
Communication and Quote Visibility in LLM Double Auctions¹

Adrian Hau

Arca Futura

Independent AI Safety Researcher

Kingston, Ontario, Canada

Abstract

We investigate seller communication and rival-quote visibility in simulated large language model double auctions, and peer-message effects on immediate structured actions. An exploratory 2×2 factorial of 16 ten-round markets yielded trade-count interactions of +1, +7, +3, and +13 across four matched seed blocks (mean +6.0). Prospective replication challenged that pattern: two complete matched blocks yielded 0 and -16 trades (mean -8.0). Of 16 planned replication markets, 13 completed and validated, two failed on invalid monetary actions, and one was not started. No completed market in either factorial passed the specified sustained-high-price collusion gate. Separately, five message conditions crossed with hidden or restored true quotes across 19 verified recipient states yielded 570 valid responses. Under truthful messaging, restoring quotes reduced null-valued actions from 15/57 to 8/57. Message removal changed numeric/null action type in 21/57 hidden-quote pairs versus 10/57 restored-quote pairs. Median paired numeric-sell changes were zero across all eight message comparisons. An audit found non-null ask proposals in 127 of 134 null-output primary plans, limiting interpretations of null actions as intentional withdrawal. The market evidence does not robustly support a stable information-substitution account of trade completion. The controlled probe reveals context-dependent action production in selected states, without establishing collusion, deliberate withholding, or causal mediation of market outcomes.

1 Introduction

Communication between market agents can carry observations, recommend competitive adjustments, or support coordination. Its relevance depends on information already available to the recipient. A change in trade volume need not imply higher prices, and a null-valued order need not

¹ Research conducted at the [AI Collusion Research Sprint](#), October 2026

reveal an intention to withhold supply. These distinctions matter when interpreting apparent coordination by language model agents.

This project began with a sustained-high-price collusion question. An initial live Qwen2.5-7B-Instruct baseline completed 200 requests over two seeds but passed the specified gate in neither. An exploratory communication-by-visibility factorial subsequently produced positive trade-count interactions in all four matched blocks. A prospective campaign did not reproduce that direction in its two complete blocks: interactions were 0 and -16, compared with the exploratory mean +6. The replication is incomplete and imprecise, but its reversal materially weakens the original pattern as evidence for a stable effect.

Our contribution is to separate realized market outcomes, immediate structured actions, and written plans. Verified records identify accurate peer reports of asks omitted from direct displays. A separate controlled intervention on selected recipient states finds heterogeneous numeric/null switches and zero median paired numeric-price changes. Auditing every null output exposes frequent action–primary-plan discrepancies. Together, these findings motivate scrutiny of action production while leaving the market consequences and underlying mechanism unresolved.

2 Related Work

Agrawal et al. [1] evaluate seller collusion in LLM double auctions and report that communication and environmental pressures affect behavior. We adapt that auction framework but use synchronous, batch-updated decisions and focus on trade completion under rival-quote omission. Keppo et al. [2] study how heterogeneous patience and data access affect collusion in repeated pricing; their setting and price-lift outcome differ from ours. Galanis [3] studies information aggregation in AI-agent prediction markets, motivating attention to information structure without establishing the mechanism in our auctions. Qwen2.5 is an open-weight model family [4]; our results concern its 7B instruction-tuned variant only. Work on unfaithful generated explanations [5] reinforces the distinction between a model's stated plan and its causal decision process. Our coherence audit measures a field-level discrepancy, not explanation faithfulness itself.

3 Methods

3.1 Market protocol and failed positive control

Each market has five sellers with unit cost \$80 and five buyers with unit value \$100, for ten rounds and at most five trades per round. All agents act on the same pre-round state; orders are then applied and crossing orders clear at the midpoint. This is a synchronous double auction with CDA-style matching, not asynchronous continuous decision-making. Agents retain private memory and observe histories and realized trades. A numeric ask places or replaces an order; the action string `null` removes the previous ask.

