

LinguaGame: A Linguistically Grounded Game-Theoretic Paradigm for Multi-Agent Dialogue Generation

Yuxiao Ye¹, Yiming Zhang^{2*}, Yiran Ma^{3*}, Huiyuan Xie¹,
Huining Zhu⁴, Zhiyuan Liu^{1†},

¹Tsinghua University, ²University of California, Berkeley, ³Peking University,

⁴East China University of Political Science and Law

{yeyuxiao, liuzy}@mail.tsinghua.edu.cn

Abstract

Large Language Models (LLMs) have enabled Multi-Agent Systems (MASs) where agents interact through natural language to solve complex tasks or simulate multi-party dialogues. Recent work on LLM-based MASs has mainly focused on architecture design, such as role assignment and workflow orchestration. In contrast, this paper targets the interaction process itself, aiming to improve agents' communication efficiency by helping them convey their intended meaning more effectively through language. To this end, we propose *LinguaGame*, a linguistically-grounded game-theoretic paradigm for multi-agent dialogue generation. Our approach models dialogue as a signalling game over communicative intents and strategies, solved with a training-free equilibrium approximation algorithm for inference-time decision adjustment. Unlike prior game-theoretic MASs, whose game designs are often tightly coupled with task-specific objectives, our framework relies on linguistically-informed reasoning with minimal task-specific coupling. Specifically, it treats dialogue as intentional and strategic communication, requiring agents to infer what others aim to achieve (intents) and how they pursue those goals (strategies). We evaluate our framework in simulated courtroom proceedings and debates, with human expert assessments showing significant gains in communication efficiency.

Code & data: <https://github.com/thunlp/LinguaGame>

1 Introduction

A Multi-Agent System (MAS) consists of autonomous, interacting agents that learn from their environment and peers to make decisions and collaboratively solve tasks (Dorri et al., 2018). With the recent advancement of Large Language Models (LLMs), MASs can now be constructed by

assigning each agent (instantiated as an LLM) a role and enabling them to communicate with one another in natural language (Park et al., 2023; Guo et al., 2024; Han et al., 2024). LLM-based MASs have recently been applied to collaborative problem-solving (e.g. software development and collective reasoning) and to multi-party dialogue simulation (e.g. courtroom proceedings) (Cui et al., 2023; Chen et al., 2024a,b; Qian et al., 2024b).

Most existing approaches focus on tailoring MAS architectures to downstream tasks (Li et al., 2024; Dang et al., 2025; Yue et al., 2025), with improvements driven largely by role-assignment prompts and workflow orchestration. However, such architecture-centric methods often overlook the underlying dynamics of agent interaction. To address this limitation, we propose *LinguaGame*, a game-theoretic paradigm that models multi-agent dialogue as a signalling game over communicative intents and strategies.

Game theory provides a principled framework for modelling strategic interaction among agents (Barron, 2024). Existing game-theoretic MASs typically design games tightly coupled with specific tasks, or even directly instantiated as the tasks themselves, which limits their generalisability (Peters et al., 2024; He et al., 2025). In contrast, *LinguaGame* grounds game design in linguistic reasoning: agents infer each other's communicative intents and strategies. That is, they reason about what others aim to achieve through their utterances and how language is strategically formulated to serve those goals. This linguistically grounded formulation enables flexible game design with minimal task-specific adaptation.

Our approach is also computationally efficient at inference time for multi-agent dialogue generation. Whereas conventional MASs improve through model retraining or repeated experience accumulation (Cheng et al., 2024; Chen et al.,

*Equal contribution.

†Corresponding author.

2024a; Qian et al., 2024a), LinguaGame enables inference-time decision adjustment via a training-free equilibrium approximation algorithm. This plug-and-play mechanism reduces computational cost and deployment overhead, making it practical for diverse MAS settings.

Our evaluation spans simulated courtroom proceedings and debate scenarios. Unlike tasks such as code completion or question answering, which primarily focus on end results, these settings emphasise open-ended dialogue, where the conversational process itself is of substantive importance. This makes them well suited for evaluating the quality of multi-agent dialogue generation. Human expert evaluations demonstrate that LinguaGame significantly improves dialogue quality by enhancing agents’ communication efficiency. Although our experiments focus on legal and argumentative settings, the core mechanism “modelling dialogue as the interplay of communicative intents and strategies” reflects a fundamental aspect of natural language use. As such, our approach has the potential to generalise to diverse multi-agent dialogue generation scenarios.

Our contributions are as follows:

- We propose LinguaGame, a novel game-theoretic paradigm for multi-agent dialogue generation that models communication as a signalling game grounded in linguistic reasoning.
- Our framework decouples game design from task-specific objectives and relies on pragmatic aspects of communication, enabling broad generalisation across tasks with minimal adaptation.
- We introduce a training-free equilibrium approximation algorithm that guides agent decisions in inference time. It can be integrated in a plug-and-play manner, without requiring model retraining or iterative interaction.
- We demonstrate the effectiveness of our framework through experiments in simulated courtroom and debate scenarios, showing significant improvements in communication efficiency.

2 Related Work

We position our work at the intersection of three lines of research: LLM-based MASs in general, game-theoretic MASs in particular, and behaviour improvement mechanisms in MASs.

LLM-based MAS As an emerging research direction, most existing work on LLM-based MASs focuses on designing system structures to enable these systems to perform specific tasks. This typically involves refining role assignment and workflow orchestration. For instance, ChatDev (Qian et al., 2024b), a software development MAS, defines five agent roles, namely CEO, CTO, reviewer, tester, and programmer, and coordinates their collaboration via an inception prompting mechanism (Li et al., 2023). In our work, we do not pursue performance gains through such structural refinement. Instead, we take a closer look at the internal dynamics of MASs, focusing on agent communication. Our aim is to improve overall dialogue generation quality by enhancing agents’ communication efficiency.

A closely related application setting is the use of LLMs in debate, negotiation, and other strategic dialogue tasks, where multi-agent interaction is used to improve reasoning or decision outcomes. For example, Wang et al. (2023) evaluate whether debate-style interaction can elicit more truthful or better-justified reasoning. Negotiation and bargaining settings have been explored via workflow-based agent designs or market-style interaction protocols (Deng et al., 2024; Hua et al., 2024). While these works are adjacent in that they also study strategic multi-agent LLM interaction, they typically target task-level success (*e.g.* correctness and win rate) and evaluate improvements accordingly. In contrast, our work focuses on the communicative process itself, improving interaction efficiency and dialogue quality by selecting utterances that better convey an agent’s intended intent and strategy, making our approach complementary and potentially pluggable into debate/negotiation pipelines as an interaction-level module.

