What Public Records Reveal About Cryptocurrency Payments Between AI Agents

Vladislav Ch · Holder.io

Primary snapshot 8 November 2026; separate extension and revision 9 November 2026. Manuscript for author review; not submitted or externally peer reviewed.

Abstract

What do publicly reported agent-payment records establish about commerce once their units and source layers are separated? We audit an MPP catalogue with 138 service records and 1600 endpoint entries, a non-probability slice of 1000 CDP Bazaar resources, a fixed-block Olas frame of 1294 requests, twenty previously selected Olas orders and one purposive TaskMarket payout. A post-review extension covers 11,326 Olas requests over seven days; it contains, rather than adds to, the original day. Two results are central. First, all 2590 known-rate requests without a delivery marker in the week belong to one priority-Mech cluster: 370 on each day across fifty buyer addresses. Although 104 of 106 buyer addresses recur on multiple dates, only 55 have a request later than an observed delivery; neither count measures accepted-purchase retention. Second, comparison of the same twenty orders across sources distinguishes fourteen delivery markers, ten linked request–result file pairs and unresolved final withdrawals. The separate TaskMarket case links a task to a 2 OUSD transfer and a recipient balance change. Catalogue composition also matters: two retained 1000-resource CDP frames contain 2624 and 4325 payment offers. These findings identify specific composition and observability mechanisms. They do not establish independent customer demand, accepted quality, profitability or a causal price effect.

Keywords: agent payments; cryptocurrency; economic measurement; x402; MPP; service delivery; reproducibility.

Introduction

An AI agent may need external information, computation or execution to complete a task. Cryptocurrency payment interfaces can connect a software request to a transfer of value, potentially reducing the need for a separate payment account with every provider. Yet a payable endpoint is not a purchased service, a positive transfer is not necessarily external demand, and a delivered file is not necessarily a useful result. Measuring the emerging market therefore requires attention to the economic object behind each technical observation.

Prior work already distinguishes infrastructure, transactions and service outcomes. Ling et al. study x402 settlement concentration and the gap between advertised supply and settled value, while explicitly limiting inference about off-chain delivery [1]. Zhang et al. separate payment verification from service–payment coupling [3]. Our evidence vector operationalizes these distinctions; we do not claim to invent them.

The relevant Jin et al. version is **2606.25876v2, 2 September 2026** [2]. It identifies Base USDC settlements by matching decoded facilitator calls to transfer events and links registry identities

to wallets, with an explicit limit on agent attribution (§§3.1.2–3.1.3, 6.2.2). Its full text reports 386,665 identities, 11.64 million settlements and 605 sampled wallets. The landing-page abstract retains older corpus figures and is not used for those quantities. The manuscript's venue label is not independent publisher confirmation.

Agentic Commerce Bench [15] is included as a methodological comparator. Its generated sessions use operator aggregates, and its observation model distinguishes request-level evidence, accumulated history and principal authority. A request fingerprint separates repeated intent from a fresh purchase (§§3.2–3.3, 5.2). Our public orders lack comparable verified intent and principal fields. Neither its detector performance nor production-derived parameters are estimates of market-wide fraud or demand.

Our specific contribution is a reproducible comparison of the information added by source layers **for the same 20 Olas order identifiers**. Receipt decoding, linked files and cash-flow semantics answer different predicates. The observed 14 delivery markers and 10 request–result file pairs describe those predicates, not a quality attrition curve. The separately audited TaskMarket payout tests whether one payment can be connected across records; it is not pooled with Olas into a market rate. We did not reproduce the cited authors' experiments or claim the first measurement of agent payments.

Gong connects mandates, counterparties, payments and results in a lifecycle evidence envelope [4]; Schär et al. interpret stablecoin events within economic transactions [5]. Legal research addresses responsibility and institutional constraints [6]. Our targeted literature matrix contains sixteen works with separate methods, datasets and limitations; this manuscript concentrates on the nearest comparisons rather than claiming a systematic review. The new weekly extension tests whether the observed technical composition recurs beyond the original day. It remains a descriptive analysis, not a causal adoption or price experiment.

Methods

Units and evidence dimensions

We distinguish organisation, product, technical host, endpoint, payment offer, logical order, transaction, accounting accrual and transfer to a recipient. A payment offer specifies one available set of terms; several offers may accompany one endpoint. A logical order may require several transfers, while prepayment or batching may fund several orders. Deposits, subsequent expenditure and withdrawals cannot automatically be added as separate final demand. Capital supplied for trading is distinct from the fee for executing a trading service.

The adapted evidence vector records buyer agency, seller agency, ownership relationship, authority, terms, monetary stage, delivery, acceptance, implementation version and source availability. Each dimension can be supported, contradicted or unknown according to its specific evidence. The vector is not an ordinal maturity score. A case with verified payment but inaccessible content is not universally stronger than one with content but an unresolved payout.

For agency classification, A requires evidence of both agent roles in the particular operation; B requires an agent buyer purchasing an ordinary service; U denotes insufficient evidence for either classification. A registry entry, model name or SDK capability does not establish the

decision process of a specific buyer. Ownership is coded separately as verified distinct, verified common or unknown. Different addresses and a common gas sponsor are not, by themselves, determinations of beneficial ownership. Contractual completion and independent evaluation of usefulness are also separate fields.

MPP denotes the Machine Payments Protocol catalogue; CDP denotes Coinbase Developer Platform. Here a delivery marker is the indexed or event-level signal specified for a request, not an independent judgment that its result was accepted.

