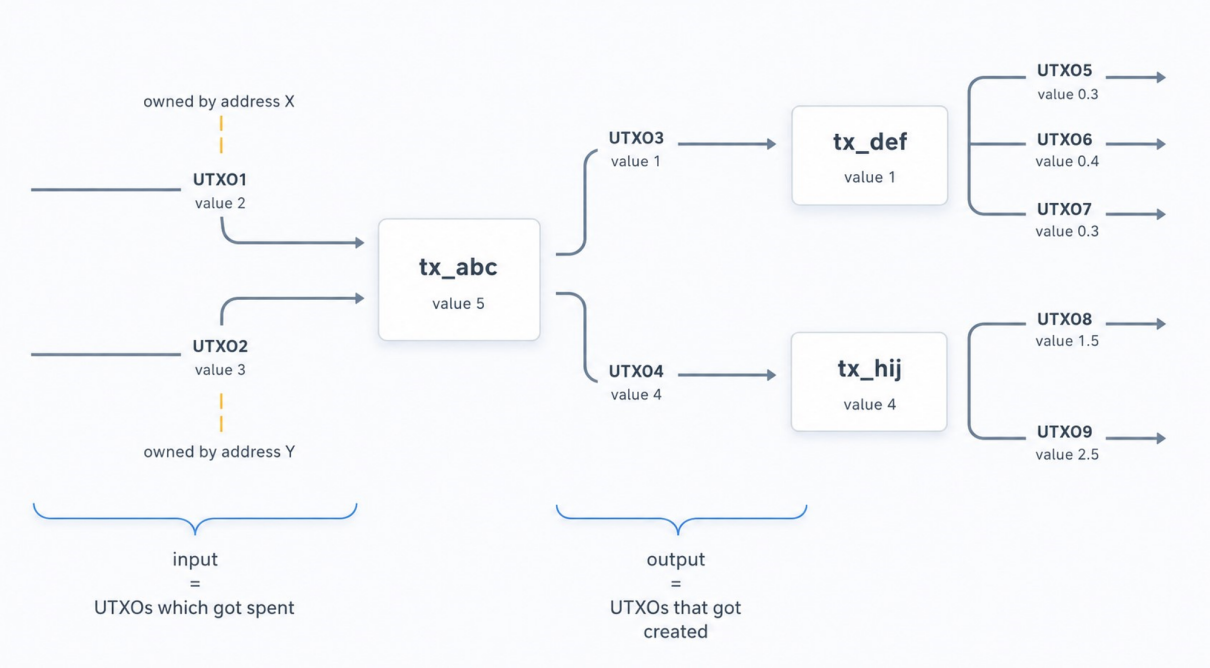

Note: The figure demonstrates the UTXO (Unspent Transaction Output) model, showing how inputs (UTXOs being spent) are consumed in a transaction (e.g., transaction tx abc) to generate outputs (newly created UTXOs), which can later serve as inputs for subsequent transactions (e.g., tx def and tx hij).

and economic behaviour within the Bitcoin ecosystem. Data are linked via identifiers, as analysis typically involves combining information across dimensions.

An output is "unspent" if it has not yet been used as an input in a subsequent transaction. The sum of all unspent outputs at any given time reflects the total supply of coins. Each transaction consumes one or more UTXOs as inputs and creates new UTXOs as outputs.[3]

A fundamental feature of the UTXO model is that the entire value of each transaction input must be consumed, leading to the aggregation challenge: if a UTXO worth 4 BTC is used to transfer 1.5 BTC, the transaction creates two outputs, one of 1.5 BTC to the recipient and one returning 2.5 BTC to the sender as change. Critically, the change output is not identifiable from the transaction data alone unless implied by the fact that change is returned to the same sender address, referred to as address reuse. Nakamoto (2008) recommends generating a new key pair for each transaction precisely to prevent linking of addresses, which has the side effect of making change outputs indistinguishable from the intended transfer. As a consequence, the transaction would only show the input and outputs without a way to distinguish transfer value from change.

This design choice has direct measurement consequences. Aggregate transaction volumes calculated from UTXO data include change outputs that do not represent economic transfers (Cole et al., 2022). Identifying change outputs requires the use of heuristics based on assumptions about user behaviour that can significantly alter estimates. One case that ameliorates the aggregation challenge arises when change is returned to a sender address. This practice has become increasingly common, as discussed in Section 4. A further measurement implication of the UTXO model is that it permits analysis of the age and velocity of individual outputs. This type of breakdown is not available to the same degree in account-based systems that pool balances.

[3]The value of transaction outputs equals the value of total inputs minus the transaction fee. These transaction fees form part of miner rewards. Miners validate new blocks under Bitcoin's proof-of-work consensus mechanism and receive the fees paid on all transactions in the block and newly generated Bitcoin via a "coinbase" transaction, the first transactions in a block.

3.2 The programmability challenge: the case of Ethereum

The Ethereum chain employs an account-based model in which addresses hold balances in the native asset, Ether. Transactions update these account balances directly as highlighted on the left side of Figure 2, avoiding the change-output ambiguity inherent to UTXO systems. Transfers in Ether include a sender, receiver and other transaction details such as transfer value, fees and block number. Ethereum charges fees based on computational work times the price. Computational work is measured in a unit called gas. The gas price times the gas used determines the transaction fee. The gas price reflects factors like network demand and congestion. Users can set higher maximum fees to prioritise their transactions.

The distinguishing feature of Ethereum is its Turing-complete smart contract layer. Besides externally owned accounts (EOAs) that hold Ether balances, smart contract accounts (SMAs) can hold Ether balances and store executable code. The code stored at smart contract addresses can, in theory, execute any type of computation (Turing complete) given sufficient computational resources.

When a transaction calls an SMA, the Ethereum Virtual Machine (EVM) executes the specified functions using the parameters from the transaction input data, as shown in the right panel of Figure 2. In addition, smart contracts can interact with each other, enabling composable and interconnected applications. A single transaction may trigger multiple smart contracts, generating multiple on-chain events. Thus, smart contracts can facilitate conditional and programmable transactions, enabling a wide array of DeFi applications that offer financial services across multiple DeFi protocols, akin to "financial lego" (for example, Auer et al., 2023b). Such applications include lending platforms, decentralised exchanges and token issuance.

This flexibility in programmability complicates the analysis of smart contracts and introduces fundamental measurement challenges. Since an SMA can store arbitrary execution logic, the economic signal may be buried in execution noise. Therefore, deciphering transaction logic and economic meaning can be hampered by complexity as a direct result of freedom in programmability. In addition, the same economic activity may appear across different implementations and different data objects such as transaction inputs, event logs or execution traces.

Community standards help reduce this fragmentation and facilitate smart contract analysis at scale. The ERC-20 (Ethereum Request for Comment 20)[4] standard for fungible tokens defines a common interface that includes functions for transferring tokens, approving allowances and retrieving balances. This standardisation facilitates systematic analysis across tokens issued via smart contracts.[5] Major stablecoins including USDT and USDC follow this standard. However, adherence is voluntary, and smart contracts may adopt ERC-20 function signatures without implementing their intended semantics.

Tokens are not inherently valuable; their on-chain issuance may be unrelated to any economic rationale. Instead, they can represent a diverse range of value propositions, such as asset-backed stablecoins or tokenised (wrapped) versions of native cryptoassets that facilitate the integration of native cryptoassets in DeFi protocols. Token rights, value and utility usually depend on off-chain arrangements. For instance, although issuance is visible on-chain, fiat-backed stablecoins are backed by off-chain custodial reserves.

Native asset transfers and smart contracts data is captured in various parts of transaction processes. Transactions can include transfer of Ether but also smart contract interactions. To invoke a smart contract, the transaction's receiver is the contract's address, and the input data field encodes the function selector and arguments (e.g., amount and recipient for a token transfer). The contract code at the receiver address executes the requested operation and subsequent operations depending on the smart contract logic.

After transaction processing and completion of execution by the EVM, transaction data are visible in receipts. Transaction receipts include the status (failure/success) and transaction details such as the block number the transaction was included in, input data and the Ether amount transacted. Transaction

[4]ERC stands for Ethereum Request for Comment. ERCs are a class of Ethereum Improvement Proposals (EIP) that specify application-level interfaces and behaviours, enabling contracts, wallets and apps to interoperate.

[5]Non-fungible tokens (NFTs) mostly follow ERC-721. For simplicity, we use fungible tokens and tokens interchangeably and refer to NFTs explicitly.

