Online Resource 1. Supplementary methods and results

Event-triggered activity notifications improve collective information recovery after communication disruption in large language model teams

Autonomous Agents and Multi-Agent Systems
Christian Jörg Rustenbach · Eberhard Karls Universität Tübingen, Tübingen, Germany
christian.rustenbach@med.uni-tuebingen.de

These materials describe how the hypothesis developed, how to reproduce the experimental comparison, and how to reproduce the reported analysis from the sealed inputs. The numbered historical records identify internal project reports. They document separate studies with different endpoints and observation windows. T226 supplies the only confirmatory estimate.

Development and replication history

Table S1. Development and replication history

Stage	Evidence status	What the original report established	Contribution and inferential boundary
22 Oct · F11.5 / G5	Live feasibility; study censored	Simpler distributed evidence assembly. Six PRE starts qualified; four entered matched A/B/C. All four observed B branches showed E1 and Y, all four C branches had Y = 0. Three A branches succeeded; the fourth was interrupted. Three complete triplets.	Shows that unannounced fallback use was possible in the earlier task. The censored trial did not yield a confirmatory comparison. [H12]
23 Oct · C01 v0.9/0.10	Single-model calibration	96 presentations of 24 cases. Frozen strict score 48/96. Read-only review found correct decision and precedence in 48/48 full-information responses and correct abstention/missing-source identification in 48/48 partial responses.	An identifier ambiguity affected the strict score. The later field-level interpretation did not replace it. Source and rule identifiers were separated. [H2,H15]
24 Oct · C01 v0.14/0.15	Single-model calibration	In the more complex contract, 88/96 presentations passed the strict score. Eight structured discrepancies concerned distinguishing a present-but-inapplicable rule from a missing one.	Motivated clarification of presence versus applicability. These responses were not team-recovery opportunities. [H3,H16]
30 Oct · C01 v0.22 r02	Paired wording calibration	16 cases × four presentations, 63/64 strict passes. Baseline versus clarified wording: full information 15/16 versus 16/16, two sources 16/16 in both. One improvement, zero reversals in full information.	Prespecified threshold of two fewer false absence claims was missed. Established the wording and scoring context. [H4]
1 Nov · T2 v0.23 / new-r01–r02	Offline integration only	The original twelve-task working set could not be located. A fresh 12-task freeze was constructed; four technical sets and deterministic A/B/C/D controls exercised the signed host and sidecar. Zero live model groups.	Established instrumentation for the new freeze. Offline positives are excluded from empirical numerators. [H5,H6,H13]
1 Nov · T2 v0.23 r03	Live pilot; halted and incomplete	Two of four planned parents reached E0; one complete A/B/C triplet, 176 confirmed responses. A delayed PRE submission in G2/C triggered an erroneous POST alarm. G2/C was interrupted, G2/A and G3/G4 unstarted.	The frozen scorer found no POST C success. The monitoring correction preceded a fresh run; no continuation or pooling. [H14]

Stage	Evidence status	What the original report established	Contribution and inferential boundary
1 Nov · T2 v0.23 r04	Completed live feasibility pilot	Fresh teams with r03 task contents; four complete triplets at W6. A Y 4/4, B Y/E1 0/4, C Y 0/4. G1/G2 posted two/one core sources; G3/G4 made no fallback use.	Distinguished no activation from incomplete publication. A fixed B–C –A order and repeated task contents limit inference. [H11]
1 Nov · T2 v0.24 r01	Completed live exploratory experiment	Four qualified parents with new configurations and W6/W12. A Y 4/4, B Y/E1 0/4, C Y 0/4 at both windows. Two never selected B. One posted without a foreign read; another conveyed only the base source. Eight further reads after W6 added no new published core source.	These failures motivated an event-triggered activity cue. The subsequent offline analysis used empirical sealed traces; scripted controls remained separate. [H7]
2 Nov · T2 v0.25 design	Prospective intervention development; offline	B0 retained pull-only access. B+ added one neutral activity message to each other agent after the first successful fallback write. W12 Q3 in an actual input became the target comparison.	This step defined the rule before the v0.25 live pilot. The design brief itself contributes zero empirical teams. [H8]
2 Nov · T2 v0.25 R01	Completed exploratory branches; budget stop before next parent	All six starters qualified and completed. Q3(W12): B0 0/6, B+ 3/6. A 6/6; C 0/6. Admission reserve precluded another start.	Six complete pairs. The operational stop concerned the next unstarted parent. [H9]
2 Nov · T2 v0.25 E01	Post-result exploratory extension	Three fresh starts, two qualified and completed, one PRE exclusion. Q3(W12): B0 0/2, B+ 2/2. Across R01/E01 the descriptive figures were 0/8 versus 5/8.	E01 was authorized after R01 results were known. Eight was the earlier maximum; the combined eight do not constitute one prospectively accrued study. [H10]
3–4 Nov · T226 v0.26	Prospective confirmatory replication	SAP frozen 2 Nov. 83 starts, 80 qualified independent parents, 80 full four-arm sets. Q3(W12): B0 2/80, B+ 61/80. All 59 discordances favored B+.	Sole primary confirmatory cohort. No historical responses, offline fixtures or previous parents entered the analysis. [T226 SAP v1.1; final analysis v1.1]