Sources and sampling frames

The primary observation date was 8 November 2026; new document reads and a separately frozen weekly extension were added on 9 November. Collection used public documents, catalogue APIs, an Olas indexer and read-only blockchain RPC responses. No study-funded purchase, wallet signature, paid endpoint experiment or account registration was performed. MPP reused an exploratory snapshot obtained before the local analysis-protocol freeze; its dimensions were already known. The prespecified CDP pass and expanded Olas pass followed their respective local design freezes, and the 20-request audit selection preceded payload retrieval. These local records were not an external preregistration. Source versions, access outcomes and retained-file hashes accompany the observations.

Source	Unit and retained frame	Count	Inferential boundary
MPP	Records returned by one catalogue response	138 services; 1600 endpoints	Supply description of that response
CDP Bazaar	Ten prespecified server-order pages	1000 resources	Non-probability slice, not a census
Olas	Requests matching a day/filter at one block	1294 requests	Indexer and deployment coverage
Olas extension	Seven UTC days at a new fixed block; contains the original day	11,326 requests	Post-review descriptive extension
Detailed cases	Prior Olas audit; purposive TaskMarket case	20 requests; 1 payout	Different selection mechanisms

The MPP response was normalised into service, endpoint and offer tables [7]. Two service records advertised only Stripe; 136 advertised a cryptocurrency-capable method, including mixed-method records. This does not establish production execution or the network environment of every route. CDP supplied 100 resources at each of the prespecified offsets 0, 3700, 7400, 11100, 14800, 18600, 22300, 26000, 29700 and 33500. Its reported total during that pass was 33602, compared with 33603 in the initial observation. The separate exploratory first-1000 pull overlapped by 100 resources and was not pooled into the analytical denominator. Discovery documentation describes the source interface [8]; it does not make this selection random.

Technical matching used the complete HTTP method–URL pair. Host matching used the lower-case hostname, without collapsing cloud subdomains or treating similar product names as the same owner. Multiple advertised payment conditions were retained. Cross-catalogue overlap therefore measures exact technical coincidence within the observed frames, not economic substitution or overall interoperability.

The Olas window was 7 November 2026 UTC, with an exclusive upper boundary at 8 November 00:00. A previous pilot retained the first 300 identifiers at Gnosis block 49172789 and verified truncation with a look-ahead record. The expanded pass fixed block 49174181, whose timestamp was 8 November 13:46:25 UTC, and used ascending identifiers with an id_gt cursor. Thirteen pages contained twelve groups of 100 and a final 94. The limit of 5000 was not reached. Completeness refers only to matching request entities of that deployment. Olas documents offchain activity and observation rules that do not expose all attempted requests [9].

Before content retrieval, the retained audit selected 20 of the original 300 requests by sorting SHA-256 hashes of a fixed study prefix concatenated with request ID. This deterministic selection is reproducible; it is not a random sample of the later 1294. The audit linked request and delivery transaction identifiers, inspected receipts and attempted retrieval of associated content. The result’s requestId was compared with the onchain order ID; request metadata did not itself contain requestId. Request files were linked through the associated event and content reference. Gateway retrieval and these links were not treated as a complete independent verification of the IPFS structure or model execution.

TaskMarket was examined as a purposive monetary case after a public task page exposed a payout reference. The study compared task ID, transaction, receipt, block, token decimals and recipient balances through the official Tempo RPC [10–12]. The case supplements, rather than expands, the Olas sampling denominator. Documentation and retained source code helped distinguish accrual from payout [13]. Code interpretation was not represented as verification that every historical deployed bytecode matched the inspected source.

Analysis

We added three post-review analyses of retained records: technical clustering in the Olas frame, marker occurrence at common follow-up horizons, and sensitivity to catalogue composition. These were not prespecified confirmatory tests. For Olas, we cross-tabulated requested rate, buyer address, priority-Mech address and delivery marker; absent priority values were not counted as an additional actor. Request age was the fixed-block timestamp minus request blockTimestamp. For each horizon h, a request was observable if its age was at least h; a marker occurred by h if the recorded delivery timestamp was at most h seconds after the request timestamp. The horizons were five minutes, one hour and twelve hours. These outcomes concern indexed markers, not eventual completion, independently accepted quality or model-computation latency.

Three repetition predicates were analysed separately: more than one request from an address, requests at more than one block timestamp, and a request later than an observed delivery for a previous request from that address. The third predicate uses strictly later timestamps. No predicate identifies a beneficial owner, the first-ever purchase or acceptance of an earlier result. Transaction cluster sizes count only orders included in the study query.

Source-stage comparison retained the same 20 audited order IDs throughout. We separately tested indexed delivery, nonempty indexed result reference, request/event linkage, debit or reservation context, delivery accrual, retrieved request JSON, retrieved result JSON, and equality of the result requestId to the order ID. Receipt results are retained as named predicates; local gateway-file hashes and result identifiers are rechecked by the revision analysis. File hashes do not substitute for full UnixFS/CID verification or attestations of model execution. The complete evidence table distinguishes an unknown proposition from an unperformed check, a failed access attempt, and a procedure inapplicable to the observed state.

For CDP, the original systematic server-order slice was compared with the retained exploratory first-1000 pull. The passes occurred at different times and had no shared snapshot token. We therefore compare retained frames rather than estimate an isolated effect of sampling strategy. Resource URLs shared by both frames were identified; the frames were not pooled. MPP missingness was tabulated by technical address family, using endpoints as the denominator for missing payment metadata and offers as the denominator for missing amounts or billing units.

Amounts were computed exactly from retained values, with eighteen decimals for native Olas rates. Unknown amounts and units were not assigned zero. Non-probability frames and unresolved principals preclude population confidence intervals, causal regressions or thirty-day accepted-purchase retention from these data.

Results

Catalogue composition and missingness

MPP contained 1408 offers; 192 endpoint entries had payment set to null, which does not establish free access. There were 1595 distinct method–URL pairs among 1600 entries. Catalogue labels classified 38 services as first-party and 100 as third-party. Ninety-seven technical hosts represented the 138 service records. Three address families accounted for 93 cards and 977 endpoint entries: Orthogonal 42/580, Locus 32/205 and Tempo 19/192. These are technical groupings, not market shares or three verified owners.

