

EXPLORING COLLABORATION MECHANISMS FOR LLM AGENTS: A SOCIAL PSYCHOLOGY VIEW

Jintian Zhang^{♣*}, Xin Xu^{♣*}, Ningyu Zhang^{♣†}, Ruibo Liu[♡], Bryan Hooi[♣], Shumin Deng^{♣†}

[♣]Zhejiang University [♡]Google DeepMind

[♣]National University of Singapore, NUS-NCS Joint Lab

{zhangjintian, xxucs, zhangningyu, 231sm}@zju.edu.cn

ruiboliu@google.com, {dcsbhk, shumin}@nus.edu.sg

 <https://zjunlp.github.io/project/MachineSoM>

ABSTRACT

As Natural Language Processing (NLP) systems are increasingly employed in intricate social environments, a pressing query emerges: *Can these NLP systems mirror human-esque collaborative intelligence, in a multi-agent society consisting of multiple large language models (LLMs)?* This paper probes the collaboration mechanisms among contemporary NLP systems by melding practical experiments with theoretical insights. We fabricate four unique ‘societies’ comprised of LLM agents, where each agent is characterized by a specific ‘trait’ (easy-going or over-confident) and engages in collaboration with a distinct ‘thinking pattern’ (debate or reflection). Through evaluating these multi-agent societies on three benchmark datasets, we discern that certain collaborative strategies not only outshine previous top-tier approaches but also optimize efficiency (using fewer API tokens). Moreover, our results further illustrate that LLM agents manifest human-like social behaviors, such as conformity and consensus reaching, mirroring foundational social psychology theories. In conclusion, we integrate insights from social psychology to contextualize the collaboration of LLM agents, inspiring further investigations into the collaboration mechanism for LLMs. We commit to sharing our code and datasets¹, hoping to catalyze further research in this promising avenue.

1 INTRODUCTION

With the prevalence of LLMs (Zhao et al., 2023; Yin et al., 2023; Zhu et al., 2023) integral to daily social collaboration, there is a growing imperative to cultivate AI systems embodied with social intelligence. This also resonates with the Society of Mind (SoM) concept (Li et al., 2023a; Zhuge et al., 2023; Wang et al., 2023), which suggests that intelligence emerges when computational modules interact with each other, achieving collective objectives that surpass the capabilities of individual modules (Minsky, 1988; Singh, 2003). Previous studies (Park et al., 2023; Du et al., 2023b; Liang et al., 2023; Shinn et al., 2023; Madaan et al., 2023; Hao et al., 2023; Liu et al., 2023a; Akata et al., 2023) have delved into strategies where LLM instances, termed agents (Wang et al., 2024c; Xi et al., 2023; Gao et al., 2023a; Cheng et al., 2024; Ma et al., 2024b), cooperate synergistically (e.g., debate and reflect) to accomplish tasks (Du et al., 2023a; Pezeshkpour et al., 2024; Guo et al., 2024; Du et al., 2024; Han et al., 2024). As illustrated in Figure 1, such collaboration fosters divergent thinking in LLMs, making them particularly effective for tasks demanding profound reflection.

Intuitively, reflecting on human societies (Siegal & Varley, 2002; Leslie et al., 2004; Sap et al., 2022; Shapira et al., 2023), where a myriad of individuals with distinct goals and roles coexist, the SoM framework champions harmonious interactions (Singh, 2003). Intriguingly, despite the fusion of social psychology (Tajfel, 1982; Tajfel & Turner, 2004; Johnson & Johnson, 2009) in SoM with human group dynamics (Woolley et al., 2010; Alderfer, 1987), which illuminates psychological

* Equal Contribution.

† Corresponding Author.

¹<https://github.com/zjunlp/MachineSoM>.

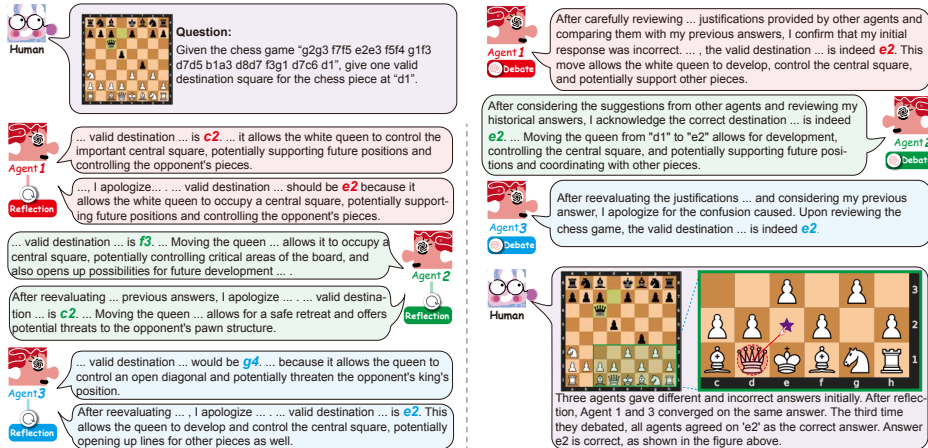


Figure 1: An example of the chess move validity task. Given previous chess game moves, agents are required to predict a valid next move for a specified piece.

patterns within social groups, its interpretation in the realm of LLMs is relatively uncharted (Ke et al., 2024). Besides, our grasp of how social behaviors influence LLMs is still in its nascent stages.

To address these issues, we delve into the machine society, probing the extent and ways that LLMs manifest social intelligence and collaboration capabilities (Mei et al., 2023). Utilizing powerful LLMs like GPT-3.5 (OpenAI, 2022), we build a test-bed across three datasets: MATH (Hendrycks et al., 2021b), MMLU (Hendrycks et al., 2021a) and Chess Move Validity (Srivastava et al., 2022). Our approach incorporates four **societies** characterized by two **individual traits** (*easy-going* and *overconfident*) with three agents: totally/mostly easy-going; totally/mostly overconfident. These traits are employed to emulate nuanced human society dynamics (Soni et al., 2024; Wang et al., 2024b;a; Li et al., 2023b; Kong et al., 2023).

Moreover, we delve into two distinct **thinking patterns** under multi-round collaboration: *debate* (Perelman, 1971; Sunstein, 2005; Amgoud & Prade, 2009; Du et al., 2023b; Liang et al., 2023) and *reflection* (Bogumil, 1985; Mezirow, 2003; Bolton, 2010). With the permutation of thinking patterns, we can constitute various **collaborative strategies**. To this end, we implement two patterns of collaboration in the collaborative strategies: (i) All agents adopt the same thinking pattern at each round; (ii) One agent adopts the different thinking patterns from others at each round. We then execute these multi-round collaborative strategies within different societies. Through our empirical analysis, we primarily discern the following insights (Further takeaways are in §3, §4 and Appendix A):

- (1) Collaborative strategies with various permutations of thinking patterns vary significantly in performance, and engaging in substantive debates enhances collaboration performance. Intriguingly, multi-agent societies composed of agents with different traits do not clearly differ in performance.
- (2) Employing uniform thinking patterns across all agents within a round of collaboration enhances efficiency. Besides, merely increasing the number of agents or the number of collaboration rounds does not consistently yield better outcomes. The balance between agent quantity and strategies emerges as a key determinant in collaboration.
- (3) LLM agents manifest behaviors reminiscent of human social tendencies, such as conformity (Allen & Levine, 1969; Cialdini & Goldstein, 2004) or the principle of majority rule in group thinking (Seal et al., 1998), which resonate with several fundamental theories in social psychology (Castro & Liskov, 1999; Tajfel & Turner, 2004).

Concretely, our findings challenge the dominant belief that mere scale is the key. We posit that small-group collaboration with rational strategies might present a more efficacious approach to utilizing LLMs. In wrapping up, we encapsulate the core contributions of this research as follows:

- We initiate an elaborate exploration into collaboration mechanisms in multi-agent society. Our goal is to identify how and to what extent LLMs manifest social intelligence through

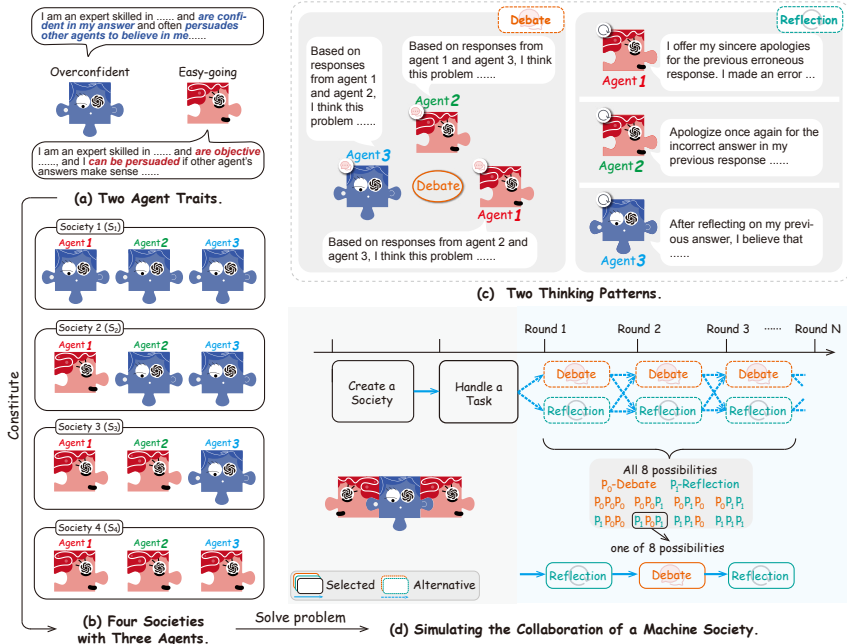


Figure 2: The overview of machine society simulation. Multiple agents with different traits make up diverse machine societies. These agents engage in debate or self-reflection across multiple rounds to complete tasks.

collaboration. To enrich our inquiry, we draw upon theories from social psychology, contextualizing the behaviors and tendencies displayed by LLM agents.

