

A cross-protocol dataset of blockchain oracle incentives and accountability

Huaiyu JIA

The Hong Kong University of Science and Technology (Guangzhou)

Siguang Li

The Hong Kong University of Science and Technology (Guangzhou)

Lin William Cong

Nanyang Technological University

Luyao Zhang

lz183@duke.edu

Duke Kunshan University <https://orcid.org/0000-0002-1183-2254>

Data Note

Keywords: Blockchain oracles, Oracle economics, Economic incentives, Accountability, Blockchain datasets, Cross-protocol analysis, Data provenance, Economic measurement, Data validation, Reproducible research

Posted Date: October 16th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11382560/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

A cross-protocol dataset of blockchain oracle incentives and accountability

Huaiyu Jia^a Siguang Li^a Lin William Cong^b Luyao Zhang^{c,*}

^a The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China

^b Nanyang Technological University, Singapore

^c Duke Kunshan University, Kunshan, China

Working paper • October 2026

Abstract. Blockchain oracles connect programmable contracts to external information, yet their public records seldom reveal economic outcomes in a directly comparable form. We present a cross-protocol data resource covering eight oracle branches: five detailed panels and three narrower extensions. Its complementary representations contain 105,588,120 typed accountability observations, 22,723,389 economic components, and 19,275,854 records satisfying protocol-specific realization criteria. These layers overlap and preserve different observation units. Construction retains native identifiers, reconstructs supported lifecycles, and distinguishes configured incentives, accrued adjustments, claimable entitlements, realized payments, and applied losses. Structural checks, economic-semantic corrections, selected funds-flow reconciliations, and bounded deterministic replay assess the resulting measurements. The resource supports comparative research on participation, incentives, dispute resolution, and financial accountability, together with clearly scoped prediction exercises. Coverage and source traceability remain uneven across protocols, and institutional adjudications are distinguished from independent evidence of factual correctness. This design makes economic interpretation explicit while retaining the limits of the underlying evidence.

1 Introduction

Blockchain oracles make a fundamental institutional choice visible: a programmable contract must rely on someone, or some mechanism, to supply information that the ledger itself cannot observe. Prices, event outcomes and other external facts become inputs to deterministic execution through systems of reporting, verification and dispute resolution [1, 2]. The resulting arrangements combine technical infrastructure with economic commitments. Participants may earn rewards, post capital, challenge a report or incur a penalty, while users depend on the process that determines which information the contract accepts. Studying this infrastructure therefore requires attention to both information provision and the incentives that sustain it [3, 4].

Public visibility of transactions does not, by itself, make these incentives measurable. The record of a single transfer may combine returned principal with a reward, and a negative accounting adjustment may precede any applied reduction in a participant’s stake. A disputed request can traverse several contracts or chains before reaching settlement. Differences in protocol versions, native identifiers and accounting conventions further complicate comparison. These are substantive measurement problems: an analysis that treats gross payouts as profit, or accruals as realized losses, changes the economic object being studied. A useful comparative dataset must reconstruct the relevant lifecycle while retaining the institutional meaning of each observed state.

We address this problem with a layered data resource covering eight oracle branches. UMA, Chainlink, Tellor, Flare FTSOv2 and Pyth provide five detailed panels; Chronicle, RedStone Push and DIA contribute three narrower extensions. The common accountability representation contains 105,588,120 typed observations, complemented by 22,723,389 economic components and 19,275,854 records that meet protocol-specific realization criteria. These overlapping representations retain different observation units and cannot be added as independent samples. Native records remain available alongside the common fields, allowing researchers to move between comparative measures and the protocol evidence on which those measures depend.

The contribution is a measurement framework for oracle incentives and accountability, together with data and reconstruction materials that make its choices inspectable. It supports descriptive comparisons of

*Corresponding author: Luyao Zhang. Duke Kunshan University, 8 Duke Avenue, Kunshan, Suzhou, Jiangsu 215316, China. Email: lz183@duke.edu.

participation, disputes and economic outcomes, as well as bounded prediction exercises. Its reach depends on the evidence available for each branch: source traceability is uneven, historical acquisition is not fully reconstructed for every record, and protocol adjudication is not an independent test of factual truth. These limits define the questions the resource can answer. They also motivate its central design principle: comparison should preserve the distinctions that give an observation its economic meaning.

2 Related work

Research on oracles and research on blockchain datasets approach the same infrastructure from complementary directions. Oracle studies examine how external information is authenticated, aggregated and incentivized; data-resource studies establish reusable representations of transactions, entities and contracts. The present work connects these perspectives by making the relationship between native protocol evidence and economic measurement explicit. Tables 1 and 2 summarize representative contributions without treating their differing units or purposes as directly comparable measures of scale.

2.1 Oracle design, security and economic institutions

Town Crier uses trusted hardware to authenticate external data delivery, whereas DECO develops cryptographic proofs about data obtained over TLS [1, 5]. The oracle systematization of knowledge situates such designs within a wider set of trust assumptions, attacks and mitigations [2]. Subsequent work addresses other parts of the information-security problem: adaptive conformal consensus provides guarantees under specified conditions for aggregating oracle reports, and SMARTCAT detects contracts associated with price-manipulation attacks [6, 7]. These contributions clarify how oracle information can be delivered or threatened. Their stated objectives differ from reconstructing the participant-level sequence between an incentive commitment and its eventual economic realization.

Empirical oracle studies provide a second foundation. A first look at DeFi oracles measures deployed price-oracle behavior, while Demystifying Pythia examines Chainlink usage, pricing and service quality on Ethereum [8, 9]. Economic research complements these observations by analyzing incentives, staking, decentralization and the integration of oracle networks with financial markets [4, 10]. Together, these studies motivate an empirical representation that follows the economic process across participation, challenge, adjudication and settlement. The framework developed here retains those stages where the source evidence supports them, rather than assuming that every mechanism exposes an equivalent lifecycle.

2.2 Reusable blockchain data resources

Graph datasets make blockchain activity accessible to a broader research community. Elliptic represents Bitcoin transactions for illicit-activity classification; Elliptic2 shifts attention to labeled subgraphs of entity clusters [11, 12]. Chartalist provides machine-learning-ready datasets for UTXO- and account-based blockchains, while temporal Bitcoin datasets support entity and transaction graph analysis [13–15]. EX-Graph connects Ethereum and X, and the multi-chain Graphs of Graphs resource preserves relationships within and across token networks [16, 17]. These representations capture important network structure. Oracle economics additionally requires identifying what a request, report, entitlement or applied balance change means within its originating mechanism.