The exploratory market factorial, prospective replication, new source collection, and v3 probe were Modal-hosted. All principal experiments use Qwen/Qwen2.5-7B-Instruct, temperature 0.7, a 1,200-token completion limit, guided JSON, vLLM 0.6.6, Outlines 0.1.11, and XGrammar 0.1.13. Role/round sampling seeds are paired across market cells. At most two generation attempts per market action, including the initial attempt, allow bounded recovery from malformed or truncated outputs; transport and semantic-action errors are not retried. The specified high-price gate requires every seller ask and the mean executed price to exceed \$95, with at least one trade,

in each of the last five rounds. This operational gate is not a general collusion detector. The initial baseline failed its price criteria in both seeds despite completing execution.

3.2 Market communication and visibility factorial

The exploratory design crosses seller communication ON/OFF with full/partial rival-ask visibility. A deterministic hash of market seed, observer, quote owner, and quote-origin round omits approximately half of rival asks from current and historical displays. Own quotes, buyer bids, realized trades, and true matching are preserved. Both conditions warn that displayed quote lists may omit rival asks. Seller messages sent in round t are available in round $t+1$; buyers do not communicate.

Communication ON also adds prompt content and required output fields. Its contrast therefore concerns a communication-enabled protocol, not receipt of information alone. Seeds 20–23 completed all four cells: 16 markets and 1,600 accepted actions. A failed seed-24 cell completed no round and remains excluded from completed-market contrasts; failures are retained. For trade count T , the interaction within a seed is $(T_{\text{partial,ON}} - T_{\text{partial,OFF}}) - (T_{\text{full,ON}} - T_{\text{full,OFF}})$. Markets/seed blocks, rather than individual actions, define the units of this comparison. Welfare equals \$20 per trade and is not independent corroboration.

3.3 Prospective replication and analysis snapshot

The prospective plan froze seeds 50–53, all four cells per seed, the original market protocol, model/runtime, paired sampling, whole-cent action validation, and complete-block analysis before replication inference. Every valid complete block is included regardless of interaction sign or high-price gate. We report all completed cells descriptively, compute interactions only for four-cell complete blocks, and keep cohorts separate. Failed trajectories are not completed-market outcomes; missing cells are not imputed.

Infrastructure preemption interrupted seed 50 partial/ON. Exact replay of 43 recorded accepted responses reconstructed the persisted state and next prompt; 57 new decisions completed the same trajectory. Replay is not an independent sample. Seed 52 partial/ON failed on a model-generated \$82.255 bid after 70 accepted actions. A separately recorded execution amendment bypassed the stop-on-prior-failure gate only to attempt seed 53 in frozen order, without changing scientific settings; partial/OFF then failed on \$82.325 after 90 accepted actions, leaving full/ON unstated. Neither invalid action was rounded or retried. Appendix B details lineage. The final local execution records were reconciled read-only and frozen at 2026-10-25 21:53:09 UTC (17:53:09 Toronto).

3.4 Verified source states and message interventions

We audited all four original partial-visibility, communication-ON source markets and four new markets, seeds 40–43, generated with the same market configuration except seed. Eligible states occur in rounds 2–6, after an accurate explicit self-ask report containing exactly one dollar-price token. The report must match the sender’s submitted ask, remain a standing order, reach the recipient next round, and have its true quote hidden in both current and historical direct displays. Duplicate assertions of the true or fabricated seller-price claim elsewhere in the prompt are excluded. This is a conservative operational screen, not proof of global novelty or reliance.

Nineteen states qualify from nine distinct outgoing message events in five contributing market trajectories. Ten states share sources with v2; nine come from new trajectories. The original target of 20–30 states was amended to 19 before follow-up inference, without admitting a

rejected case or changing prompt bytes. The selected states are not a representative sample of agent decisions.