Figure 2: Native and token transfer on Ethereum

(a) Native ETH transaction (b) Example token transaction via smart contract

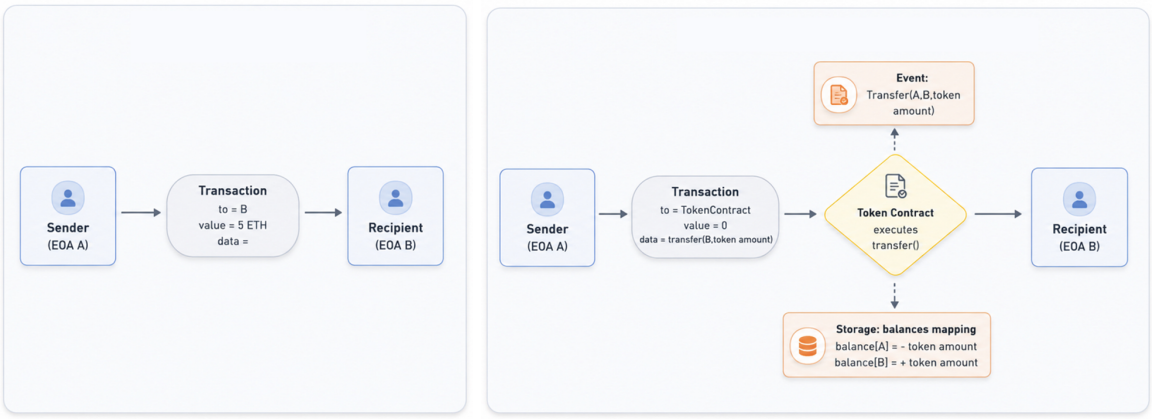

Note: The left panel illustrates the transfer of a native cryptoasset (Ether) between two externally owned accounts (EOAs). In contrast, the right panel depicts the example of an EOA initiating a token transfer by sending inputs to a smart contract account (SMA), which contains the logic governing the token transfer. The token receiver is provided to the SMA through the input data. Upon a successful transfer, the SMA updates its storage to reflect the new token balances. Token transfers can also be implemented in a different way, depending on the smart contract's logic.

receipts also hold an array of all emitted logs of the transaction. Logs capture events emitted by smart contracts, with their presence and structure determined by the contract's code specification. They are commonly used in DeFi analytics and for integration purposes, for example service providers using logs to update token balances. Community standards like ERC-20 or ERC-721 facilitate reproducible and extensive analysis across DeFi protocols. Logs are non-authoritative and may be incomplete, inconsistent or misleading. Nevertheless, code transparency and verifiability substantially constrain manipulation and developers often have a natural incentive to provide comprehensive logs for verification and promoting usage.

In contrast to smart contract-defined data, traces capture the execution steps. This includes invoking a smart contract function and any downstream calls to other contracts. Analysing traces often requires deciphering bytecode and validating events through direct state differences (state diffs). State diffs record the net effect of a transaction on account balances and smart contract data (e.g. token holdings). Unlike logs, which are defined by developers, state diffs reflect the actual on-chain outcome that is observable through the net change in token balances. While traces provide a granular view of execution, this granularity comes at the cost of significant complexity in data extraction and analysis, highlighting the challenges of systematically measuring programmable systems. Although programmability introduces complexity and the risk of data misspecifications, standardised log data offers a key advantage by enabling straightforward and scalable measurement verifiable via traces data.

3.3 The comparability challenge: the case of Tron

Tron adopts an account-based model and supports smart contract execution, similar to Ethereum. In its aim for EVM compatibility, Tron closely mirrors Ethereum's design and implementation, reflecting similar programmability challenges in measurement. At the same time, Tron diverges in several key aspects. The chain is designed to offer high throughput and low-cost transaction processing. Like Ethereum, all state-changing events are initiated by transactions, the fundamental units through which data are written to the distributed ledger. Tron transactions follow a similar structure, including a unique transaction ID, sender and recipient addresses, transfer value and execution metadata. However, rather than using gas fees, Tron employs a dual-resource system consisting of bandwidth and energy. Regular transactions consume bandwidth, while smart contract executions consume energy. Both resources can be obtained by staking the native TRX tokens—which locks up tokens in exchange for access to network capacity and reduces reliance on variable fees. While this structure may effectively lead to lower fees for some users,

fees depend on network congestion and user behaviour.

Economically, these design differences imply that ostensibly identical tokens, issued on different chains, serve different use cases. While the nature of the token may be identical, the underlying infrastructure on which it resides may create a comparability challenge. On Tron, the fee regime could make high-frequency, low-value stablecoin transfers suitable while Ethereum's variable fees and rich smart contract ecosystem may attract different transaction types. Aggregating token activity across chains thus could conflate payment-like turnover on Tron with liquidity provision and collateralisation on Ethereum. While cross-chain analytics is essential for comprehensive analysis, disaggregating data by technical architecture can provide valuable insights into specific use cases.

The use cases of smart contracts may also vary based on architectural differences. A distinctive feature of the Tron ecosystem compared to Ethereum is the TRC-10 token standard. TRC-10 tokens are issued and managed by native system contracts, without requiring execution within the TVM (Tron Virtual Machine). System contracts provide predefined functionality determined by the protocol level. Any account may issue a TRC-10 token upon payment of an issuance fee. During issuance, the creator defines a set of core attributes, including the token name, total supply, exchange rate relative to native TRX, circulation period, description and staking parameters. The data extraction logic remains similar to Ethereum.

In terms of measurement, the Tron protocol explicitly classifies transactions by their operational role. The network currently supports approximately 60 distinct transaction types. Among these, only two—smart contract creation and smart contract invocation—are executed by the TVM and directly associated with user-deployed smart contracts. Other transaction types invoke built-in system functionalities such as account management, token issuance, voting and resource delegation These functional differences change the scope of the programmability challenge due to design and actual usage.

4 Results

To identify and address the aforementioned measurement challenges, we employ a simplified classification for DeFi data layers, based on Schär (2021) and Auer et al. (2023b). Bitcoin operates with a single-layer architecture dedicated to native asset transfers (base layer). In contrast, Ethereum and similar blockchain networks adopt a more complex two-layer architecture. The base layer facilitates native cryptoasset transfers while simultaneously triggering the smart contract layer that enables DeFi applications through smart contract functionality. The smart contract data layer encompasses a broad spectrum of activities, including lending pool operations and token issuance.

We dissect smart contract activity and stablecoin transfers by leveraging a structured data set from project Mercurius that encompasses 100 billion records spanning various dimensions, including output data and block-level information.[6] The significant volume of transactions across Bitcoin, Ethereum, and Tron underscores their activity levels, while also illustrating the functional distinction between base layer and smart contract layer, as depicted in Figure 3. Although Bitcoin boasts the highest overall market capitalisation and has processed more than one billion transactions, Ethereum records more than twice as many transactions, and Tron nearly ten times as many. Furthermore, records related to smart contract activity on Ethereum and Tron collectively amount to more than 30 billion entries. Careful interpretation of these figures is warranted, given the overlap between base layer and smart contract layer data, as well as within data at the smart contract level. Despite these overlaps, the magnitude of smart contract data underscores the rapid evolution of the ecosystem, moving beyond the simple transfer of native cryptoassets and reflecting the increasing complexity and functionality across blockchain networks.

Within DeFi applications, the data cover the majority of stablecoin activity. The most widely used stablecoins are pegged to fiat currencies and, so far, are predominantly tied to the US dollar. Tether

[6]Data includes transactions and master data such as addresses. Note that certain data points are stored multiple times across various stages of the data processing pipeline. The processing framework adheres to a medallion architecture, wherein the *bronze* layer retains raw, unaltered data in tables, the *silver* layer processes the data to facilitate analysis and the *gold* layer consolidates the information into aggregated indicators. The number therefore indicates total data points, but not a count of distinct blockchain events or empirical observations.

Figure 3: Overview of data records (in billion)

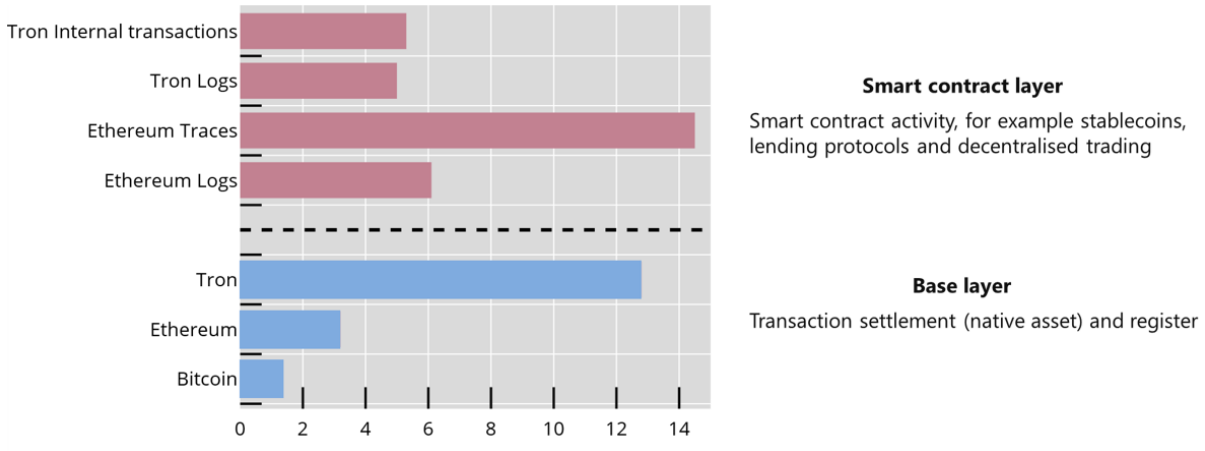

Note: There is overlap between Ethereum and Tron data across the base layer and the smart contract layer, as well as within smart contract layer data. Source: Authors' calculations.

(USDT) has the largest market capitalisation, followed by Circle's USDC, albeit with a significant gap. Issuance is highly concentrated across Ethereum and Tron (Figure 4). Our analysis leverages transaction-level data from Ethereum and Tron, capturing 82% of the stablecoin market with near-complete coverage of USDT. Reflecting the high concentration within DeFi, USDT's total supply of roughly USD 180 billion is issued almost exclusively on Ethereum and Tron, with other blockchains accounting for only close to 3% of its circulation. Similarly, the issuance of USDC, with a supply of USD 75 billion, is predominantly issued on Ethereum, which accounts for nearly two-thirds of its total issuance. By comparison, the issuance of other stablecoins, such as Sky Dollar (USDS) with a supply of around USD 8 billion, is significantly smaller.

Data across multiple blockchains highlight the challenges of comparability in measurement, as transfers of the same stablecoin on different infrastructures enable abstracting asset-specific use cases from the underlying technological framework.