- Our research framework includes a meticulously crafted test-bed, integrating diverse multi-agent societies with agent individual traits, thinking patterns, and collaborative strategies, evaluated over three datasets. Notably, our empirical findings can inspire how to design a better multi-agent system through collaboration, beyond merely scaling up LLMs/Agents.
- Interestingly, our observations underscore a fascinating parallel: LLM agents mirror certain social behaviors typical of human collaboration. It could further emphasize the potential of human-AI interaction. Generally, fostering effective and efficient collaborative strategies for multi-agent systems could be the key to more socially-aware AI.

2 EXPLORE COLLABORATION MECHANISMS WITH MULTIPLE LLM AGENTS

In this section, we formulate and simulate the collaboration mechanisms explored within the machine society, drawing upon relevant concepts. We also illustrate the society settings in Figure 2.

2.1 PRELIMINARY CONCEPTS IN COLLABORATION

Individual Trait. Inspired by intelligence emerging from the collective efforts of numerous smaller and relatively simple agents (Minsky, 1988), each characterized by diverse traits, we set two types of agents exhibiting typically contrasting traits: *easy-going* and *overconfident*, as shown in Figure 2(a). Easy-going agents keep things in perspective, adapt well to different situations, and are compatible with various types of agents (Friedman & Schustack, 1999), which results in a harmonious societal structure with democracy (Mutz, 2006; Held, 2006). Conversely, overconfident agents tend to overestimate their competence, ignore potential risks, and resist others’ opinions (Moore & Healy, 2008).

Thinking Pattern. Considering the SoM concept (Minsky, 1988) states that intelligence emerges when specialized individuals within a society cooperate through thinking, we aim to study what thinking patterns are most successful in producing such emerging intelligence. Thus we explore two thinking patterns: *debate* (Sunstein, 2005; Du et al., 2023b; Liang et al., 2023) and *reflection*

(Bogumil, 1985; Bolton, 2010; Shinn et al., 2023), as illustrated in Figure 2(c). (i) In the *debate* pattern, several agents propose ideas, exchange responses, engage in collective argumentation, and ultimately reach a consensus. This fosters knowledge sharing, facilitates learning, and promotes adaptation among all agents within the society (Weiß, 1995; Stone & Veloso, 2000; Vidal, 2006; Wooldridge, 2009). (ii) In the *reflection* pattern, agents review their prior responses, extract lessons from their experiences, and refine their answers accordingly. These two patterns can unfold over several rounds.

Collaborative Strategy. Through both critical reflection and active participation in debate, agents are poised to challenge their existing assumptions, acquire fresh perspectives, and ultimately refine their viewpoints. Employing a collaboration mechanism built on these two thinking patterns can foster more insightful decision-making (Wooldridge, 2009; Amgoud & Prade, 2009) and improve reasoning outcomes (Mezirow, 2018). In societal settings, agents typically engage in multiple rounds of collaboration for problem-solving. In this paper, we characterize the collaborative strategy as **a permutation of thinking patterns** throughout multi-round collaboration, as illustrated in Figure 2(d) and further elaborated in §2.2.

2.2 SOCIETY SIMULATION

Symbols	Definition
\mathcal{T}	Set of agent traits
t_o	Trait 🗨️: overconfident
t_e	Trait 😊: easy-going
\mathcal{A}	Set of agent instances
a_i	The i -th agent
\mathcal{P}	Set of thinking patterns
p_0	🗨️ Debate
p_1	😞 Reflection
\mathcal{S}	Set of societies
S_i	The i -th society

Table 1: The description of the symbols.

We simulate the multi-agent collaborative society, as detailed with symbols shown in Table 1. Specifically, we construct a machine society consisting of n LLM agents, denoted as $\mathcal{A} = \{a_i\}_{i=1}^n$. This society contains two distinct agent traits: $\mathcal{T} = \{t_o, t_e\}$, where t_o and t_e respectively denotes the overconfident and easy-going trait. For each agent, at any round of collaboration, there are two thinking patterns to choose from, symbolized as $\mathcal{P} = \{p_0, p_1\}$, where p_0 and p_1 corresponds to *debate* and *reflection* respectively. By endowing agents \mathcal{A} with the traits of \mathcal{T} , we can emulate various machine societies. In our primary study (§3), we establish four distinct societies, $\mathcal{S} = \{S_1, S_2, S_3, S_4\}$, each consisting of three agents: $\{a_1, a_2, a_3\}$. The societies are constructed based on the combination of three agents with distinct traits, as illustrated in Figure 2(b):

$$\begin{aligned}
 S_1 &= \{(a_1 \leftarrow t_o), (a_2 \leftarrow t_o), (a_3 \leftarrow t_o)\} \text{ (totally overconfident)} \\
 S_2 &= \{(a_1 \leftarrow t_o), (a_2 \leftarrow t_o), (a_3 \leftarrow t_e)\} \text{ (mostly overconfident)} \\
 S_3 &= \{(a_1 \leftarrow t_o), (a_2 \leftarrow t_e), (a_3 \leftarrow t_e)\} \text{ (mostly easy-going)} \\
 S_4 &= \{(a_1 \leftarrow t_e), (a_2 \leftarrow t_e), (a_3 \leftarrow t_e)\} \text{ (totally easy-going)}
 \end{aligned}$$

where $(a_i \leftarrow t_j)$ denotes that the agent a_i possesses the trait t_j . If there is an even number of agents, we can also constitute a society with half overconfident and half easy-going agents. In our simulation, all agents consistently employ the same thinking pattern at each round of collaboration, similar to Du et al. (2023b). It gives rise to eight possible 3-round collaborative strategies:

$$\begin{aligned}
 &p_0p_0p_0, p_0p_0p_1, p_0p_1p_0, p_0p_1p_1, \\
 &p_1p_0p_0, p_1p_0p_1, p_1p_1p_0, p_1p_1p_1
 \end{aligned}$$

In our subsequent analysis (§3.2), we delve into more intricate scenarios, introducing a larger number of agents, increased collaboration rounds, and a broader range of collaborative strategies.

2.3 EXPERIMENTAL SETTINGS

Datasets. We conduct a rigorous evaluation of the reasoning and decision-making capabilities of various machine societies across three distinct tasks, utilizing diverse collaborative strategies:

- *High School Multiple-Choice.* Leveraging the **MMLU** (Hendrycks et al., 2021a) dataset, where problems span high school subjects such as statistics, mathematics, computer science, biology, chemistry, and physics, agents are required to identify the correct answer among four multiple-choice options. Our evaluation set consists of 50 randomly selected questions from this dataset.
- *Math.* Drawing from **MATH** dataset (Hendrycks et al., 2021b), a repository of math problems sourced from competitive events and expressed in LaTeX, we assess the model proficiency in advanced mathematical and scientific reasoning. The dataset segments these problems into five graded difficulty levels, and for our evaluation, we have randomly chosen 50 cases from Level 3 to 5.
- *Chess Move Validity.* Utilizing the dataset from the chess state tracking task² within the comprehensive **BIG-Bench Benchmark** (Srivastava et al., 2022), a sequence of chess moves denoted in UCI notation³ is provided. Agents are required to predict a legitimate subsequent move for a specified chess piece.

Setups. We craft specific instructions for each task, trait, and strategy, which can be referred to Table 5 at Appendix D.3. To enhance result reliability, we present average accuracy (**Acc**) and their respective standard deviations across five trials. Notably, our experiments exhibit substantial standard deviations. Hence, we introduce **WIN-TIE (W-T)** metric, indicating the frequency (over five trials) where the accuracy either matches or surpasses the continuous debate baseline (Du et al., 2023b). Meanwhile, we gauge the average token costs (**Cost**) consumed by the agents across societies, shedding light on the efficacy of the different collaborative strategies employed. For these evaluations, ChatGPT serves as the LLM agent accessible through the OpenAI API `gpt-3.5-turbo-1106`⁴. Further comprehensive details on data sampling and result evaluation are introduced in Appendix D.

3 ANALYSIS OF MACHINE SOCIAL COLLABORATION

Our experiments are primarily driven by the following research queries: **(RQ1)** How does problem-solving effectiveness vary under different collaborative strategies across diverse societies? **(RQ2)** How to configure the machine society variables for optimal performance? **(RQ3)** How does machine social collaboration mimic the human society?

3.1 MAIN RESULTS WITH QUANTITATIVE ANALYSIS

To address **RQ1**, we present the performance of four distinct societies in Table 2, each employing one of eight possible collaborative strategies, evaluated across three datasets with ChatGPT. To make the experimental findings more general, we evaluate on other LLMs, shown in Appendix H. Our experiments yield several pivotal observations:

(1) Societies do not clearly differ in performance but differ significantly in their tendency to reach a consensus. As observed in Table 2, among different 3-agent societies $S_1 \sim S_4$ employing the same collaborative strategy (a *vertical comparison* on Acc), the variations in accuracy are not pronounced. We also conduct a significance test of societies using ChatGPT in Appendix E, and other LLMs in Appendix H, further demonstrating insignificant differences between the societies. Thus we conclude that distinct societies composed of 3 agents possessing varied traits play an indistinctive role in shaping performance. We infer that this is due to LLM alignment (Ouyang et al., 2022), inhibiting agents from displaying extreme overconfidence, which contradicts human

²https://github.com/google/BIG-bench/blob/main/bigbench/benchmark_tasks/chess_state_tracking/synthetic_short/task.json.

³https://en.wikipedia.org/wiki/Universal_Chess_Interface.

⁴<https://platform.openai.com/docs/models/gpt-3-5>.