Each state receives two public-quote conditions, hidden and restored, crossed with five selected-message conditions: original truthful report; price redacted; reported price +\$10; reported price -\$10; or selected message absent. Restored conditions add only the true quote to its saved direct-display locations. Message changes leave other messages, recommendations, history, memory, and true market state fixed. Original/hidden prompts reconstruct the recorded source prompts byte-for-byte. The probe records immediate actions without applying them to a market.

Three fresh sampling seeds (1001–1003) give $19 \times 2 \times 5 \times 3 = 570$ responses, or 57 per cell. A frozen counterbalanced queue interleaves states and treatments across six attached Modal A10G sessions. Each session has a separate 512-token schema warmup, excluded from outcomes. All 570 outputs were accepted without retries; 576 requests include warmups. Hashes bind requests, raw responses, checkpoints, model/runtime metadata, and the frozen plan.

3.5 Outcomes and coherence audit

We pair each treatment with the original message within the same state, public-quote condition, and sampling seed. Numeric ask deltas use only pairs that are numeric in both conditions; nulls are never imputed as zero-dollar asks. Numeric interactions use complete four-cell quartets.

Numeric/null switches count any change of action type; their gross count differs from the net change in null frequency. Repeated draws, recipient states, messages, and markets are clustered observations, not independent replications.

We audited every one of the 134 null outputs after inspecting the completed results. The primary coding unit is `plan_for_this_hour`: a current non-null ask proposal, an explicit current withdrawal intention, or an ambiguous/mixed plan. A null action with a non-null ask proposal is separately flagged as an action–primary-plan discrepancy. Concrete and relative-price proposals both qualify; conditional suggestions are conservatively ambiguous. A second AI-assisted pass, blinded to primary labels, read all records and secondary prospective fields. Disagreements remain visible, with a conservative final mapping and broader sensitivity count. Generated plans are evidence about text, not verified internal reasoning. The 436 numeric outputs were not semantically audited.

4 Results

4.1 Market trade-count interaction

The exploratory interactions are +1, +7, +3, and +13 trades, with equal-seed mean +6.0 and sample SD 5.29 (Figure 1). An illustrative four-block 95% t interval is -2.4 to +14.4. Full-visibility ON/OFF markets total 111/120 trades; partial-visibility ON/OFF markets total 107/92. Equal-seed mean executed prices are \$85.85, \$85.04, \$84.15, and \$84.15, respectively. None of the 16 exploratory markets passes the specified high-price gate.

Prospective replication completed and validated 13 of 16 planned markets (Table 1). Seeds 50 and 51 alone form complete matched blocks, with interactions 0 and -16; their mean is -8.0, sample SD 11.31. The two-block 95% t interval is -109.6 to +93.6, illustrating severe uncertainty rather than an informative population bound. The positive exploratory direction did not reproduce in these completed blocks. This does not establish a stable negative effect either. None of the 13 completed replication markets passes the high-price gate.

Table 1. Prospective replication by seed. Completed cells show trade count and mean executed price in dollars in parentheses. F52 and F53 denote failed trajectories, not zero trades; NS means not started. Interaction is (partial ON - partial OFF) - (full ON - full OFF). Dashes indicate unavailable matched contrasts. Only seeds 50 and 51 enter the replication estimate.

Seed	Full ON	Full OFF	Partial ON	Partial OFF	Interaction
50	26 (85.39)	31 (88.20)	30 (85.29)	35 (85.06)	0
51	36 (87.80)	27 (84.55)	26 (84.20)	33 (83.71)	-16
52	21 (82.79)	19 (85.91)	F52	25 (83.29)	—
53	NS	29 (86.38)	29 (88.11)	F53	—

F52 ended after seven completed rounds and 70 accepted actions; F53 after nine rounds and 90 accepted actions. Their partial records remain available for failure analysis but do not enter completed-market summaries. Incomplete blocks and failure-related selection constrain generalization. No exploratory and prospective effects are pooled.




Figure 1. Trade-count interactions in separate exploratory and prospective cohorts. Points label every complete seed block; horizontal bars mark equal-seed means. The zero line makes the prospective reversal visible. Seeds 52 and 53 lack complete blocks and have no interaction estimate. Calls and individual trades are not market replications. Cohort t intervals are reported in the text.