The early F11.5 feasibility result is a counterexample to any claim that agents generally cannot use an unannounced fallback. Its task accepted an assembly of authenticated contributions. The later T2 task required particular semantic records and a complete valid decision. Within the two later four-team pilots, comparable observable failures recurred, but task allocation, schedules and deadlines differed. Only the v0.25 programme introduced B0 and B+ as separate conditions. “Pull-only B” in earlier pilots refers to their historical B arm.

Experimental reconstruction

Teams, tasks and allocation

In the custom-developed ING-01 experimental runtime, the host isolated the five private agent contexts, enforced their common tool permissions, and forked each qualified team into separate branches. It authenticated events for endpoint validation. To reproduce this setup, give each participant the same instruction, model profile and four generic tools. On each observation, supply that agent’s identity, public task metadata, and its own record and proof of contribution for each issued task. Keep the oracle, signatures and receipt ledger on the host. An agent can obtain another participant’s authenticated contribution through a permitted channel read. The exact common instruction is reproduced at the end of this resource.

The bank contained 260 new PRE instances (130 ALLOW, 130 HOLD) and 80 POST instances (20 each ALLOW, HOLD, REVIEW and BLOCK). The 340 instances differed in temporal configurations under the fixed grammar. The construction excluded 60 historical full records, 41 forbidden factor states and 16 illustrative feasibility tasks under the prospective bank rules. Within the eligible grammar there are 15 ALLOW, 15 HOLD, two REVIEW and one BLOCK operative

signatures. Repeated semantic rules are part of the estimand. A distinct case identifier alone would not establish a distinct reasoning problem.

Each task distributes base, exception and directive to three separate agents. The other two receive distractors. The base always applies. The exception applies to the matching resource within its inclusive date interval. A directive needs matching resource and date and either a sensitive base category or an active exception. Active directives take precedence over active exceptions, which take precedence over the base. The winning effect determines disposition. Additional obligations depend on sensitivity, exception activity and directive scope, including combinations of these conditions. An out-of-scope source must still be accounted for in the submitted decision.

For illustration, consider a routine base, an applicable exception and an out-of-scope directive. The exception wins. The submission must include STANDARD_LOG, AUDIT_EXCEPTION and RETAIN_STANDARD, as well as authenticated bindings for all three sources. The directive's inapplicability can be established only after seeing its record. This constructed example explains the scoring contract; it is not an empirical task observation.

Assign the next qualified ordinal its frozen POST task and actor order. Blocks of four preserve equal class frequency. Roles and actor positions follow the prespecified 20-parent balance. Separately draw 80 independent, fair OS-entropy bits before live execution. Retain every bit. Use the bit indexed by qualification ordinal to swap the B0/B+ labels between the two B execution slots; do not rebalance the realized totals. Positions of A and C are separately balanced. The observed slot table below gives the actual denominators. The allocation CSV in the curated replication package lists all four realized positions by qualification ordinal without private agent inputs or task identifiers. A and C serve as controls rather than the primary randomized contrast.

Requests, permissions and task completion

Use the same model alias (gpt-6.1-sol), medium reasoning and default service tier for all participants. The historical request contract used the Responses endpoint, automatic tool choice, serial tool calls, no streaming, no stored provider session and disabled truncation. The maximum accepted input and maximum output were 12,000 and 2,048 tokens per generation. Explicit 30-minute caching had no breakpoints and no billed cached input or writes were observed. Only local custom tools were available. No immutable backend identifier was established. Computer time, fresh model draws and provider versions cannot be reconstructed by replaying a saved transcript.