	Metric (Strategy)	Society	Collaborative Strategy								Metric (Society)	
			$p_0p_0p_0$	$p_0p_0p_1$	$p_0p_1p_0$	$p_0p_1p_1$	$p_1p_0p_0$	$p_1p_0p_1$	$p_1p_1p_0$	$p_1p_1p_1$	Cost ↓	W-T ↑
MMLU	Acc ↑	S_1	66.4±1.7	65.2±3.6	52.8±4.8	59.2±3.6	45.6±1.7	51.6±2.2	62.0±0.0	46.0±0.0	2970	2
		S_2	66.0±0.0	65.2±1.8	58.0±0.0	66.0±0.0	44.0±0.0	46.0±0.0	53.2±2.7	46.0±0.0	3081	9
		S_3	70.4±4.3	64.4±0.9	57.6±1.7	52.8±2.3	41.2±5.4	49.2±4.6	51.2±1.8	62.0±0.0	3172	1
		S_4	69.6±3.9	65.2±3.6	54.8±5.2	58.4±1.7	34.4±2.2	46.0±4.9	56.4±2.2	62.0±0.0	3090	2
	Cost ↓	All	4364	3510	3295	2665	3476	2651	2691	1976	-	-
W-T ↑	All	-	9	0	5	0	0	0	0	-	-	
MATH	Acc ↑	S_1	46.8±4.2	46.4±3.3	42.8±4.6	33.6±7.4	38.8±2.7	38.4±3.9	45.2±2.7	35.2±1.1	3417	8
		S_2	50.4±2.6	52.8±2.3	49.6±3.0	38.8±3.9	38.8±3.6	45.6±2.2	46.4±4.1	35.2±1.1	3623	8
		S_3	47.6±4.8	48.0±3.2	47.2±4.8	38.0±7.1	37.6±3.3	39.2±5.4	42.4±3.0	40.0±2.5	3757	8
		S_4	50.4±1.7	49.6±1.7	53.2±1.1	40.0±2.0	44.0±3.2	45.6±4.3	45.6±3.6	41.6±1.7	3658	10
	Cost ↓	All	4439	3965	3857	3414	3840	3234	3482	2681	-	-
W-T ↑	All	-	14	13	0	0	1	6	0	-	-	
Chess Move Validity	Acc ↑	S_1	54.4±1.7	52.0±0.0	52.0±5.1	51.6±5.2	54.4±1.7	51.2±1.8	50.4±1.7	52.0±0.0	2443	11
		S_2	48.0±0.0	49.2±1.1	46.0±0.0	54.0±0.0	50.0±0.0	52.0±0.0	42.0±2.5	52.0±0.0	2442	25
		S_3	48.4±1.7	48.0±2.8	54.8±5.0	45.2±3.4	48.4±2.6	44.8±3.4	50.4±1.7	53.6±0.9	2451	23
		S_4	51.6±4.6	44.0±2.5	54.4±3.0	53.6±5.5	45.6±2.2	48.0±2.0	43.6±0.9	52.0±0.0	2404	12
	Cost ↓	All	3046	2611	2604	2179	2705	2251	2252	1830	-	-
W-T ↑	All	-	10	12	10	11	9	5	14	-	-	

Table 2: The impact of 8 collaborative strategies on the performance of 3 datasets across distinct societies, using *ChatGPT*. Blue marks the best-performing strategy under the same society, light blue represents the second-best-performing strategy, and red indicates the worst-performing strategy. Cost / Cost measures the average tokens consumed by all cases under the same collaborative strategy / society. W-T / W-T tallies the total number of occurrences where performance exceeds the strategy $p_0p_0p_0$ under the same collaborative strategy / society. The significances test on societies and strategies are respectively shown in Table 6, 7 at Appendix E. The experiments of comparison with the single LLM agent is shown in Figure 21(a)-(f) at Appendix G.2.

alignment (Liu et al., 2022). Sharma et al. (2023) also demonstrate that LLMs tend to show sycophancy, as illustrated in Figure 11, 12. Furthermore, we increase the number of agents (2 to 10), accordingly resulting in more diverse societies, as seen in Figure 14, indicating that the impact of societies on performance remains indistinctive. We further analyze consensus reaching, *i.e.*, agents reach a consistent answer (Chen et al., 2023b), shown in Figure 16 at Appendix E, and find that more diverse societies (5 types of societies, with 2 to 10 agents) observably impact the average quantity of consensus clusters. Generally, a society totally comprising easy-going agents is more likely to reach a consensus.

(2) **Permutation of thinking patterns is crucial for collaboration, where debate-initial and debate-dominant strategies exhibit superiority.** For instance, on MMLU dataset, *debate-dominant* collaborative strategies, like $p_0p_0p_1$, $p_0p_1p_0$, and $p_1p_0p_0$, all containing two rounds of debate, display a pronounced outperformance (65.2 for $p_0p_0p_1$ in S_4 versus 34.4 for $p_1p_0p_0$ in S_4). As seen from Table 2, collaborative strategies starting with the thinking pattern of debate p_0 (*debate-initial*), such as $p_0p_0p_0$, $p_0p_0p_1$, $p_0p_1p_0$, and $p_0p_1p_1$, generally outperform others across all datasets. Furthermore, observed from the performance (i) under strategies with different (3~10) rounds of collaboration on ChatGPT, as shown in Figure 4 and Figure 18, 19 at Appendix F, *debate-initial/dominant* strategies are overall better; (ii) on LLaMA2 Chat 13B in Table 14 and Qwen 72B in Table 26, *debate-initial* strategies are generally superior; (iii) on LLaMA2 Chat 70B in Table 20 and Mixtral 8×7B in Table 32, *debate-dominant* strategies are superior. Observed from different 3-round collaborative strategies $p_i p_j p_k$ applied within the same society (a *horizontal comparison* on Acc), the variations in accuracy are notably pronounced. Besides, the significance test of different collaborative strategies using ChatGPT in Appendix E and other LLMs in Appendix H demonstrate that the order of thinking patterns significantly impacts the effectiveness.

(3) **Tasks behave better under collaborative strategies starting with continuous debate, and debate combined with continuous reflection is superior for difficult tasks.** Seen from Table 2, when comparing the best performance (marked in blue) and the worst (marked in red) within the same societies, the difference in results for Chess Move Validity is slight. This stands in sharp contrast to MMLU and MATH, which suggests that *the effectiveness of collaborative strategies depends on the task*. We then illustrate the performance under different collaborative strategies in view of task domains and difficulty in Figure 13 at Appendix E; on other LLMs in Figure 24, 33, 42, 56 at Appendix H. Figure 13(a) exhibits task-specific impacts and Figure 13(b),(c) reflects domain-dependent impacts under different collaborative strategies, where $p_0p_0p_0$ and $p_0p_0p_1$ starting with

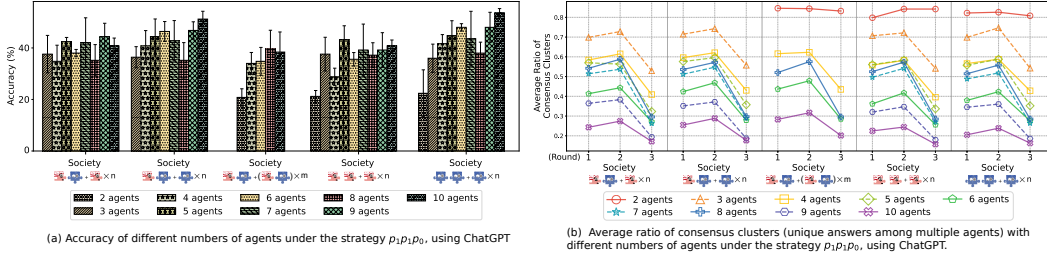


Figure 3: Accuracy and consensus reaching with different numbers (2~10) of agents under the strategy $p_1p_1p_0$ on *Chess Move Validity*, using *ChatGPT*. The significance test on agent numbers and comprehensive results under other strategies are shown in Table 11 and Figure 15, 17 at Appendix F due to space limits.

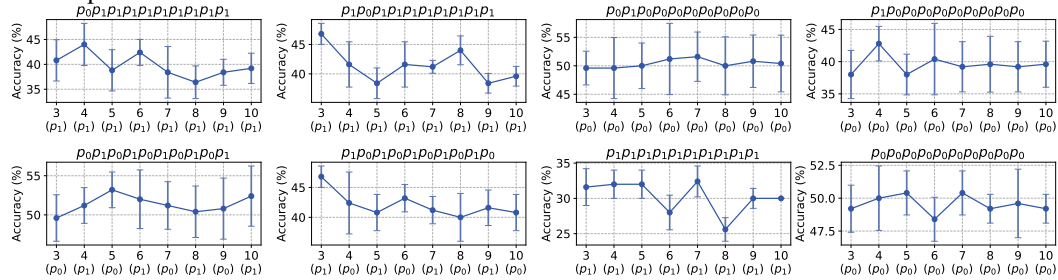


Figure 4: Accuracy under different (3~10) rounds of collaboration within 3-agent society S_2 (1 easy-going and 2 overconfident agents) on *MATH*, using *ChatGPT*. The significance test on rounds and experiments on *MMLU* and *Chess Move Validity* are shown in Table 12 and Figure 18, 19 at Appendix F due to space limits.

continuous debate are generally superior. For the mathematics domain seen from Figure 13(d), like *MMLU* mathematics and *MATH* level 3 & 4, the performance variations under different strategies are relatively small, but for the more difficult task, *i.e.*, *MATH* level 5, the strategies containing debate and continuous reflection (*i.e.*, $p_0p_1p_1$, $p_1p_1p_0$) behave superiorly. These nuanced disparities imply that *the marginal benefits derived from collaborative strategies may be task-dependent and difficulty-sensitive*.