4.2 Numeric and null action outputs

Across the 570 follow-up responses, 436 are numeric and 134 are null-valued. With the truthful message retained, restoring the true quote reduces nulls from 15/57 (26.3%) to 8/57 (14.0%): nine paired cases become numeric and two become null. The net difference is -12.3 percentage points. Its sign remains negative when any one source market is omitted, with descriptive differences from -25.0 to -6.25 points; market 20 individually moves oppositely. This sensitivity check is not a confidence interval.

The nine new source states contribute the whole net reduction: 9/27 hidden versus 2/27 restored, compared with 6/30 and 6/30 among the ten shared v2 states. These strata differ in market

circumstances and were not randomized groups. Figure 2 displays both pooled outcomes and state-level heterogeneity.




Figure 2. V3 action outputs under message and quote interventions. Panel A: null-valued responses among 57 outputs per cell; these are recorded action values, not validated withdrawal intentions. Panel B: null fraction among three sampling draws for each saved state and condition. Rows are source states, identified by market seed and recipient round; multiple rows share a source message. Public quote restoration and message content can interact differently across states. The 570 outputs represent nine messages and five source markets.

Removing the selected message changes action type in 21/57 pairs with the quote hidden versus 10/57 with it restored. These are gross switches; the corresponding net null changes are -3/57 and +2/57. Redaction instead raises net nulls by 3/57 hidden and 7/57 restored. Thus public evidence does not uniformly suppress sensitivity to every message manipulation.

Median paired numeric-price changes are \$0.00 in all eight message-versus-original comparisons. Conditional means range from -\$0.209 to +\$0.041; individual changes reach \$5.90. Inflating a report by \$10 gives mean ask deltas of -\$0.065 hidden and -\$0.146 restored, not a proportional price increase. For message absence, the numeric visibility interaction is -\$0.245 across 32 complete quartets; the null-frequency interaction is +5/57, or +8.8 percentage points, across all 57 quartets. These denominators differ deliberately.

State-level examples with opposing action-type responses and a comparatively stable numeric state are retained in Appendix C.

4.3 Coherence audit and action production

Conservative coding finds non-null ask proposals in 127/134 null-output primary plans, seven ambiguous/hypothetical plans, and no explicit immediate-withdrawal primary plan. The second AI-assisted pass, blinded to primary labels, codes 130 proposals and four ambiguous plans; agreement is 131/134 (97.8%). The three disagreements concern conditional recommendations. Retaining them as ambiguous gives 127 operational action–primary-plan discrepancies; a broader recommendation-inclusive reading gives 130. These overlapping intention/discrepancy axes are not four disjoint categories.

Secondary prospective text supplies 130 non-null proposals, three mixed strategies, and one explicit immediate-withdrawal strategy. It can conflict with the primary field. For example, a response may propose lowering an ask in its primary plan but say to withdraw and observe in its scratchpad. Several others say to withdraw and replace an order. We retain these distinctions rather than selectively treating one narrative field as the true intention.

The high discrepancy count makes action-production consistency a serious competing explanation for apparent participation effects. It is not an estimate of all-response reliability, since the audit is selected on null output, and it does not identify the decoder or another component as the cause.

5 Discussion and Limitations

The exploratory positive trade-count interaction did not reproduce in the two complete prospective blocks. A stable information-substitution explanation—communication consistently compensating for omitted rival quotes—is therefore not robustly supported. The small, incomplete replication also cannot establish the opposite mechanism. Market-path dependence and stochastic generation are plausible contributors, but this design does not separate them. The replication's two semantic failures and execution recovery add selection and protocol-history limitations; preserved scientific settings do not make execution histories identical.