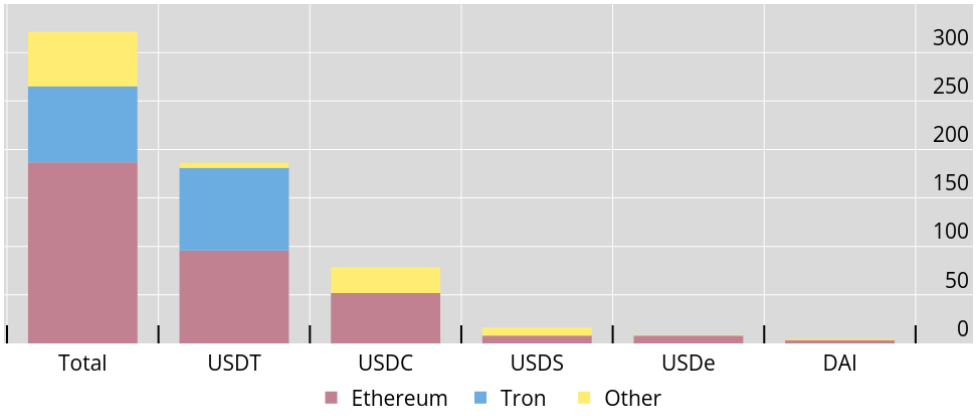

Figure 4: Stablecoin supply by blockchain (in billions of USD)

Sources: Tether Transparency, Circle Transparency, and DefiLlama (snapshot 16 April 2026); authors' calculations.

4.1 Bitcoin metrics

We start by examining the impact of different methodological choices on indicators of Bitcoin, highlighting the aggregation challenge in measurement. We leverage data covering the period from 2009 to 2026. The

dataset comprises approximately 1.3 billion transactions and a total of 3.6 billion transaction outputs.

First, we analyse transfer values in the Bitcoin network and the role of change; see Section 4.1.1. Next, we examine measures of illiquid coins, with prolonged dormancy serving as a proxy for lost coins; see Section 4.1.2. Finally, we leverage characteristics of the UTXO structure to estimate different measures of Bitcoin market capitalisation; see Section 4.1.3.

4.1.1 Transaction values

Typically, a Bitcoin transaction consists of at least one output directed to the recipient and another output returned as change to the sender. Heuristics aim to identify the change output in a Bitcoin transaction. We investigate boundary assumptions by estimating an upper bound, a lower bound and an adjusted value for transfer volumes, employing relatively simple methodological approaches. The upper bound represents the total value of transfers within the network without any adjustments to the data, calculated as the sum of all UTXOs of transactions i at time t:

$$Total\ Volume_t = \sum_{i=1}^{n} UTXO_{i,t} \tag{1}$$

The naive underlying assumption posits that no change outputs exist in transactions, thereby treating all outputs as indicative of transfer activity. A more meaningful approach excludes self-transfers—transactions with outputs returned to a sending addresses—to derive an adjusted estimate. Self-transfers are most likely change from a transaction and unlikely to hold economic significance by themselves. However, self-transfers allow separating transfer values from change outputs. The adjusted transfer volume at time t calculates the total UTXOs of transaction i, minus the self-transfers across n transactions:

$$Adjusted\ Bound\ Volume_t = \sum_{i=1}^{n} (UTXO_{i,t} - SelftransferUTXO_{i,t}) \tag{2}$$

Given the uncertainty in identifying change, we estimate a naive lower bound using a (over)simplistic change heuristic. The approach excludes either the self-transfer output or the highest output of a transaction, based on the assumption that actual transfer values represent the smallest output(s) of a transaction, thereby strictly avoiding overestimation.[7] For each transaction i, either the value of the output related to a self-transfer or the UTXO with the highest value is treated as change:

$$Lower\ Bound\ Volume_t = \sum_{i=1}^{n} (UTXO_{i,t} - (SelftransferUTXO_{i,t} + I_{i,t} \cdot maxUTXO_{i,t})) \tag{3}$$

where $I_{i,t}$ is an indicator function that equals 1 if there are no self transfers in a transaction, and 0 otherwise. The underlying assumptions heavily impact the aggregate measure of on-chain transactions. Figure 5 illustrates that on-chain flows excluding self-transfers substantially reduce flows, up to a factor of six. This is striking, given that the upper bound is widely reported as the transfer value. The significant reduction in transfer values caused by self-transfers reflects both the characteristics of the Bitcoin protocol and user behaviour. The value excluded due to self-transfers increases markedly from March 2016 onwards, driven by a rise in address reuse. In the early stages of Bitcoin adoption, users appeared to prioritise privacy and security over convenience. However, as Bitcoin gained broader appeal among casual users and service providers, there has likely been a shift in priorities, with privacy concerns

[7] Note that in special cases like Coinjoins, there could be more than one change address. To account for this, only the lowest transaction output could be considered. This alternative specification barely changes results. The estimated lower bound could even be further adjusted. For instance, analysis might exclude transactions that involve mixer services, spam activities or intermediary transfers within a specific time frame. Given the uncertainty in identifying such transactions, we refrain from further adjustments other than those described above. Due to these exceptions, our lower bound should be regarded as a conservative heuristic rather than a mathematically guaranteed bound.

Figure 5: Bounded transaction value estimates

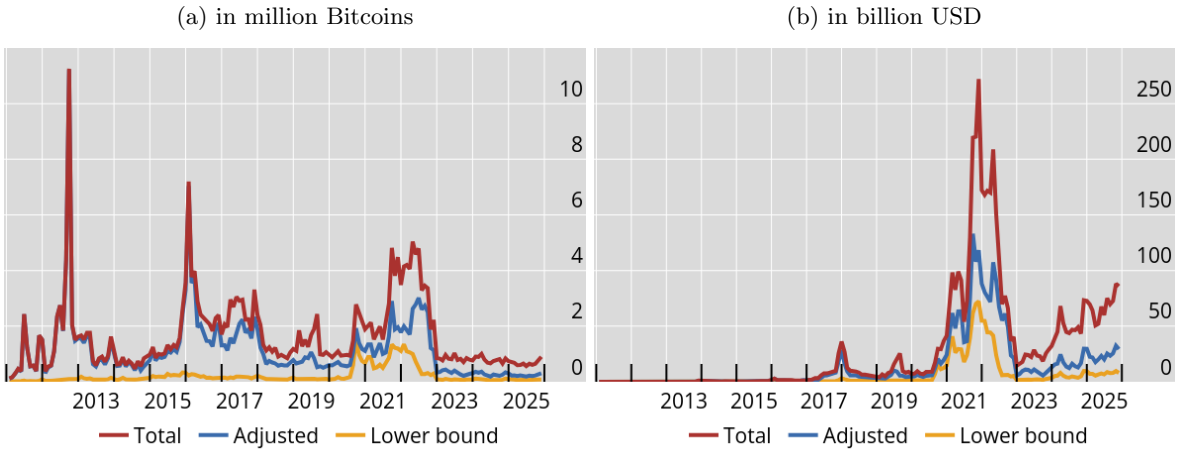

Note: The figures show monthly averages of daily transaction values. The left panel depicts transferred Bitcoins in millions, the right panel shows transferred Bitcoins in USD. The total value consists of all on-chain transactions on a given day. The adjusted value excludes self-transfers, i.e. outputs returned to the sender address. The lower bound value either subtracts self-transfers (if applicable) or removes the highest output of each transaction instead.

giving way to convenience. While some users value a higher level of privacy, crypto service providers may even deliberately opt to signal ownership rather than obfuscate transactions. Interestingly, the lower bound remains quite stable at a low level over time, likely reflecting the overly cautious assumptions in identifying change.

4.1.2 Lost and illiquid coins

Beyond the challenge of identifying transaction values with certainty, the loss of private keys—and thus access to Bitcoins—complicates metrics like market capitalisation and the stock of usable coins. In contrast to the traditional financial system, recovering lost access to self-hosted wallets is not possible. While dormant assets do not imply loss of access in traditional banking, Bitcoins that are dormant for an extended period of time often reflect the fact that a user or entity lost access. Although the exact amount of lost coins is unknown, we use time-based assumptions to identify UTXOs where private keys are likely permanently lost. Besides, dust coins which are UTXOs smaller than transaction fees are uneconomical to transfer. The Annex provides a brief overview of dust coins. Given the relatively small amount and the dependence on network fees, we focus on lost coins.

To gauge lost coins, we first examine supply age and output-velocity (lifespan) distributions. At any point in time, the total supply of Bitcoin is equal to the sum of all UTXOs. The UTXO age is defined as the time an unspent transaction output has remained unspent. To quantify the velocity of transaction outputs, we measure the time interval between their creation and spending. This metric applies exclusively to spent outputs, as unspent outputs have no defined velocity (yet).

The current UTXO age and velocity of outputs are depicted in Figure 6. The largest fraction of supply, over 8 million Bitcoin, was moved less than a year ago. At the same time, the majority of Bitcoin supply was moved over one year ago, with 6 million Bitcoin moving over 5 years ago. Around 3.5 million Bitcoin have remained dormant for over 10 years and 1.8 million Bitcoin have not been moved for over 15 years. The anonymous creator(s) of Bitcoin is believed to hold over 1 million of these. Given that in most cases they have not been moved since they were mined, we consider them lost.

Due to Bitcoin's asymptotic supply cap, with a halving of block rewards roughly every 4 years, relatively few, high-value outputs drive supply age. The velocity of outputs indicates that the majority of outputs are spent quickly. Specifically, 99% of the outputs were spent within one year. The majority of the daily transfer values consist of outputs moved on the same day or one day after creation. Despite these trends, a small fraction, i.e. roughly 24 thousand Bitcoin, was spent after being dormant for more

than 10 years, with almost 3,000 Bitcoin being spent after being dormant for 15 years. This underscores the uncertainty of time-based heuristics to identify lost coins. Even so, UTXOs older than 15 years rarely move, accounting for only about 0.1% of the supply older than 15 years.