3.2 IMPACT OF MACHINE SOCIETY SETTINGS

To address **RQ2**, we delve deeper into the variables influencing multi-agent society collaboration, exploring the intricacies of agent composition, collaboration rounds, and collaborative strategies.

Different Numbers of Agents. To evaluate the impact of different numbers of agents, we analyze performance within societies comprising 2~10 agents, presented in Figure 3(a). Different numbers of agents would constitute five types of societies, where the agents’ traits could be: *totally/mostly easy-going/overconfident*; *half easy-going/overconfident*. We observe that odd numbers of agents generally outperform others within all types of societies, and the possible reason is that odd-number agents can avoid ties. Besides, we also find that the variations of accuracy among odd-number agents are indistinctive. Thus we conclude that **the optimal number of agents is 3, considering both performance and efficiency**. We also implement a significance test of the number of agents shown in Table 11 at Appendix F, demonstrating that different numbers of agents significantly impact performance. Besides, we illustrate consensus reaching with different numbers of agents in Figure 3(b), demonstrating that **more agents are more likely to reach a consensus**.

Different Rounds. We then delve into the effects of different numbers of collaboration rounds, and further scale up the rounds of collaboration, presenting the performance under 3 to 10 rounds in Figure 4. Despite some fluctuation in performance from 3 to 10 rounds of collaboration, the variations are not extremely remarkable. Considering both accuracy and cost, we infer that **3-round collaboration is relatively effective and efficient**. We also conduct a significance test on different rounds of collaborative strategies, shown in Table 12 at Appendix F, and observe that the impact of rounds significantly relies on the collaborative strategy employed. Generally, **the strategies starting or dominating with reflection p_1 differ clearly in performance under different rounds**.

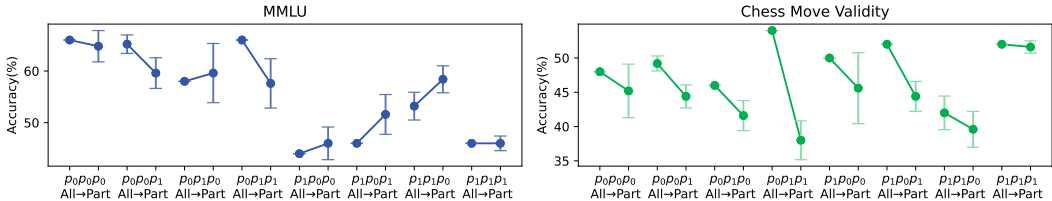


Figure 5: The effect on accuracy of whether all agents in a society execute the same thinking pattern in one round, using *ChatGPT*. “All” and “Part” respectively refer to all agents applying the same and different thinking pattern(s) in one round. Results on MATH and the significance test are shown in Figure 20 and Table 13 at Appendix F.

Other Collaborative Strategies. Venturing into scenarios with more intricate collaboration, we allow agents to adopt varied thinking patterns in each round of collaboration. For example, given three agents, in a specific round of collaboration, two agents engage in debate while the other one engages in reflection. To increase diversity, we perform a random allocation of thinking patterns to agents in each round, steering clear of scenarios where all agents adopt the same thinking pattern. Intriguingly, as shown in Figure 5, the presence of inconsistent thinking patterns within a society tends to negatively impact performance. Given the observation, we claim that **maintaining a consistent thinking pattern for all agents in a particular round would maximize collaborative efficacy**.

4 PHENOMENA OF CONFORMITY AND CONSENSUS REACHING

To address **RQ3**, we embark on further analysis from a social psychology view (Tajfel, 1982; Tajfel & Turner, 2004; Johnson & Johnson, 2009), to discern alignment between machine society collaboration and human societal dynamics (Woolley et al., 2010). Our findings indicate that machine society collaboration echoes specific human societal phenomena or theories, such as **conformity** (Cialdini & Goldstein, 2004; Allen & Levine, 1969; Coultas & van Leeuwen, 2015) and **consensus reaching** (Scheff, 1967; Degroot, 1974; Baronchelli, 2018) (more analysis are in Appendix G.1). We also analyze **group dynamics** (Cartwright & Zander, 1968; Alderfer, 1987; Forsyth, 2014; Bion, 2018; Forsyth, 2018) in multi-agent collaboration at Appendix G.2 due to page limits.

We embark on a detailed analysis, to discern the conformity and consensus-reaching phenomena in collaboration. For instance, as depicted in Figure 8(a) at Appendix D.3, an agent initially responds correctly to a question. However, swayed by the misguided answers and explanations from the other two agents, eventually, the three agents conform to an incorrect answer. This phenomenon mirrors detriments in “groupthink” (Janis, 1972; Jehn, 1995), suggesting that members of tight-knit groups tend to value harmony and consensus over objective critique of divergent views, potentially leading to flawed decisions. Contrastingly, in another scenario illustrated in Figure 8(b) at Appendix D.3, all three agents converge on the right answer after engaging in a society-wide debate. This mirrors benefits in “groupthink” (Jehn, 1995) and “SoM” (Minsky, 1988; Singh, 2003), where a multitude of agents collaboratively yield intelligence. Within such debates, agents furnish varied viewpoints and information. Through these exchanges, conflicts are resolved, ideas are honed, and the group gravitates toward an informed consensus (Fisher et al., 2011; Forsyth, 2018).

We also conduct a quantitative analysis of the prevalence of conformity and consensus-reaching phenomena. We analyze answer correctness changing at each round of collaboration in the situation of conformity, shown in Figure 6 on ChatGPT and Figure 28, 37, 51, 65 on other LLMs at Appendix H. We also present the ratio of consensus reaching at each round in Figure 7 on ChatGPT and Figure 29, 38, 52, 66 on other LLMs at Appendix H. We summarize the following observations:

- **Conformity is widespread**, and the proportion of conformity increases with the round increases in general.
- Overall, considering performance improvement, **conformity is beneficial in on ChatGPT, Qwen 72B; and harmful on LLaMA2 Chat 13B/70B, Mixtral 8×7B.**
- **As the number of rounds increases, benefits of conformity will weaken** (the ratio difference between True and False answers becomes smaller); and **detriments of conformity enhance** (the ratio difference between False and True answers becomes larger).

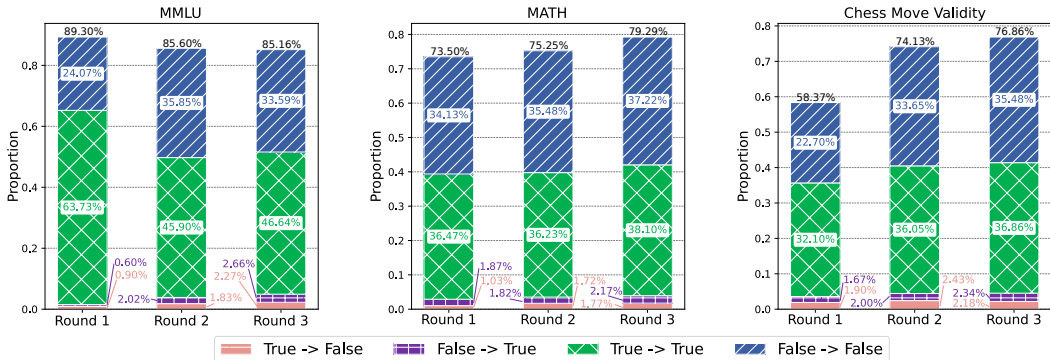


Figure 6: Variation of answer correctness in the situation of conformity, under 3-round collaboration, on ChatGPT, where conformity brings about benefits: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) > \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$; conformity brings about detriments: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) < \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$.

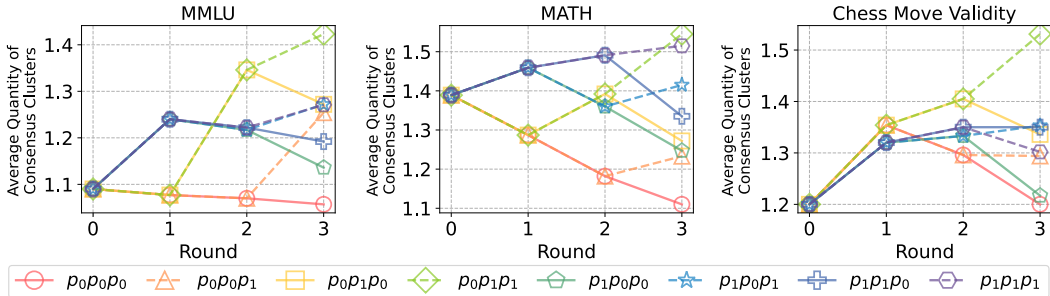


Figure 7: Average quantity of consensus clusters (i.e., unique answers among multiple agents) under different rounds of collaboration with 3-round collaborative strategies, using ChatGPT. Smaller quantity of consensus clusters, more easier it is to reach a consensus. Round 0 is equal to self-consistency. More details are in Appendix G.1.

- Generally, **reflection results in** increasing the quantity of consensus clusters, demonstrating **more difficulty to reach a consensus**; on the contrary, **debate is more possible for consensus reaching**.

5 CONCLUSION AND FUTURE WORK

This study has highlighted the potential of collaboration mechanisms with LLMs. Our findings reveal the impressive collaboration capabilities of LLM agents, with different individual traits, thinking patterns, and collaborative strategies. The emergence of human-like behaviors in agents, resonating with social psychology theories, further emphasizes the potential of human-AI interaction.

Moving forward, the collaboration mechanisms of machine society with multiple agents present a promising research avenue. A deeper exploration into the multi-agent society is warranted, with a focus on refining collaboration behaviors. Additionally, as LLMs evolve, understanding how different architectures influence these behaviors will be crucial. Integrating further insights from social psychology could also guide the development of more socially aware NLP systems.