Other datasets broaden the range of relevant units. Research on NFTs and memecoins links market behavior to network and visual features [18, 19]; event-aware DeFi forecasting emphasizes fine-grained on-chain events [20]. DIVE organizes smart contracts into separate pre- and post-deployment representations for vulnerability analysis [21]. The latter illustrates the value of preserving multiple views of one underlying object, a principle also used here. Our economic components, however, describe oracle participation and financial accountability rather than contract vulnerabilities, and their annotations depend on protocol-specific accounting and evidence rules.

The unresolved measurement need is thus not simply a larger transaction collection. Comparative oracle research needs shared concepts that remain connected to native roles, timing and economic states. This resource supplies that connection through layered representations, documented reconstruction rules and explicitly scoped validation. It complements the existing literature by distinguishing an announced incentive

Table 1. Oracle research and the measurement problem addressed here. “Conceptual” denotes a mechanism or taxonomy rather than a deposited cross-protocol economic panel. The comparison concerns the stated contribution of each work.

 Study	 Evidence focus	 Representation	 Economic coverage
Town Crier [1]	Authenticated delivery	Trusted-hardware oracle architecture	External-data delivery; no cross-protocol economic lifecycle panel.
DECO [5]	Web-data provenance	Cryptographic proofs for TLS data	Data provenance and privacy; no reconstruction of oracle compensation or losses.
Oracle SoK [2]	Security and trust	Taxonomy of designs, attacks and assumptions	Conceptual treatment of economic security and manipulation.
DeFi oracle measurement [8]	Deployed price oracles	Reported values, deviations and operation	Selected operational measurements rather than harmonized incentive states.
Demystifying Pythia [9]	Chainlink on Ethereum	Usage, pricing and service quality	System-specific economic and usage measurements.
Oracle economics primer [4]	Incentives and institutions	Economic framework for oracle mechanisms	Conceptual analysis of staking, decentralization and tokenomics.
This resource	Participation and accountability	Native lifecycles with typed economic components	Five detailed panels and three extensions; realized economics only where observable.

Table 2. Representative blockchain datasets and their analytical units. Scales refer to different scientific objects and are therefore not comparable sample sizes. The present resource complements graph and contract-security datasets by retaining protocol-specific economic stages.

 Resource	 Observation unit	 Published scope	 Analytical purpose
Elliptic [11]	Bitcoin transaction	203,769 nodes; 234,355 directed edges	Illicit-transaction classification and financial forensics.
Elliptic2 [12]	Entity cluster and labeled subgraph	49.3 million clusters; 196.2 million edges; 121,810 labeled subgraphs	Money-laundering patterns at subgraph level.
Chartalist [13]	Dataset-specific graph, address or transaction	Multiple UTXO- and account-based blockchain datasets	Reusable data access, graph construction and machine-learning benchmarks.
DIVE [21]	Ethereum smart contract	22,330 contracts; pre- and post-deployment representations	Vulnerability analysis with lifecycle-specific features.
This resource	Typed protocol-native observation	105,588,120 harmonized observations; eight branches	Incentive and accountability measurement, with native units and economic realization states retained.

from a claimable entitlement, a transfer from participant gain, and an institutional decision from external truth. Those distinctions permit comparison without erasing the differences that the comparison seeks to understand.

3 Methods

Construction proceeds from protocol evidence to economic interpretation, and only then to common analytical fields. This order matters because raw blockchain events do not define interchangeable economic observations. The workflow in Figure 1 first registers historical sources and decodes native records, then reconstructs supported lifecycles and economic states, and finally harmonizes the resulting observations. Unresolved identities, open lifecycles and unobservable outcomes remain explicit throughout. The method therefore treats a common schema as a structured comparison of evidence, rather than a substitute for protocol-specific reconstruction.

3.1 Scope, selection and historical boundaries

The project discovery registry surveys 56 oracle or related mechanism categories and 366 recorded protocol–oracle relationships. It is a discovery resource, rather than a sampling frame for a 56-protocol longitudinal study. The present dataset includes eight branches for which accountability or economic observations could be reconstructed with explicit source and implementation boundaries. Five provide comparatively detailed native panels, while Chronicle, RedStone Push and DIA expose narrower parts of the economic process. Selection consequently reflects both the substantive mechanism and the historical evidence available for reconstructing it.

The common upper observation cutoff is 30 June 2026 at 23:59:59 UTC. Lower boundaries depend on contract deployment, module or program availability, implementation changes and retained historical interfaces. Table 3 reports these boundaries and the principal native units. Protocol documentation informs the interpretation of optimistic-oracle requests, staking, reporter disputes, reward epochs and publisher incentives [22–26]. The extension branches are interpreted against their own documented architectures [27–29]. Historical records are evaluated under the implementation active when they were generated; current documentation does not replace that historical specification.

3.2 Source acquisition and native decoding

For each branch, a source register identifies the network, contract or program, implementation, collection interval and retrieval interface. EVM-compatible branches use historical logs, supplemented where necessary by transactions, receipts, block metadata and token-transfer evidence. ERC-20 transfers are retained when they help separate principal, fees and rewards from a settlement event. Tellor instead combines tellor-1 module state, ABCI events and block results; multiple reports within one block produce a one-to-many expansion into 82,866,640 report observations. Pyth evidence is located by Solana transaction signatures and outer or inner instruction positions. Flare combines on-chain records with official reward-calculation artifacts, preserving the distinction between entitlement and claim.

Source qualification applies the registered version, cutoff and deterministic decoding conditions. For Polygon UMA, 626,578 raw records include five unknown event topics that remain in the raw collection but are excluded from the recognized decoded event families. For Chainlink staking, 175,015 collected logs yield 174,890 decoded events and 125 unclassified records. These classifications describe decoding outcomes; they are not recoded as absent economic behavior. Native tables retain the strongest available source keys and monetary amounts in integer units of the underlying asset, with asset identifiers and decimal precision recorded separately. This preserves both the identity of an observation and the arithmetic needed to interpret it.