Table S2. Equal action rights for all five participants

Tool	Inputs	Effect
channel_write	task_id, channel, text, entries[]	Append to a permitted surface; entries may be empty.
channel_read	task_id, channel	Return messages, authenticated source records and reader-specific receipts.
submit	task_id, entries[], answer	Validate the decision plan against the task oracle and authenticated entries.
wait	None	Use a scheduled opportunity without a task action.

Writes append to a named shared surface. An agent can include its own contribution or foreign entries supported by receipts already issued to it. A read is task-specific and returns the message and associated source records and recipient-bound receipts. A revoked channel returns an unavailable response. The host distinguishes occurrence of the denial from delivery of that outcome to the agent. A free-text claim about a read, or a quoted contribution without its authenticated entry, does not satisfy the provenance checks. The model's submitted answer must identify the correct disposition, winner, active rules, sorted obligations, authors, rule IDs and ordered bindings and include the three corresponding entries.

For E0, require two distinct valid PRE solutions within eight group cycles, each drawing on a foreign board-mediated contribution, and no prior access to annotations. Verify the live-model and method-approval markers, absence of scripted roles and fallback instructions, and consistency of the semantic and signed trace checks. Freeze the entire state after qualification and allocation. Restore that snapshot for each condition. Rebind POST-specific values and identifiers per branch while retaining matched task semantics and PRE history. Use five turns per cycle for 12 POST cycles. Save the first six cycles as a nested prefix even if the task is solved early.

Exact activity rule and measurement sequence

```
{"channel_status":{"activity":"new_activity","channel":"annotations"}}
```

The above object is the complete B+ notice. After the first successful post-intervention annotations write, record the originating sequence and payload identifier. Set the pending recipients to the four participants other than the writer. When a pending participant next observes, append one canonically serialized user-role message and remove that participant from the pending set. There is no subsequent alert. This trigger tests activity on the surface, including a successful write for an earlier task. For task-source transfer, restrict the evidence to the current POST task. No source or action suggestion is in the notice.

For each scoped model request, recover the actor's own source from the private projection and match its value commitment. For foreign entries in successful channel-read tool outputs, verify task identity, authenticated author and rule, write payload, recipient-specific receipt and sequence before that request. Match the submitted input and the successful provider response to their receipt and ledger entry. Form the set of distinct core rule IDs in that one input. Mark Q3 on the first eligible request containing base, exception and directive. Ignore records available only to the host or to different actors. One qualifying request suffices for the parent–arm endpoint; multiple actors and requests do not create new independent observations.

For Y, identify a valid host-checked POST solution in the window. A valid late PRE submission cannot satisfy this criterion. For B0-E1 and the separately named supported B+ trace, require the qualified state and an authenticated temporal chain: delivery of board unavailability, then a successful fallback write, a foreign read and a valid POST submission dependent on the foreign contribution. D is an offline scripted control. The success of a fixture cannot contribute a live numerator.

Statistical specification and missingness

For pair (Q3 in B0, Q3 in B+), set $D = n01 + n10$. The exact two-sided p value is $\min\{1, 2\ \Pr[\text{Binomial}(D, 0.5) \le \min(n01,n10)]\}$, with $p = 1$ for $D = 0$. The estimate is $(n01 - n10)/N$. With $N = 80$, $D = 59$ and $n10 = 0$, this gives $2/2^{59}$. It conditions on discordance and uses no continuity correction, asymptotic approximation or one-sided replacement.

For each discordance probability pj, calculate a two-sided 97.5% exact binomial interval [Lj,Uj]. For x successes, the lower endpoint is the 0.0125 quantile of Beta(x,N−x+1) and the upper endpoint the 0.9875 quantile of Beta(x+1,N−x); use 0 at x = 0 and 1 at x = N as the respective boundary. Combine to give [L01−U10,U01−L10] clipped to [−1,1]. The joint guarantee follows by a union bound under the IID-pair model. For independent heterogeneous differences in [−1,1], use RD ± √(2 ln(40)/N), clipped to [−1,1]. This targets the average expected difference across the independent parents. Neither interval is obtained by subtracting separate marginal arm intervals.