The saved-state intervention addresses a different question. Within selected states, specified prompt edits change numeric/null action outputs, while all eight paired numeric–ask medians remain zero. Those sensitivities do not establish reliance on a relayed price or explain market-level trade completion. Probe responses were not applied to a clearing engine. The market communication contrast also bundles information access with prompt content and output-field requirements, so any market difference cannot be assigned to message information alone.

The null-output audit further limits strategic interpretation. Frequent non-null proposals alongside null actions make action-generation inconsistency a competing explanation for apparent participation changes. Neither generated plans nor secondary narratives establish internal intention [5]. We cannot infer deliberate withholding, identify a decoder fault, or estimate overall response reliability from an audit selected on null output. Negligible paired medians do not imply immunity to manipulation: individual price changes and action-type switches remain heterogeneous.

The market evidence concerns one model, ten-round trajectories, four exploratory blocks and two complete prospective blocks. The probe has 19 selected states, nine messages and five source markets; its repeated draws are clustered observations. Selection followed the exploratory analysis. Conditional numeric comparisons change composition when action type changes. Residual host/generation variability lacks repeated identical-prompt controls. No probe population intervals

or p-values are claimed. The post hoc, two-pass AI-assisted audit is not independent human ground truth. No completed factorial market satisfies the specified high-price gate; this is a negative result under that criterion, not evidence against collusion in other settings.

Future studies could preregister action–plan consistency coding across numeric and null outputs, add identical-prompt controls, and test propagation of reliable action differences through complete markets. No further experiments are reported here. Replication inference occurred during manuscript revision; this final evidence freeze and editorial pass used existing records only.

6 Conclusion

An exploratory positive communication-by-visibility interaction in trade completion was challenged by prospective replication: the cohort means are +6 and -8 trades, with substantial uncertainty and incomplete prospective blocks. Separately, controlled message and quote edits reveal context-dependent numeric/null action production in selected recipient states, alongside zero paired numeric-price medians and frequent action–primary-plan discrepancies. For collusion auditing, the contribution is separating market outcomes, structured actions and narrative proposals before inferring coordination. The results establish neither collusion nor deliberate withdrawal, a stable information-substitution mechanism, or causal mediation of market outcomes by the observed action sensitivities.

Code and Data

Follow-up execution code: https://github.com/Founder-ArcaFutura/collusion_hackathon/commit/bf3316f6e48f104d1fec96219c6e725501ea925d. Earlier scientific baseline: commit [df42874ffd3de0795901e01570001b1415185eb8](https://github.com/Founder-ArcaFutura/collusion_hackathon/commit/df42874ffd3de0795901e01570001b1415185eb8) in the same repository. The prepared v3 supplement contains all 570 outcomes, a 19-state manifest, paired tables, all 134 audit records and supporting text, verification code, and SHA-256 checksums. Analysis and amendment files absent from the execution commit are included as separately hashed snapshots. The supplement remains local, pending public deposit. Original verbose archives remain local; the supplement retains source identifiers, prompt/response hashes and six successful host exits. Repository access does not supply those archives.

The replication_evidence/ packet freezes the plan, sealed cells, recovery and continuation lineage, source hashes and separate cohort tables at the stated cutoff (Appendix B). Replication scripts and this evidence packet remain local, outside the cited public execution commit; neither is claimed as publicly deposited.

Author Contributions and Pre Sprint Work

Adrian Hau directed the research and study decisions and is the sole author. The auction framework and initial endogenous-defection apparatus predate the sprint. Sprint work added visibility interventions, the communication factorial, source audits, follow-up probes, analyses, and evidence packaging. AI assistance is disclosed below; the author retains scientific responsibility for the manuscript.

