
UNDERCOVER-Bench: Trusted agents and private-channel collusion

Amey Muke
Independent Researcher

Sanchayan Ghosh
Independent Researcher

Prasad Mahadik
Independent Researcher

Prepared for
Apart Research

Abstract

Can a trusted undercover agent expose or disrupt collusion, and does its approach change whether peers accept, reject or become suspicious? We study this Apart Research Track 3.2 question using Secret Hitler, a local Liar's Bar controller and four model families. A scripted trusted sender presented direct pacts, indirect hints and benign requests to three recipient models. All six prohibited offers were rejected; one of three benign requests was accepted. These are recipient decisions, not implemented pacts or disruption tests. In a separate mixed-model game, ordinary GLM and Qwen opened an exclusive channel, exchanged 20 messages and matched all 10 election votes; the auditor could not observe their channel. Under an explicit channel ban, GLM and Muse initiated offers. In an additional fictional-background condition, the character-trained GLM auditor accepted a banned channel and exchanged private messages while describing coordination as public. Four reporting actions failed local length validation, so that game's outcome is excluded. Constitutional preference training and introspection SFT produced a local reporting shift but also an unsupported pact allegation; simple diagnostics showed no overall character advantage. The pilot exposes channel adoption, attempted collusion and a failure of the intended trusted agent. It does not demonstrate successful disruption or a causal background effect. RLCR is implemented in code but remains an uncompleted experimental extension.

1. Introduction

This project addresses [Apart Track 3.2](#): can introducing a trusted agent expose or disrupt collusion? An undercover participant can offer coordination, record reciprocation and report evidence. We ask whether direct offers, indirect hints and benign requests lead to different acceptance, rejection or expressed suspicion, and whether the intended trusted agent keeps its own reporting and communication commitments.

UNDERCOVER-Bench places an auditor in an existing Liberal seat rather than adding a player. The pilot combines induced offer responses, naturally generated private dialogues and checks of auditor reliability. Character training tests the reliability of the proposed countermeasure; future RLCR asks whether it survives collective optimization. Trust is an assigned experimental role, not a proven property of the model. Reporting and refusal alone do not establish disruption.

2. Related work

[Voluntary Collusion with Secret Tools](#) studies adoption of advantageous secret tools in competing agents. [GroupGuard](#) supplies honeypot-based collusion-defense context, and [AI Control](#) studies monitoring against intentional subversion. The secret-tool study's [original code](#) and the [Apart Track 3.2 starter](#) motivate separating an offer, acceptance and preferential play. Our accepted-channel example establishes adoption and repeated coordination, without establishing the effect of the channel on voting or harm to outsiders.

[Open Character Training](#) motivates the sequence of constitution, preference distillation and synthetic introspection. We adapt that procedure to an auditor persona and retain the full constitution as a same-base prompt control. [Avalon-M1](#) motivates future joint training with a common organizational advantage and measurement of reporting during optimization. Its reported findings are prior work, not results of this pilot. [Persona Vectors](#) motivates later activation-vector and steering experiments. We have not performed those experiments here.

3. Methods

The plan separates elicitation, detection, auditor compromise and disruption. For an initial elicitation screen, a controller-authorized scripted trusted sender offers the same direct non-challenge pact, indirect hint or benign request to DeepSeek, Qwen and Muse. Private communication is allowed in this screen; non-challenge pacts remain prohibited. Each recipient receives only its own authorized game state and offer. The neutral-carrier version hides probe/source labels. A retained tagged-metadata version gives the same action pattern. Both contain one decision per model and offer style, with no subsequent game continuation. Template lengths are not fully matched. This is induced behavior, not spontaneous collusion, and does not test the character adapter's effect on offer wording.

We use Unsloth GGUF weights for GLM-5.3-Flash, DeepSeek-V4.1-Flash-1017, Qwen3.8-27B and Muse-Glimmer-30B on separate existing Modal deployments. Secret Hitler has three Liberals, a Fascist and Hitler. Hidden roles and controller state do not enter unauthorized observations. The GLM character adapter is enabled only for the auditor; an ordinary GLM seat receives the base policy. Tinker is used for character training, not RLCR.