Figure 6: Supply age and output velocity (percent of total)

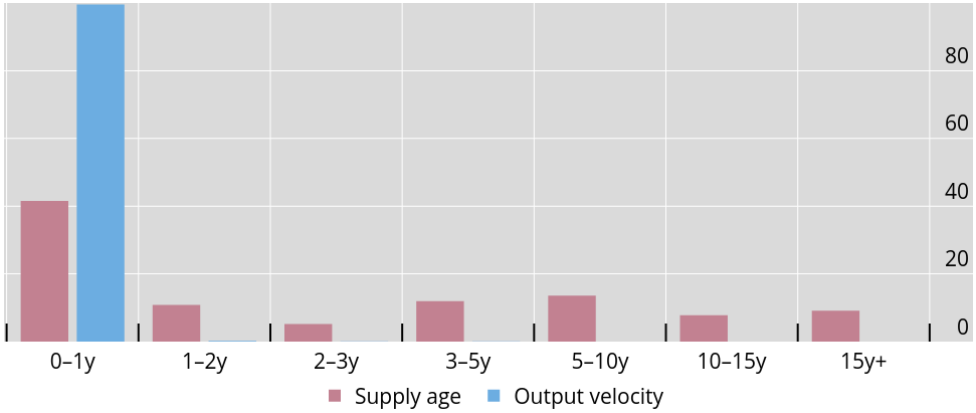

Note: The figure shows the composition of the Bitcoin supply by age bands. The velocity of output values in the Bitcoin network is calculated as the time interval between the creation and spending of an output. The supply age includes unspent transaction outputs, whereas velocity only includes spent outputs. Percentages are calculated by age band's share of total, measured in Bitcoins. Total supply is roughly 20 million Bitcoins, total outputs spent over 9 billion Bitcoin. Data as of March 2026.

4.1.3 Market capitalisation

Building on the findings, we qualify the simple measure of market capitalisation with two adjustments. Often-reported figures calculate overall market capitalisation at time t ($MC_{total,t}$) by multiplying total supply, meaning all UTXOs, at the current price:

$$MC_{total,t} = \sum_{i=1}^{n} UTXO_{i,t} \cdot Price_t \tag{4}$$

To reflect likely lost coins, we calculate market capitalisation without outputs older than 15 years.[8] This amounts to roughly over 1.8 million Bitcoin which do not factor into market capitalisation.

$$MC_{excl\_LostCoins,t} = \left(\sum_{i=1}^{n} UTXO_{i,t} - \text{LostCoins} \right) \cdot Price_t \tag{5}$$

This reflects a hindsight adjustment, identified only by looking back to historical cut-off points, and it may change if coins move again,

As second adjustment, we use an alternate measure which implicitly accounts for UTXO age. It serves neither as a substitute nor a strict lower bound for current market value. This realised market capitalisation (RMC) is defined by valuing each output i at the price of its creation c rather than the current price, meaning when the UTXO was last spent in a transaction.

$$RMC_t = \sum_{i=1}^{n} UTXO_{i,t} \cdot Price_{c_i} \tag{6}$$

Realised market cap reflects price swings and is continuously updated by transaction flows, unlike a uniform spot-price valuation. Outputs that were created early but not spent are weighted at a price

[8] However, excluding Bitcoin not moved for 15 years likely represents a lower bound estimate since the loss of private keys presents a frequent issue.(for example, Badev and Chen, 2014).

significantly lower than in recent years. Therefore, older supply have little effect on the realised cap. Conversely, outputs last moved at a higher price do not decrease market capitalisation when prices go down.

Realised market cap typically lies below the simple market capitalisation measure due to an overall increase in price levels and the decrease of new Bitcoin supply minted in each block. Assumed lost coins, created shortly after the inception of the network and not moved since, factor in at low prices. However, when outputs were spent at higher prices and are not moved when prices decrease, realised market cap is higher than the simple measure.

The three measures are displayed in Figure 7, showcasing a range of estimates. The overall market cap shows highly volatile growth over time, with peaks during 2017, 2021 and starting in 2024. The realised market cap lies well below the simple cap measure most of the time, in particular during times of price surges. The simple cap is up to four times the realised cap. The realised market cap grows steadily over time, showing a smoother trajectory compared to the other metrics. At the same time, realised market cap lies above the overall market cap when Bitcoin prices fell steeply in 2022. This pattern is driven by holder behaviour, rather than short-term price volatility.

The analysis shows that the measures are highly sensitive to the underlying assumptions and methodologies. Rather than relying on single figures, a range of estimates offers better measures of economic significance and trends.

Figure 7: Measures of market capitalisation (in billion USD)

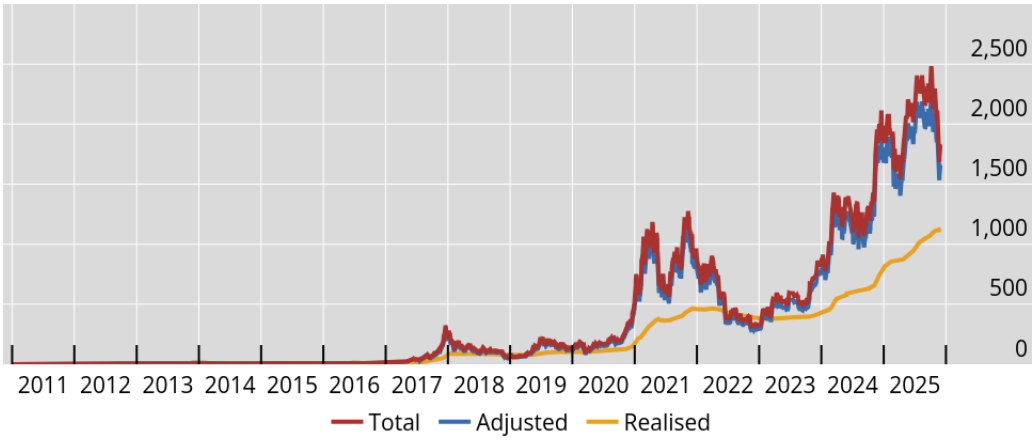

Note: The graph depicts different measures of market capitalisation in billion USD at a daily frequency. Total market capitalisation is the supply of Bitcoin at a given point in time multiplied by the current price. Adjusted market capitalisation deducts lost coins that were not moved in over 15 years and are considered lost over the full time period. The realised capitalisation smooths volatility by valuing coins at the price when they were last moved.

4.2 DeFi metrics

Next, we turn to the programmability and comparability challenges. We exploit the smart contract layer, which underpins a wide range of financial and non-financial applications. In particular, we focus on stablecoins—due to their potentially significant economic implications.

In order to provide an overview of deployed smart contracts, we begin by classifying them based on bytecode and adherence to community standards (see Section 4.2.1). This serves as the starting point for deeper analysis of smart contract types. Next, we leverage text analysis on contracts issuing tokens to uncover the dominance of liquidity pool tokens and highlight the challenges of name-based classification due to imitation and noise (see Section 4.2.2). Using granular decentralised exchange trading data, we demonstrate how economic activity in DeFi is heavily concentrated around stablecoins which act as a gravitational centre for the ecosystem (see Section 4.2.3). We then analyse stablecoin holdings by account type across Ethereum and Tron, revealing that the same stablecoin can serve distinct economic functions

on different chains—distinctions that aggregate stablecoin metrics often obscure (see Section 4.3). Finally, we document the proliferation of spurious stablecoin contracts, shedding light on the dynamics that complicate on-chain measurement frameworks (see Section 4.2.4).

4.2.1 Smart contract classification

While the underlying logic of smart contracts is highly flexible, allowing developers to implement virtually any functionality, community-driven standards have emerged to ensure interoperability and consistency across the ecosystem. These standards facilitate the technical classification of smart contracts.[9] Ethereum Improvement Proposals (EIPs) include ERC-20 for fungible tokens and ERC-721 for non-fungible tokens (NFTs).[10] Beyond token standards, proxy contracts have emerged as a prominent architectural feature. A proxy contract holds state and delegates calls to a separate implementation contract, which contains the executable logic. This design enables upgrades and bug fixes without altering the contract address or requiring state migration.

Deployment of smart contracts occurs in two ways: by sending a transaction containing compiled bytecode with an empty recipient field, or internally via factory contracts that invoke CREATE opcodes. The former is observable as a transaction, while the latter appears only as an internal call within execution traces. We identify smart contract deployments by inspecting execution traces for both null-address transactions and internal CREATE/CREATE2 operations, which together capture the full population of deployed contracts.

By applying a systematic technical identification approach, we classify smart contracts, enabling a more granular understanding of their role and functionality within the Ethereum ecosystem. The classification of smart contract types is performed deterministically using bytecode. A contract is labelled as ERC-20 if its deployed bytecode contains a set of functions defined by the standard. Similarly, a contract is classified as ERC-721 if functions corresponding to ERC-721 standards are detected. The proxy contract category matches one of several standards, such as the ERC-1967 Transparent Upgradeable Proxy Standard, the ERC-1167 Minimal Proxy Standard or the ERC-1967 Beacon Proxy Standard.