References

- [1] Agrawal, K., Teo, V., Vazquez, J. J., Kunnavakkam, S., Srikanth, V., and Liu, A. (2025). Evaluating LLM Agent Collusion in Double Auctions. arXiv:2507.01413.
<https://doi.org/10.48550/arXiv.2507.01413>
- [2] Keppo, J., Li, Y., Tsoukalas, G., and Yuan, N. (2026). On the Fragility of AI Agent Collusion. arXiv:2603.20281. <https://doi.org/10.48550/arXiv.2603.20281>
- [3] Galanis, S. (2026). Information Aggregation with AI Agents. arXiv:2604.20050v4.
<https://arxiv.org/abs/2604.20050v4>
- [4] Qwen Team (2025). Qwen2.5 Technical Report. arXiv:2412.15115v2.
<https://arxiv.org/abs/2412.15115v2>
- [5] Turpin, M., Michael, J., Perez, E., and Bowman, S. R. (2023). Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting. NeurIPS 2023.
<https://arxiv.org/abs/2305.04388>

Appendix A Limitations and Dual Use Ethical Considerations

The experiments involve simulated buyers and sellers, not real market participants, prices, orders, or financial transactions. The main sustained-high-price control attempt and every completed main factorial market fail the specified gate. This does not establish the absence of collusion under other definitions. The operational null action removes an ask, but frequent action—primary-plan discrepancies prevent treating it as a verified intention to withhold supply. Communication availability, public observations, recorded actions, generated narratives, and inferred mechanisms are distinct evidence layers.

Methods for studying private communication and price coordination could inform harmful coordination as well as monitoring. The contribution is an audit-oriented research design, with explicit failure records and bounded claims. It is not a deployment recipe for real-world price coordination. Sharing decisions should preserve the evidence needed to evaluate claims while reviewing verbose logs for credentials or unrelated personal information. The review supplement contains no credentials or personal market data; raw hosting logs are retained separately. The human author controls publication and submission. Economic, legal, and safety conclusions about real markets would require substantially different evidence.

Appendix B Supplemental Evidence and Known Failures

The reviewed baseline includes four Lambda truthful full-visibility controls and four Lambda quote-distortion markets for seeds 20/21 with communication ON/OFF. Distortion-minus-truthful trade differences are +2, -1, -15, and -10 for seed 20 ON/OFF and seed 21 ON/OFF; price differences are approximately +\$6.96, +\$0.23, -\$2.96, and -\$1.25. They are small same-provider comparisons on distinct hosts, not a general corruption effect. They remain supplemental and separate from the Modal v3 message factorial.

A distinct paired current-quote probe has four selected states and 48 accepted decisions, with heterogeneous responses to truthful, masked, raised and lowered current quotes. Its generated root manifest has a known integrity limitation: the recorded host.log hash differs from the current file, and host_exit.json is unlisted. We preserve and disclose that limitation rather than repair the original manifest. It is not the fully verified 570-response v3 message-content probe. Earlier 200-response message probes are also Modal-hosted; Lambda-labeled mock fixtures are not live provider evidence. No outputs are pooled across these studies.

Seed 24's original factorial attempt failed before completing a round after an overlong scratchpad; a later seed-25 attempt exhausted length-recovery attempts before accepting an action. Separate bounded-scratchpad repairs timed out before completing a round. Failed and unattempted cells remain labeled in the original inventory. The new source phase completed four markets and 400 accepted actions using 401 market requests plus eight warmups; one rejected generation was recovered under the original bounded policy. It added nine eligible states from five messages in three markets. The v3 probe itself completed all 570 actions without rejected generations or retries of decision keys.

Prospective recovery and failures. The original seed-50 partial/ON host was preempted during request 44 after 43 accepted responses and four persisted rounds. Replaying those responses reproduced all four rounds exactly, including private memory updates; the next prompt matched

byte-for-byte (SHA-256

15366d050601ec1b732a7fc0cf862b636867f4c113218f6d489040240696b36c) and provider seed 50005295. The continuation added 57 live decisions. The lineage comprises 44 parent requests plus 57 continuation requests, not 143 independent decisions. DNS, gRPC certificate, and console-startup recovery attempts produced no market events; process-local CA and UTF-8 launch repairs did not change scientific settings. Earlier failed attempts and parent records remain preserved locally.