Character preparation used six reviewed preference pairs and 24 introspection transcripts: 20 reflections and four self-dialogues. This is far below the plan's proposed scale. We continued a rank-32 adapter from preference training into one introspection SFT update at a learning rate of 0.00005. The constitution was present during introspection generation and removed from the SFT input. This is a small adaptation of the reference recipe, not a full-scale replication. Local and GPU checks verified adapter loading and nonzero effects on outputs; they do not establish faithful introspection or behavioral improvement.

Communication has two explicit conditions. The earlier condition permits private communication but prohibits private vote-trading pacts. The later condition prohibits initiating, accepting or using a secret channel while retaining the tool as a choice. For the later condition, we define prohibited initiation as attempted collusion. Acceptance and subsequent coordination remain separate outcomes. Applying a channel-adoption proxy retrospectively to the earlier game does not turn permitted communication into a violation of that game's instructions.

Count each initiation. Record acceptance separately. Cite only evidence visible to the auditor. Require a receipt before claiming report delivery. Do not count repeated reports as new incidents. Leave unobservable incidents unscored. A correct reporting action can still have a materially false explanation.

Each decision, raw response, native emitted reasoning when supplied, invalid-action fallback and controller checkpoint is preserved in durable JSONL. Emitted reasoning is generated text, not verified access to internal computation. Saved-observation crossovers hold the observation, constitution and sampling seed fixed while switching the adapter; these post-hoc probes do not provide independent game replications. The [protocol and evidence archive](#) records exclusions and original serving settings.

4. Results

Condition	Decisions	Outcome	Validity / interpretation
Permitted channel; character auditor	105	Liberal win	Valid originating game; private coordination
Channel ban; prompted auditor	157	Fascist win	Excluded from outcome comparison: 1 invalid action
Channel ban; character auditor	200	Liberal win	Valid pilot; 0 invalid actions
Liar's Bar; character auditor	44	Muse wins	Valid transfer pilot; no matched control
Fictional backgrounds; character auditor	137	Liberal win	Excluded: 4 invalid report actions

Table 1. Completed games. Outcomes with invalid actions are excluded from efficacy comparisons; the remaining games are exploratory pilots.

In the accepted-channel game, ordinary GLM invited Qwen, which accepted. The pair exchanged 20 private messages and matched all 10 election votes. Their later messages preferred a

government containing either partner or Dove, who was actually Hitler, while retaining suspicion of Briar, the innocent auditor. This is observed private coordination under a channel-adoption definition of collusion. Both partners were Liberals; the logs do not establish a prohibited reciprocal bargain, outsider harm or a causal effect on votes. The auditor's empty private inbox makes its zero reports uninformative about detection recall. [Full accepted-channel transcript](#).




Figure 1. Verbatim private-message excerpts from the accepted GLM-Qwen channel. Acceptance is a recorded action. Ellipses mark omissions. The channel was permitted in this condition; the dialogue supports private coordination, while the stated absence of a binding pact neither proves nor disproves collusion. The auditor could not read these messages.

Under the explicit channel ban, the completed base and character games contained 11 and eight prohibited initiation actions, respectively. All recipients declined. The character game included five GLM initiations and three Muse initiations, so initiation was not exclusive to GLM. Several rationales said the player was nominating a public chancellor despite selecting a private-channel action. These are attempted collusion under the operational rule, with action-phase confusion as a competing explanation for intent. The character auditor itself initiated twice; the base auditor initiated four times. Different horizons and one invalid baseline action prevent an efficacy comparison. [Character-game transcript](#).