The resulting classification of smart contracts on Ethereum in Table 1 highlights that most deployed contracts do not adhere to the predefined technical standards. Out of a total of 67.5 million deployed and active contracts on Ethereum, over 54 million remain uncategorised. Proxy contracts make up the largest classified category, with close to 12 million contracts, while fungible tokens account for 1.4 million and NFTs for 0.1 million contracts. This underscores the programmability challenge in measurement with most smart contracts exhibiting unclear use cases based on technical standards.

At the same time, the volume of ERC-20 tokens—including major stablecoins—highlights the rapid proliferation of token creation. Within these technical benchmarks we leverage softer criteria, such as naming conventions to derive economic meaning.

4.2.2 Token text analysis

ERC-20 tokens carry optional metadata for name, symbol and decimals.[11] We apply text analysis to token names and symbols to map the token landscape based on metadata.

Word frequencies in Table 2 show that liquidity pool tokens from major decentralised exchanges (DEXs) dominate the corpus. The symbol `uni-v2` alone appears in roughly one third of all token contracts, together with names linked to PancakeSwap and Sushiswap. This pattern reflects the logic of automated market makers: each token pair is traded in a dedicated pool. Liquidity providers receive

[9] Adherence to standards does not guarantee the functional or technical soundness of a smart contract.

[10] On the Tron blockchain, smart contract logic follows a similar logic. Tron supports Solidity as well as other programming languages, enabling developers to create contracts with functionality analogous to Ethereum's standards. In particular, Tron's TRC-20 and TRC-721 standards mirror Ethereum's ERC-20 and ERC-721, providing compatibility for fungible and non-fungible tokens, respectively.

[11] Decimals define the smallest divisible unit of a token, determining how many decimal places it can be divided into—analogous to how the US dollar divides into cents. The default is 18.

Table 1: Technical classification of smart contracts on Ethereum (in millions)

Category	**Active**	**Total**
Non-fungible tokens (ERC-721)	0.1	0.1
Fungible tokens (ERC-20)	1.4	1.5
Proxy contracts	11.8	11.9
Uncategorised contracts	54.2	78.5
Total deployed contracts	**67.5**	**92.0**

Note: Categorisation is based on identified contracts from deployment transactions and the inclusion of functions defined by the respective standards in the bytecode. Data as of March 2026.

tokens representing their share in these pools. Because liquidity pool tokens follow uniform naming conventions, their counts proxy the number of traded token pairs on a DEX.

At the same time, metadata are an imperfect basis for classification. There is no central registry for token names or symbols and any contract can adopt any label. Generic terms such as *token*, *coin* and *protocol*, alongside incentive-oriented words like *claims* and *rewards*, appear frequently. Frequent terms such as *inu* and *ai* align with popular thematic trends and speculative narratives, suggesting that token naming conventions may be influenced by prevailing market sentiment or promotional strategies rather than substantive economic function. This adds noise to name-based classification and inflates counts of economically distinct assets.

A further challenge arises from imitation. Frequently used words cluster around well-known stablecoins such as USDT, USDC and DAI, even though the authentic versions are issued from single contracts. Contracts may even replicate bytecode, making imitations technically distinguishable only by their unique address. Consequently, name- and symbol-based metrics overstate the variety of economically relevant tokens and blur the boundary between actual and potentially spurious deployments. Token naming can be used maliciously to mimic real tokens and in some cases these tokens likely relate to scamming. However, symbol reuse alone establishes neither economic irrelevance nor deceptive intent. These findings underscore the measurement challenge of programmability, as token metadata reflects DeFi's dynamism and fragmentation that hinder systematic classification and reliable issuer identification.

From a regulator's perspective, identifying assets to monitor should rely on high activity rather than token symbols and names. For effective regulation, linking the assets in public networks to the legal entities that issue them is essential. This linkage needs to rely on off-chain information.

Table 2: Top 15 most used names, symbols and words in ERC-20 tokens

Token Symbol	**N**	**Token Name**	**N**	**Word**	**N**
uni-v2	465,576	uniswap v2	465,535	v2	467,974
usdt	6,867	pancake lps	6,725	uniswap	466,231
cake-lp	6,730	sushiswap lp token	4,432	token	61,725
slp	4,704	balancer pool token	3,313	coin	25,683
bpt	3,401	invisicoin	3,253	inu	22,273
inv	3,326	tether usd	2,793	ai	14,283
test	2,914	test	1,469	pepe	13,072
eth	2,914	jh	1,469	claim	9,880
pepe	1,884	thetoken	1,423	lp	9,758
$	1,557	$	1,420	rewards	9,749
ee	1,522	usdt	1,408	website	9,615
usdc	1,512	vista	1,089	lps	7,277
dai	1,302	ethereum	1,025	pancake	6,757
vista-lp	1,073	shibaswap lp token	1,025	protocol	6,428
sslp	1,044	tether	1,023	usd	6,414

4.2.3 Stablecoins and decentralised exchange trading activity

Text analysis shows that tokens from major DEX platforms dominate ERC-20 metadata. Based on this result, we extract liquidity pool data from the Uniswap V2 LP protocol and measure which assets are traded most intensively in these pools. Uniswap is an automated market maker where users trade against liquidity pools containing token pairs rather than an order book. Thus, swap counts provide a clean proxy for how intensively pools are used.

Uniswap V2 swap activity is highly concentrated and anchored by stablecoin markets. Rankings in Figure 8 indicate that WETH accounts for most swaps and the most active counterparts are stablecoins. WETH likely plays a prominent role as the tokenised form of Ethereum's native cryptoasset that facilitates use in DeFi protocols.[12] Swap activity is concentrated in a few pools—most notably the pairs WETH–USDC and WETH–USDT. This links the metadata patterns to economic use: liquidity pool tokens proliferate on DEXs based on the number of available pairs, yet trading and price discovery are anchored in a small set of pools. The pattern highlights the role of stablecoins as the primary quote and settlement assets. It also implies that the resilience of DeFi trading depends on major stablecoins.

Figure 8: Token swap activity based on total swap counts (in millions)

(a) Token ranking

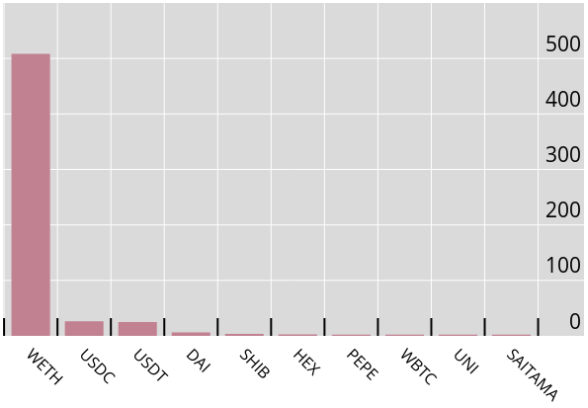

(b) Token pair ranking

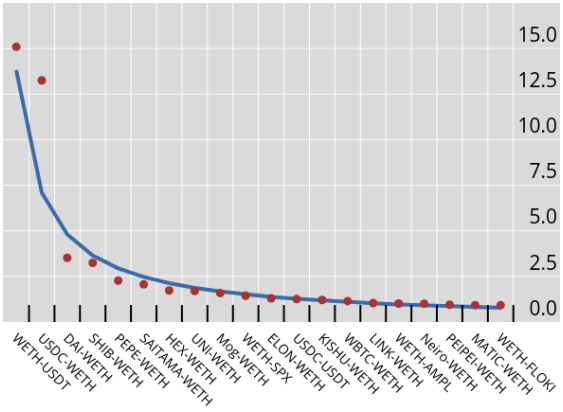

Note: Token tickers taken from respective smart contract metadata. The left panel displays the top 10 tokens based on the number of swap transactions involving the respective token expressed in millions, based on granular Uniswap V2 data until end 2025. The right panel plots the top 20 token pairs, based on swap transactions of the pair expressed in millions, based on granular Uniswap V2 data until end 2025 and a fitted power law distribution.

4.2.4 Spurious activity

The expansion of major stablecoins has also attracted spurious activity, imitating metadata of well-known stablecoin smart contracts. Such activity encompasses a range of opportunistic behaviours, including scam activity that induces users to transfer cryptoassets. We identify potentially spurious USDT contracts by filtering tokens whose symbol is "USDT". We exclude the actual Tether USDT contract itself from the spurious sample based on its publicly known address.

Spurious contract activity in Figure 9 is plotted against USDT activity. Spurious activity increases sharply in 2020. By the end of the sample, the cumulative issuance of tokens reaches roughly 100 billion. There is a sharp increase in spurious issuance in 2020 when stablecoins gained traction. The scale and persistence of activity of spurious USDT tokens over the recent years indicates that these tokens generate substantial transaction traffic on the Ethereum chain. Higher transaction volumes in USDT correlate positively with spurious activity shown by quarterly averages of USDT transfer volumes and transfers

[12] Users can receive (redeem) WETH (Ether) by sending Ether (WETH) to the smart contract address at a ratio of 1:1. The increasing role of smart contract activity is also observable from transfer values. Starting in 2021, WETH transfer volumes exceeded native Ether transfer volumes. In 2025, WETH transfers were around 7 times as high as transaction volume in native Ether transactions.