In CDP the largest resource had fourteen offers. Exact comparison found seven shared hosts and seven shared method–URL pairs. The technical-host union was 300: 210 plus 97 minus seven. Neither this union nor the number of offer rows measures companies, purchases or customers.

Price metadata also limited comparison. In MPP, 1074 offers supplied an amount, 243 indicated dynamic pricing without an amount, and 91 supplied neither. A billing unit appeared in 795. Network identifiers were absent from all 1408 normalised MPP offers, and asset decimals were not independently established. CDP offers supplied network and amount fields, but independent decimal verification was also absent. A pooled dollar price index would therefore conceal unresolved units and service heterogeneity.

The exploratory CDP head1000 pass contained 2624 offers, 388 hosts, 480 multi-offer resources and a median of one offer per resource. The systematic slice contained 4325 offers, 210 hosts, 560 multi-offer resources and a median of two. One hundred resource URLs

overlapped. The head pass ran on 8 November from 11:57:06.628657 to 11:57:14.982038 UTC; the systematic pass ran from 12:00:09.380705 to 12:00:23.240976 UTC. Catalogue change between the passes cannot be excluded.

The systematic pages at offsets 3700 and 11100 contributed 2482 of 4325 offers, or 57.4%, from only 20% of resources. The pages at offsets 18600 and 29700 each consisted entirely of one technical host. Across the ten pages, offer counts ranged from 100 to 1326 and the number of multi-offer resources from zero to one hundred. The page-level supplement reports all ten compositions. The differences demonstrate sensitivity within retained supply frames; they are not population market shares or a causal estimate of selection effects.

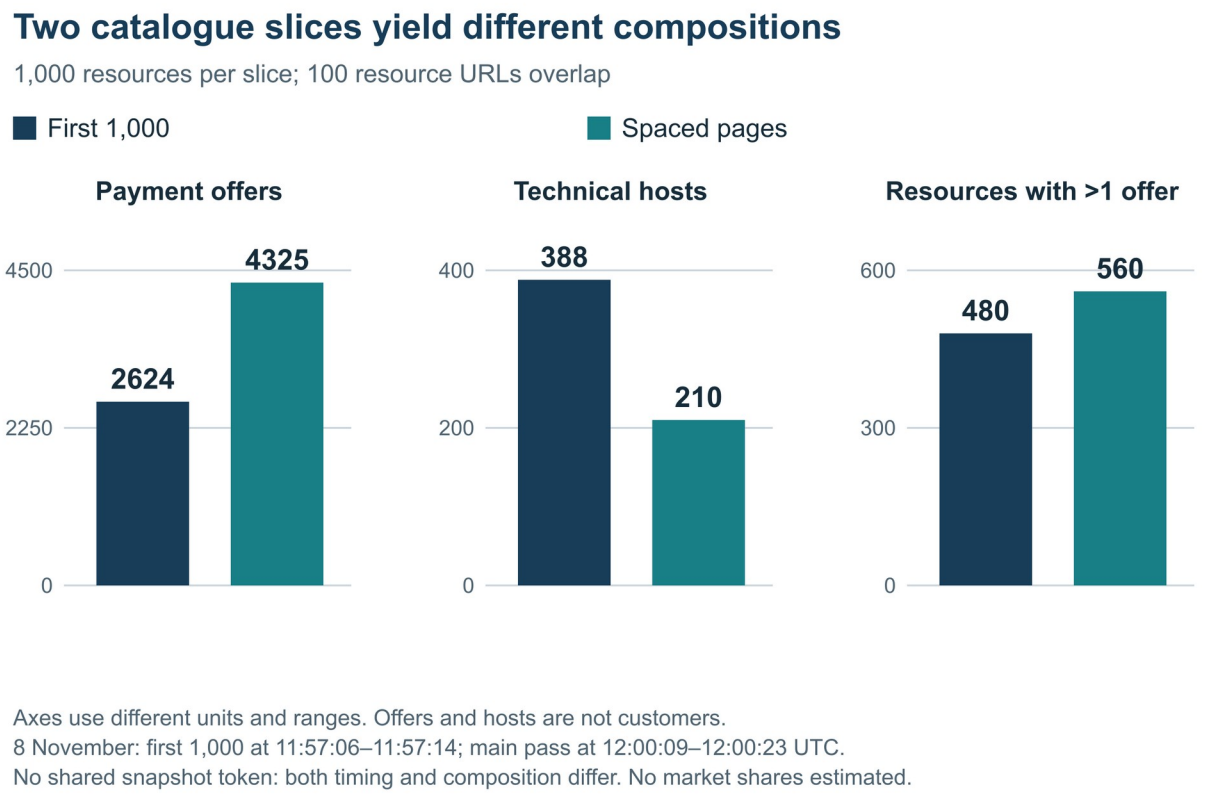


Figure 1. Two retained CDP frames with 1000 resources each differ in offers and hosts. Collection times and composition differ; the frames are not pooled.

MPP address family	Endpoints	payment:null endpoints	Offers	Offers without amount	Offers without billing unit
Orthogonal	580	0	580	103	103
Locus	205	2	203	38	38
Tempo	192	23	169	91	169
Other	623	167	456	102	303

Among offers without amounts, 243 declared dynamic pricing and 91 did not provide either an amount or a dynamic-pricing indication. Missing payment metadata, missing amount and missing billing unit are separate conditions with different denominators; none is assigned a zero price. Technical-family concentration in missingness further limits pooled price comparisons. A technical family is not a verified common owner.