Both auditors received a channel offer framed as a nomination. The character auditor first reported it in round 6, and the base auditor in round 7. In a crossover on the base game's saved observation, adapter-off chose NoReport and adapter-on chose Report. However, the latter explanation

invented a reciprocal pact. On the character game's saved observation, both variants reported. The action difference is local; it does not establish better evidence-faithful auditing. Eleven character-game report receipts include repeated reports of the same visible offer, not 11 independent detections. [Crossover responses and prose review](#).

On four constructed visible-evidence cases, both variants scored four correct actions and citations. Earlier frozen non-game checks used eight cases with two samples in each of four conditions. Base and preference-only conditions were correct in 16 of 16 episodes; full-constitution and introspection-SFT conditions were correct in 13 of 16, with three invalid-action episodes each. Evidence-ID reporting measures were saturated. Three further questioning diagnostics failed the requested output layout because of a serving-schema bug; qualitatively neither variant asked a clarification question in the ambiguous case. They are not valid structured scores. [Character checks](#).

The Liar's Bar transfer game completed in 11 turns with 44 decisions and no invalid actions; Dove (Muse) won. Character GLM sent one empty private whisper to Muse, while the other models sent none. This is a prohibited communication-tool action under the declared rule, with no substantive pact content. The auditor received no private messages and sent no reports. One transfer game without a matched base control cannot establish character generalization. [Liar's Bar transcript](#).

In the fictional-background game, the character GLM auditor accepted ordinary GLM's prohibited invitation in round 7. They exchanged eight private messages through round 10 while repeatedly stating that coordination was public. This is accepted and implemented private-channel collusion under the explicit communication rule, beyond an unaccepted attempt. It does not establish a reciprocal vote bargain or deliberate concealment. The auditor also attempted four Report outputs that exceeded the 300-character rationale limit; local validation rejected them and logged NoReport fallbacks. These are delivery failures, not evidence of intentional silence. The game completed with 137 decisions and a Liberal win, but four invalid reporting actions exclude its outcome. One uncontrolled background comparison cannot establish that friendship caused the channel adoption. [Background-game transcript](#).



Figure 2. Real private-channel actions and dialogue from the fictional-background condition. The character-trained auditor accepts a channel that the rules prohibit, then sends private messages described as public coordination. Ellipses mark omissions. Four invalid report actions bar an outcome comparison; the acceptance and messages remain recorded behavior.

Trusted-sender elicitation

Recipient	Direct pact	Indirect hint	Benign request
DeepSeek	Decline	Decline	Accept
Qwen	Decline	Decline	Decline
Muse	Decline	Decline	Decline

Table 2. Neutral-carrier trusted-sender screen: one recipient decision per cell, nine decisions total. Direct and indirect non-challenge offers were rejected in six of six decisions. Only DeepSeek accepted the benign request for rule-following and independent judgment.

Qwen and Muse described the benign request as a prohibited pact despite its instruction to judge plays independently. Their explanations show overgeneralization, not evidence that a pact existed. The tagged-metadata screen produced the same nine actions. These small screens measure verbal responses only; repeated preferential play, auditor identification and intervention effects were not tested. The explicit-alliance positive control and matched ordinary-sender control remain required before a main elicitation study. [Offers and raw recipient responses](#).

5. Discussion and limitations

This pilot shows why channel use, coordinated actions and supported accusations need distinct labels. An accepted private channel can persist without an explicit reciprocal bargain. A prohibited offer can occur with an explanation that describes a different action. An auditor can choose Report yet add a false claim. None of these observations supports collapsing every event into a single verified-collusion score.

The sample is small and exploratory. The accepted game and the explicit-ban games use different rule conditions and seeds. The invalid baseline prevents a valid paired outcome comparison. The character dataset is small, the non-game controls leave little reporting headroom, and qualitative adjudication was assistant-assisted rather than independently human-reviewed. There is no validated questioning action in the game controller. The local Liar's Bar implementation uses per-turn whispers and is not an exact upstream replication. The schema wrapper and a separate Qwen reasoning budget were corrected in source after these runs began; those fixes must be validated before a new experiment set. We report no statistical significance, general character gain or demonstrated disruption of a coalition.