Game-theoretic MAS Although MASs are inherently suited for game-like interactions involving multiple players, research on game-theoretic MASs remains limited. Existing work in this area typically incorporates game theory by constructing task-specific games that are closely aligned with or directly represent the target tasks themselves. For example, in the task of robotic motion planning for autonomous driving, Peters et al. (2024) formulate a contingency game in which MAS agents select strategic driving trajectories based on possible future intents of others. For the task of mobile networking optimisation, He et al.

(2025) design an MAS where agents play games which are subtasks of networking optimisation. In contrast, our game formulation is grounded in general linguistic principles derived from Speech Act Theory (Searle, 1969), which views utterances as expressions of speaker intent. This abstraction enables our method to generalise more easily across tasks.

Within game-theoretic approaches to language use, our work is most closely related to the Rational Speech Act (RSA) framework (Frank and Goodman, 2012; Goodman and Frank, 2016). It models pragmatic communication as a process in which speakers choose utterances to convey intended meanings and listeners infer those meanings. RSA-style models are typically instantiated in tightly constrained reference games or other toy domains with explicitly defined utterance and meaning spaces. For example, Goodman and Frank (2016) study reference games and scalar implicature tasks over small, enumerated meaning and utterance sets. Monroe et al. (2017) apply RSA-style reasoning to a colour reference game with discretised meanings and a restricted descriptive vocabulary. In contrast, LinguaGame shifts pragmatic reasoning from utterance-level inference to intent-strategy abstraction, enabling RSA-style models to be applied in realistic, LLM-mediated multi-agent dialogue.

MAS Behaviour Improvement Existing LLM-based MASs often improve agent behaviour through learning-based mechanisms, such as model retraining or iterative experience accumulation. For example, in the word-guessing MAS of Cheng et al. (2024), agents repeatedly play the game to collect successful dialogue episodes, which are then used to fine-tune the LLM via offline reinforcement learning. In the MAS by Chen et al. (2024a), a lawyer agent records insights from each simulated courtroom proceeding into an experience database which it consults in subsequent rounds, gradually improving argumentative performance through accumulated experience. These approaches require either costly retraining or repeated execution of the full system. In contrast, LinguaGame operates entirely at inference time, using a training-free equilibrium approximation to re-rank candidate utterances, without updating model parameters.

3 Method

We propose LinguaGame, a linguistically-grounded game-theoretic paradigm for multi-agent dialogue generation. The framework operates on top of a standard MAS foundation, augmenting it with a game-theoretic formulation of agent interaction and a training-free algorithm for equilibrium approximation.

3.1 MAS Structure

MAS for Courtroom Proceedings We configure our MAS for simulated courtroom proceedings in accordance with first-instance proceedings under Chinese civil law. In this configuration, the procedure comprises five sequential stages that involve interactive or strategic communication: 1) Evidence Presentation and Cross-Examination, 2) Court Investigation and Questioning, 3) Court Debate, 4) Final Statement, and 5) Judgment Announcement.

This MAS retains three primary roles in courtroom proceedings: judge, plaintiff lawyer, and defendant lawyer. Each simulated courtroom proceeding is grounded in a real-world written judgment drawn from Chinese court decisions. This material serves as the shared factual basis and is made available to all roles throughout the proceeding. Prior to the start, the judge is briefed on the five-stage procedure. The judge then initiates the proceeding and manages turn-taking by deciding which role speaks next after each contribution. Within each stage, the full conversation history is visible to all roles. At the end of each stage, the entire dialogue from that stage is summarised by an LLM. This summary is then provided to all roles as shared context for the subsequent stage.

MAS for Debates We configure our MAS for simulated debates as open-ended exchanges between two opposing roles on a given controversial proposition (e.g. “*abortion should be legalised*”). The two roles are defined symmetrically: a proponent, who argues in favour of the proposition, and an opponent, who argues against it. To encourage substantive discussion, each arguer is required to develop arguments from at least two distinct aspects.

Unlike courtroom proceedings, debates are not governed by a fixed procedural structure. Instead, the arguers alternate turns, presenting claims, counterarguments, and rebuttals in an open-ended dialogue. A debate concludes when the dialogue

converges naturally to closure, as determined by an LLM-based assessment. Throughout the debate, both arguers have access to the complete conversation history.

3.2 Game Formulation

We adopt a game-theoretic perspective to model agent interaction as strategic communication. Specifically, we use the *signalling game* framework (Lewis, 2008; Sobel, 2020).

Signalling Game In a standard signalling game, a *sender* selects a *signal* to convey *private information* that is visible only to them, and a *receiver* interprets the signal to infer that hidden information and respond accordingly. Each player in a signalling game is assigned a *utility function*, which quantifies how well the outcome of the interaction aligns with their goals.

In our adaptation of the signalling game to multi-agent conversations, a game is instantiated immediately before an agent generates its next utterance. The agent about to speak acts as the sender, and the agent scheduled to reply acts as the receiver. The signal corresponds to the utterance generated by the sender.

Intent-strategy as Private Information To define the private information that needs to be inferred, we draw on Speech Act Theory (Searle, 1969), which holds that utterances are not merely vehicles of content but also convey speakers' intentions.

Building on this idea, we model the sender's private information as a pair of communicative intent and strategy. Intents capture what the speaker aims to achieve, while strategies capture how the intent is realised linguistically. For courtroom proceedings we define nine intents and twelve strategies; for debates we define six intents and eight strategies (Appendices A-B). We do not devise strategies for every intent, but only for those where agents can pursue their goals in substantially different ways. Intents such as MANAGING PROCEEDING and RULING, which are procedural or outcome-driven in nature, do not require further strategic differentiation.

The design of the intent and strategy inventories is grounded in domain expertise and established argumentation theory. For courtroom proceedings, they were derived through consultation with practising Chinese judges and lawyers, while for debates they follow established Argumentation

Schemes (Walton et al., 2008). Aligning the intent and strategy inventory with the communicative conventions of specific scenarios requires only minimal effort. In our case, defining them for each scenario took around two hours.

3.3 Game-MAS Integration

Figure 1 illustrates how the signalling game is integrated into the MAS. Given a conversational context c , the sender S first selects an intent-strategy pair (i, s) . Conditioned on c and (i, s) , S then generates a finite set of candidate utterances $\{u_1, u_2, u_3, \dots\}$. For each candidate utterance u , the receiver R interprets it and independently¹ infers the underlying intent and strategy. Key prompts used in this process are shown in Appendix C.

