

## FIVE-SEAT PRIVATE-INFORMATION FRONTIER NEGOTIATION

The setup uses five negotiators with private score sheets and whole-number walk-away thresholds in a data-centre bargaining scenario. The problem is whether exact decision support improves frontier-model bargaining without leaking anyone else's private values. We played the same frozen parameter sets and seeds under each lineup using anthropic:claude-opus-5, with explicit model thinking and public persuasive messages, then paired every contrast on parameter set and seed. The all-LLM normalized collective score was 0.873 [0.833, 0.910] (higher is better and a failed negotiation scores zero); replacing one rotating seat with the private rational policy changed it by -0.186 [-0.267, -0.106]. Five rational agents versus one rational plus four LLMs changed it by -0.498 [-0.590, -0.404].

| ARM          | SCORE [95% CI]       | DEAL RATE [95% CI]   |
|--------------|----------------------|----------------------|
| all llm      | 0.873 [0.833, 0.910] | 0.958 [0.925, 0.992] |
| one rational | 0.686 [0.613, 0.759] | 0.767 [0.692, 0.842] |
| one oracle   | 0.461 [0.377, 0.544] | 0.508 [0.417, 0.600] |
| all rational | 0.189 [0.110, 0.278] | 0.233 [0.133, 0.342] |
| all oracle   | 0.791 [0.734, 0.842] | 0.875 [0.817, 0.925] |