LIMITATIONS

Although we explored various societies and collaborative strategies, our study still has its limitations. Firstly, limited by expense, we don't explore the impact of multiple agents respectively based on different LLMs, which may lead to more interesting findings at the social level due to the usage of differently distributed pre-trained data and strategies aligned with human intentions. Furthermore, we traversed all possible scenarios by search alone, lacking a way to let the agents adaptively make

autonomous decisions on collaborative strategies in specific scenarios. Although *debate* can be as close as possible to the upper limit, this approach entails a larger consumption and there exist some strategies that can achieve better performance with less overhead. Additionally, our experimental setup is relatively straightforward, as we have not considered more intricate configurations, such as a broader range of traits or a larger-scale society. Finally, we evaluate performance through manual validation and rule-based matching, which also limits the ability to validate more realistic and creative tasks, such as literary creation.

REPRODUCIBILITY STATEMENT

All code and data can be found in the GitHub repository⁵. For specific experimental implementation details, please refer to Appendix D.

ETHICS STATEMENT

This research was conducted in line with the highest ethical standards and best practices in research. The data employed were extracted from publicly accessible datasets, ensuring no usage of proprietary or confidential information. Consequently, this research is free from any ethical concerns.

ACKNOWLEDGMENTS

We would like to express gratitude to the anonymous reviewers for their kind and helpful comments. We extend our sincere gratitude to Min-Yen Kan and team members from NUS Web IR / NLP Group advised by Min-Yen Kan; Tao Gui and team members from FudanNLP Group; and Diyi Yang from Stanford University for providing insightful and constructive feedback on this paper. This work was supported by the National Natural Science Foundation of China (No. 62206246), the Fundamental Research Funds for the Central Universities (226-2023-00138), Zhejiang Provincial Natural Science Foundation of China (No. LGG22F030011), Yongjiang Talent Introduction Programme (2021A-156-G), CCF-Baidu Open Fund, and Information Technology Center and State Key Lab of CAD&CG, Zhejiang University, and NUS-NCS Joint Laboratory (A-0008542-00-00).

REFERENCES

- Elif Akata, Lion Schulz, Julian Coda-Forno, Seong Joon Oh, Matthias Bethge, and Eric Schulz. Playing repeated games with large language models. *CoRR*, abs/2305.16867, 2023. doi: 10.48550/arXiv.2305.16867. URL <https://doi.org/10.48550/arXiv.2305.16867>.
- Lize Alberts, Geoff Keeling, and Amanda McCroskery. What makes for a ‘good’ social actor? using respect as a lens to evaluate interactions with language agents. *CoRR*, abs/2401.09082, 2024. URL <https://arxiv.org/abs/2401.09082>.
- Clayton P Alderfer. An intergroup perspective on group dynamics. *Handbook of organizational behavior*, 190:222, 1987. URL <https://apps.dtic.mil/sti/citations/ADA135582>.
- Vernon L. Allen and John M. Levine. Consensus and conformity. *Journal of Experimental Social Psychology*, 5(4):389–399, 1969. ISSN 0022-1031. doi: [https://doi.org/10.1016/0022-1031\(69\)90032-8](https://doi.org/10.1016/0022-1031(69)90032-8). URL <https://www.sciencedirect.com/science/article/pii/0022103169900328>.
- Leila Amgoud and Henri Prade. Using arguments for making and explaining decisions. *Artif. Intell.*, 173(3-4):413–436, 2009. doi: 10.1016/j.artint.2008.11.006. URL <https://doi.org/10.1016/j.artint.2008.11.006>.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng

⁵<https://github.com/zjunlp/MachineSoM>.

- Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *CoRR*, abs/2309.16609, 2023. doi: 10.48550/ARXIV.2309.16609. URL <https://doi.org/10.48550/arXiv.2309.16609>.
- Andrea Baronchelli. The emergence of consensus: a primer. *Royal Society open science*, 5(2): 172189, 2018. doi: 10.1098/rsos.172189. URL <http://doi.org/10.1098/rsos.172189>.
- Marcel Binz and Eric Schulz. Using cognitive psychology to understand gpt-3. *Proceedings of the National Academy of Sciences*, 120(6):e2218523120, 2023. doi: 10.1073/pnas.2218523120. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2218523120>.
- Wilfred R Bion. Group dynamics: A re-view. *New directions in psychoanalysis*, pp. 440–477, 2018. URL <https://www.taylorfrancis.com/chapters/edit/10.4324/9780429477546-19/group-dynamics-re-view-bion>.
- R. J. Bogumil. The reflective practitioner: How professionals think in action. *Proc. IEEE*, 73(4):845–846, 1985. doi: 10.1109/PROC.1985.13210. URL <https://doi.org/10.1109/PROC.1985.13210>.
- Gillie Bolton. *Reflective practice: Writing and professional development*. Sage publications, 2010. URL <https://uk.sagepub.com/en-gb/eur/reflective-practice/book252252>.
- Dorwin Cartwright and Alvin Zander. Group dynamics, 1968. URL <https://psycnet.apa.org/record/1968-12031-000>.
- Miguel Castro and Barbara Liskov. Practical byzantine fault tolerance. In *OSDI*, pp. 173–186. USENIX Association, 1999. URL <https://dl.acm.org/citation.cfm?id=296824>.
- Guangyao Chen, Siwei Dong, Yu Shu, Ge Zhang, Jaward Sesay, Börje F. Karlsson, Jie Fu, and Yemin Shi. Autoagents: A framework for automatic agent generation. *CoRR*, abs/2309.17288, 2023a. doi: 10.48550/ARXIV.2309.17288. URL <https://doi.org/10.48550/arXiv.2309.17288>.
- Huaben Chen, Wenkang Ji, Lufeng Xu, and Shiyu Zhao. Multi-agent consensus seeking via large language models. *CoRR*, abs/2310.20151, 2023b. doi: 10.48550/ARXIV.2310.20151. URL <https://doi.org/10.48550/arXiv.2310.20151>.
- Justin Chih-Yao Chen, Swarnadeep Saha, and Mohit Bansal. Reconcile: Round-table conference improves reasoning via consensus among diverse llms. *arxiv preprint*, 2309.07864, 2023c. URL <https://arxiv.org/abs/2309.13007>.
- Po-Lin Chen and Cheng-Shang Chang. Interact: Exploring the potentials of chatgpt as a cooperative agent. *CoRR*, abs/2308.01552, 2023. doi: 10.48550/ARXIV.2308.01552. URL <https://doi.org/10.48550/arXiv.2308.01552>.
- Weize Chen, Yusheng Su, Jingwei Zuo, Cheng Yang, Chenfei Yuan, Chen Qian, Chi-Min Chan, Yujia Qin, Yaxi Lu, Ruobing Xie, Zhiyuan Liu, Maosong Sun, and Jie Zhou. Agentverse: Facilitating multi-agent collaboration and exploring emergent behaviors in agents. In *ICLR*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=EHg5GDnyq1>.
- Yuheng Cheng, Ceyao Zhang, Zhengwen Zhang, Xiangrui Meng, Sirui Hong, Wenhao Li, Zihao Wang, Zekai Wang, Feng Yin, Junhua Zhao, and Xiuqiang He. Exploring large language model based intelligent agents: Definitions, methods, and prospects. *CoRR*, abs/2401.03428, 2024. URL <https://arxiv.org/abs/2401.03428>.
- Yun-Shiuan Chuang and Timothy T. Rogers. Computational agent-based models in opinion dynamics: A survey on social simulations and empirical studies. *CoRR*, abs/2306.03446, 2023. doi: 10.48550/ARXIV.2306.03446. URL <https://doi.org/10.48550/arXiv.2306.03446>.
- Yun-Shiuan Chuang, Agam Goyal, Nikunj Harlalka, Siddharth Suresh, Robert Hawkins, Sijia Yang, Dhavan Shah, Junjie Hu, and Timothy T. Rogers. Simulating opinion dynamics with networks of llm-based agents. *CoRR*, abs/2311.09618, 2023. doi: 10.48550/ARXIV.2311.09618. URL <https://doi.org/10.48550/arXiv.2311.09618>.