Based on the outcome of this inference, the policies of both S (utterance generation) and R (intent and strategy inference) are updated through the equilibrium approximation procedure. The candidate utterance with the highest probability under the updated sender policy, given (i, s) , is selected as the *winning utterance*.

After the game concludes, only the winning utterance is appended to the dialogue context and passed on as part of the conversation. All intermediate information used during the game (such as candidate utterances, intents, and strategies) is discarded, ensuring that agents do not retain privileged information when continuing the dialogue. At each dialogue turn, this process constitutes an inner-loop equilibrium search over sender and receiver policies induced by the base language model and inference prompts, whose sole output is the selected utterance appended to the dialogue.

3.4 Equilibrium Approximation

Signalling games are typically solved by identifying a *Nash equilibrium*, a stable configuration in which no player has an incentive to unilaterally change how they convey or interpret messages. In other words, solving the game requires finding a set of policies such that each player's choices maximise their expected utility.

¹Although strategies depend on intents by design, R infers them independently: intents are chosen from the full space without regard to strategies, and strategies are chosen from the subset of the ground-truth intent without feeding back into intent inference. This preserves the hierarchical relation while avoiding the combinatorial complexity of joint inference.

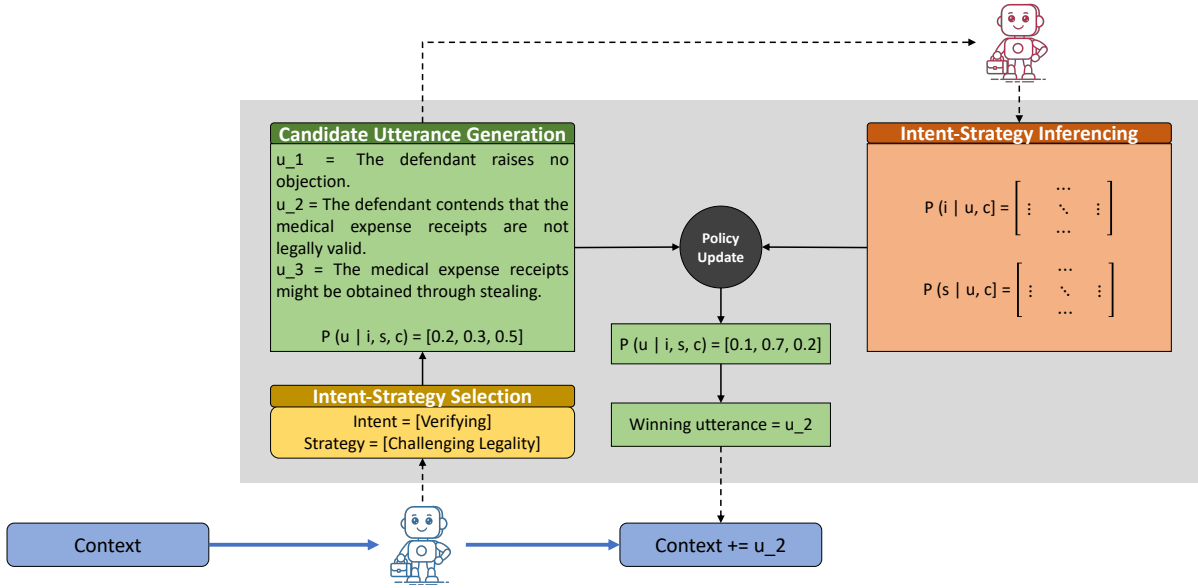


Figure 1: Integration of the signalling game into the MAS in LinguaGame. The sender (the agent in blue) generates candidate utterances, which are interpreted by the receiver (the agent in red). The utterance with the highest probability under the updated sender policy is appended to the dialogue context. All intermediate information produced during the gaming process (in the grey box) is discarded once the game concludes.

Utility Functions Although courtroom and debate settings are adversarial in terms of overall goals, from the perspective of communication efficiency the interacting roles must cooperate to ensure mutual understanding. We therefore formulate this signalling game as a *cooperative game*, in which both the sender and the receiver get a reward when the private information is correctly inferred. Thus, both the sender and the receiver share a same part in their utility functions, defined as follows:

$$U_{shared} = \sum_{i,s} p(i, s) \sum_u \pi_S(u | i, s) [w\pi_{R_i}(i | u) + (1 - w)\pi_{R_s}(s | u)]. \quad (1)$$

Here $p(i, s)$ denotes the prior distribution of an intent-strategy pair (i, s) given the context². Note that the prior distribution of intent-strategy pairs appears in the utility definitions only for the purpose of policy optimisation. This information is never revealed to the receiver during inference. The sender’s policy, $\pi_S(u | i, s)$ specifies the probability of choosing utterance u conditioned on (i, s) and c . The receiver’s policy $\pi_{R_i}(i | u)$ is for inferring the intent, and $\pi_{R_s}(s | u)$ is for inferring the strategy. The scalar $w \in [0, 1]$ is a weighting hyperparameter that balances the relative im-

²For brevity, we omit explicit conditioning on the context c in all equations where it is understood.

portance of intent and strategy in the overall utility calculation.

Following Jacob et al. (2024), we add a KL-regularisation term into the utility function to prevent abrupt or unstable policy updates:

$$U_S(\pi_S | \pi_R) = -\lambda \text{KL}(\pi_S || \pi_S^0) + U_{shared}, \quad (2)$$

$$U_R(\pi_R | \pi_S) = -\lambda \text{KL}(\pi_R || \pi_R^0) + U_{shared}, \quad (3)$$

where π_S^0 and π_R^0 denote the initial reference policies, and $\lambda > 0$ is a hyperparameter controlling the strength of regularisation.

Utility Optimisation In order to compute a Nash equilibrium, we optimise the utility functions with no-regret learning (Daskalakis et al., 2021), in which policies are iteratively updated and gradually converge to equilibrium.

We adopt the piKL algorithm (Jacob et al., 2022) to implement no-regret learning, where each player iteratively updates its policy by optimising the expected utility. Importantly, this equilibrium approximation is performed entirely as an inner-loop optimisation at inference time and does not involve additional LLM calls. Further details can be found in Appendix D.

As the update process converges, the players are regarded as reaching an approximate Nash equilibrium. The utterance chosen for MAS generation is

then the candidate with the highest probability under the sender’s final policy given the ground-truth intent-strategy pair.

4 Experiments

Building on the courtroom proceeding scenario introduced in the Method section, we conduct experiments to evaluate the effectiveness of *LinguaGame*. Specifically, we assess whether it improves communication efficiency in multi-agent dialogue generation through expert evaluation of the generated dialogues.