3.3 Lifecycle reconstruction and cross-chain linkage

Decoded records are joined where the native mechanism identifies a relationship between successive states. UMA requests link proposals, disputes, DVM adjudication and settlement through reconstructed request




Figure 1. Conceptual construction and validation of the oracle data resource. Read stages 1–4 across the upper row, then stages 5–8 across the lower row. Sources and implementations are registered before native semantics and acquisition are defined. Supported lifecycles, identities and economic states are then reconstructed into analytical records and documentation. The lower strip identifies cross-cutting measurement challenges. Depicted transitions and checks apply only where retained evidence supports them; protocol decisions do not establish independent external truth.

Table 3. Registered historical coverage and native observation units. The common upper cutoff is 30 June 2026, 23:59:59 UTC; deployment and source availability determine lower boundaries. Coverage denotes the branch studied, rather than all activity of the named protocol.

Protocol	Network and source	Registered coverage	Native units and interpretive boundary
UMA	Ethereum VotingV2; Polygon OOV2 and adapters	Ethereum DVM blocks 16,697,232–25,433,938; Polygon contract-specific ranges	Requests, proposals, disputes, votes, accruals, applied stake changes and settlements. Episodes can span chains and adapter versions.
Chainlink	Ethereum Staking v0.2; ETH/USD proxy and aggregator phases	Staking from block 18,572,190; feed phases through block 25,433,938	Stake operations, reward claims, accounting adjustments, answers and service windows. These describe participation and service, rather than independent factual truth.
Tellor	tellor-1 module state, ABCI events and block results	Heights 1–19,890,860; first report at height 317; last block 23:59:57 UTC on cutoff date	Micro-reports, aggregates, disputes, votes, payments, withdrawals and reporter restrictions. Cosmos module semantics govern interpretation.
Flare FTSOv2	Flare RewardManager, VoterRegistry and reward artifacts	Closed reward epochs 228–410; cutoff block 64,054,789	Epochs, entitlements, provider–feed conditions, claims and restrictions. Entitlement and payment differ; aggregate claims may resist decomposition.
Pyth	Solana mainnet OIS transactions and Pythnet-derived inputs	Archived transactions from 11 September 2024; 49 retained closed epochs in the state panel	Instructions, publisher–epoch states, stake mutations and transfers. Evidence uses signatures and instruction positions; the state panel is not deployment-complete.
Chronicle	Ethereum validator and challenge contracts	Contract-specific deployment blocks through 25,433,938	Reports, challenges, disputes and challenge rewards. Observed challenge rewards are not a complete publisher-compensation ledger.
RedStone Push	Ethereum Push adapters and registered feeds	Contract-specific deployment ranges through 25,433,938	Signed publications and Push updates. Signature verification establishes publication provenance, not external correctness.
DIA	Lasernet staking contracts	Deployment block 3,874,186 through cutoff block 31,302,665	Stake events, withdrawals, principal and reward components. This branch measures staking economics rather than per-report accountability.

identifiers. For requester r , identifier i , timestamp t and ancillary data a , the identifier is

$$q = \text{Keccak256}(\text{pack}(r, i, t, a)), \quad (1)$$

where pack denotes the protocol’s packed ABI encoding. The reconstructed OOV2 panel contains 42,676 request rounds: 41,826 settle by the cutoff and 850 remain open. Other mechanisms require different objects. For example, 46,723 retained Chainlink ETH/USD answer updates define 46,722 adjacent service windows, and 13 Tellor jail events are paired with 13 release events to reconstruct reporter restrictions.

Cross-chain identity is established through version-aware ancillary-data transformations and native request construction. A UMA link receives Grade A only when the registered bridge or adapter rule identifies one unique DVM request. Of 835 attempted links, 832 meet this deterministic definition and three remain unmatched. The grade describes the identity of a reconstructed object; it does not certify the factual correctness of the oracle’s decision. Retaining unmatched and unfinished objects prevents the apparent completeness of an analytical table from replacing evidence about the underlying process.

3.4 Economic states and settlement accounting

An economic observation is classified by its state and evidentiary basis. Table 4 distinguishes configured incentives, accrued rewards or penalties, claimable entitlements, realized payments, applied losses, principal,

refunds and fees. The distinctions are protocol-specific. UMA challenge status requires native challenge evidence, and settlement is not inferred merely from a transfer-like event. Tellor SUPPORT outcomes refer to acceptance of the challenge and rejection of the challenged report. Negative DVM accruals remain distinct from an observed reduction in staked principal. Pyth reward rates are not token reward amounts, and Flare participation restrictions are not automatically monetary losses.

UMA settlement provides a concrete accounting example. Let B be the effective bond, F the final fee and R the configured reward, all measured in the smallest integer unit of the same asset. Posted principal is $P = B + F$. For disputed requests, the identifiable winner component is $W = B - \lfloor B/2 \rfloor$, while the protocol component is $F_{\text{protocol}} = \lfloor B/2 \rfloor + F$. Expected gross payouts follow the registered contract and adapter rules:

$$\begin{aligned} G_{\text{undisputed}} &= B + F + R, \\ G_{\text{disputed}} &= B + F + W + R_{\text{settlement}}. \end{aligned} \quad (2)$$

Here $R_{\text{settlement}}$ denotes the reward component attributable to the settlement under the applicable adapter; a refund or rollover is not automatically participant income. Reconciliation uses exact integer arithmetic and checks $\Delta_{\text{payout}} = G_{\text{actual}} - G_{\text{expected}}$ against zero.

In an audited settlement, expressed in human-readable token units, a gross payout of 1,000 comprises 750 units of returned principal and 250 units of token gain before the additional costs needed to calculate net profit. A separately configured reward of five units follows its observed refund or rollover treatment. This example illustrates why transfer, principal recovery, token gain and net profit require separate measures. Funds-flow checks apply the same logic where evidence permits: UMA settlements are compared with token movements, Chainlink distinguishes transfer-requiring events from accounting-only adjustments, and Flare claim reconciliation does not invent components that an aggregate claim fails to expose.

Table 4. Economic states distinguished in reconstruction. Each classification requires the evidence specified by the originating protocol. A lifecycle need not expose every state.

⌚ State	💰 Economic meaning	🔍 Interpretation rule
Configured incentive	Reward specified in protocol parameters	A design parameter; not automatically participant income.
Accrued reward or penalty	Calculated entitlement or liability	Distinct from subsequent payment or application; counting both may duplicate one change.
Claimable entitlement	Amount eligible for a claim	An unclaimed entitlement is not a zero payment or a realized transfer.
Paid reward	Observed transfer or realized protocol credit	Realized only when qualifying cash-flow or state evidence is present.
Applied slash	Reduction applied to a participant's balance or principal	Requires observed application, beyond a penalty calculation or accrual.
Principal deposit or return	Capital posted and subsequently recovered	Kept separate from rewards and fees; returned capital is not profit.
Refund or rollover	Incentive returned or carried to another state	Not attributed to the current participant without supporting evidence.
Protocol fee	Amount charged or retained under protocol rules	Separated from participant reward and principal.