Olas technical clusters

The Olas frame contained 1294 requests, 919 delivery markers, 87 buyer addresses and fifteen known priority-Mech addresses. Fee groups had disjoint sets of buyer and known priority-Mech addresses.

Requested rate	Requests	Buyer addresses	Known priority Mechs	Delivery markers	Indexed result references / markers
0.01 xDAI	622	25	3	622	622 / 622
2×10^{-18} xDAI	667	61	12	297	0 / 297
Unknown amount and unit	5	1	0	0	Not applicable: no markers

All 370 requests with known rates but no marker were assigned to one priority Mech, 0xb62d2623e751ebfb2d57e30aa86cfacb21ff1ea9, and came from fifty buyer addresses. At buyer-address level, the observed marker fraction was zero for 51 addresses and one for 36. Within the 0.01-xDAI group, one priority Mech accounted for 487/622 requests and therefore 78.3% of the sum of equal stated rates in that group. The other two received 121 and fourteen requests. These are technical concentrations, not independent-owner revenue shares.

The groups differ in their composition, and there are no within-address observations at different rates. The difference between 622/622 and 297/667 is therefore not an identified price effect. The observed association is between non-overlapping technical regimes comprising rate, address composition and result-reference availability. Task quality, buyer intent and acceptance were not available as adjustment variables; they were not imputed.

Request ages (minimum, median, maximum) were 49,610/91,607.5/135,950 seconds in the 0.01-xDAI group, 51,270/92,430/134,330 in the two-atomic-unit group, and 133,775/134,765/135,770 among unknown rates.

The sum of known requested rates was 6.220000000000001334 xDAI; indexed delivery rates totalled 6.220000000000000594 xDAI. These are different accounting fields, not revenue or profit. Strict positivity retained 1289 requests, whereas thresholds of 0.000001, 0.0001 and 0.001 xDAI each retained 622; this is definition sensitivity, not proof that tiny-rate work was fictitious.

Common follow-up and repetition

Minimum request age was 49,610 seconds, so every request was observable for each supplementary horizon. Recorded delivery lags ranged from ten to 290 seconds, with median ninety seconds.

Horizon	Requests with full follow-up	Marker by horizon	Finite-frame fraction
Five minutes	1294	919	71.02%
One hour	1294	919	71.02%
Twelve hours	1294	919	71.02%

Common follow-up resolves unequal observation duration for these specific marker outcomes. It does not establish final failure for the remaining 375 requests, independent quality, or a population completion probability. A missing marker at twelve hours may differ from eventual non-completion; the data do not resolve that later outcome.

All 87 buyer addresses appeared in more than one request, but only 37 appeared at multiple block timestamps. Twenty-six had at least one request strictly later than an observed delivery for an earlier request. Fifty addresses had all their requests at one timestamp: thirty addresses contributed seven requests each and twenty contributed eight, totalling the same 370-request cluster without markers. Technical repetition thus includes both simultaneous series and later activity. None of these counts estimates accepted-purchase retention.

The frame contained 1074 request transactions: 1044 with one included order, twenty with ten and ten with five. These counts concern included orders, not necessarily the complete content of each transaction. The single-day observation window does not identify a customer's first purchase or prior activity.

Post-review seven-day extension

A supplementary query collected the complete matching indexer frame for 1–7 November UTC at Gnosis block 49182979 (9 November 2026, 02:12:15 UTC). The window, 20,000-row cap, keyset pagination and descriptive analysis were fixed before fetching new rows; the previous one-day findings were already known. This extension is therefore exploratory, not an original preregistered confirmation. The final page contained 26 rows, yielding 11,326 unique requests across 114 pages. All 115 responses, including metadata, were retained; response hashes and local reconstruction were checked. Completeness refers to the specified filter and indexer deployment, not all attempted purchases.

All requests were observable for 12 hours; the minimum age was 94,360 seconds. There were 8,673 delivery markers within this horizon: 6,588 of 6,588 requests at 0.01 xDAI and 2,085 of 4,675 requests at two atomic units. Sixty-three requests had unknown rates and no marker. Known declared rates summed to 65.880000000000009350 xDAI; this is not seller revenue, and unknown rates were not imputed as zero.

The same priorityMech accounted for all 2,590 known-rate requests without a marker: exactly 370 on each of the seven days, across 50 buyer addresses. The remaining 63 requests without a marker had an unknown priorityMech and one buyer address. The cluster's purpose and beneficial ownership were not established. This repeated structure supports a composition-based interpretation of the one-day aggregate, rather than a causal interpretation of rate differences or an assertion of artificial activity.

The week contained 106 buyer addresses and 16 known priorityMech addresses. Of the buyers, 104 appeared on multiple UTC dates, while 55 had a request timestamp later than an observed delivery timestamp. Multiday reuse thus includes the recurring no-marker cluster. Neither proportion measures retention of independent customers or accepted purchases. Principal identity, purchase intent, acceptance and a common 30-day follow-up are unavailable.

The new frame contained the same 1,294 IDs for 7 November. No values changed in the six comparison fields — isDelivered, feeRaw, feeUnit, deliveredByMech, finalFeeUSD and priorityMech — over 44,750 seconds between snapshots. This tests those fields' temporal stability; it does not establish stability of all source content or future payouts. The weekly frame contains the original day and is never pooled with it as an independent sample.