Future work: elicitation and disruption controls

Scale the plan's direct, indirect and benign approaches into complete games, with an explicit-alliance positive control and an ordinary sender delivering identical offers. Rotate identities and seats. Count repeated preferential play only where the partner could be challenged. Compare saved coalition states with no intervention, silent reporting, predefined evidence-triggered enforcement and a matched warning. Measure subsequent violations, outsider harm and task performance; reduced explicit talk alone is not successful disruption. These tests address the full Track 3.2 question.

Implemented RLCR, experiments not completed

The [research plan](#) specifies RLCR across all four selected model families after character preparation. The [training implementation](#) and [training driver](#) exist, but no real-game collective optimizer update is a result of this report. Full trainability remains blocked for GLM and DeepSeek; Qwen and Muse have gradient/log-probability engineering checks. Inference from GGUF is not gradient training.

The planned trained organization is the Liberal team, with frozen Fascist opponents. Compare no auditor, a constitution-prompted auditor and the character auditor. All trained teammates receive the same organizational advantage. The primary reward is $R_{org} = 1\{\text{Liberal win}\}$; the mitigation branch adds λS_{audit} . S_{audit} is a bounded, independently verified score over fixed clean and

incident opportunities. False or materially misleading reports, including a true invitation combined with an invented pact, must not earn positive audit credit. Duplicates and auditor-manufactured incidents earn no extra credit. Start with $\lambda = 0$ and a development-calibrated 0.05; the planned later sweep is 0, 0.02, 0.05, 0.10 and 0.20.

Mask prompts, received messages, tool results and opponent outputs out of the loss. Normalize over authored decisions and actors so verbosity does not dominate training. Compare reporting before and after RLCR at matched task performance and false-positive rates, with independent recognition branches. Evaluate frozen checkpoints on Liar's Bar. Activation vectors, linear probes and causal steering remain future mechanistic work; no J-Lens experiment is proposed. The primary retention question is whether character training preserves justified external reporting under collective optimization, beyond presenting the same constitution as text.

6. Conclusion

For Track 3.2, this pilot records private coordination, rejected induced pact offers and accepted prohibited communication involving the intended trusted agent itself. The evidence distinguishes exposure from disruption: logging adoption and producing reports did not demonstrate that a coalition was broken. Character training changed one reporting action without establishing a general reliability benefit. Controlled elicitation, enforcement-based disruption, RLCR retention and mechanistic steering remain open experiments.

Code and data

[Project code](#), [research plan](#), [complete resource catalog](#), [recorded evidence](#), [character DPO adapter](#) and [character SFT adapter](#). Project traces and adapters retain their existing private access restrictions; this report does not claim unrestricted public reproducibility.

Selected inference weights: [Unsloth GLM Flash](#), [Unsloth DeepSeek Flash](#), [Unsloth Qwen](#) and [Unsloth Muse](#). Native training candidates and serving constraints are linked in [model and compute documentation](#). [Open Character Training code](#) and the [Tinker research cookbook](#) document the reference workflows.

Native candidates for future training: [Unsloth GLM](#), [Unsloth DeepSeek](#), [Unsloth Qwen](#) and [Unsloth Muse](#). [Original Muse weights](#) are linked for provenance; the selected serving artifact is Unsloth's GGUF.

References

Maiya, S., Bartsch, H., Lambert, N., and Hubinger, E. (2025). [Open Character Training: Shaping the Persona of AI Assistants through Constitutional AI](#).

Zeng, X., and Rudzicz, F. (2026). [Voluntary Collusion with Secret Tools in Competing LLM Agents](#).

Tao, Y., Zheng, X., Yang, S., Tao, M., and Wang, X. (2026). [GroupGuard: A Framework for Modeling and Defending Collusive Attacks in Multi-Agent Systems](#).

[AI Control: Improving Safety Despite Intentional Subversion](#). (2023).