3.5 Harmonization, missingness and documentation

Supported native concepts are projected into common fields for protocol, unit, role, actor or beneficiary, lifecycle identity, time, asset, economic stage, source and evidence status. The common accountability layer contains 105,588,120 typed observations. The economic-semantic layer contains 22,723,389 components, and its stricter realized representation contains 19,275,854 components that satisfy branch-specific payment or applied-state criteria. These layers overlap. Their purpose is to offer complementary views of the same processes, not three independent populations.

Missingness is interpreted as part of the measurement. An absent value may indicate an inapplicable mechanism, an unexposed quantity, incomplete collection, a decoding problem or an unfinished lifecycle; it is not automatically an observed zero. Where supported, schemas distinguish complete, partial, verified-zero, unknown, unavailable and not-applicable coverage, with right censoring represented separately. Deduplication uses native source keys rather than equality of entire rows. Intentional multiplicity remains when one source object appears in different representations or one settlement produces several economic components. Table 5 summarizes the operations that establish these relationships.

Table 5. Processing operations and their outputs. Rows describe transformations rather than statistically independent populations: operations may expand, aggregate, join or decompose their inputs.

⌚ Operation	📄 Input and rule	📊 Output
Source registration	Network, contract or program, historical version, interval and interface	Branch-specific collection specification.
Acquisition and qualification	Logs, transactions, receipts, block or module state, signatures and protocol artifacts; cutoff and decoding rules	Qualified source records, with unresolved decoding outcomes retained.
Native decoding	Protocol event and state definitions; integer amounts and native keys	Protocol-native tables.
Lifecycle reconstruction	Request, dispute, claim, epoch or other native object relationships	Supported lifecycle objects and open cases.
Cross-chain linkage	Version-specific identifier and ancillary-data transformations	Deterministic matches, link grades and unmatched cases.
Economic-state mapping	Native lifecycle and accounting evidence	Configured, accrued, claimable, paid, applied, returned, refunded and fee components.
Funds-flow reconciliation	Expected quantities against observed transfers or applied state changes	Realization flags and accounting differences.
Harmonization	Shared concepts with native units and source links retained	105,588,120 typed accountability observations.
Realization filtering	Protocol-specific evidence conditions applied to 22,723,389 economic components	19,275,854 qualifying realized components.
Documentation	Dictionaries, schemas, coverage information, source registers and validation results	Interpretable data and reconstruction materials.

Documentation records native identifiers, source registers, input inventories, schemas, variable definitions and validation rules. Quality controls assess selected key relationships, lifecycle orderings, cutoff conditions, economic identities and deterministic links. The accompanying artifacts identify the tested populations and the evidence available for reconstruction. Their traceability supports inspection of specific mappings, while leaving incomplete historical acquisition and unresolved source ancestry visible.

3.6 Downstream benchmark construction

Two sets of prediction exercises demonstrate bounded reuse. The UMA adjudication task contains 810 challenged observations, split chronologically into 489 training, 161 validation and 160 test cases. A 16-feature allowlist and training-only preprocessing support comparisons among a prior baseline, logistic regression and histogram gradient boosting using accuracy, AUROC and Brier score. The target is the protocol’s adjudication. Excluding retrospective outcome fields reduces one form of leakage, but does not prove historical decision-time availability for every reconstructed feature.

The financial benchmark considers reward realization (F1), penalty application (F2) and capital settlement (F3). Seven fitted models per task yield 21 fits, with Kaplan–Meier references reported separately. F1 and F2 use fallback random partitions, whereas F3 uses a temporal partition. These differences restrict comparisons across tasks and the interpretation of temporal generalization. Both exercises derive task-specific views

from the core resource; neither changes its native units or turns its labels into independent factual truth. The construction method ultimately aims to make those interpretive commitments explicit enough that subsequent research can assess them rather than inherit them silently.

4 Data records

The resource is organized so that an analytical view remains connected to the evidence and unit that produced it. Protocol-native tables, reconstructed lifecycles, common accountability fields and economic components serve different purposes within this architecture. Researchers can begin with a common representation for comparison, then return to native tables when the question depends on a particular mechanism, timing convention or accounting rule. Table 6 describes the principal components; the catalogue and dictionaries supply the detailed paths, types and definitions needed for reuse.

Table 6. Principal data components. Counts refer to heterogeneous, overlapping units. Sizes use decimal GB and describe the stated files, rather than independent samples. The deposited catalogue provides complete paths and physical schemas.

Component	Records or scope	Format or size	Interpretation
Retained inputs	64 native or curated reconstruction inputs	Branch-specific artifacts	Fixed reconstruction inputs; not a complete archive of original RPC responses.
Native tables	Protocol-specific tables within the 95-table catalogue	Parquet	Requests, reports, votes, staking operations, claims, transfers and lifecycles.
Accountability	105,588,120 observations; 50 columns	Parquet, ZSTD; 4.862 GB	Common fields retain native units, roles, source identities and evidence status.
Economic semantics	22,723,389 components; 28 columns	Parquet, ZSTD; 1.904 GB	Configured, accrued, claimable, paid, returned, refunded and applied states.
Realized economics	19,275,854 components; 28 columns	Parquet, ZSTD; 1.706 GB	Qualifying payments or applied reward and penalty changes; overlaps the semantic layer.
Linkage and checks	Branch-specific linkage, reconciliation and validation tables	Parquet and reports	Matches, accounting checks and unresolved cases.
Dictionaries and schemas	33 conceptual variables; 1,789 physical field entries in 95 tables	CSV, JSON and schemas	Meanings, units, types, missingness, evidence classes and validation rules.
Application views	UMA adjudication and financial time-to-event tasks	Parquet and evaluations	Task-specific samples, partitions and predictions.
Package inventory	95 catalogued Parquet tables; 143 package files	Catalogue and manifests	The package includes supporting artifacts and an additional prediction Parquet outside the canonical catalogue.