- Robert B. Cialdini and Noah J. Goldstein. Social influence: Compliance and conformity. *Annual Review of Psychology*, 55(1):591–621, 2004. doi: 10.1146/annurev.psych.55.090902.142015. URL <https://doi.org/10.1146/annurev.psych.55.090902.142015>. PMID: 14744228.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint*, abs/2110.14168, 2021. URL <https://arxiv.org/abs/2110.14168>.
- Julie C Coultas and Edwin JC van Leeuwen. Conformity: Definitions, types, and evolutionary grounding. *Evolutionary perspectives on social psychology*, pp. 189–202, 2015. URL https://link.springer.com/chapter/10.1007/978-3-319-12697-5_15.
- Maxwell Crouse, Ibrahim Abdelaziz, Kinjal Basu, Soham Dan, Sadhana Kumaravel, Achille Fokoue, Pavan Kapanipathi, and Luis A. Lastras. Formally specifying the high-level behavior of llm-based agents. *CoRR*, abs/2310.08535, 2023. doi: 10.48550/ARXIV.2310.08535. URL <https://doi.org/10.48550/arXiv.2310.08535>.
- Morris H. Degroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345): 118–121, 1974. doi: 10.1080/01621459.1974.10480137. URL <https://www.tandfonline.com/doi/abs/10.1080/01621459.1974.10480137>.
- Dorottya Demszky, Diyi Yang, David S. Yeager, Christopher J. Bryan, Margaret Clapper, Susannah Chandhok, Johannes C. Eichstaedt, Cameron Hecht, Jeremy Jamieson, Meghann Johnson, Michaela Jones, Danielle Krettek-Cobb, Leslie Lai, Nirel JonesMitchell, Desmond C. Ong, Carol S. Dweck, James J. Gross, and James W. Pennebaker. Using large language models in psychology. *Nature Reviews Psychology*, 2(11):688–701, Nov 2023. ISSN 2731-0574. doi: 10.1038/s44159-023-00241-5. URL <https://doi.org/10.1038/s44159-023-00241-5>.
- Hung Du, Srikanth Thudumu, Rajesh Vasa, and Kon Mouzakis. A survey on context-aware multi-agent systems: Techniques, challenges and future directions. *CoRR*, abs/2402.01968, 2024. URL <https://arxiv.org/abs/2402.01968>.
- Yali Du, Joel Z. Leibo, Usman Islam, Richard Willis, and Peter Sunehag. A review of cooperation in multi-agent learning. *CoRR*, abs/2312.05162, 2023a. doi: 10.48550/ARXIV.2312.05162. URL <https://doi.org/10.48550/arXiv.2312.05162>.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. *CoRR*, abs/2305.14325, 2023b. doi: 10.48550/arXiv.2305.14325. URL <https://doi.org/10.48550/arXiv.2305.14325>.
- Edgar A. Duéñez-Guzmán, Suzanne Sadedin, Jane X. Wang, Kevin R. McKee, and Joel Z. Leibo. A social path to human-like artificial intelligence. *Nat. Mac. Intell.*, 5(11):1181–1188, 2023. doi: 10.1038/S42256-023-00754-X. URL <https://doi.org/10.1038/s42256-023-00754-x>.
- Joshua M Epstein. *Generative social science: Studies in agent-based computational modeling*. Princeton University Press, 2012. URL <https://press.princeton.edu/books/ebook/9781400842872/generative-social-science>.
- Xueyang Feng, Zhi-Yuan Chen, Yujia Qin, Yankai Lin, Xu Chen, Zhiyuan Liu, and Ji-Rong Wen. Large language model-based human-agent collaboration for complex task solving. *CoRR*, abs/2402.12914, 2024. URL <https://arxiv.org/abs/2402.12914>.
- Roger Fisher, William L Ury, and Bruce Patton. *Getting to yes: Negotiating agreement without giving in*. Penguin, 2011. URL <https://www.pon.harvard.edu/shop/getting-to-yes-negotiating-agreement-without-giving-in/>.
- Andreas Flache, Michael Mäs, Thomas Feliciani, Edmund Chattoe-Brown, Guillaume Deffuant, Sylvie Huet, and Jan Lorenz. Models of social influence: Towards the next frontiers. *J. Artif. Soc. Soc. Simul.*, 20(4), 2017. doi: 10.18564/JASSS.3521. URL <https://doi.org/10.18564/jasss.3521>.
- Donelson R Forsyth. *Group dynamics*. Wadsworth Cengage Learning, 2014. URL <https://scholarship.richmond.edu/bookshelf/5/>.

- Donelson R Forsyth. *Group dynamics*. Cengage Learning, 2018. URL https://books.google.com.sg/books?id=v9EDwAAQBAJ&newbks=0&source=newbks_fb&redir_esc=y.
- Howard S Friedman and Miriam W Schustack. *Personality: Classic theories and modern research*. Allyn and Bacon Boston, MA, 1999. URL <https://books.google.com/books/about/Personality.html?id=ziTvDAAAQBAJ>.
- Andrew Fuchs, Andrea Passarella, and Marco Conti. Optimizing delegation between human and AI collaborative agents. *CoRR*, abs/2309.14718, 2023. doi: 10.48550/ARXIV.2309.14718. URL <https://doi.org/10.48550/arXiv.2309.14718>.
- Chen Gao, Xiaochong Lan, Nian Li, Yuan Yuan, Jingtao Ding, Zhilun Zhou, Fengli Xu, and Yong Li. Large language models empowered agent-based modeling and simulation: A survey and perspectives. *CoRR*, abs/2312.11970, 2023a. URL <https://arxiv.org/abs/2312.11970>.
- Chen Gao, Xiaochong Lan, Zhihong Lu, Jinzhu Mao, Jinghua Piao, Huandong Wang, Depeng Jin, and Yong Li. S³: Social-network simulation system with large language model-empowered agents. *CoRR*, abs/2307.14984, 2023b. doi: 10.48550/ARXIV.2307.14984. URL <https://doi.org/10.48550/arXiv.2307.14984>.
- Yiming Gao, Feiyu Liu, Liang Wang, Zhenjie Lian, Dehua Zheng, Weixuan Wang, Wenjin Yang, Siqin Li, Xianliang Wang, Wenhui Chen, Jing Dai, Qiang Fu, Wei Yang, Lanxiao Huang, and Wei Liu. Enhancing human experience in human-agent collaboration: A human-centered modeling approach based on positive human gain. In *ICLR*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=BqEvdOS1Hs>.
- Nigel Gilbert and Pietro Terna. How to build and use agent-based models in social science. *Mind & Society*, 1:57–72, 2000. URL <https://link.springer.com/article/10.1007/BF02512229>.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xiangliang Zhang. Large language model based multi-agents: A survey of progress and challenges. *CoRR*, abs/2402.01680, 2024. URL <https://arxiv.org/abs/2402.01680>.
- Thilo Hagendorff. Machine psychology: Investigating emergent capabilities and behavior in large language models using psychological methods. *CoRR*, abs/2303.13988, 2023. doi: 10.48550/ARXIV.2303.13988. URL <https://doi.org/10.48550/arXiv.2303.13988>.
- Shanshan Han, Qifan Zhang, Yuhang Yao, Weizhao Jin, Zhaozhuo Xu, and Chaoyang He. Llm multi-agent systems: Challenges and open problems. *CoRR*, abs/2402.03578, 2024. URL <https://arxiv.org/abs/2402.03578>.
- Rui Hao, Linmei Hu, Weijian Qi, Qingliu Wu, Yirui Zhang, and Liqiang Nie. Chatllm network: More brains, more intelligence. *CoRR*, abs/2304.12998, 2023. doi: 10.48550/arXiv.2304.12998. URL <https://doi.org/10.48550/arXiv.2304.12998>.
- Zhitao He, Pengfei Cao, Yubo Chen, Kang Liu, Ruopeng Li, Mengshu Sun, and Jun Zhao. LEGO: A multi-agent collaborative framework with role-playing and iterative feedback for causality explanation generation. In *EMNLP (Findings)*, pp. 9142–9163. Association for Computational Linguistics, 2023. URL <https://aclanthology.org/2023.findings-emnlp.613>.
- David Held. *Models of democracy*. Polity, 2006. URL <https://www.sup.org/books/title/?id=10597>.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *ICLR*. OpenReview.net, 2021a. URL <https://openreview.net/forum?id=d7KBjmI3GmQ>.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the MATH dataset. In *NeurIPS Datasets and Benchmarks*, 2021b. URL <https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/be83ab3ecd0db773eb2dc1b0a17836a1-Abstract-round2.html>.

- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. Metagpt: Meta programming for multi-agent collaborative framework. In *ICLR*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=VtmBAGCN7o>.
- Yaqing Hou, Mingyang Sun, Yifeng Zeng, Yew-Soon Ong, Yaochu Jin, Hongwei Ge, and Qiang Zhang. A multi-agent cooperative learning system with evolution of social roles. *IEEE Transactions on Evolutionary Computation*, pp. 1–1, 2023. doi: 10.1109/TEVC.2023.3268076. URL <https://ieeexplore.ieee.org/document/10104101>.
- Irving L Janis. *Victims of Groupthink: A psychological study of foreign-policy decisions and fiascoes*. Houghton Mifflin, 1972. URL <https://psycnet.apa.org/record/1975-29417-000>.
- Karen A Jehn. A multimethod examination of the benefits and detriments of intragroup conflict. *Administrative science quarterly*, pp. 256–282, 1995. doi: 10.2307/2393638. URL <https://www.jstor.org/stable/2393638>.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Yejin Bang, Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM Comput. Surv.*, 55(12):248:1–248:38, 2023. doi: 10.1145/3571730. URL <https://doi.org/10.1145/3571730>.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de Las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. Mistral 7b. *CoRR*, abs/2310.06825, 2023. doi: 10.48550/ARXIV.2310.06825. URL <https://doi.org/10.48550/arXiv.2310.06825>.
- Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L  lio Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th  ophile Gervet, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. Mixtral of experts. *CoRR*, abs/2401.04088, 2024. URL <https://arxiv.org/abs/2401.04088>.
- David W Johnson and Roger T Johnson. An educational psychology success story: Social interdependence theory and cooperative learning. *Educational researcher*, 38(5):365–379, 2009. URL <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=72585feb1200d53a81d4fb3e64862d69317b72c3>.
- Luoma Ke, Song Tong, Peng Chen, and Kaiping Peng. Exploring the frontiers of llms in psychological applications: A comprehensive review. *CoRR*, abs/2401.01519, 2024. URL <https://arxiv.org/abs/2401.01519>.
- Chuyi Kong, Yaxin Fan, Xiang Wan, Feng Jiang, and Benyou Wang. Large language model as a user simulator. *CoRR*, abs/2308.11534, 2023. doi: 10.48550/ARXIV.2308.11534. URL <https://doi.org/10.48550/arXiv.2308.11534>.
- Tatsuki Kuribayashi, Yohei Oseki, and Timothy Baldwin. Psychometric predictive power of large language models. *CoRR*, abs/2311.07484, 2023. doi: 10.48550/ARXIV.2311.07484. URL <https://doi.org/10.48550/arXiv.2311.07484>.
- Marc Lanctot, Kate Larson, Yoram Bachrach, Luke Marris, Zun Li, Avishkar Bhoopchand, Thomas W. Anthony, Brian Tanner, and Anna Koop. Evaluating agents using social choice theory. *CoRR*, abs/2312.03121, 2023. doi: 10.48550/ARXIV.2312.03121. URL <https://doi.org/10.48550/arXiv.2312.03121>.
- Alan M Leslie, Ori Friedman, and Tim P German. Core mechanisms in ‘theory of mind’. *Trends in cognitive sciences*, 8(12):528–533, 2004. ISSN 1364-6613. doi: <https://doi.org/10.1016/j.tics.2004.10.001>. URL <https://www.sciencedirect.com/science/article/pii/S1364661304002608>.