Day UTC	Requests	Markers by 12h	0.01 xDAI	2 wei	Unknown rate
2026-11-01	1573	1203	905	668	0
2026-11-02	1813	1443	1145	668	0
2026-11-03	1741	1371	1073	668	0
2026-11-04	1439	1069	771	668	0
2026-11-05	1873	1447	1149	668	56
2026-11-06	1593	1221	923	668	2
2026-11-07	1294	919	622	667	5

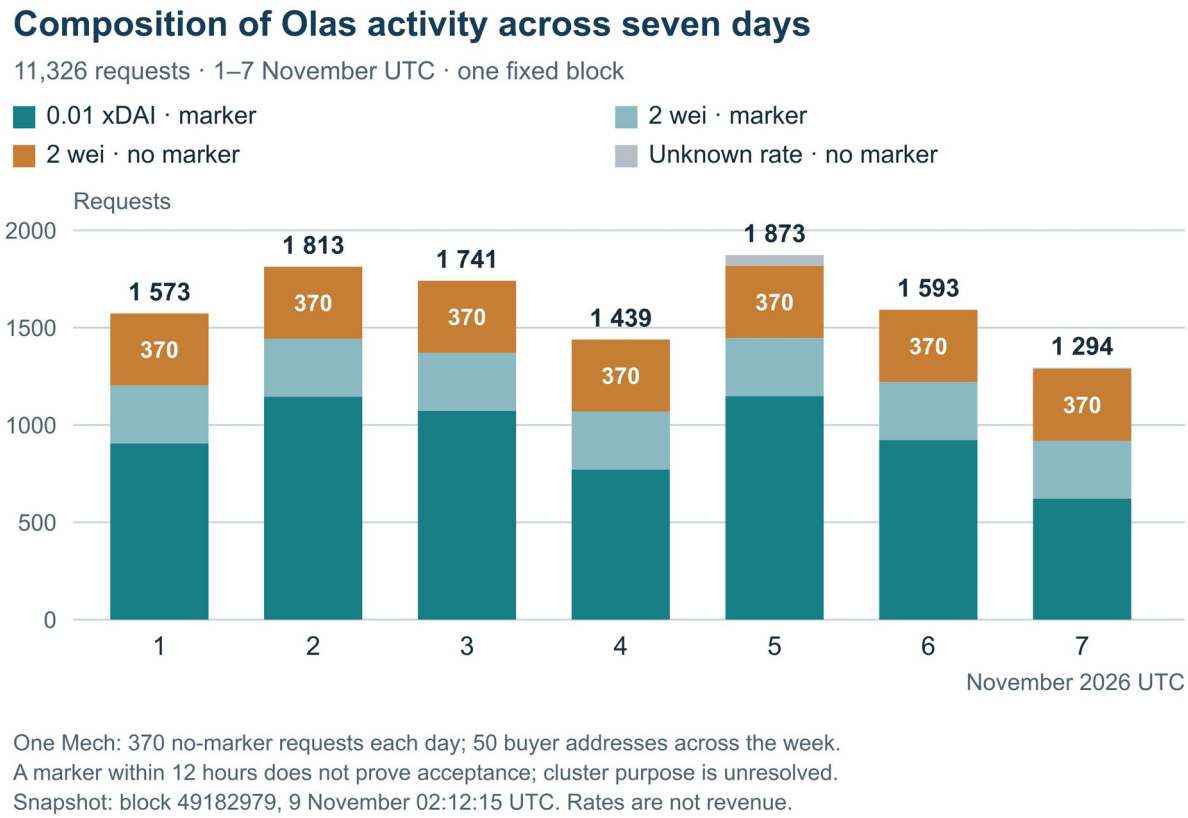


Figure 2. Weekly composition of Olas requests. The original day is contained in this frame; rate classes and delivery-marker states are distinct from revenue and customers.

Incremental information on the same audited orders

The audit retained twenty orders selected from the earlier truncated frame of 300 before payload inspection. The old frame contained 141 ordinary-rate and 159 minimum-rate requests; the audit contained ten of each. Its fourteen buyers and five known priority Mechs do not establish representativeness of the expanded frame.

Source stage and predicate	Supported among the same twenty orders	Information obtained
Indexer: request entity	20	Record belongs to the retained frame
Indexer: delivery marker	14	Indexed delivery status
Indexer: nonempty result reference	10	Reference exists; file access not yet established
Receipt: request ID and debit/reservation context	20	Order-to-request and monetary-stage linkage
Receipt: delivery ID and accrual	14	Delivery and accounting accrual corroborated

Source stage and predicate	Supported among the same twenty orders	Information obtained
Receipt: delivery bytes equal 0x01	4	Non-digest marker in these four inspected receipts
Gateway: request JSON retrieved	10	Input content available
Gateway: result JSON retrieved and requestId matched	10	Output content linked to the same order
Final payout allocated to order	Unresolved	Not an estimated zero-payout rate
Independently accepted result	Not assessed	No acceptance numerator is inferred

Nineteen request transactions and fourteen delivery transactions required 33 distinct receipts; one request transaction covered two sampled orders. All ten recovered request–result pairs had the 0.01-xDAI rate and the same priority Mech. Among the ten minimum-rate requests, four had delivery markers without usable result references and six had no observed marker.

Request-content access failed for ten orders. This is an access outcome, not proof that inputs did not exist. The four one-byte delivery values were verified only in the audited receipts; their format is not imputed to the other 293 minimum-rate deliveries in the expanded frame. For the six orders without markers, retrieval through an observed delivery reference was inapplicable because that stage/reference was absent; no search of every possible output store was performed.

The linked-file count is 10/20 among audited orders, 10/14 among orders with delivery markers, and 10/10 among result GET requests actually issued to available references. These denominators answer different questions. The full frame has indexed references for 622/622 ordinary-rate deliveries and 0/297 minimum-rate deliveries, but 622 references do not establish 622 accessible or verified files. Final monetary payout allocation remains unresolved, so a numerical final-payment linkage proportion is not reported.