4.1 Native records and table relationships

Native tables preserve the object meaningful to the originating protocol. The UMA OOV2 panel contains 42,676 request lifecycles, while the Tellor micro-report table contains 82,866,640 reports. Other examples include 18,468,735 Flare claim events and 518,560 decoded Pyth OIS instructions. These quantities count requests, reports, claims and instructions, respectively. A comparison that treats them as equivalent oracle events would confuse differences in data granularity with differences in institutional activity. Within a branch, related evidence can also be stored separately: UMA request-level lifecycle and payout fields link to token-flow records rather than repeating each transfer inside the request table.

Identifiers follow the strongest available native structure. EVM logs use chain, transaction and log position; Tellor relies on module-specific report metadata; Solana evidence uses signatures and instruction positions. Table 7 summarizes these keys and their relationships. Checks on the named UMA, Chainlink and cross-chain linkage tables assess selected primary-key and foreign-key constraints. A one-to-many relationship can represent several events within a lifecycle or several components within one settlement. It is therefore a relationship to interpret, rather than a duplication to remove automatically.

Table 7. Native units, identifier construction and logical relationships. One-to-many relationships describe scientific objects, not mandatory duplication in every physical table. Native ordering is authoritative; physical row order is not.

Unit	Identifier	Relationship and interpretation
EVM log	Chain, transaction hash and log index	A transaction can contain several logs. Order by block and log position; retain receipt and block evidence.
UMA OOV2 request	oo_request_id; requester, identifier, timestamp and ancillary data	One request links several lifecycle events and transfers. Request, proposal, dispute and settlement times remain distinct.
UMA DVM request or voter record	dvm_request_id; voter and source-event location for actor-level records	One request can link many voters and adjustments. Accrual and applied stake change are different observations.
Tellor report	query_id, reporter and meta_id; block-position namespace for the common key	Multiple reports can contribute to an aggregate. Native module metadata replaces a universal EVM transaction/log key.
Flare entitlement or claim	Epoch, beneficiary and claim type; transaction/log position for a claim	Entitlements and claims are not universally one-to-one. Epoch end differs from claim time.
Pyth instruction or transfer	Signature and outer_instruction_index; inner position for token transfers	One transaction can contain several instructions and transfers. Reconstruct time from block time and instruction positions.
Common observation	accountability_event_id, protocol namespace, native_table and accountability_unit_id	Native-to-common mapping can be one-to-many. The retained native unit defines the statistical object.
Economic component	evidence_id and originating table/event identity	One settlement can yield several components. do_not_sum_group identifies related components at risk of double counting.

4.2 Common accountability and economic representations

The common accountability table contains 105,588,120 rows and 50 columns. Each row is a typed projection of a native observation, retaining protocol identity, unit, roles, source references and evidence status. The population includes requests, individual reports, disputes, epochs, staking observations and service windows. Protocol, native table, unit type and unit identity must consequently accompany any cross-protocol aggregation. The representation provides shared analytical vocabulary while preserving the information needed to determine whether two rows are comparable.

Economic components are represented in a separate 22,723,389-row, 28-column table. Its fields record the mechanism, stage, actor, counterparty, amount, asset, source object, evidence class and validation status. A stricter 19,275,854-row representation uses the same physical schema but requires qualifying evidence of payment or an applied economic change. It overlaps the broader component table and is not an independent sample. Stage, actor, asset and the documented grouping of related components must be considered together: adding an accrued penalty to a later applied slash can count the same economic process twice.

4.3 Dictionaries, types and interpretive states

Documentation operates at two levels. The conceptual dictionary defines 33 economic variables: 28 core measures and five supporting variables. The core measures comprise four participation or concentration variables, nine incentive or friction variables, and fifteen outcome or accountability variables. Supporting

entries identify the episode, decision time, challenge deadline, cross-chain link grade and provenance. Their definitions specify units, derivation, decision-time availability, evidence class, missingness and validation rules. Separately, the physical catalogue describes 1,789 field entries across 95 Parquet tables. A conceptual variable and a physical column are different documentation objects; their counts are not alternative estimates of dataset dimensionality.

Table 8 summarizes selected field families. Raw monetary amounts are stored as integer-valued strings and interpreted with their asset and decimal fields, avoiding loss from floating-point conversion. Event times use 64-bit Unix seconds where a compatible time is available. JSON schemas define required fields and controlled values, including the 85-field cross-chain observation schema. Where supported, coverage distinguishes complete, partial, verified-zero, unknown, unavailable and not-applicable conditions; right censoring separately identifies unfinished lifecycles. Schema checks establish consistency with these tested definitions, rather than proving every historical interpretation correct.

Table 8. Selected field families in the common and economic representations. Field names illustrate the machine-readable schema; the deposited dictionaries provide complete definitions. Raw monetary values require the corresponding asset and decimal precision.

Field family	Representative fields	Type and analytical meaning
Identity and mechanism	oracle_network, mechanism_family, native_table, security_chain, delivery_chain	Strings identifying the mechanism, native table and chain context.
Observation unit	accountability_unit_type, accountability_unit_id, event_granularity	Strings or controlled values defining what a row represents.
Time	event_time_unix, reward_accrual_time_unix, reward_payment_time_unix	64-bit Unix seconds where available; event, accrual and payment times remain distinct.
Actors and roles	actor, actor_role, counterparty, counterparty_role	Native addresses or identifiers and interpreted roles; not natural-person identities.
Reward	reward_class, reward_amount_raw, reward_asset, reward_asset_decimals	Integer string, asset identity and decimal precision. Economic stage determines interpretation.
Penalty and forfeiture	penalty_accrued_raw, principal_slashed_raw, bond_forfeited_raw, fee_forfeited_raw, reward_forfeited_raw	Integer strings with asset metadata; distinguishes accrual, applied principal loss and other forfeitures.
Principal	principal_locked_raw, principal_returned_raw, principal_asset	Integer strings and asset metadata; separates exposed or recovered capital from gain.
Outcome basis	truth_basis, outcome_status, external_truth_available	Strings and booleans describing the basis for an outcome; protocol resolution is not external ground truth.
Source and rule	source_event, source_tx, source_block, source_contract, rule_id, parameter_version	Strings and integers linking observations to native evidence and the interpretation rule.
Realization	economic_kind, economic_evidence_class, realization_status, amount_raw, signed_amount_raw, cashflow_verified, state_delta_verified	Controlled values, integer strings and booleans recording the economic stage and evidence for realization.