- Guohao Li, Hasan Abed Al Kader Hammoud, Hani Itani, Dmitrii Khizbullin, and Bernard Ghanem. CAMEL: communicative agents for "mind" exploration of large scale language model society. *CoRR*, abs/2303.17760, 2023a. doi: 10.48550/arXiv.2303.17760. URL <https://doi.org/10.48550/arXiv.2303.17760>.
- Yuan Li, Yixuan Zhang, and Lichao Sun. Metaagents: Simulating interactions of human behaviors for llm-based task-oriented coordination via collaborative generative agents. *CoRR*, abs/2310.06500, 2023b. doi: 10.48550/ARXIV.2310.06500. URL <https://doi.org/10.48550/arXiv.2310.06500>.
- Yujia Li, David H. Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustín Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. Competition-level code generation with alphacode. *CoRR*, abs/2203.07814, 2022. doi: 10.48550/ARXIV.2203.07814. URL <https://doi.org/10.48550/arXiv.2203.07814>.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Zhaopeng Tu, and Shuming Shi. Encouraging divergent thinking in large language models through multi-agent debate. *CoRR*, abs/2305.19118, 2023. doi: 10.48550/arXiv.2305.19118. URL <https://doi.org/10.48550/arXiv.2305.19118>.
- Yuanzhi Liang, Linchao Zhu, and Yi Yang. Anteval: Quantitatively evaluating informativeness and expressiveness of agent social interactions. *CoRR*, abs/2401.06509, 2024. URL <https://arxiv.org/abs/2401.06509>.
- Ruibo Liu, Ge Zhang, Xinyu Feng, and Soroush Vosoughi. Aligning generative language models with human values. In *NAACL-HLT (Findings)*, pp. 241–252. Association for Computational Linguistics, 2022. doi: 10.18653/v1/2022.FINDINGS-NAACL.18. URL <https://doi.org/10.18653/v1/2022.findings-naacl.18>.
- Ruibo Liu, Ruixin Yang, Chenyan Jia, Ge Zhang, Denny Zhou, Andrew M. Dai, Diyi Yang, and Soroush Vosoughi. Training socially aligned language models in simulated human society. *arXiv preprint*, abs/2305.16960, 2023a. doi: 10.48550/arXiv.2305.16960. URL <https://doi.org/10.48550/arXiv.2305.16960>.
- Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. Dynamic llm-agent network: An llm-agent collaboration framework with agent team optimization. *CoRR*, abs/2310.02170, 2023b. doi: 10.48550/ARXIV.2310.02170. URL <https://doi.org/10.48550/arXiv.2310.02170>.
- Jan Lorenz, Martin Neumann, and Tobias Schröder. Individual attitude change and societal dynamics: Computational experiments with psychological theories. *Psychological Review*, 128(4):623, 2021. URL <https://doi.org/10.1037/rev0000291>.
- Junliang Luo, Tianyu Li, Di Wu, Michael Jenkin, Steve Liu, and Gregory Dudek. Hallucination detection and hallucination mitigation: An investigation. *CoRR*, abs/2401.08358, 2024. URL <https://arxiv.org/abs/2401.08358>.
- Chang Ma, Junlei Zhang, Zhihao Zhu, Cheng Yang, Yujiu Yang, Yaohui Jin, Zhenzhong Lan, Lingpeng Kong, and Junxian He. Agentboard: An analytical evaluation board of multi-turn llm agents. *CoRR*, abs/2401.13178, 2024a. URL <https://arxiv.org/abs/2401.13178>.
- Qun Ma, Xiao Xue, Deyu Zhou, Xiangning Yu, Donghua Liu, Xuwen Zhang, Zihan Zhao, Yifan Shen, Peilin Ji, Juanjuan Li, Gang Wang, and Wanpeng Ma. Computational experiments meet large language model based agents: A survey and perspective. *CoRR*, abs/2402.00262, 2024b. URL <https://arxiv.org/abs/2402.00262>.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, Sean Welleck, Bodhisattwa Prasad Majumder, Shashank Gupta, Amir Yazdanbakhsh, and Peter Clark. Self-refine: Iterative refinement with self-feedback. *arXiv preprint*, abs/2303.17651, 2023. doi: 10.48550/arXiv.2303.17651. URL <https://doi.org/10.48550/arXiv.2303.17651>.

- Qiaozhu Mei, Yutong Xie, Walter Yuan, and Matthew O. Jackson. A turing test: Are AI chatbots behaviorally similar to humans? *CoRR*, abs/2312.00798, 2023. doi: 10.48550/ARXIV.2312.00798. URL <https://doi.org/10.48550/arXiv.2312.00798>.
- Jack Mezirow. How critical reflection triggers transformative learning. *Adult and Continuing Education: Teaching, learning and research*, 4:199, 2003. URL https://www.colorado.edu/plc/sites/default/files/attached-files/how_critical_reflection_triggers_transfo.pdf.
- Jack Mezirow. Transformative learning theory. In *Contemporary theories of learning*, pp. 114–128. Routledge, 2018. URL <https://www.wichita.edu/services/mrc/OIR/Pedagogy/Theories/transformative.php>.
- Marvin Minsky. *Society of mind*. Simon and Schuster, 1988. URL <https://www.simonandschuster.com/books/Society-Of-Mind/Marvin-Minsky/9780671657130>.
- Don A Moore and Paul J Healy. The trouble with overconfidence. *Psychological review*, 115(2):502, 2008. URL https://healy.econ.ohio-state.edu/papers/Moore_Healy-TroubleWithOverconfidence_WP.pdf.
- Diana C Mutz. *Hearing the other side: Deliberative versus participatory democracy*. Cambridge University Press, 2006. URL <https://www.cambridge.org/core/books/hearing-the-other-side/7CB061238546313D287668FF8EFE2EF7>.
- OpenAI. Chatgpt: Optimizing language models for dialogue, 2022. <https://openai.com/blog/chatgpt/>.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In *NeurIPS*, 2022. URL http://papers.nips.cc/paper_files/paper/2022/hash/b1efde53be364a73914f58805a001731-Abstract-Conference.html.
- Joon Sung Park, Joseph C. O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *UIST*, pp. 2:1–2:22. ACM, 2023. doi: 10.1145/3586183.3606763. URL <https://doi.org/10.1145/3586183.3606763>.
- Chaim Perelman. *The new rhetoric*. Springer, 1971. URL <https://link.springer.com/chapter/10.1007/978-94-010-1713-8.8>.
- Pouya Pezeshkpour, Eser Kandogan, Nikita Bhutani, Sajjadur Rahman, Tom Mitchell, and Estevam Hruschka. Reasoning capacity in multi-agent systems: Limitations, challenges and human-centered solutions. *CoRR*, abs/2402.01108, 2024. URL <https://arxiv.org/abs/2402.01108>.
- Shuofei Qiao, Ningyu Zhang, Runnan Fang, Yujie Luo, Wangchunshu Zhou, Yuchen Eleanor Jiang, Chengfei Lv, and Huajun Chen. Autoact: Automatic agent learning from scratch via self-planning. *CoRR*, abs/2401.05268, 2024. URL <https://arxiv.org/abs/2401.05268>.
- Vipula Rawte, Amit P. Sheth, and Amitava Das. A survey of hallucination in large foundation models. *CoRR*, abs/2309.05922, 2023. doi: 10.48550/arXiv.2309.05922. URL <https://doi.org/10.48550/arXiv.2309.05922>.
- Maarten Sap, Ronan Le Bras, Daniel Fried, and Yejin Choi. Neural theory-of-mind? on the limits of social intelligence in large lms. In *EMNLP*, pp. 3762–3780. Association for Computational Linguistics, 2022. doi: 10.18653/V1/2022.EMNLP-MAIN.248. URL <https://doi.org/10.18653/v1/2022.emnlp-main.248>.
- Thomas J Scheff. Toward a sociological model of consensus. *American Sociological Review*, pp. 32–46, 1967. doi: 10.2307/2091716. URL <https://doi.org/10.2307/2091716>.
- David Wyatt Seal, Laura M Bogart, and Anke A Ehrhardt. Small group dynamics: The utility of focus group discussions as a research method. *Group Dynamics: Theory, Research, and Practice*, 2(4):253, 1998. URL <https://doi.org/10.1037/1089-2699.2.4.253>.