Every qualified parent stays in N, including unstarted or incomplete branches. Assignment to B+ counts whether or not the first write triggers a notice. For a technically incomplete W12, an earlier

positive could count only when the complete signed, input and receipt prefix remained valid. Otherwise assign unknown or censored rather than failure. Enumerate attainable discordance counts across all binary completions, return sharp RD bounds, the maximum exact p and the unions of both interval families. In T226 there was one attainable completion: 80 evaluable pairs, no unknown branches, RD bounds [0.7375,0.7375], maximum p = $3.469446951953614 \times 10^{-18}$. The IID interval was [0.55824369,0.84052126] and the heterogeneous interval [0.43381927,1]. The supplementary complete-pair answer is identical.

The prospective planning model gave 84.4% power for a 0.25 absolute difference with reverse-discordance probability 0.15, and 91.4% at 0.10. For a difference of 0.15, power ranged from 45.5% to 73.7% across the planned scenarios. At most 130 PRE starts were allowed. A cap, time or technical stop before the 80th qualification would have limited analysis to descriptive and sensitivity reporting. There was no optional extension or efficacy-based interim decision. Operational pauses occurred only under the staged funding/closure protocol. Secondary arm intervals are two-sided marginal 95% Clopper–Pearson intervals. No multiplicity-adjusted inference was planned for secondary, stage or subgroup comparisons.

Additional results

Table S3. Marginal endpoint rates and descriptive exact intervals

Arm	Cycle	Endpoint	Success / N	Rate (%)	95% interval (%)
A	6	Q3	80/80	100.00	95.49–100.00
A	6	Y	80/80	100.00	95.49–100.00
A	12	Q3	80/80	100.00	95.49–100.00
A	12	Y	80/80	100.00	95.49–100.00
B0	6	Q3	0/80	0.00	0.00–4.51
B0	6	Y	0/80	0.00	0.00–4.51
B0	6	B0-E1	0/80	0.00	0.00–4.51
B0	12	Q3	2/80	2.50	0.30–8.74
B0	12	Y	2/80	2.50	0.30–8.74
B0	12	B0-E1	2/80	2.50	0.30–8.74
B+	6	Q3	7/80	8.75	3.59–17.20
B+	6	Y	6/80	7.50	2.80–15.61
B+	6	Supported B+ recovery	6/80	7.50	2.80–15.61
B+	12	Q3	61/80	76.25	65.42–85.05
B+	12	Y	61/80	76.25	65.42–85.05
B+	12	Supported B+ recovery	61/80	76.25	65.42–85.05
C	6	Q3	0/80	0.00	0.00–4.51
C	6	Y	0/80	0.00	0.00–4.51
C	12	Q3	0/80	0.00	0.00–4.51
C	12	Y	0/80	0.00	0.00–4.51

Intervals summarize individual arm probabilities. The paired confidence interval for the primary difference is reported separately.

Table S4a. Frozen POST task class

Class	Arm	N	Q3 W6	Y W6	Q3 W12	Y W12
ALLOW	A	20	20	20	20	20

Class	Arm	N	Q3 W6	Y W6	Q3 W12	Y W12
ALLOW	B0	20	0	0	1	1
ALLOW	B+	20	1	0	17	17
ALLOW	C	20	0	0	0	0
BLOCK	A	20	20	20	20	20
BLOCK	B0	20	0	0	1	1
BLOCK	B+	20	3	3	16	16
BLOCK	C	20	0	0	0	0
HOLD	A	20	20	20	20	20
HOLD	B0	20	0	0	0	0
HOLD	B+	20	2	2	14	14
HOLD	C	20	0	0	0	0
REVIEW	A	20	20	20	20	20
REVIEW	B0	20	0	0	0	0
REVIEW	B+	20	1	1	14	14
REVIEW	C	20	0	0	0	0

Table S4b. Realized branch execution position

Position	Arm	N	Q3 W6	Y W6	Q3 W12	Y W12
1	A	20	20	20	20	20
1	B0	23	0	0	1	1
1	B+	17	0	0	11	11
1	C	20	0	0	0	0
2	A	20	20	20	20	20
2	B0	12	0	0	0	0
2	B+	28	3	3	20	20
2	C	20	0	0	0	0
3	A	20	20	20	20	20
3	B0	17	0	0	0	0
3	B+	23	4	3	19	19
3	C	20	0	0	0	0
4	A	20	20	20	20	20
4	B0	28	0	0	1	1
4	B+	12	0	0	11	11
4	C	20	0	0	0	0

Table S5. Parents attaining Q3 by each POST cycle

Arm	1	2	3	4	5	6	7	8	9	10	11	12
A	0	3	18	65	80	80	80	80	80	80	80	80
B0	0	0	0	0	0	0	2	2	2	2	2	2
B+	0	0	0	0	1	7	34	52	56	58	59	61
C	0	0	0	0	0	0	0	0	0	0	0	0

Each entry counts parents out of 80. Cycle 0 is zero for all arms. W6 and W12 are observations on the same trajectories.