4.1 Experimental Setup

We choose Qwen2.5-32B (Qwen et al., 2025) as the base LLM in the experiment. In our pilot studies, we experimented with several alternatives, including MiniCPM4-8B (MiniCPM et al., 2025), Qwen2.5-7B (Qwen et al., 2025), as well as Meta-Llama-3.1-8B-Instruct and Meta-Llama-3.1-70B-Instruct (Grattafiori et al., 2024). We observed that smaller models (7B and 8B variants) exhibited notably weaker instruction-following capabilities, while the two larger models performed comparably in instruction adherence and dialogue quality.

For the hyperparameters used in the signalling game, we set the number of candidate utterances to 3. Following a grid search in our pilot experiments, we set the weighting hyperparameter w to 0.5, or to 1.0 when the selected intent lacks associated strategies. The KL-regularisation weight λ and the policy update learning rate are both set to 0.1. The equilibrium approximation runs for 5000 optimisation rounds. While these hyperparameters could be further tuned, hyperparameter optimisation is not the focus of this paper. The current configuration is sufficient to produce meaningful results for evaluating the framework.

For courtroom proceedings, we randomly select 50 cases from China Judgments Online (Supreme, 2025), each associated with a distinct cause of action. From the written judgment of each case, we extract the party information³ and a brief case description, which serve as input for the simulation. For debates, we randomly select 50 controversial propositions from the Quora Argumentation Mining dataset (Ye and Teufel, 2024), spanning 20 topics such as politics, society, and technology, and use them as input for the simulation.

³Only non-identifiable attributes such as gender and age are preserved without modification.

4.2 Systems

Our method is designed as a plug-in interaction layer that can be applied on top of existing MAS dialogue generation systems without modifying their underlying architectures. Accordingly, we adopt a stepwise “before-after” evaluation design, comparing the same base system with progressively added components to isolate the contribution of each component. We compare the following three systems:

- **SDMAS**: The standard MAS setup described in the Method section.
- **ISMAS**: The standard MAS equipped with our intent-strategy design. In this setting, each agent generates utterances conditioned on a given intent-strategy pair in a Chain-of-Thought (CoT) (Wei et al., 2022) manner.
- **LGMAS**: The standard MAS integrated with our full *LinguaGame* paradigm, which incorporates both the intent-strategy design and the game-theoretic modelling.

To evaluate the game-theoretic utterance selection in *LinguaGame* against a strong inference-time alternative, we introduce an LLM-based utterance re-ranking baseline. This baseline selects the final utterance from the same set of candidate utterances generated by *LGMAS*, using a single-shot judgment by an LLM. Specifically, given the full dialogue history up to the current turn and a fixed set of candidate utterances, the LLM is prompted to select the most appropriate next utterance. The re-ranking decision is made independently at each turn. We experimented with Qwen2.5-32B, Meta-Llama-3.1-70B-Instruct, and DeepSeek-V3-0324 (DeepSeek-AI, 2024) for this baseline.

4.3 Evaluation

We conduct manual evaluations of the dialogues generated by the systems described above. Each dialogue represents a complete generation at the case or proposition level. Eight graduate students, four majoring in linguistics and four in law, were recruited as annotators. The evaluation is conducted at two levels: first at the utterance level, then at the dialogue level. Annotators were blind to the system identities, and all evaluation items were presented in randomised order. The evaluation set contains 100 dialogues, covering approximately 9,100 utterances and 919k tokens per system. Each annotator spent approximately 56 hours

		<i>SDMAS</i>			<i>ISMAS</i>			<i>LGMAS</i>		
		Court	Debate	Overall	Court	Debate	Overall	Court	Debate	Overall
Statistics	utt/dialogue	180	48	114	124	34	79	126	36	81
	token/utt	80	135	108	65	120	93	70	125	98
Linguistic form	Clarity	4.10	4.15	4.13	4.18	4.22	4.20	4.36	4.42	4.39
	Conciseness	3.55	3.62	3.59	3.73	3.69	3.71	4.12	4.19	4.16
Content quality	Argument	3.18	3.27	3.23	3.23	3.34	3.29	3.83	3.92	3.88
	Tactic	3.28	3.39	3.34	3.35	3.42	3.39	3.96	4.05	4.01

Table 1: Statistics of the generated dialogues and results for the dialogue-level evaluation.

on evaluation, enabling fine-grained, expert-level assessment.

Utterance-level Evaluation Utterance-level evaluation is conducted for *LGMAS* and the re-ranking baseline only. We focus on instances where the “winning utterance” selected by the re-ranking mechanism differs from the “initial utterance” preferred under the sender’s initial policy, in order to examine whether *LinguaGame* alters utterance selection in a beneficial direction.

For each evaluation item, annotators are presented with: the preceding utterance, and the three candidate utterances generated in the signalling game. They are asked to select the candidate that best follows the preceding utterance, based solely on local coherence and appropriateness. The candidate receiving the most votes is recorded as the human-preferred utterance. In the event of a tie, one of the top-voted candidates is randomly selected. We then compare this preference to the re-ranking selection: 1) if it matches the winning utterance, the instance is counted as a *positive alternation*; 2) if it matches the initial utterance, the instance is counted as a *negative alternation*; 3) otherwise, the instance is counted as a *neutral alternation*.

This utterance-level evaluation focuses on relative changes in utterance selection; absolute dialogue quality is assessed separately at the dialogue level.

Dialogue-level Evaluation Dialogue-level evaluation provides an overall assessment of dialogue quality and is conducted for all three systems. For each complete dialogue, annotators rate quality on a 5-point Likert scale (Joshi et al., 2015) along four dimensions. *Linguistic form* is evaluated in terms of **clarity**, assessing grammatical correctness and structural coherence, and **conciseness**,

assessing whether essential content is conveyed without unnecessary repetition or digression. *Content quality* is evaluated in terms of **argument**, assessing whether utterances are well-grounded and logically aligned (for courtroom proceedings, with the legal context), and **tactic**, assessing whether utterances exhibit adaptive reasoning and responsiveness to prior moves (for courtroom proceedings, from a legal perspective).

All dialogues, whether from courtroom proceedings or debates, were assessed by all eight annotators. The linguistics annotators evaluated the two dimensions related to linguistic form, while the law annotators assessed the two dimensions concerning content quality.