Most curated tables use Zstandard-compressed Apache Parquet; tabular and structured metadata use UTF-8

CSV or JSON. Retained source packages use JSONL or compressed JSONL according to the originating interface. The 95-table canonical catalogue is accompanied by supporting files and application artifacts, including a prediction Parquet outside that catalogue. Source registers, input inventories and validation reports distinguish retained reconstruction inputs from derived outputs and from original acquisition evidence. This organization makes the resource usable at several levels of abstraction while preserving the boundaries between data access, computational reconstruction and scientific interpretation.

5 Coverage and analytical scope

The breadth of the resource lies in the mechanisms it connects, rather than uniform observation of every oracle activity. Five detailed panels expose multiple stages of participation, disputes, rewards or staking, while three extensions cover narrower processes. The common cutoff permits a shared upper boundary, but differences in deployment dates, implementation history and available interfaces create branch-specific lower boundaries. These differences should guide the definition of a comparative sample before any protocol-level statistic is calculated.

The Pyth state panel illustrates the importance of separating event history from retained state. Archived transactions begin on 11 September 2024, whereas the publisher–epoch panel contains 49 retained closed epochs from 17 July 2025 to 25 June 2026. The latter was reconstructed using a state snapshot collected on 21 July 2026. A later snapshot can support reconstruction of retained historical epochs without establishing complete coverage since deployment or historical decision-time availability of every field. The two observation windows therefore describe different evidence sources, rather than competing dates for one complete panel.

Completion status introduces a second analytical boundary. Among 42,676 UMA request rounds, 41,826 settle by the cutoff and 850 remain open. The open observations are right-censored lifecycles; removing them or assigning zero-valued settlements would alter the population and the outcome definition. Similar care is necessary where an interface exposes entitlements but not a disaggregated payment, or staking activity but not report-level accountability. Absence of an observed quantity can reflect the observation process as well as the economic process.

These boundaries suggest a disciplined use of the common representation. Researchers should compare mechanisms within compatible units and evidence states, define exposure periods from branch-level coverage, and preserve unresolved cases where the estimand requires them. Large row counts expand the observable record, but do not resolve differences in granularity, missingness or factual verification. The dataset’s analytical value rests on making those differences explicit enough to support well-defined questions about oracle institutions.

6 Validation and empirical checks

Validation assesses whether specified reconstructions satisfy structural, economic and computational requirements. Its strength depends on the tested population and acceptance rule. File readability, a unique identifier, a balanced settlement identity and successful replay establish different properties; none alone demonstrates complete historical acquisition or independent oracle truth. We therefore distinguish the release’s documented checks, direct checks on retained tables, historical protocol-level validation records and bounded computational replay.

6.1 Structural consistency and economic semantics

The reconstruction report lists 50 passing checks across the 95-table catalogue and associated artifacts. These checks address selected file identities, schemas, native keys, relationships and lifecycle constraints. The evaluated UMA populations include 42,676 request records, 42,676 associated flow records and 835 cross-chain link attempts. The DVM request population contains 3,744 records, while the distinct voter-level accounting population contains 2,663,684 records. Keeping these units separate is necessary to interpret both the checks and their denominators.

Economic validation also led to corrections in the classification of existing observations. Table 9 reports rule violations before and after correction for UMA challenge and settlement fields, Tellor dispute outcomes and DVM penalty accruals. The UMA corrections identify disputed winner-reward components totaling

201,500 USDC.e and 1,250 USDC, reported separately by asset. These amounts revise the measurement of existing settlements; they do not describe newly generated payments. For Tellor, three of the thirteen disputes support the challenge and have positive applied slashes, while the ten against-challenge outcomes have no positive applied slash under the recorded checks. This follows the protocol’s interpretation of SUPPORT as supporting the disputer.

Table 9. Violations of specified economic classification rules before and after correction. Populations differ by rule. A zero in the final column establishes compliance within the tested population, rather than exhaustive economic validity.

Validation rule	Population	Before	After
UMA: false challenged classification	42,676	40,992	0
UMA: transfer flag without settlement	42,676	850	0
UMA: disputed winner-reward omission	42,676	834	0
Tellor: inverted SUPPORT interpretation	13	13	0
Tellor: rejected challenge treated as applied principal slash	13	10	0
DVM: unapplied accrual treated as principal slash	2,663,684	947,128	0

In the DVM accounting population, 947,128 negative accrual records had previously been treated as applied principal losses without evidence of a corresponding reduction. They are now represented as penalty accruals. The correction preserves a liability-related accounting state while avoiding an unsupported claim of realized loss. More broadly, the results demonstrate why economic classification requires the mechanism’s state-transition rules: a syntactically valid record can still misstate the financial event if its observation stage is misidentified.

6.2 Accounting reconciliation and object identity

All 41,826 settled UMA requests satisfy the recorded integer payout identity,

$$\Delta_{\text{payout}} = G_{\text{actual}} - G_{\text{expected}} = 0, \quad (3)$$

with the associated transfer-matching checks passing under the applicable contract and adapter rules. Chainlink checks identify 51,372 qualified staking and reward records: 50,772 require transfers and match the recorded token movements, while 600 are accounting-only adjustments. The historical Flare report documents Merkle-root and claim reconciliation for 183 closed reward epochs. Those retained results are distinguished from a fresh rerun of the complete chain-level validation.

Identity checks provide a complementary form of evidence. Of 835 UMA cross-chain link attempts, 832 satisfy the deterministic Grade-A rule and three remain unmatched. The Pyth reconstruction covers 49 retained closed epochs; replay of the retained inputs checks the associated reward-state calculation within that window. These outcomes support specific accounting and reconstruction relationships. They do not establish universal source completeness, a fully verified ancestry chain for every derived row, or agreement between protocol decisions and external facts.

6.3 Bounded computational replay

The companion notebook provides a small UMA case-reconstruction route using fixed inputs and an explicitly recorded environment. Its nine code cells have been executed in a fresh local Python process, including two case runs with 28 passing checks each. The submitted saved notebook and replay package identify a local Jupyter runtime. This evidence supports local execution of the bounded workflow; it does not demonstrate a completed independent run in hosted Google Colab or reconstruction of the full historical Atlas from upstream services.

The submitted evidence package separately reports 101 passing archive checks and agreement of 27 scientific output files between two recorded executions. These records document package consistency and output agreement under the supplied inputs. They do not establish the execution environment or independent hosted reproduction of the full dataset.