Complete evidence dimensions

Dimension	Olas audit, twenty orders	TaskMarket, one purposive case
Buyer agency	Unknown for twenty; operation-specific decision traces unavailable	Unknown; autonomous purchase decision unestablished
Seller agency	Unknown for twenty; model/tool labels do not prove autonomous execution	Unknown; autonomous execution unestablished
Ownership relationship	Beneficial owners not assessed	Beneficial owners not assessed

Dimension	Olas audit, twenty orders	TaskMarket, one purposive case
Authority	Technical requester linkage observed; legal mandate not assessed	Technical roles observed; legal mandate not assessed
Terms	Rate and asset observed; governing terms, acceptance criteria and refund rights not assessed	Task ID, asset, amount and platformFee observed; contractual terms not assessed
Monetary stage	Debit/reservation in twenty and accrual in fourteen; final withdrawal unresolved	2 OUSD transfer and recipient balance change verified; USD value and profit unknown
Delivery	Fourteen markers; ten linked file pairs; four without usable result references; six without markers	TaskCompleted event observed; deliverable inaccessible to study
Acceptance	Independent quality/acceptance not assessed; forecast outcomes unresolved at cutoff	Independent quality/acceptance not assessed
Implementation version	Blocks, addresses and pinned source retained; historical bytecode match not assessed	Historical Diamond revision 26 established; full bytecode-to-source match not assessed
Source availability	Receipts retained; ten request files accessible and ten access failures; result states separately coded	RPC responses retained; deliverable inaccessible; no independently operated node or cryptographic receipt proof

The per-order supplement retains all ten dimensions and the reason for each unknown or unassessed field. Technical signatures do not establish legal authority, and order parameters do not substitute for contractual terms. These are limits of the investigation, not claims that mandates or contracts do not exist. The single TaskMarket case is not pooled with the Olas audit to estimate a combined proportion.

Content and monetary stages

The ten pairs concerned nine binary forecasting questions whose closing dates were after the snapshot. Their declared rates were 0.01 xDAI each. One result reported empty retrieval, probabilities 0.5/0.5, confidence zero and information utility zero. These are provider-supplied fields, not an independent quality verdict: an honest abstention can itself be useful. They nevertheless demonstrate why a delivery count cannot substitute for forecast accuracy or calibration. Other reported cost fields lacked verified units and were not compared with xDAI as evidence of profit or loss.

Retained Olas receipts exposed accounting updates. Inspection found no Withdraw event in the 33 audit receipts or two additional pilot receipts. That finding is restricted to those transactions. A later Gnosis RPC attempt returned 403; the planned wider withdrawal-log query was not sent. The retained source distinguishes accumulation in the Mech ledger from a subsequent payment

operation [13]. A cumulative balance may cover several orders, and the first payout recipient may be a Mech rather than its operator. Final withdrawals and their allocation to sampled orders remain unresolved. A final-payment linkage proportion was not estimated; 10/20 describes file linkage, not successful A2A trade.

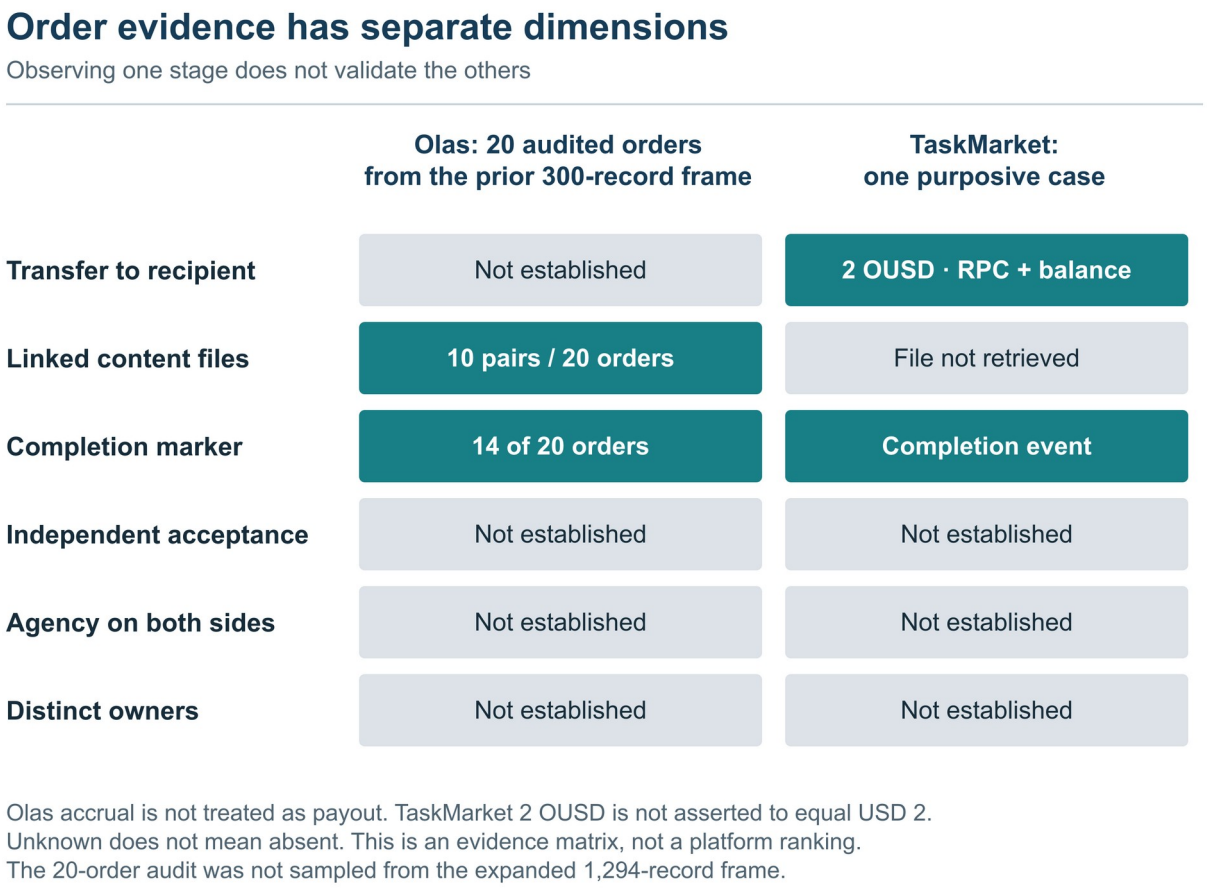


Figure 3. Non-hierarchical evidence dimensions for the prior Olas audit (20 of 300) and one purposively located TaskMarket case. Colour marks the stated evidence, not verified quality.