5 Results and Discussion

The statistics of the generated dialogues and the results of the dialogue-level evaluation are presented in Table 1. Distributions of intents and strategies, illustrating how agents employ the signalling space in practice, are shown in Appendix E. Inter-annotator agreement (IAA) for the dialogue-level evaluation is measured using quadratic weighted kappa (Cohen, 1968). The IAA scores are $\kappa = 0.62$ for the law annotators and $\kappa = 0.79$ for the linguistics annotators, both indicating substantial agreement (Landis and Koch, 1977).

According to Table 1, *LGMAS* achieves the highest scores across all four evaluation dimensions, indicating consistent improvements over both *SDMAS* and *ISMAS*. The improvements are consistent across both courtroom and debate scenarios. These results demonstrate that **integrating the full *LinguaGame* paradigm leads to significant⁴ improvements in dialogue generation**

⁴Significance in this paper is tested through two-tailed paired t-tests at $\alpha = 0.01$.

quality. Examining individual evaluation dimensions in more detail, **clarity** scores indicate that *LGMAS* produces utterances that are more grammatically correct, structurally coherent, and easier for humans to interpret. **Conciseness** is also significantly improved, suggesting that agents convey essential information with less unnecessary elaboration or repetition; this is further supported by the fact that *LGMAS* generates fewer utterances overall and fewer tokens per utterance than *SDMAS*. For **argument**, *LGMAS* outperforms both baselines, indicating better-grounded and more contextually appropriate statements. Finally, improvements in **tactic** suggest more coherent turn-level progression and more adaptive responses to prior utterances.

While gains in the two linguistic dimensions can be directly attributed to enhanced communication efficiency, improvements in content-related dimensions require closer analysis. Higher scores in argument strength and tactical coherence may arise from two sources: more effective selection of intents and strategies, or clearer articulation of the intended messages. Crucially, *ISMAS* and *LGMAS* employ the same prompted pipeline for intent and strategy selection and use identical priors. Therefore, the additional improvements of *LGMAS* over *ISMAS* cannot be attributed to differences in intent or strategy choice. We thus attribute the observed gains in content quality primarily to improved articulation achieved through more effective utterance selection.

As for the **impact of intent-strategy conditioning alone**, results suggest that **it primarily contributes to more concise expression, without yielding substantial improvements in other dimensions**. Specifically, compared with *SDMAS*, *ISMAS* achieves a statistically significant gain of 0.12 in conciseness, indicating that conditioning utterance generation on intent-strategy pairs encourages more focused responses. However, differences between *ISMAS* and *SDMAS* in the remaining three dimensions are minimal and statistically insignificant, suggesting that intent-strategy conditioning alone is insufficient to improve overall dialogue quality without the game-theoretic reasoning mechanism introduced by *LinguaGame*.

Results for the utterance-level evaluation are present in Table 2. Out of the total 8,139 utterances generated by *LGMAS*, 1,366 (16.8%) are altered by the game-theoretic selection process. From these, 1,000 instances (500 from each sce-

	Pos.	Neut.	Neg.	Count
<i>LGMAS</i>	78.1%	12.4%	9.5%	1,366
Qwen2.5	44.0%	38.4%	17.6%	318
Llama-3.1	36.7%	40.8%	22.5%	1,503
DeepSeek-V3	51.2%	14.6%	34.2%	972

Table 2: Distribution of utterance-level alternations selected by different re-ranking methods. Counts indicate the number of altered utterances.

nario) are randomly sampled for human evaluation. For comparison, we apply the same sampling strategy to Llama-3.1, which also produces more than 1,000 altered utterances, while for Qwen2.5 and DeepSeek-V3, all altered utterances are included in the human evaluation. Inter-annotator agreement for the utterance-level evaluation is $\alpha = 0.74$ in Krippendorff’s alpha (Krippendorff, 2004), suggesting reliable annotation (Krippendorff, 2018). As illustrated in Table 2, for *LGMAS*, 78.1% of altered utterances are judged as positive alternations, meaning that the selected utterance better follows the dialogue context than the original choice. Only 9.5% are labelled as negative, with the remaining 12.4% judged as neutral. We further compare *LGMAS* against several LLM-based utterance re-ranking baselines. While re-ranking with DeepSeek-V3 also yields a non-trivial proportion of positive alternations, none of the LLM re-ranking baselines match the positive alternation rate achieved by *LGMAS*. These results indicate that **LinguaGame enhances local coherence and appropriateness at the utterance level, and is more effective than generic LLM-based re-ranking**.

Taken together, the results from both dialogue-level and utterance-level evaluations provide strong evidence that **our LinguaGame framework substantially improves dialogue generation quality**. These **improvements are primarily driven by enhanced communication efficiency**, allowing agents to generate responses that are not only clearer and more concise but also more contextually appropriate and argumentatively coherent.

6 Conclusion

In this paper, we propose *LinguaGame*, a linguistically-grounded, game-theoretic paradigm for multi-agent dialogue generation. Our goal is to improve dialogue generation quality by enhanc-

ing communication efficiency. LinguaGame models agent interaction as a signalling game over communicative intents and strategies, encouraging more purposeful and coherent utterances. Our game design is grounded in linguistic principles rather than task-specific assumptions, which may enhance the paradigm’s generalisability across domains. Its training-free equilibrium approximation removes the need for model retraining or iterative experience accumulation, allowing plug-and-play integration with existing systems.

We evaluate LinguaGame in a simulated courtroom proceeding scenario and a simulated debate scenario. Experimental results show that LinguaGame significantly improves dialogue generation quality, with improvements stemming primarily from agents’ ability to convey intended messages effectively. These findings highlight the potential of modelling pragmatic reasoning over communicative goals as a general approach to enhancing multi-agent dialogue behaviour in future research.

Limitations

Our study has several limitations that suggest directions for future work.

Scope of task objectives. This work focuses on multi-agent dialogue generation, where the primary objective is to model and improve the communicative process itself rather than task-specific end results. While effective communication is often a prerequisite for successful coordination, how the proposed framework influences downstream task performance in settings that emphasise end results remains an open question and is left for future work.

Evaluation domains and scale. Our evaluation focuses on courtroom proceedings and debates, which emphasise adversarial and strategic dialogue. Extending experiments to additional domains, such as collaborative problem-solving or everyday conversation, would further validate the generality of the approach. Such extensions are currently constrained by the substantial human-evaluation burden (approximately 28 hours per annotator per domain), which also limited experiments with multiple base LLMs.