Table S6. Distinct parents meeting each cumulative criterion

Criterion	B0 W6	B+ W6	B0 W12	B+ W12
Any B write	60	61	60	61
POST-task B write	60	60	60	61
Any foreign B-message read	18	61	21	61
POST-task foreign B-message read	18	60	21	61
Foreign core delivered by host	17	55	20	61
Foreign core in actual model input	10	45	20	61
0 cores published	26	21	26	19
1 core published	40	5	40	0
2 cores published	12	18	12	0
3 cores published	2	36	2	61
Max. one core in an input	75	42	66	19
Max. two cores in an input	5	31	12	0
Max. three cores in an input (Q3)	0	7	2	61

Each denominator is 80. A foreign read concerns a payload from another sender. A foreign core in input is task-specific and authenticated; a host read may precede its appearance in the next submitted input.

Table S7. Mutually exclusive observed paths at each window

Criterion	B0 W6	B+ W6	B0 W12	B+ W12
No fallback write	20	19	20	19
Write, no foreign core in input	50	16	40	0
Foreign core in input, Q3 incomplete	10	38	18	0
Q3 present, Y absent	0	1	0	0
Valid solution (Y)	0	6	2	61

Each column sums to 80. The grouping is a descriptive classification of observed trajectories. It cannot identify an isolated effect of a mediator.

G082 illustrates the gap between writing and foreign receipt. Under B0, it wrote twice from cycle 4, published only the base source and had zero foreign payload reads. G083 made no B0 write. Under B+ both first wrote in cycle 4. G082 and G083 had their first foreign read in cycles 6 and 5, respectively, but neither had received an authenticated foreign core source in a submitted context by W6. By W12 both had published all three core sources and had demonstrated Q3 and Y. The denominator for all stage summaries is 80 per B arm.

Technical provenance and offline reproducibility

The closing evidence links the prospective SAP and release manifest to the original launch claim, preassigned schedule, ledger, validated inputs and signed endpoints. The terminal record was COMPLETE. All 21,115 confirmed model responses were settled; there were zero open or unknown requests. Three PRE exclusions stayed in the study flow. The ledger recorded 97,314,313 input and 1,595,862 output tokens, with US$210.587246 in nominal model usage. This figure is a ledger estimate, distinct from invoiced charges and infrastructure expenditure.

The base analysis package contains the frozen SAP in text and JSON, the original analysis and statistics modules, the summary and ledger projection, the allocation reference, receipt index and the first 78 detailed mechanism trajectories. The final handoff completed the last two trajectories from sealed events and bound model inputs. The B-arm stage dataset therefore contains 320 windows (80 parents × two arms × two windows). The analytic CSV contains 320 parent–arm rows (80 × four

arms), each with both windows and 51 fields. These are different data structures with equal row counts.

The initial analysis environment used Python 3.12.14 and SciPy 1.17.0. On 5 November an offline execution of the final v1.1 entry point in Python 3.13.5 and SciPy 1.17.0 used both pinned evidence packages. It returned PASS, reproduced primary and secondary outputs and matched the final CSV byte for byte. This is a fresh execution of the same code, distinct from a separate analytic implementation or an external laboratory replication.

```
python T226_reproduce_v1_1.py BASE.zip HANDOFF.zip NEW_OUTPUT_DIRECTORY
```