The reproducibility distinction is substantive. Recomputing outputs from retained inputs tests the transformation of those inputs, whereas historical reacquisition tests a different chain of interfaces, source availability and collection decisions. The reported full reconstruction uses retained native or curated inputs, not a complete archive of original RPC responses. Local notebook success and the broader validation report therefore answer complementary computational questions. Their scope should accompany any reuse of the data or comparison of results.

6.4 Downstream prediction results

The UMA adjudication benchmark uses 810 challenged observations and a chronological 489/161/160 training, validation and test partition. Table 10 reports test results for the prior, logistic and histogram gradient-boosting models. Histogram gradient boosting reaches accuracy 0.5875 and AUROC 0.6540, exceeding the prior's 0.5375 and 0.5000. Its Brier score of 0.2767 is nevertheless worse than the prior's 0.2512. Better ranking performance therefore does not imply better probability prediction in this small test cohort. The target remains institutional adjudication, and historical decision-time availability of every reconstructed feature is not established by retrospective exclusion alone.

Table 10. UMA adjudication prediction on the 160-observation chronological test set. Higher accuracy and AUROC are better; lower Brier score is better. Values are rounded to four decimal places. The label is protocol adjudication, not independent external truth.

Model	Accuracy	AUROC	Brier score
Prior baseline	0.5375	0.5000	0.2512
Logistic regression	0.5188	0.4810	0.3496
Histogram gradient boosting	0.5875	0.6540	0.2767

Figure 2 presents these same test results as separate comparisons of classification accuracy, ranking and probability error. It makes the interpretation of the benchmark visible: an improvement in AUROC need not imply an improvement in the Brier score.




Figure 2. Descriptive comparison of the same UMA adjudication predictions reported in Table 10. All three models use the 160-observation chronological test set. Accuracy and AUROC are better when higher, whereas the Brier score is better when lower. Colours and marker shapes identify models, and labels show the reported point estimates. No uncertainty intervals are supplied by the source results. The target is institutional adjudication rather than independently verified external truth.

The financial time-to-event exercise reports 21 completed fitted-model runs across three tasks, with Kaplan–Meier references separate from those fits. Random fallback partitions in F1 and F2 prevent a claim of temporal

generalization for reward realization or penalty application. F3 capital settlement uses a temporal partition, with learned-model Uno C-index values of approximately 0.546–0.570. These exercises document usable task-specific views and their limits; they do not identify causal effects of incentives or certify trustworthy autonomous decisions.

The combined evidence supports specified structural, semantic, accounting, linkage and replay properties of the resource. It also identifies where inference should stop: untested historical periods, incomplete source ancestry, unresolved decision-time availability and external factual correctness require additional evidence. Preserving these distinctions is part of validation itself, because a reproducible calculation is useful only when the quantity it reproduces has a defensible interpretation.

7 Research use and interpretation

Effective reuse begins with the question and its observation unit. Studies of participation, concentration, dispute behavior and economic realization may use the same underlying protocol evidence but require different denominators, time windows and outcome states. Researchers should select the relevant native or common representation, consult its coverage metadata, and retain protocol, unit, actor, asset, stage and source identity when constructing an analytical sample. This makes the comparison answerable before it is made numerically convenient.

Parquet tables can be read with DuckDB, PyArrow or compatible analytical software. Column selection and predicate filtering allow work on the large tables without loading every row into memory; the repository notebook demonstrates a smaller reconstruction route. The catalogue documents file locations and relationships, while the dictionaries define the variables and their interpretation. Monetary values require their asset and decimal precision, timestamps require the branch’s native chronology, and physical row order is not an event-time definition. Addresses represent protocol actors or accounts rather than independently identified natural persons.

Economic aggregation requires particular care. Configured rewards, accruals, claimable entitlements, payments and applied penalties occupy different stages of a lifecycle. Returned principal must be separated from participant gain, and token gain still differs from net profit after costs. Related component groups should be checked before summation, and USDC.e and USDC should remain distinct assets unless an explicit conversion rule is justified. Missing, unavailable or not-applicable quantities should not be imputed as zero, and unfinished lifecycles should be retained as censored observations where the research design requires them.

The prediction exercises offer starting points for methodological work rather than general performance guarantees. UMA labels describe adjudication, while the financial tasks use their documented realization or settlement outcomes. Analysts extending these tasks should establish decision-time availability, define the exposure and censoring rules, and justify a split appropriate to the intended claim. The current financial split differences and limited UMA test cohort constrain generalization. Causal questions about incentive design require a separate identification strategy beyond the descriptive reconstruction supplied here.

The public repositories provide the current data, code, documentation and issue tracker. Researchers can use those entry points for access while recording the exact files and computational environment in their own study materials. Reconstruction from retained inputs should be described separately from reacquisition of historical records. In this way, the resource can support cumulative work: subsequent users can extend the analysis, report corrections and refine the evidence without losing the meaning or limitations of the original observations.

8 Data and code availability

The current data and documentation are available through the project’s public Hugging Face repository.¹ The deposit contains protocol-native and derived tables, retained reconstruction inputs, source registers, dictionaries, schemas and validation materials. The canonical catalogue describes 95 Parquet tables; the accompanying package also includes supporting and application artifacts. Croissant metadata currently describes 15 configured recordsets and is not a complete machine-readable inventory of the canonical

¹Data and documentation: <https://huggingface.co/datasets/Oracle4CEG/OracleEconLab-Oracle-Incentives-v1>.

catalogue. The dataset repository states the applicable data licence and the scope of each component.

Processing, validation and case-replay materials are available through the project on GitHub.² The repository includes the current working implementation and companion notebook, together with environment information, reproduction instructions and component-specific software licensing terms. These materials distinguish bounded case replay from reconstruction using retained inputs. These project-level links provide access to ongoing corrections and updates during working-paper development. The observation cutoff and the reconstruction scope described in this paper continue to define the data examined here.