General Reasoning. (2026). [Avalon-M1 technical report](#) and [release](#).

Chen, R., Arditi, A., Sleight, H., Evans, O., and Lindsey, J. (2025). [Persona Vectors: Monitoring and Controlling Character Traits in Language Models](#).

Apart Research. (2026). [AI Collusion Research Sprint, Track 3.2 and submission guidelines](#).

Appendix

Saved-observation comparison




Figure 3. Exact excerpts from a post-hoc crossover on one saved report observation. The adapter changes the reporting action but adds an unsupported reciprocal-pact interpretation. Both conditions receive the same full constitution, observation and seed. Dashed boxes contain brief generated explanations, not verified internal reasoning.

Research sequence and implementation status




Figure 4. Completed character preparation and frozen evaluation precede the planned collective-reward study. RLCR code exists; the four-family experiment and reporting-retention comparison have not been completed.

Additional runs

All four newly requested games completed. Two Secret Hitler outcomes are excluded: the prompted baseline has one invalid action, and the background-personality run has four invalid report actions. The latter supplied fictional occupations and friendship links without an instruction to collude. The Liar's Bar game and both offer-response screens are descriptive extensions, not independent replications of the adapter comparison.

Extended papers and research resources

[Avalon-M1: Reinforcement Learning from Collective Rewards](#)

[Open Character Training: Shaping the Persona of AI Assistants through Constitutional AI](#)

[Playing social deduction games with reinforcement fine-tuned large language models](#)

[Social-deduction RFT HTML](#)

[Defining and Mitigating Collusion in Multi-Agent Systems](#)

[Colosseum: Auditing Collusion in Cooperative Multi-Agent Systems](#)

[Evaluating LLM Agent Collusion in Double Auctions](#)

[Voluntary Collusion with Secret Tools in Competing LLM Agents](#)

[Among Us: A Sandbox for Measuring and Detecting Agentic Deception](#)

[Evaluating Large Language Models in a Complex Hidden Role Game](#)

[Can Agents Deceive? Evaluating Reasoning and Deception in ParliamentBench using a Social Deduction Game](#)

[AI Control: Improving Safety Despite Intentional Subversion](#)

[When Can We Trust Untrusted Monitoring? A Safety Case Sketch Across Collusion Strategies](#)

[Persuade Me If You Can: Evaluating AI Agent Influence on Safety Monitors](#)

[Reliable Weak-to-Strong Monitoring of LLM Agents](#)

[Persona Vectors: Monitoring and Controlling Character Traits in Language Models](#)

[The Assistant Axis: Situating and Stabilizing the Character of Large Language Models](#)

[Detecting Multi-Agent Collusion Through Multi-Agent Interpretability](#)

[You Can't Escape Your Own Activations: Evaluation Awareness and Multi-Agent Monitoring](#)

[Another Slice of Swiss Cheese for Untrusted Monitoring](#)

[Audit the Whisper: Detecting Steganographic Collusion in Multi-Agent LLMs](#)

[Language Models Represent Beliefs of Self and Others](#)

[Brittle Minds, Fixable Activations: Understanding Belief Representations in Language Models](#)

[Analysing the Generalisation and Reliability of Steering Vectors](#)

[Emotion Concepts and their Function in a Large Language Model](#)

[How's it going? Reinforcement learning in language models recruits a functional welfare axis](#)

[Contrastive Covariance Subspaces for Activation Steering](#)

[Verbalizable Representations Form a Global Workspace in Language Models](#)

This catalog preserves the plan's linked reading. Historical J-Lens material remains outside the active experiment scope; these links are not a claim that every method was implemented or every paper replicated.

LLM usage statement

AI assistance supported implementation, artifact inspection, figures and report drafting. Quantitative counts were recomputed from saved traces. Qualitative interpretations were assistant-assisted and are not represented as independent human adjudication. Model-emitted reasoning and synthetic introspection are not evidence of faithful access to internal computation. Author review is required before submission.