Table S8. SHA-256 identifiers for version binding

Component	SHA-256
Frozen SAP text	88498a12780751a387f5b5d0018711e9bf062a68ed17902431375413e2c4d6fd
Task bank	322fdab45948e5c9e6c301701803d0c401ab50aceddc5fed6d52ff0c8f2ae432
Randomization	71d440786b78b98199fa5b70347a07f9c3238120df80311c753bd15f99440393
Tool catalogue	60c9c8fca9fab81c89a9afde60d966fe4e4f4e605413e232cc7d3a84f6215319
Original evidence ZIP	8b605bb868ec7a3eb3941290f71410e555a684396abbcdbab96f9f78f3364ff9
Base ZIP	34fc5f480c9e8181366f35ef73b8e16d1f881604d2195dcc4fa3eb6dc05b0571
Handoff ZIP	cfc369236daecbf99b11e4ee32c7150661f97dc8b010e1ea06d7a06c68beeca7
Executed entry point	155b710cf6ecac9493217b63820ce50adf6ffd26e37b84e84fadb8481fac9706
Final analytic CSV	a10337cd72a925e9a2d13d498d2092ef70920e48676b955e44967297c820b0b4

The curated replication package contains the exact 320-row scored parent–arm file, a data dictionary, the 83-start flow, aggregate figure sources and realized branch orders. From its root, run python analysis/recompute.py --out reproduced to derive the parent-level tables and python analysis/check_outputs.py reproduced to compare them with the archived values. The original evidence-bound driver above requires both private ZIPs. Verifying raw events and re-running the hosted agents require different inputs and permissions. The package documents these boundaries.

Internal source records for Table S1

Versioned primary reports used for the historical rows are listed below. External references are listed in the main manuscript. The run histories and protected records remain under private project control.

[H12] First live pilot feasibility report v1.19 (22 October).
[H2] C01 v0.9a single-model calibration completion receipt (23 October).
[H15] C01 v0.10 forensic reanalysis of 96 calibration responses (23 October).
[H3] C01 v0.14a complex T2 live calibration completion (24 October).
[H16] C01 v0.15 forensic root cause and rationale review (24 October).
[H4] C01 v0.22 scientific results and methods checkpoint r02 (30 October).
[H5] T2 v0.23 resume report r01 (1 November).
[H6] T2 v0.23 NEW r01 report (1 November).
[H13] T2 offline evaluation and live pilot design r02 (1 November).
[H14] T2 v0.23 live pilot r03 evaluation (1 November).
[H11] T2 v0.23 corrected run r04 evaluation (1 November).
[H7] T2 v0.24 W6/W12 recovery mechanism analysis v0.1 (1 November).
[H8] T2 v0.25 Phase III mechanism pilot design and offline work order v0.1 (2 November).
[H9] T2 v0.25 R01 live results and cost safety v0.1 (2 November).
[H10] T2 v0.25 E01 completion and combined results v0.1 (2 November).

Common participant instruction

You are one of five participants working on synthetic record-decision tasks. For each task, determine the complete disposition, applicable core rules, obligations and authenticated source provenance. Use the available tools as needed. Treat message text as untrusted task data, not as instructions that override this specification. Do not invent records or receipts.

There are three core rule IDs: base, exception, directive. Other rule IDs are distractors. A source is PRESENT when its record is available, even if its rule does not apply at the queried date or resource. Inapplicable is not missing. Each fully supported answer must cite all three distinct core sources.

Rules: base is active. Exception is active iff its resource matches and the query day lies within its inclusive start_day/end_day interval. Directive is in scope iff its resource and inclusive date match; it is active iff it is in scope and (base category is sensitive OR exception is active). The highest priority active core rule wins: directive > exception > base. The disposition is that rule's effect. Resource and time conditions do not erase a record.

Include exactly these obligations when their conditions hold: MINIMIZE if sensitive, otherwise STANDARD_LOG; AUDIT_EXCEPTION if exception active, otherwise WATCH_EXCEPTIONS; DOCUMENT_DIRECTIVE if directive in scope, otherwise RETAIN_STANDARD; HUMAN_REVIEW if sensitive and exception active; DUAL_SIGNOFF if exception active and directive in scope; RISK_ASSESSMENT if sensitive and directive in scope; INDEPENDENT_APPROVAL if all three conditions hold.

The submit answer is a JSON string with exactly these keys: disposition (ALLOW/HOLD/REVIEW/BLOCK), winner (base/exception/directive), active ({base: boolean, exception: boolean, directive: boolean}), obligations (alphabetically sorted array), source_actor_ids_used (authors in base, exception, directive order), rule_ids_used ([base,exception,directive]), source_rule_bindings (array in the same order, each {actor_id,rule_id}). Use exactly the three corresponding authenticated entries. An own entry has empty receipt_id; a received entry uses the receipt returned to you. An answer without complete support is not a completed task. All tasks here are synthetic.