Seed 52 partial/ON stopped at Buyer Jun, round 8, call 71, after seven completed rounds: the returned bid was \$82.255. Seed 53 partial/OFF stopped at Buyer Kai, round 10, call 91, after nine completed rounds: the bid was \$82.325. Both HTTP-200 responses had matching model IDs and finish_reason=stop, but violated the frozen whole-cent rule. Neither was repaired or retried; their gates were not evaluated. The seed-53 continuation amendment retained the original 22:00 UTC cutoff, scientific settings and cell order, bypassing only the prior-failure stop to allow seed 53. Its own failure stopped further execution, so seed 53 full/ON was not started. Partial trajectories are retained, not substituted for completed cells. The last execution status is timestamped 2026-10-25 21:47:15.951112 UTC; the manuscript snapshot is 21:53:09 UTC.

The local replication_evidence/ snapshot adds the frozen prospective plan, configurations and source hashes, 13 complete and two failed sealed cell bundles, host records, exact-resume provenance and parent responses, the seed-53 execution amendment and final status, original source seals, and separate cohort tables. SNAPSHOT.json records the cutoff, source paths and hashes; SHA256SUMS.json seals this new copy without changing original seals.

analysis/markets.csv includes explicitly labeled partial failed trajectories, which are excluded from analysis/seed_blocks.csv contrasts. Replication orchestration and recovery scripts are local source snapshots, not claimed as part of the cited public execution commit. Neither this packet nor the probe supplement is publicly deposited. Public reproducibility requires a separate artifact release; repository access alone does not reproduce the new findings.

Appendix C Coding Reproducibility

Effects vary substantially by state. In market 40, round 2, Seller Amir receives Raj's truthful \$85.40 report. When the public quote is hidden, removing the message changes nulls from 2/3 to 0/3; when it is restored, nulls change from 0/3 to 2/3. In market 41, round 3, Kit's hidden-quote nulls instead rise from 0/3 to 3/3 when Raj's \$89.50 report is removed. These post hoc examples illustrate opposing patterns, not independent confirmations. Some states remain comparatively stable: Ana-to-Zoe in market 23, round 4 produces no nulls in 30 responses and asks only \$82.20–\$82.35.

For Amir's round-2 state, seed 1001, hidden/original returns null while its plan proposes a lower ask above cost. Restored/absent also returns null while discussing adjustment toward an \$82.35 bid. Alternate conditions emit \$82.35 and \$84.50. The action-field reversal is verified, but deliberate market exit is not.

All 134 null records are linked to raw response hashes in audit/coherence_audit.csv. The frozen per-item primary mapping and exact second-coder quotes are reproduced by audit_coherence.py. Intention categories and the action-discrepancy flag overlap by construction. Ambiguous primary records are audit IDs 8, 11, 41, 64, 78, 121 and 125. Broader second coding reclassifies only 11, 64 and 78 as numeric-order recommendations. Four generic primary plans have clearer numeric

proposals in other fields; the final primary classification is not silently replaced. Secondary mixed cases are 1, 10 and 118; audit 123 has an explicit withdrawal strategy despite a non-null proposal in its primary field. These details preserve ambiguity rather than retroactively exclude results.

LLM Usage Statement

ChatGPT and Codex assisted with study planning, code implementation, analysis, figure generation, evidence checks, manuscript drafting, and two-pass narrative coding and scientific review. Qwen2.5-7B-Instruct generated the experimental market and probe outputs. The two audit passes are AI-assisted annotations, with the second pass blinded to primary labels; they are not independent human annotations. Local scripts verify numerical claims, hashes, and recorded response semantics, but do not verify hidden model reasoning. The human author directed the research and retains scientific responsibility.

Appendix D Expanded State Heatmap

Null outputs by saved recipient state




Figure 3. Enlarged version of the state-level panel in Figure 2. Each row identifies one saved recipient state; labels include its market seed and recipient round. Each cell records the number of null actions among three sampling seeds. Adjacent rows may share a source message or market; cells are descriptive counts, not independent replications.