Figure 9: USDT and spurious token activity (in billions)

(a) USDT supply and spurious issuance

(b) USDT and spurious transaction volumes

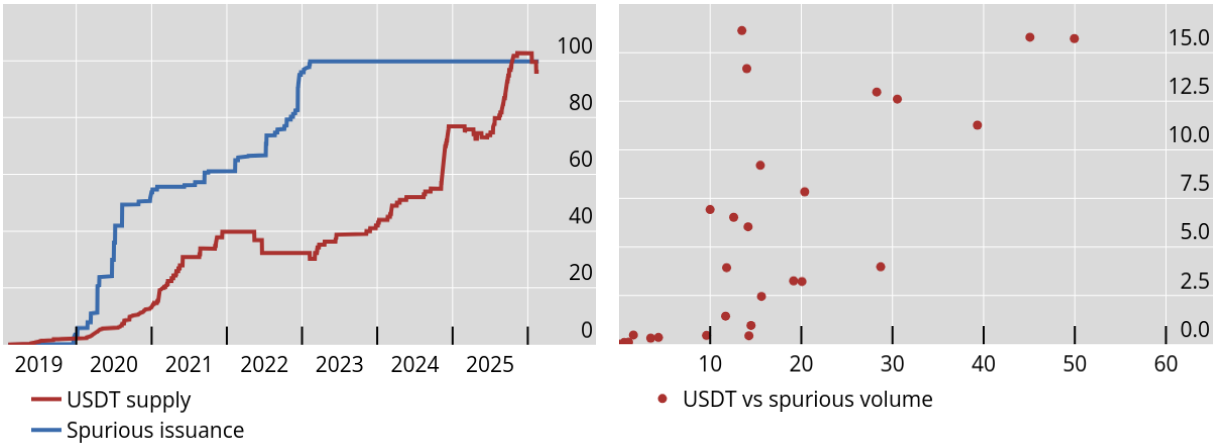

Note: The left panel plots USDT supply and spurious issuance token units. Spurious tokens are identified by aggregating smart contracts with the token symbol USDT (other than the USDT contract issued by Tether). Note that spurious tokens may not be pegged to the USD and units may not be meaningful. Therefore, extreme outliers are removed. The right panel plots USDT (x axis) and spurious contracts (y axis) quarterly average transaction volumes. Outliers are removed for spurious token activity.

associated with spurious tokens. While spurious activity is linked to arbitrary token values, unlikely to be backed by USD reserves, transfers are substantial and reach up to over 15 billion.[13]

Spurious activity highlights another aspect of the programmability challenge. Technically similar or even identical smart contracts represent different economic motives. Technical classification offers only limited mitigation. Specifically, the current DeFi ecosystem does not link issuing entities and institutions to on-chain assets. Registries or technical verification could mitigate these regulatory challenges and could also help prevent fraudulent activity.

4.3 Comparability of stablecoin use cases across blockchains

To evaluate use case dynamics of the largest stablecoin USDT across different blockchains, we examine transaction data extracted from smart contract event logs across Ethereum and Tron. USDT on the Tron blockchain operates as a TRC-20 token and equivalently as ERC-20 on Ethereum. All transactions related to the respective smart contracts are recorded on-chain via event logs that are emitted during contract execution.[14]

We break down holdings by externally owned accounts (EOAs) and smart contract accounts (SMAs). While EOAs are typically controlled by private keys and represent users or entities, SMAs store code and execute predefined logic. To classify accounts as EOA or SMA, we identify SMAs from smart contract deployments. By breaking down token holdings, we can disentangle the economic roles of stablecoins within the DeFi ecosystem. A high concentration of assets in SMAs likely reflects active use in DeFi applications, such as staking, lending or liquidity provision. In contrast, a significant share of assets in EOAs suggests a preference for holding tokens as a store of value or for transactional purposes such as payments or remittances. Data across blockchains unveil the role of infrastructure in determining use cases, rather than token identity alone. One caveat is that technical address types are imperfect proxies for economic roles: smart-contract balances may reflect bridges, wrappers, or custodial arrangements, and EOAs balances cannot distinguish payments, remittances, store-of-value demand or exchange custody.

[13] In order to avoid a small number of transactions with extreme values driving results, the sample excludes transactions with a transfer value above 100 billion.

[14] Transaction logs are decoded to extract the sending and receiving addresses, as well as the transferred amounts, which are normalised using the token's decimal precision. For robustness, we cross-reference USDT supply by extracting issuance (mint) events and correcting for redemption (burn) events and tokens destroyed from blacklisted addresses.

An analogous concept to our approach is the widely cited metric of total value locked (TVL) in DeFi, which measures the aggregate value of assets deposited into DeFi protocols. TVL provides a snapshot of the scale and adoption of DeFi ecosystems, reflecting liquidity held by smart contracts. TVL is typically calculated by summing up the holdings of SMAs associated with specific DeFi applications. However, as shown by Saggese et al. (2025), TVL is prone to inconsistent methodologies and even manipulation to inflate values. In contrast to aggregating token balances by DeFi protocols via querying their address balances, we aggregate token holdings via token transactions. This provides an alternative measure to standard TVL approaches that strictly avoids double-counting of tokens. While TVL can provide a full picture of token holdings at a protocol level, our approach yields a comprehensive view of holdings on a token level.[15]

Stablecoin activity has seen rapid growth since 2020, with USDT supply surges in 2021 and 2023, reaching over USD 180 billion in early 2026. The distribution of USDT holdings across EOAs and SMAs across Ethereum and Tron in Figure 10 indicates that the SMA share for Ethereum rose steadily through 2022. USDT supply decreased during the 2022 crypto market downturn, while the SMA share rose above 20%, hovering between 15–20% until late 2024. Coinciding with the US election and a more favourable policy outlook, stablecoin issuance shot up markedly. However, this uptick was not accompanied by an increase in smart contract holdings, suggesting it did not come with an increased usage of DeFi protocols. Since late 2024, the share of USDT smart contract holdings on Ethereum has declined to 10-15%. Similarly, the passage of the GENIUS Act, a stablecoin regulation in the US, led to increased stablecoin issuance but did not result in a sustained rise in smart contract holdings, pointing to a decoupling between stablecoin issuance and DeFi usage.

Figure 10: USDT holdings and smart-contract shares: Ethereum and Tron

(a) Holdings of total supply by account type (in billion)

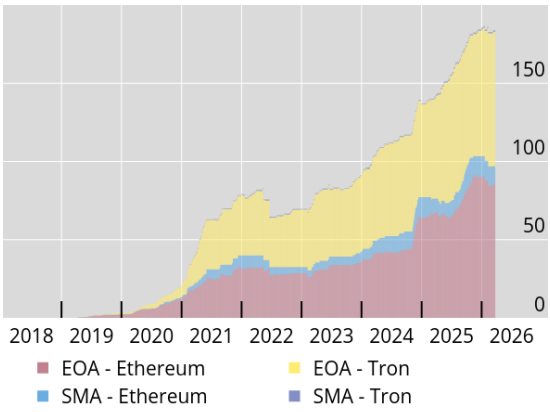

(b) Smart contract share of total holdings by chain (in percent)

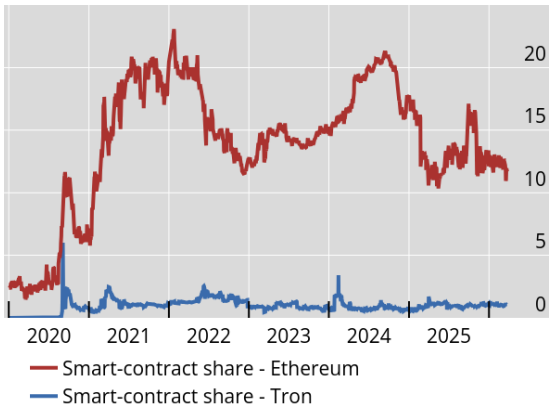

Note: The left panel plots externally owned account (EOA) and smart contract account (SMA) holdings of USDT on Ethereum and Tron (billions). The right panel plots the smart contract holdings as share of total holdings on each chain (percent).

While substantial amounts of stablecoins on Ethereum are held within DeFi protocols, this pattern does not extend to Tron where smart contract holdings account for only 1% most of the time. This stark contrast illustrates the comparability challenge as it hints at distinct economic behaviour across chains. USDT on Ethereum often functions as collateral and a liquidity provision instrument. USDT on Tron, appears to serve predominantly transactional and store-of-value motives, operating to a larger degree outside smart contract ecosystems. The same token, issued by the same entity, thus serves different economic roles depending on the infrastructure on which it is held. Treating USDT as a monolithic instrument could conceal its multifaceted role across different infrastructures.

[15]While both approaches theoretically yield the same results, there are flaws in verifiability and transparency in TVL computation. While event log analysis offers superior granularity compared to periodic balance queries, our approach requires the full transaction history and thus suffers from challenges in terms of scalability and completeness of token coverage. In addition, some smart contracts holdings may not relate to DeFi activity.

5 Conclusions

This paper demonstrates that, despite the transparency of public blockchains, widely used indicators of cryptoasset and DeFi activity are highly dependent on methodological choices and underlying assumptions. Metrics such as transaction volumes, market capitalisation and total value locked often suggest a degree of accuracy that is not supported by the nature of the underlying data. Using granular information from Bitcoin, Ethereum and Tron, we show that even modest adjustments in data treatment can lead to substantial differences in inferred economic activity.

Three structural sources of measurement divergences underpin these results. First, Bitcoin's UTXO architecture introduces an *aggregation challenge*: unless change outputs, lost coins and valuation conventions are accounted for, both transaction volumes and market capitalisation are systematically overstated. Second, Ethereum's smart contract environment gives rise to a *programmability challenge*, as economically relevant activity is dispersed across transactions, logs, execution traces and state changes, while extensive token proliferation obscures meaningful signals. Third, the issuance of stablecoins across multiple blockchains creates a *comparability challenge*. Identical tokens can fulfil distinct economic functions depending on the underlying infrastructure. Our analysis of USDT illustrates this point: on Ethereum, stablecoins are more closely intertwined with DeFi intermediation, whereas on Tron they are more likely used for payment-like or store-of-value purposes.