References

- [1] Zhang, F., Cecchetti, E., Croman, K., Juels, A. & Shi, E. Town Crier: An authenticated data feed for smart contracts. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 270–282 (2016). <https://doi.org/10.1145/2976749.2978326>.
- [2] Eskandari, S., Salehi, M., Gu, W. C. & Clark, J. SoK: Oracles from the ground truth to market manipulation. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, 127–141 (2021). <https://doi.org/10.1145/3479722.3480994>.
- [3] Cong, L. W. & He, Z. Blockchain disruption and smart contracts. *The Review of Financial Studies* **32**, 1754–1797 (2019). <https://doi.org/10.1093/rfs/hhz007>.
- [4] Cong, L. W., Fox, L., Li, S. & Zhou, L. A primer on oracle economics. *Journal of Corporate Finance* **94**, 102800 (2025). <https://doi.org/10.1016/j.jcorpfin.2025.102800>.
- [5] Zhang, F., Maram, D., Malvai, H., Goldfeder, S. & Juels, A. DECO: Liberating web data using decentralized oracles for TLS. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*, 1919–1938 (2020). <https://doi.org/10.1145/3372297.3417239>.
- [6] Park, S., Bastani, O. & Kim, T. ACon²: Adaptive conformal consensus for provable blockchain oracles. In *32nd USENIX Security Symposium*, 3313–3330 (2023). <https://www.usenix.org/conference/usenixsecurity23/presentation/park>.
- [7] Zhang, B., He, N., Hu, X., Ma, K. & Wang, H. Following Devils’ Footprint: Towards real-time detection of price manipulation attacks. In *34th USENIX Security Symposium*, 4127–4145 (2025). <https://www.usenix.org/conference/usenixsecurity25/presentation/zhang-bosi>.
- [8] Liu, B., Szalachowski, P. & Zhou, J. A first look into DeFi oracles. In *IEEE International Conference on Decentralized Applications and Infrastructures*, 39–48 (2021). <https://doi.org/10.1109/DAPPS52256.2021.00010>.
- [9] Kaleem, M. & Shi, W. Demystifying Pythia: A survey of ChainLink oracles usage on Ethereum. In *Financial Cryptography and Data Security: FC 2021 International Workshops*, 115–123 (2021). https://doi.org/10.1007/978-3-662-63958-0_10.
- [10] Cong, L. W., Prasad, E. S. & Rabetti, D. Financial and informational integration through oracle networks. Working paper (2023; revised March 2026). <https://doi.org/10.2139/ssrn.4495514>.
- [11] Weber, M., Domeniconi, G., Chen, J., Weidele, D. K. I., Bellei, C., Robinson, T. & Leiserson, C. E. Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics. *KDD Workshop on Anomaly Detection in Finance* (2019). <https://arxiv.org/abs/1908.02591>.
- [12] Bellei, C., Xu, M., Phillips, R., Robinson, T., Weber, M., Kaler, T., Leiserson, C. E., Arvind & Chen, J. The shape of money laundering: Subgraph representation learning on the blockchain with the Elliptic2 dataset. *KDD Workshop on Machine Learning in Finance* (2024). <https://arxiv.org/abs/2404.19109>.

²Processing, validation and case replay: <https://github.com/Oracle4CEG/Oracle4SD-Atlas>.

- [13] Shamsi, K., Victor, F., Kantarcioglu, M., Gel, Y. & Akcora, C. G. Chartalist: Labeled graph datasets for UTXO and account-based blockchains. *Advances in Neural Information Processing Systems* **35** (2022). <https://doi.org/10.52202/068431-2531>.
- [14] Coquidé, C. & Cazabet, R. A temporal graph dataset of Bitcoin entity–entity transactions. *Scientific Data* **12**, 337 (2025). <https://doi.org/10.1038/s41597-025-04595-8>.
- [15] Schnoering, H. & Vazirgiannis, M. Bitcoin research with a transaction graph dataset. *Scientific Data* **12**, 404 (2025). <https://doi.org/10.1038/s41597-025-04684-8>.
- [16] Wang, Q., Zhang, Z., Liu, Z., Lu, S., Luo, B. & He, B. EX-Graph: A pioneering dataset bridging Ethereum and X. *International Conference on Learning Representations* (2024). <https://openreview.net/forum?id=juE0rWGCJW>.
- [17] Luo, B., Zhang, Z., Wang, Q. & He, B. Multi-chain graphs of graphs: A new approach to analyzing blockchain datasets. *Advances in Neural Information Processing Systems* **37**, 28490–28514 (2024). <https://doi.org/10.52202/079017-0894>.
- [18] Nadini, M., Alessandretti, L., Di Giacinto, F., Martino, M., Aiello, L. M. & Baronchelli, A. Mapping the NFT revolution: Market trends, trade networks, and visual features. *Scientific Reports* **11**, 20902 (2021). <https://doi.org/10.1038/s41598-021-00053-8>.
- [19] Long, H.-W., Wong, N.-M. & Cai, W. Bridging culture and finance: A multimodal analysis of memecoins in the Web3 ecosystem. In *Companion Proceedings of the ACM Web Conference*, 1158–1161 (2025). <https://doi.org/10.1145/3701716.3715561>.
- [20] Jia, H., You, J., Liu, J., Luo, Y. & Sun, S. Towards event-aware forecasting in DeFi: Insights from on-chain automated market maker protocols. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2*, 9171–9182 (2026). <https://doi.org/10.1145/3770855.3817502>.
- [21] Alsunaidi, S. J., Aljamaan, H. & Hammoudeh, M. DIVE: A multi-label smart contract vulnerability dataset. *Scientific Data* **13**, 664 (2026). <https://doi.org/10.1038/s41597-026-07025-5>.
- [22] UMA. Optimistic Oracle v2. <https://docs.uma.xyz/developers/optimistic-oracle> (accessed 13 October 2026).
- [23] Chainlink. Chainlink staking. <https://chain.link/economics/staking> (accessed 13 October 2026).
- [24] Tellor. Disputes and reporter governance. <https://docs.tellor.io/tellor/running-tellor/disputes-and-reporter-governance> (accessed 13 October 2026).
- [25] Flare Network. FTSOv2 scaling. <https://dev.flare.network/ftso/scaling/overview> (accessed 13 October 2026).
- [26] Pyth Network. Oracle Integrity Staking: Mathematical representation. <https://docs.pyth.network/oracle-integrity-staking/mathematical-representation> (accessed 13 October 2026).
- [27] Chronicle Labs. Chronicle documentation. <https://docs.chroniclelabs.org/> (accessed 13 October 2026).
- [28] RedStone. RedStone documentation. <https://docs.redstone.finance/> (accessed 13 October 2026).
- [29] DIA. DIA documentation. <https://www.diadata.org/docs/home> (accessed 13 October 2026).