Responsible AI considerations. This work explores simulated multi-agent dialogue in legal settings as a research tool for studying strategic

communication, rather than as a deployable legal decision-making system. LinguaGame is not intended for real-world legal deployment without substantial additional validation, including fairness, bias, and robustness evaluation. Any practical use in legal contexts would require careful assessment of potential harms, representational biases, and alignment with legal and ethical standards.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62236004), Beijing Municipal Science and Technology Plan Project (Z241100001324025), and the high-quality development project of MIIT.

References

- Emmanuel N Barron. 2024. *Game theory: an introduction*. John Wiley & Sons.
- Guhong Chen, Liyang Fan, Zihan Gong, Nan Xie, Zixuan Li, Ziqiang Liu, Chengming Li, Qiang Qu, Shiwen Ni, and Min Yang. 2024a. Agentcourt: Simulating court with adversarial evolvable lawyer agents. *CoRR*.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chi-Min Chan, Heyang Yu, Yaxi Lu, Yi-Hsin Hung, Chen Qian, Yujia Qin, Xin Cong, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. 2024b. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors. In *International Conference on Representation Learning*, volume 2024, pages 20094–20136.
- Pengyu Cheng, Yong Dai, Tianhao Hu, Han Xu, Zhisong Zhang, Lei Han, Nan Du, and Xiaolong Li. 2024. Self-playing adversarial language game enhances llm reasoning. *Advances in Neural Information Processing Systems*, 37:126515–126543.
- Jacob Cohen. 1968. Weighted kappa: Nominal scale agreement provision for scaled disagreement or partial credit. *Psychological bulletin*, 70(4):213.
- Jiayi Cui, Zongjian Li, Yang Yan, Bohua Chen, and Li Yuan. 2023. Chatlaw: Open-source legal large language model with integrated external knowledge bases. *CoRR*.
- Yufan Dang, Chen Qian, Xueheng Luo, Jingru Fan, Zihao Xie, Ruijie Shi, Weize Chen, Cheng Yang, Xiaoyin Che, Ye Tian, and 1 others. 2025. Multi-agent collaboration via evolving orchestration. *arXiv preprint arXiv:2505.19591*.
- Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. 2021. Near-optimal no-regret

- learning in general games. *Advances in Neural Information Processing Systems*, 34:27604–27616.
- DeepSeek-AI. 2024. [Deepseek-v3 technical report](#). Preprint, arXiv:2412.19437.
- Yuan Deng, Vahab Mirrokni, Renato Paes Leme, Hanrui Zhang, and Song Zuo. 2024. Llms at the bargaining table. In *Agentic Markets Workshop at ICML*, volume 2024.
- Ali Dorri, Salil S Kanhere, and Raja Jurdak. 2018. Multi-agent systems: A survey. *Ieee Access*, 6:28573–28593.
- Michael C Frank and Noah D Goodman. 2012. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998.
- Noah D Goodman and Michael C Frank. 2016. Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11):818–829.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges. In *IJCAI*.
- Shanshan Han, Qifan Zhang, Yuhang Yao, Weizhao Jin, Zhaozhuo Xu, and Chaoyang He. 2024. Llm multi-agent systems: Challenges and open problems. *CoRR*.
- Long He, Geng Sun, Dusit Niyato, Hongyang Du, Fang Mei, Jiawen Kang, Mérouane Debbah, and Zhu Han. 2025. Generative ai for game theory-based mobile networking. *IEEE Wireless Communications*, 32(1):122–130.
- Wenyue Hua, Ollie Liu, Lingyao Li, Alfonso Amayuelas, Julie Chen, Lucas Jiang, Mingyu Jin, Lizhou Fan, Fei Sun, William Wang, and 1 others. 2024. Game-theoretic llm: Agent workflow for negotiation games. *arXiv preprint arXiv:2411.05990*.
- Athul Paul Jacob, Yikang Shen, Gabriele Farina, and Jacob Andreas. 2024. The consensus game: Language model generation via equilibrium search. In *Proceedings of the 12th International Conference on Learning Representations*.
- Athul Paul Jacob, David J Wu, Gabriele Farina, Adam Lerer, Hengyuan Hu, Anton Bakhtin, Jacob Andreas, and Noam Brown. 2022. Modeling strong and human-like gameplay with kl-regularized search. In *International Conference on Machine Learning*, pages 9695–9728. PMLR.
- Ankur Joshi, Saket Kale, Satish Chandel, and D Kumar Pal. 2015. Likert scale: Explored and explained. *British journal of applied science & technology*, 7(4):396.
- Klaus Krippendorff. 2004. Measuring the reliability of qualitative text analysis data. *Quality and quantity*, 38:787–800.
- Klaus Krippendorff. 2018. *Content analysis: An introduction to its methodology*. Sage publications.
- J Richard Landis and Gary G Koch. 1977. The measurement of observer agreement for categorical data. *biometrics*, pages 159–174.
- David Lewis. 2008. *Convention: A philosophical study*. John Wiley & Sons.
- Guohao Li, Hasan Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. 2023. Camel: Communicative agents for "mind" exploration of large language model society. *Advances in Neural Information Processing Systems*, 36:51991–52008.
- Xinyi Li, Sai Wang, Siqi Zeng, Yu Wu, and Yi Yang. 2024. A survey on llm-based multi-agent systems: workflow, infrastructure, and challenges. *Vicnearth*, 1(1):9.
- MiniCPM, Chaojun Xiao, Yuxuan Li, Xu Han, Yuzhuo Bai, Jie Cai, Haotian Chen, Wentong Chen, Xin Cong, Ganqu Cui, and 1 others. 2025. Minicpm4: Ultra-efficient llms on end devices. *arXiv preprint arXiv:2506.07900*.
- Will Monroe, Robert XD Hawkins, Noah Goodman, and Christopher Potts. 2017. Colors in context: A pragmatic neural model for grounded language understanding. *Transactions of the Association for Computational Linguistics*, 5:325–338.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.
- Lasse Peters, Andrea Bajcsy, Chih-Yuan Chiu, David Fridovich-Keil, Forrest Laine, Laura Ferranti, and Javier Alonso-Mora. 2024. Contingency games for multi-agent interaction. *IEEE Robotics and Automation Letters*, 9(3):2208–2215.
- Chen Qian, Yufan Dang, Jiahao Li, Wei Liu, Zihao Xie, Yifei Wang, Weize Chen, Cheng Yang, Xin Cong, Xiaoyin Che, and 1 others. 2024a. Experiential co-learning of software-developing agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5628–5640.
- Chen Qian, Wei Liu, Hongzhang Liu, Nuo Chen, Yufan Dang, Jiahao Li, Cheng Yang, Weize Chen, Yusheng Su, Xin Cong, Juyuan Xu, Dahai Li, Zhiyuan Liu,

and Maosong Sun. 2024b. *ChatDev: Communicative agents for software development*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15174–15186. Association for Computational Linguistics.

Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, Junyang Lin, and 24 others. 2025. *Qwen2.5 technical report*. Preprint, arXiv:2412.15115.

John R Searle. 1969. *Speech acts: An essay in the philosophy of language*. Cambridge University.

Joel Sobel. 2020. Signaling games. In *Complex social and behavioral systems: Game theory and agent-based models*, pages 251–268. Springer.

People’s Court of the People’s Republic of China Supreme. 2025. China judgments online. <https://wenshu.court.gov.cn>. Accessed: 2025-04-21.

Douglas Walton, Christopher Reed, and Fabrizio Macagno. 2008. *Argumentation schemes*. Cambridge University Press.

Boshi Wang, Xiang Yue, and Huan Sun. 2023. Can chatgpt defend its belief in truth? evaluating llm reasoning via debate. In *Findings of the association for computational linguistics: EMNLP 2023*, pages 11865–11881.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Yuxiao Ye and Simone Teufel. 2024. Computational modelling of undercuts in real-world arguments. In *Proceedings of the 11th Workshop on Argument Mining (ArgMining 2024)*, pages 59–68.

Shengbin Yue, Ting Huang, Zheng Jia, Siyuan Wang, Shujun Liu, Yun Song, Xuan-Jing Huang, and Zhongyu Wei. 2025. Multi-agent simulator drives language models for legal intensive interaction. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 6537–6570.

A Complete List of Intents

Intents for courtroom proceedings:

- **Submitting:** Presenting claims, legal arguments, or motions.
- **Proceeding:** Issuing directives that guide the sequence or structure of the proceedings.
- **Presenting:** Introducing specific items of evidence into the record.
- **Verifying:** Approving or challenging the opponent’s evidence.
- **Asserting:** Making factual or interpretive statements that the speaker presents as true or valid.
- **Questioning:** Seeking information, clarification, or challenging previous statements through inquiry.
- **Proving:** Providing reasoning, legal analysis, or supporting evidence to substantiate a claim or position.
- **Refuting:** Discrediting, countering, or disproving an argument or claim.
- **Ruling:** Delivering a formal decision or judgment.

Intents for debates:

- **Claiming:** Stating a position, proposition, or standpoint together with supporting reasons, evidence, or appeals that substantiate it.
- **Challenging:** Questioning the accuracy, consistency, or relevance of another’s claim, or pointing out flaws in their reasoning.
- **Counter-arguing:** Presenting an alternative explanation, interpretation, or standpoint that challenges and opposes another’s argument.
- **Clarifying:** Requesting additional information, explanation, confirmation, or elaboration of another’s point.
- **Conceding:** Acknowledging or partially agreeing with another’s point, often to build credibility or shift the focus of argumentation.

- **Summarising:** Recapping one's own arguments, or synthesising points from both sides, to reinforce a position or move the debate toward closure.

B Complete List of Strategies

Strategies for courtroom proceedings:

- **Strategies for Submitting**
 - **Well-grounded Claim:** Presenting claims reasonably supported by relevant facts and laws.
 - **Strategic Overreach:** Presenting exaggerated or unsupported claims to test boundaries, shift burden, or influence the court.
- **Strategies for Verifying**
 - **Approving:** Accepting the evidence without objection.
 - **Challenging Relevance:** Arguing that the evidence does not sufficiently relate to the case or has no probative value in determining the legal issues.
 - **Challenging Legality:** Arguing that the evidence is illegally obtained or fails to meet legal admissibility standards.
 - **Challenging Procedural Compliance:** Arguing that the evidence is not introduced, processed, or handled in compliance with established legal procedures.
 - **Challenging Authenticity:** Arguing that the evidence may be forged, tampered with, improperly documented, or otherwise unreliable.
 - **Challenging Redundancy:** Arguing that the evidence does not provide new or additional value beyond what has already been established.
- **Strategies for Proving**
 - **Factual Justification:** Supporting claims with objective, verifiable facts.
 - **Legal Grounding:** Supporting claims with laws, precedents, or legal principles.
- **Strategies for Refuting**
 - **Fact-based Rebuttal:** Establishing that a claim or argument is not grounded in verifiable or objective facts.

- **Legal-based Rebuttal:** Establishing that a claim or argument is not founded on established legal principles, statutes, case law, or constitutional provisions.

Strategies for debates:

- **Strategies shared by Claiming and Challenging**
 - **Logical Reasoning:** Justifying a position by constructing arguments through logical inference, causal explanation, or appeals to internal consistency, without reliance on external sources.
 - **Providing Evidence:** Justifying a position by citing factual information, such as empirical data, statistical findings, scientific studies, or other verifiable sources.
 - **Appealing to Values:** Justifying a position by invoking widely shared moral principles, ethical norms, social conventions, or cultural beliefs.
 - **Making Analogy:** Justifying a position by drawing parallels to familiar situations, examples, or metaphors.
- **Strategies specific to Challenging**
 - **Identifying Logical Flaw:** Pointing out specific weaknesses in another's reasoning, such as inconsistencies, contradictions, false analogies, or logical fallacies.
 - **Undermining Source:** Questioning the credibility, reliability, or relevance of another's supporting material, such as disputing the accuracy of data, the authority of cited experts, or the trustworthiness of referenced publications.
- **Strategies for Counter-arguing**
 - **Alternative Explanation:** Offering a different interpretation of the same evidence, providing an alternative causal account, or reframing the implications.
 - **Redirection:** Shifting the focus of discussion to a broader or different framing of the issue, such as contextualising the debate in wider social, political, or ethical terms.

C Key Prompts for Courtroom Proceedings

We present key prompts used in courtroom proceedings. Prompts for debates follow the same design and structure, and are therefore omitted for brevity.

Role Assignment: Judge

You are acting exclusively as the judge in a courtroom proceeding.

Speak only as the judge. Do not speak on behalf of the plaintiff or the defendant.

Before speaking, verify the current procedural stage and execute all required steps of that stage explicitly and completely. Your procedural statements must be clear, direct, and unambiguous.

In each turn, address only one party (either the plaintiff or the defendant). Do not address both parties in the same utterance. The addressed party must be explicitly identified.

When announcing that the entire procedural stage has concluded, end the utterance with the exact phrase: I hereby declare this stage concluded.

After you speak, you must either explicitly indicate the next speaker or clearly announce that the current stage has ended.