Taken together, these findings indicate that naive aggregation of on-chain data can misinform empirical research and policy assessments. Overstated indicators may exaggerate the systemic relevance of crypto markets, while poorly specified metrics may conceal emerging vulnerabilities and use cases. More broadly, the results highlight that effective analysis of cryptoassets and DeFi requires combining economic analysis with a detailed technical understanding of blockchain infrastructures. Accurately classifying assets, wallets and transactions is highly complex. Robust analysis requires a multi-pronged approach to filter and refine the data. Furthermore, regulatory initiatives could enhance data classification through a more formal identification of actors, though this remains a significant challenge in balancing the privacy features of public ledgers.

By treating commonly used indicators as assumption-dependent measures, and by advancing bounded, transparent and classification-based approaches, this paper contributes to the development of more robust analytical frameworks. As these markets continue to evolve and become more interconnected with the traditional financial system, the ability to interpret on-chain data accurately will be essential for sound policy design and timely risk identification.

The foundation for analysis of DeFi dynamics is granular data. Future research can build on the outlined data structure here to study questions, such as DeFi protocol interlinkages, decentralised exchanges, price formation mechanisms across a plethora of token pairs and interconnections across blockchains. Given their link to the traditional financial system via reserves, stablecoins present unique challenges for policymakers. Directions for future research may include studying transactional patterns across chains, stablecoin use cases across protocols and traditional actors employing DeFi architecture.

From a central bank perspective, these findings imply that on-chain indicators should be treated as noisy approximations rather than direct measures of economic activity. Effective use of blockchain data for financial stability monitoring, payments oversight and monetary analysis requires an explicit alignment between policy objectives and data constructions, together with transparent assumptions. The findings suggest the following three implications. First, point estimates should be complemented by bounded ranges reflecting protocol-specific uncertainty and retrospective revision. Second, interpreting DeFi activity requires anchoring economic meaning in technical classification combined with expert judgement, as token labels and standard logs are vulnerable to noise and imitation. Third, cross-chain analysis should distinguish infrastructure and asset identity.

Annex

Data structure

Figure 11 and Figure 12 represent the entity-relationship diagram (ERD) of the Ethereum and Tron data used in this paper, respectively. In the first stage of data processing, raw blocks, transaction receipt data and traces of all executed transactions are extracted from the Ethereum and Tron blockchain nodes. The raw data are stored in JSON format and are further transformed into more structured Delta format. This process yields three base tables—blocks, receipts and traces. Finally, from these tables, three additional tables are derived - transactions (txs), logs (derived from blocks and receipts) and smart contracts (derived from traces). The ERD illustrates the design, structure, and relationship of the resulting data model by listing the variables contained in each table together with their data types and key relational links used to connect the tables.

Figure 11: Ethereum data model

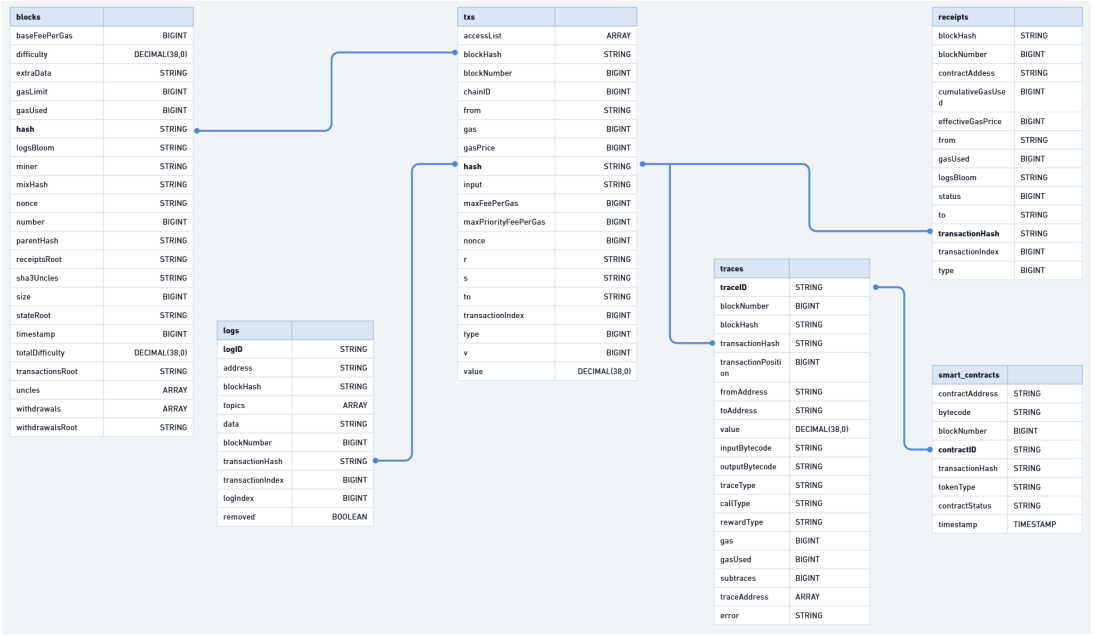

Figure 12: Tron data model

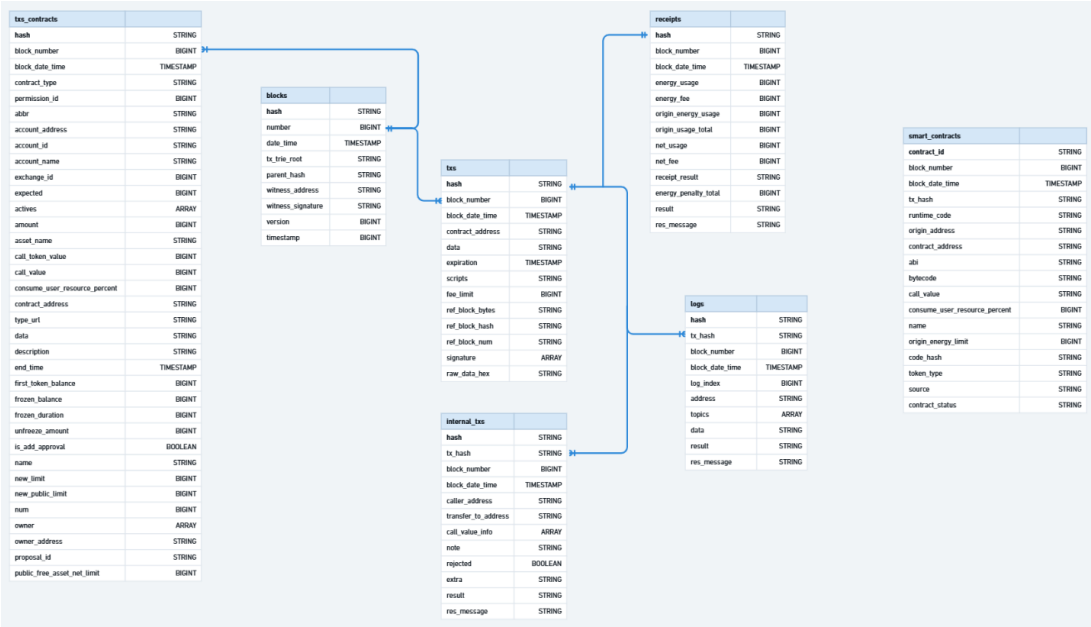

Dust coins

Identifying so-called dust coins is subject to fee levels. Fees are not constant and spike in periods of high network usage. During times with lower fee rates, low transaction outputs can still be moved, whereas in times with high fees more transaction outputs become uneconomical to move.[16]

We use a simple definition of dust coins which defines any UTXO equal or lower than 10,000 Satoshi as dust, equal to around USD 6.70 in April 2026:

$$\text{DustUTXO} \leq 0.0001 \text{ Bitcoin} \tag{7}$$

This assumption reflects the fact that fees have historically been higher than this threshold but have since declined markedly (see Figure 13b). Based on this conservative assumption, we display the current value and amount of dust in Figure 13a. In 2026, dust makes up only close to 1,500 bitcoin (0.007% of total supply). While small in terms of value, the number of UTXOs labelled dust adds up over 250 million. At the beginning of 2020, this number was about 60 million.

Note that the calculation is from today's perspective; if dust coins are spent in the future, the figure would change. Given that the value of dust coins is currently rather insignificant, even accounting for our generous definition of dust, and can potentially be spent in times of low fees, we ignore them in our analysis of market capitalisations.

[16] Also note that multiple dust coins can be aggregated into a single transaction. However, this assumes that the dust coins involved are either controlled by the same network participant or that multiple participants collaborate, combining their coins into one transaction.

Figure 13: Illiquid coins: Dust

(a) Dust coins

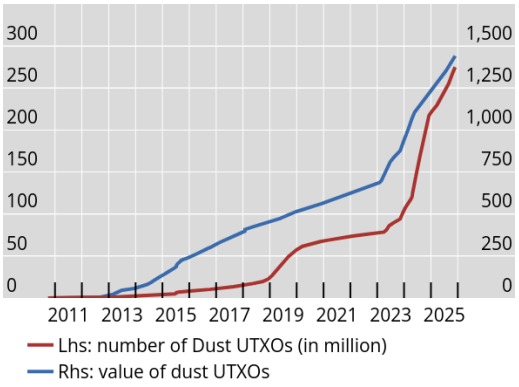

(b) Fees

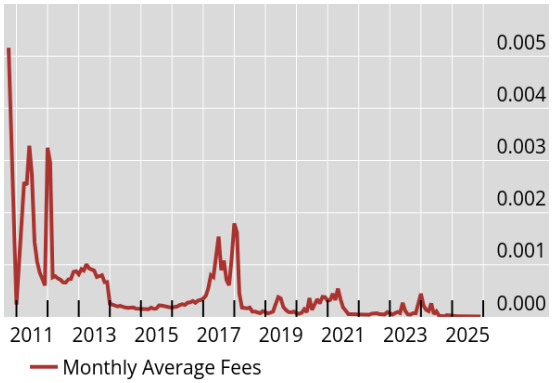

Note: The left panel displays UTXOs below 10,000 Satoshi, cumulated over time. The right panel shows monthly averages of daily transaction fees, measured in Bitcoin.