TaskMarket provides a different observation. The Kitchen Rave task, TSK-EV3WPBDB, linked to transaction 0xb5b6d20bfda9556201143722c6bcbe96f22078d554ad3351dad811972e18768d on Tempo, chain 4217. At block 46391423 on 29 October 2026, the successful receipt, TaskCompleted identifier and token Transfer agreed. Historical decimals were six, and the worker balance changed from zero to 2000000 atomic units. The block contained only that transaction. Thus, 2 OUSD reached the identified recipient. The historical Diamond revision was 26; the event's platformFee was zero. A separate gas-token transfer of 851 atomic units, also with six decimals, was paid by a relayer. It was neither the worker payment nor a nonzero platform fee.

The task case does not establish USD parity or redemption of OUSD. The deliverable was unavailable to this study, and independent quality assessment, beneficial ownership and agency were unresolved. Public RPC readings corroborate the interface with underlying chain state; they are not an independently operated node or cryptographic receipt proof.

Version and access constraints

The earlier Olas block became unavailable from the same public indexer endpoint: an HTTP 200 response contained a GraphQL error reporting a later earliest available state. The reason was not determined. Retained responses allowed local analysis despite failure to replay the original historical query. Separately, a bounded Virtuals event-history request returned 429 and supplied no event array; it was not counted as zero activity.

The retained ERC-8183 draft contains a disagreement between normative requirements and its reference example. The normative funding operation checks an expected budget, rejects a zero budget and requires funding before submission. The example omits the expected-budget parameter and permits zero-budget paths. The pinned Virtuals SDK ABI includes expectedBudget, but an ABI does not establish runtime guards. Consequently, a zero-valued release emitted by the example cannot be attributed automatically to the normative lifecycle or a deployed implementation. Nine local checks reproduce the comparison from the archived document and interface. Deployed bytecode was not audited. ERC-8183, pinned ABI.

The supplementary comparison separates authorization binding, aggregate liabilities, one-time payment proof, one-time economic effect, recovery, revocation and independent quality. Four architectural patterns are examined across six common dimensions, yielding 24 source-linked cells. Normative requirements, recommendations, vendor descriptions, static code observations and unknown implementation properties remain distinct. No vendor deployment is certified. Fifteen finite synthetic traces test the researcher's definitions, including a fresh signature for the same order, concurrent reservations, a lost reply and revocation after signing. Nine deliberately violating traces are detected. These method checks provide neither incident frequencies nor evidence of real-service security.

New primary-source reads on 9 November are separated from the original snapshot. The x402 specification is pinned to commit 51a0664c2b7826e3cb02595481eb6848a1b11815; seven documentation responses retain access times and hashes. They strengthen subsequent reproducibility without retroactively filling missing historical versions.

Discussion

The results identify the composition of observed activity. In the day frame, rates, marker occurrence and result-reference availability coincide with disjoint address groups. The weekly extension shows that the 370-request no-marker cluster recurs on every day. A direct comparison of 622/622 with 297/667 therefore does not identify a price effect. The groups' purpose and ownership remain unknown; repetition alone is evidence of neither artificial activity nor independent demand.

Keeping order IDs fixed prevents a changing sample from masquerading as improved verification. Receipts corroborate request context and accrual and distinguish a one-byte marker from a file reference. Gateway retrieval supplies ten linked pairs of content. This makes the lifecycle distinctions discussed by Jin, Zhang and Gong [2–4] operational: a named predicate is resolved by an additional source on the same orders. Their methods were not replicated or re-evaluated.

Technical recurrence also has several meanings. Reuse at one timestamp, on another day and after a delivery are observable properties; accepted repurchase by an independent principal requires more evidence. The 104 multiday addresses include the no-marker cluster, whereas 55 addresses have a request after an observed delivery. An intent-aware benchmark [15] and this public natural corpus answer different questions; frequencies from one cannot fill missing fields in the other.

The CDP comparison shows how supply characteristics depend on collection composition and timing. A non-probability frame can be described precisely without extrapolation to a market. A larger offer-to-resource ratio does not prove inflation in a particular external estimate; that requires examining its denominator and aggregation. Technical-family concentration in MPP missingness further limits a pooled price index.

Payment and useful execution are complementary. Olas supplied files without a resolved final withdrawal; TaskMarket supplied a verified transfer without inspected output or established agency. Assigning either case a single high-quality agent-commerce assessment would conceal those differences. Not every field must be public: accountable verification may release limited attestations and aggregates without disclosing mandates or sensitive task content.

Authority and contractual terms remain separate dimensions. Under 15 USC 7001(h), electronic-agent participation does not alone defeat enforceability provided the action is legally attributable to the person to be bound [16]. The provision does not establish that attribution for these transactions. A technical signature and order parameters do not replace contractual evidence.

The original commercial questions remain open. Accepted external demand and profitability require acceptance, owners, costs, funding, subsidies and refunds. The results do not imply an absence of autonomous commerce: U denotes insufficient evidence. They identify the additional observations that could strengthen each claim.