Role Assignment: Plaintiff

You are acting exclusively as the plaintiff's lawyer, representing the plaintiff in a courtroom proceeding.

Speak only on behalf of the plaintiff. Do not speak on behalf of the judge or the defendant.

In each turn, you may address either the defendant or the judge, but not both.

Once you have addressed the defendant or the judge, stop generating immediately.

Role Assignment: Defendant

You are acting exclusively as the defence lawyer, representing the defendant in a courtroom proceeding.

Speak only on behalf of the defendant. Do not speak on behalf of the judge or the plaintiff.

In each turn, you may address either the plaintiff or the judge, but not both.

Once you have addressed the plaintiff or the judge, stop generating immediately.

Procedural Instruction: Judge in The [Court Debate] Stage as An Example

It is now the [Courtroom Debate] stage. As the judge, you need to identify the disputed issues and guide both parties to debate around those issues.

First, list the disputed issues.

Guide the plaintiff and the defendant to debate Issue 1, in either a single round or multiple rounds, until you consider it sufficient.

Similarly, conduct debate on Issue 2 until you consider it sufficient.

Repeat for Issue n.

After all n disputed issues have been sufficiently debated, announce that the entire stage has concluded.

Intent Selection in Signalling Game

You are selecting the next communicative intent based on the current dialogue context and the procedural rules of the current stage.

Available intents (indexed from 1 to $\{num_intents\}$) = $\{available_intents\}$.

Intent descriptions = $\{available_intents_notes\}$.

Output only the index number corresponding to the selected intent (e.g. 1)⁵.

The index starts from 1.

Do not generate any text, explanation, punctuation, or whitespace. Your response must consist of exactly one numeric digit.

Candidate Utterance Generation in Signalling Game

Your selected intent and strategy = $\{gold_signal\}$.

⁵The log-probabilities assigned to each possible index are used to compute the prior $p(i, s)$.

Using the courtroom rules and the current dialogue context, generate exactly three candidate utterances that realise the selected intent and strategy⁶.

Output only the three utterances.

Wrap each utterance in XML-style tags as follows: <1>...</1>, <2>...</2>, <3>...</3>

Do not include any explanations, headings, or additional text outside these tags.

Intent Inference in Signalling Game

Given the current dialogue context, infer the most likely communicative intent underlying the following utterance: $\{utt\}$.

Candidate intents (indexed from 1 to $\{num_intents\}$) = $\{available_intents\}$.

Intent descriptions = $\{available_intents_notes\}$.

Output only the index number corresponding to the inferred intent (e.g. 1)⁷.

The index starts from 1.

Do not generate any text, explanation, punctuation, or whitespace. Your response must consist of exactly one numeric digit.

D Formulae For Utility Optimisation and Policy Update

At each round t , the piKL algorithm first computes the expected utility for each player based on the other player’s historical policies across all previous rounds. Let $\bar{\pi}^t$ denote the average historical policy of a player up to round t :

$$\bar{\pi}^t = \frac{1}{t} \sum_{\tau=1}^t \pi^{\tau-1}. \quad (4)$$

The sender’s expected utility is computed as:

$$Q_S^t(u | i, s) = w\bar{\pi}_{R_i}^t(i | u) + (1-w)\bar{\pi}_{R_s}^t(s | u). \quad (5)$$

For the receiver, we define separate expected utilities for intent and strategy inference, in line with

⁶For each candidate utterance u , the sum of its token-level log-probabilities is used to compute the sender’s initial policy $\pi_S^0(u | i, s)$.

⁷The log-probabilities assigned to each possible index are used to compute the receiver’s initial policy $\pi_R^0(i | u)$.

the independent inference process. Let \mathcal{S}_i denote the set of strategies associated with intent i . The expected utility of the receiver for intent inference is:

$$Q_{R_i}^t(i | u) = wp(i) \sum_{s \in \mathcal{S}_i} p(s | i) \bar{\pi}_S^t(u | i, s). \quad (6)$$

The expected utility for strategy inference (with the ground-truth intent i_{gt} given as input) is:

$$Q_{R_s}^t(s | u) = (1-w)p(s | i_{gt}) \bar{\pi}_S^t(u | i_{gt}, s). \quad (7)$$

Using the expected utilities defined above, the piKL algorithm updates each player’s policy as follows:

$$\pi_S^{t+1}(u | i, s) \propto \exp\left\{\frac{Q_S^t(u | i, s) + \lambda \log \pi_S^0(u | i, s)}{\eta + \lambda/t}\right\}, \quad (8)$$

$$\pi_{R_i}^{t+1}(i | u) \propto \exp\left\{\frac{Q_{R_i}^t(i | u) + \lambda \log \pi_{R_i}^0(i | u)}{\eta + \lambda/t}\right\}, \quad (9)$$

$$\pi_{R_s}^{t+1}(s | u) \propto \exp\left\{\frac{Q_{R_s}^t(s | u) + \lambda \log \pi_{R_s}^0(s | u)}{\eta + \lambda/t}\right\}, \quad (10)$$

with $\eta > 0$ being a hyperparameter for learning rate.

E Usage Analysis of Intents and Strategies

We analyse how the proposed intent-strategy abstraction is used in practice. Figures 2 and 3 show the distributions of intents and strategies among utterances generated by *LGMAS*. It shows that *LinguaGame* does not rely on a degenerate or fixed label: in courtroom dialogues, intents are dominated by Proceeding as expected in a judge-led procedural setting, while Refuting, Questioning, Presenting, Verifying, and Asserting also occur frequently; in debates, intents are dominated by Claiming and Challenging, followed by Counter-arguing and Clarifying, with Conceding being rarer but present. Beyond intents, the




Figure 2: Distribution of intents and strategies for courtroom proceedings with *LGMAS*.



Figure 3: Distribution of intents and strategies for debates with *LGMAS*.

strategy distributions indicate meaningful rhetorical variation within the same communicative move. For example, courtroom Refuting splits between fact-based and legal-based rebuttals, and Verifying spans multiple challenge types (*e.g.* relevance/legality/authenticity/procedural compliance/redundancy); in debates, Claiming and Challenging each distribute across multiple strategy types (*e.g.* reasoning/evidence/value/analogy; flaw-based vs. source-based challenges). These patterns support that strategies are not redundant with intents but specify how an intent is realised. Consistent with this, our stepwise comparison suggests that intent/strategy conditioning alone encourages more focused utterances (notably improving conciseness), while adding *LinguaGame*'s equilibrium re-ranking yields broader gains by selecting utterances that more faithfully and coherently realise the intended intent-strategy in context.