Example transaction of a token transfer

Transactions on Ethereum and Tron not only hold data on native cryptoassets but also smart contract calls. Smart contract data is reflected in transactions that invoke smart contracts, in the execution steps during transaction execution and after completion of a transaction. While we focus on an exemplary token transfer here, smart contract activity extends to a wide range of arbitrary programmatic operations defined by the underlying code.

A transaction to a smart contract address, depicted in Figure 14, holds the intended interaction with a contract. The transaction is sent from an EOA address to the smart contract with n Ether value of zero (0x0). The timestamp of the transaction can be inferred from the block number which does not form part of transaction instructions but is known in hindsight.

The input string includes instructions to the smart contract, which can be deconstructed into three parts. In the example, the input encodes three components: the function selector (transfer), the recipient address and the value. Convert the value from hexadecimal to a number and adjust it by the token's decimals to obtain the human-readable amount, in the example here yielding a transaction of over 1 billion USDT tokens.

The contract executes the requested operation and subsequent follow-up operations depending on the smart contract logic. After completion of execution in the EVM, the transaction produces a receipt. Transaction objects contain the input data and the Ether amount transacted, while transaction receipts record the execution outcome, including the status (Failure/Success) and emitted logs. Log data include data on developer defined events.

The execution of smart contract calls produces data across various dimensions:

1. Transactions with the inputs containing the parameters to invoke the smart contract logic.
2. Transaction receipts containing logs of smart contract activity as specified by the contract logic.
3. Traces containing the execution steps of transactions such as a call to a smart contract.
4. State differences (state diffs) containing the changes in smart contract storage before and after transactions.

Figure 15 illustrates the receipt for the same transaction as above where with the smart contract address returns the event and the token transfer is extracted from the so-called topics of the logs. The token amount transferred is returned in the field called data.

Figure 14: Example transaction call to smart contract

Note: The left-hand side includes a JSON format transaction with instructions for a smart contract address. The transaction input includes instructions for the smart contract, specifying a transfer of tokens to an address (leading '0x' inserted) and the amount in hexadecimal format. Note that token contracts specify the place where a decimal is set. In the example here the decimal value is 6.

Figure 15: Example transaction receipt with token transfer

Note: The left hand side includes a JSON format transaction receipt. The transaction receipt includes details of the transaction and the logs which specify the sender and receiver of tokens. The field data returns the token amount in hexadecimal format. Note that token contracts specify the place where a decimal is set. In the example here the decimal value is 6.

References

Alasadi, L., Bewaji, S., Gugnani, A., Gupta, T., Heijmans, R., and Van der Woerd, E. (2026). An econometric investigation on the stability of stablecoins: Are these coins stable or is their stability just a flip of the coin? *Journal of Financial Market Infrastructures*, forthcoming:0–0.

Aldasoro, I., Frost, J., and Ito, H. (2026). The impact of stablecoins on the international monetary and financial system. BIS Papers 170, Bank for International Settlements.

Altavilla, C., Boucinha, M., Burlon, L., Adalid, R., Fortes, R., and Maruhn, F. (2026). Stablecoins and monetary policy transmission. ECB Working Paper 3199, European Central Bank.

Aquilina, M., Cornelli, G., Frost, J., and Gambacorta, L. (2025). Cryptocurrencies and decentralised finance: Functions and financial stability implications. *BIS Working Paper*, 156.

Aramonte, S., Huang, W., and Schrimpf, A. (2021). Defi risks and the decentralisation illusion. *BIS Quarterly Review*.

Auer, R., Cornelli, G., Doerr, S., Frost, J., and Gambacorta, L. (2023a). Crypto trading and bitcoin prices: Evidence from a new database of retail adoption. BIS Working Papers 1049, Bank for International Settlements.

Auer, R., Haslhofer, B., Kitzler, S., Saggese, P., and Victor, F. (2023b). The technology of decentralized finance (defi). BIS Working Papers 1066, Bank for International Settlements.

Auer, R., Lewrick, U., and Paulick, J. (2025). Defiying gravity? an empirical analysis of cross-border bitcoin, ether and stablecoin flows. BIS Working Paper 1265, Bank for International Settlements.

Azzimonti, M. and Quadrini, V. (2025). Digital economy, stablecoins, and the global financial system. *National Bureau of Economic Research*, 34066.

Badev, A. and Chen, M. (2014). Bitcoin: Technical background and data analysis. Finance and Economics Discussion Series 2014-104, Board of Governors of the Federal Reserve System (U.S.).

Bewaji, S., Hamid, S., Aerts, T., Byck, S., Heijmans, R., and Van der Woerd, E. (2024). Are cryptocurrencies cryptic or a source of arbitrage? a genetic algorithm approach. *Journal of Financial Market Infrastructures*, 11(1):37–65.

Biais, B., Bisière, C., Bouvard, M., Casamatta, C., and Menkveld, A. (2023). Equilibrium bitcoin pricing. *Journal of Finance*, 78(2):967–1014.

BIS (2026). Anchoring trust in money: innovation beyond stablecoins. Annual Economic Report 2026 Chapter III, Bank for International Settlements.

Buterin, V. (2014). Ethereum whitepaper: A next-generation smart contract platform.

Böhme, R., Christin, N., Edelman, B., and Moore, T. (2015). Bitcoin: economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–38.

Cole, B., Dyhrberg, A., Foley, S., and Svec, J. (2022). Can Bitcoin be Trusted? Quantifying the economic value of blockchain transactions. *Journal of International Financial Markets, Institutions and Money*, 79(C).

Cong, L., Li, X., Tang, K., and Yang, Y. (2023). Crypto wash trading. *Management Science*, 69(11):6427–6454.

Di Casola, P., Habib, M. M., and Tercero-Lucas, D. (2023). Global and local drivers of bitcoin trading vis-à-vis fiat currencies. ECB Working Paper Series 2868, European Central Bank.

ECB Crypto-Assets Task Force (2019). Crypto-assets: Implications for financial stability, monetary policy, and payments and market infrastructures. Occasional Paper 223, European Central Bank.

Gandal, N., Hamrick, J., Moore, T., and Oberman, T. (2018). Price manipulation in the bitcoin ecosystem. *Journal of Monetary Economics*, 95:86–96.

Guo, D., Dong, J., and Wang, K. (2019). Graph structure and statistical properties of ethereum transaction relationships. *Information Sciences*, 492:58–71.

Halaburda, H., Haeringer, G., Gans, J., and Gandal, N. (2022). The microeconomics of cryptocurrencies. *Journal of Economic Literature*, 60(3):971–1013.

Hofmann, B., Mehrotra, A., and Paulick, J. (2026). Dollarisation and monetary control: what lessons for the rise of stablecoins? BIS Working Paper 1370, Bank for International Settlements.

Jahanshahloo, H., Irresberger, F., and Urquhart, A. (2023). Bitcoin under the microscope. *British Accounting Review*, 58(3).

Karau, S. (2023). Monetary policy and Bitcoin. *Journal of International Money and Finance*, 137:102880.

Kosse, A., Rice, T., Schär, F., Shirakami, T., and Siridhasanakul, J. (2026). The anatomy of stablecoin transactions. BIS Working Papers 1359, Bank for International Settlements.

Liu, A., Jahanshahloo, H., Chen, J., and Eshraghi, A. (2023). Trading patterns in the bitcoin market. *The European Journal of Finance*.

Liu, Y. and Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, 34(6):2689–2727.

Makarov, I. and Schoar, A. (2021). Blockchain analysis of the Bitcoin market. NBER Working Paper, no 29396.

Mao, Q., Wang, J., Feng, Z., and Yan, J. (2024). Unravelling stablecoin-favored ecosystem: Extracting, exploring on-chain data from tron blockchain. In *Blockchain and Web3.0 Technology Innovation and Application*, Communications in Computer and Information Science. Springer.

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., and Savage, S. (2016). A fistful of bitcoins: characterizing payments among men with no names. *Association for Computing Machinery*, 59(4):86–93.

Moeser, M. and Narayanan, A. (2022). Resurrecting address clustering in bitcoin.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf. Accessed: 2025-07-04.

Pietrzak, M. (2023). What can monetary policy tell us about bitcoin? *Annals of Finance*.

Prat, J. and Walter, B. (2021). An equilibrium model of the market for bitcoin mining. *Journal of Political Economy*, 129(8):2415–2452.

Ron, D. and Shamir, A. (2013). Quantitative analysis of the full bitcoin transaction graph. In Sadeghi, A.-R., editor, *Financial Cryptography and Data Security*, pages 6–24, Berlin, Heidelberg.

Saggese, P., Froewis, M., Kitzler, S., Haslhofer, B., and Auer, R. (2025). Towards verifiability of total value locked (tvl) in decentralized finance. BIS Working Papers 1268, Bank for International Settlements.

Schilling, L. and Uhlig, H. (2019). Some simple bitcoin economics. *Journal of Monetary Economics*, 106:16–26.

Schär, F. (2021). Decentralized finance: On blockchain- and smart contract-based financial markets. *Federal Reserve Bank of St. Louis REVIEW*, 103(2):153–174.

TRON-Foundation (2017). Tron: Advanced decentralized blockchain platform (v2.1).