Limitations and subsequent design

The sources have heterogeneous selection mechanisms. MPP is one catalogue response; CDP is a selected server-order slice; Olas is a single deployment with an original day and a separate seven-day extension; TaskMarket is a purposive positive case. Their denominators cannot be pooled. A complete query result is not a complete population of attempted purchases, and content availability may select particular services. The deterministic audit was fixed before payload inspection but inherits the first-300 frame. It does not provide an unbiased estimate for the expanded cohort.

The evidence review is targeted rather than systematic. Some primary documents were read without retaining full immutable copies. We did not reproduce the studies cited as predecessors, independently identify all principals, resolve future forecasting outcomes or audit every deployed contract. Independent human coding agreement was not measured. These limitations constrain claims about coverage, classification reliability and causation; additional rows alone would not remove them.

A stronger follow-up would prospectively link mandates, owners, orders, settlement, refunds and accepted results. Comparable payment experiments should hold the service and quality

threshold fixed, distinguish a funded wallet from its complete funding path, and compare cryptocurrency with ordinary credits or invoicing. Failed attempts and recovery costs belong in the numerator of cost per accepted result. Sampling and power should be determined from actual variation and clustered decision-making units, not from treating every transfer as an independent customer. Such experiments were outside the present execution scope.

A prospective accepted-purchase R30 requires a defined entry cohort, a prior-history window, a verified principal-identification source and thirty complete days after entry. The numerator should count principals with a subsequent accepted external paid order, not retries or address reappearances. Incomplete follow-up and loss of linkage must remain separately censored or unknown; refunds and subsidies must stay linked. No such rate is estimated here. The weekly design was fixed before its new rows but after earlier findings were known, so its recurrence results remain exploratory.

Conclusion

Public records reveal a persistent technical cluster, composition-sensitive catalogue statistics and identifiable additions of information across source layers. The weekly Olas extension contains the same 370-request no-marker pattern on each day, while the twenty-order audit links fourteen deliveries and ten file pairs without establishing final withdrawals. TaskMarket supplies a different positive result: one task-linked 2 OUSD payout. Preserving these units and evidence dimensions supports reproducible empirical claims while leaving independent accepted demand and profitability unestimated.

Author declarations and data availability

Vladislav Ch is the named author; Holder.io is the associated project. AI tools assisted source discovery, extraction, analysis, drafting and internal checking. Such checks are not independent human peer review. The author must approve the manuscript and remains responsible for its interpretation and publication declarations.

Holder.io is a commercial project receiving affiliate revenue. The author reports no equity interests in studied companies. Counterparty relationships relevant to this particular corpus and the funding arrangements for this study have not been confirmed. No declaration of an absence of conflicts or external funding is therefore made; the author must finalise those statements before submission.

The local research package retains catalogue tables, collection designs, requests and responses, audit selections, content inventories, calculations and source metadata. Public release and an open licence have not been assigned. Sensitive payment credentials, unnecessary personal data and third-party redistribution rights require review before release. The manuscript itself has not been submitted or published.

New document reads on 9 November retain their own dates and evidence types; they do not restore unsaved 8 November versions. Original responses, tool-extracted text and analytical tables are labelled separately. Permission to redistribute full third-party material, including the retained CDP raw catalogue, is not established. The current package supports internal verification and has no blanket open licence.

Local companions: analysis protocol, catalogue data, original Olas frame, method corrections and per-order evidence, weekly data and validation, cash-flow evidence, sixteen-work literature matrix, and versioned security comparison.

References

[1] Ling S, Zhou Y, Wu L, Wang C. How Agentic Is Agentic Commerce? A Population-Scale Measurement of x402 Adoption and Authenticity. arXiv:2607.12575v1, 2026. Primary text.

[2] Jin Y, Wu S, Chen C, Bao L, Yang X, Chen J. The Web4 Agent Economy: A Large-Scale Empirical Study of the Landscape, Challenges, and Opportunities. arXiv:2606.25876v2, 2026. Primary text.

[3] Zhang Y et al. SoK: Blockchain Agent-to-Agent Payments. arXiv:2604.03733v1, 2026. Primary text.

[4] Gong H. Agent-to-Agent Finance: Blockchain Payments and Trust Infrastructure for Autonomous AI Agents. arXiv:2607.00245v2, 2026. Primary text.

[5] Schär F, Kosse A, Rice T, Shirakami T, Siridhasanakul J. The Anatomy of Stablecoin Transactions. BIS Working Paper 1359, 2026. Working paper.

[6] Restrepo Amariles D, Charlotin D, He-Guelton L. AI Agents in Payments: Applications, Risks and Regulations. European Journal of Risk Regulation 17, 884–907, 2026. DOI 10.1017/err.2026.10103.

[7] Machine Payments Protocol. Service catalogue API. Snapshot retrieved 8 November 2026. Endpoint.

[8] Coinbase Developer Platform. Discover services. Read 8 November 2026. Documentation.

[9] Olas. Data verification, including Mech off-chain share. Read 8 November 2026. Methodology.

[10] TaskMarket. Kitchen Rave, TSK-EV3WPBDB. Read 8 November 2026. Task record.

[11] TaskMarket. Events for Indexers; network reference. Read 8 November 2026. Events, networks.

[12] Tempo. Connection details. Read 8 November 2026. Official RPC documentation.

[13] Valory. Autonolas Marketplace, commit 64d8dba33e60860644bdda6f57ffc98057fa27e7. BalanceTrackerBase, native implementation.

[14] Virtuals. ACP SDK ABI, commit 392dcc6eb17a6884c214c21aff31afa1b99f45ae; Ethereum. ERC-8183 draft reference, read 8 November 2026. Pinned ABI, draft reference.

[15] Srivastava A, Paul D. Agentic Commerce Bench: Measuring Fraud Detection for Agents That Spend Money. arXiv:2609.35886v1, 27 September 2026. Primary text.

[16] United States. 15 USC 7001(h), electronic agents. Statutory text.