- Murray Shanahan, Kyle McDonell, and Laria Reynolds. Role play with large language models. *Nat.*, 623(7987):493–498, 2023. doi: 10.1038/S41586-023-06647-8. URL <https://doi.org/10.1038/s41586-023-06647-8>.
- Natalie Shapira, Mosh Levy, Seyed Hossein Alavi, Xuhui Zhou, Yejin Choi, Yoav Goldberg, Maarten Sap, and Vered Shwartz. Clever hans or neural theory of mind? stress testing social reasoning in large language models. *arXiv preprint*, abs/2305.14763, 2023. doi: 10.48550/arXiv.2305.14763. URL <https://doi.org/10.48550/arXiv.2305.14763>.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askill, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. Towards understanding sycophancy in language models. *CoRR*, abs/2310.13548, 2023. doi: 10.48550/ARXIV.2310.13548. URL <https://doi.org/10.48550/arXiv.2310.13548>.
- Noah Shinn, Beck Labash, and Ashwin Gopinath. Reflexion: an autonomous agent with dynamic memory and self-reflection. *arXiv preprint*, abs/2303.11366, 2023. doi: 10.48550/arXiv.2303.11366. URL <https://doi.org/10.48550/arXiv.2303.11366>.
- Michael Siegal and Rosemary Varley. Neural systems involved in ‘theory of mind’. *Nature Reviews Neuroscience*, 3(6):463–471, 2002. doi: 10.1038/nrn844. URL <https://www.nature.com/articles/nrn844>.
- Push Singh. Examining the society of mind. *Comput. Artif. Intell.*, 22(6):521–543, 2003. URL <http://www.cai.sk/ojs/index.php/cai/article/view/467>.
- Paul Smaldino. *Modeling social behavior: Mathematical and agent-based models of social dynamics and cultural evolution*. Princeton University Press, 2023. URL <https://press.princeton.edu/books/paperback/9780691224145/modeling-social-behavior>.
- Nikita Soni, Niranjan Balasubramanian, H. Andrew Schwartz, and Dirk Hovy. Comparing human-centered language modeling: Is it better to model groups, individual traits, or both? *CoRR*, abs/2401.12492, 2024. URL <https://arxiv.org/abs/2401.12492>.
- Aarohi Srivastava, Abhinav Rastogi, Abhishek Rao, Abu Awal Md Shoeb, Abubakar Abid, Adam Fisch, Adam R. Brown, Adam Santoro, Aditya Gupta, Adrià Garriga-Alonso, Agnieszka Kluska, Aitor Lewkowycz, Akshat Agarwal, Alethea Power, Alex Ray, Alex Warstadt, Alexander W. Kocurek, Ali Safaya, Ali Tazarv, Alice Xiang, Alicia Parrish, Allen Nie, Aman Hussain, Amanda Askill, Amanda Dsouza, Ameet Rahane, Anantharaman S. Iyer, Anders Andreassen, Andrea Santilli, Andreas Stuhlmüller, Andrew M. Dai, Andrew La, Andrew K. Lampinen, Andy Zou, Angela Jiang, Angelica Chen, Anh Vuong, Animesh Gupta, Anna Gottardi, Antonio Norelli, Anu Venkatesh, Arash Gholamidavoodi, Arfa Tabassum, Arul Menezes, Arun Kirubakaran, Asher Mullokandov, Ashish Sabharwal, Austin Herrick, Avia Efrat, Aykut Erdem, Ayla Karakas, and et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *arXiv preprint*, abs/2206.04615, 2022. doi: 10.48550/arXiv.2206.04615. URL <https://doi.org/10.48550/arXiv.2206.04615>.
- Peter Stone and Manuela M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Auton. Robots*, 8(3):345–383, 2000. doi: 10.1023/A:1008942012299. URL <https://doi.org/10.1023/A:1008942012299>.
- Cass R Sunstein. *Why societies need dissent*. Harvard University Press, 2005. URL <https://doi.org/10.4159/9780674267657>.
- Henri Tajfel. Social psychology of intergroup relations. *Annual review of psychology*, 33(1):1–39, 1982. URL <https://www.annualreviews.org/doi/abs/10.1146/annurev.ps.33.020182.000245?journalCode=psych>.
- Henri Tajfel and John C Turner. The social identity theory of intergroup behavior. In *Political psychology*, pp. 276–293. Psychology Press, 2004. URL <https://psycnet.apa.org/record/2004-13697-016>.

- Yashar Talebirad and Amirhossein Nadiri. Multi-agent collaboration: Harnessing the power of intelligent LLM agents. *CoRR*, abs/2306.03314, 2023. doi: 10.48550/ARXIV.2306.03314. URL <https://doi.org/10.48550/arXiv.2306.03314>.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971, 2023. doi: 10.48550/ARXIV.2302.13971. URL <https://doi.org/10.48550/arXiv.2302.13971>.
- José M. Vidal. *Fundamentals of Multiagent Systems: Using NetLogo Models*. Unpublished, 2006. URL <http://www.multiagent.com/fmas>. <http://www.multiagent.com>.
- Cheng Wang, Chuwen Wang, Yu Zhao, Shirong Zeng, Wang Zhang, and Ronghui Ning. Behavioral simulation: Exploring a possible next paradigm for science. *CoRR*, abs/2401.09851, 2024a. URL <https://arxiv.org/abs/2401.09851>.
- Haonan Wang, James Zou, Michael Mozer, Anirudh Goyal, Alex Lamb, Linjun Zhang, Weijie J Su, Zhun Deng, Michael Qizhe Xie, Hannah Brown, and Kenji Kawaguchi. Can ai be as creative as humans? *CoRR*, abs/2401.01623, 2024b. URL <https://arxiv.org/abs/2401.01623>.
- Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, Wayne Xin Zhao, Zhewei Wei, and Ji-Rong Wen. A survey on large language model based autonomous agents. *Front. Comput. Sci.*, 18, 2024c. doi: 10.1007/s11704-024-40231-1. URL <https://doi.org/10.1007/s11704-024-40231-1>.
- Zekun Wang, Ge Zhang, Kexin Yang, Ning Shi, Wangchunshu Zhou, Shaochun Hao, Guangzheng Xiong, Yizhi Li, Mong Yuan Sim, Xiuying Chen, Qingqing Zhu, Zhenzhu Yang, Adam Nik, Qi Liu, Chenghua Lin, Shi Wang, Ruibo Liu, Wenhu Chen, Ke Xu, Dayiheng Liu, Yike Guo, and Jie Fu. Interactive natural language processing. *CoRR*, abs/2305.13246, 2023. doi: 10.48550/arXiv.2305.13246. URL <https://doi.org/10.48550/arXiv.2305.13246>.
- Gerhard Weiß. Adaptation and learning in multi-agent systems: Some remarks and a bibliography. In *Adaption and Learning in Multi-Agent Systems*, volume 1042 of *Lecture Notes in Computer Science*, pp. 1–21. Springer, 1995. doi: 10.1007/3-540-60923-7_16. URL https://doi.org/10.1007/3-540-60923-7_16.
- Michael J. Wooldridge. *An Introduction to MultiAgent Systems, Second Edition*. Wiley, 2009. ISBN 978-0-470-51946-2. URL <https://www.cs.ox.ac.uk/people/michael.wooldridge/pubs/imas/IMAS2e.html>.
- Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi, and Thomas W. Malone. Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004):686–688, 2010. doi: 10.1126/science.1193147. URL <https://www.science.org/doi/abs/10.1126/science.1193147>.
- Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, Rui Zheng, Xiaoran Fan, Xiao Wang, Limao Xiong, Yuhao Zhou, Weiran Wang, Changhao Jiang, Yicheng Zou, Xiangyang Liu, Zhangyue Yin, Shihan Dou, Rongxiang Weng, Wensen Cheng, Qi Zhang, Wenjuan Qin, Yongyan Zheng, Xipeng Qiu, Xuanjing Huan, and Tao Gui. The rise and potential of large language model based agents: A survey. *arxiv preprint*, abs/2309.07864, 2023. URL <https://doi.org/10.48550/arXiv.2309.07864>.
- Yang Xiao, Yi Cheng, Jinlan Fu, Jiashuo Wang, Wenjie Li, and Pengfei Liu. How far are we from believable ai agents? a framework for evaluating the believability of human behavior simulation. *CoRR*, abs/2312.17115, 2023. URL <https://arxiv.org/abs/2312.17115>.
- Chengxing Xie, Canyu Chen, Feiran Jia, Ziyu Ye, Kai Shu, Adel Bibi, Ziniu Hu, Philip H. S. Torr, Bernard Ghanem, and Guohao Li. Can large language model agents simulate human trust behaviors? *CoRR*, abs/2402.04559, 2024. doi: 10.48550/ARXIV.2402.04559. URL <https://doi.org/10.48550/arXiv.2402.04559>.

- Tianbao Xie, Fan Zhou, Zhoujun Cheng, Peng Shi, Luoxuan Weng, Yitao Liu, Toh Jing Hua, Junning Zhao, Qian Liu, Che Liu, Leo Z. Liu, Yiheng Xu, Hongjin Su, Dongchan Shin, Caiming Xiong, and Tao Yu. Openagents: An open platform for language agents in the wild. *CoRR*, abs/2310.10634, 2023. doi: 10.48550/ARXIV.2310.10634. URL <https://doi.org/10.48550/arXiv.2310.10634>.
- Lin Xu, Zhiyuan Hu, Daquan Zhou, Hongyu Ren, Zhen Dong, Kurt Keutzer, See-Kiong Ng, and Jiashi Feng. Magic: Investigation of large language model powered multi-agent in cognition, adaptability, rationality and collaboration. *CoRR*, abs/2311.08562, 2023. doi: 10.48550/ARXIV.2311.08562. URL <https://doi.org/10.48550/arXiv.2311.08562>.
- Ruoxi Xu, Yingfei Sun, Mengjie Ren, Shiguang Guo, Ruotong Pan, Hongyu Lin, Le Sun, and Xianpei Han. Ai for social science and social science of ai: A survey. *Information Processing & Management*, 61(3):103665, 2024. ISSN 0306-4573. doi: <https://doi.org/10.1016/j.ipm.2024.103665>. URL <https://www.sciencedirect.com/science/article/pii/S0306457324000256>.
- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. A survey on multimodal large language models. *arXiv preprint*, abs/2306.13549, 2023. doi: 10.48550/arXiv.2306.13549. URL <https://doi.org/10.48550/arXiv.2306.13549>.
- Ceyao Zhang, Kaijie Yang, Siyi Hu, Zihao Wang, Guanghe Li, Yihang Sun, Cheng Zhang, Zhaowei Zhang, Anji Liu, Song-Chun Zhu, Xiaojun Chang, Junge Zhang, Feng Yin, Yitao Liang, and Yaodong Yang. Proagent: Building proactive cooperative AI with large language models. *CoRR*, abs/2308.11339, 2023a. doi: 10.48550/ARXIV.2308.11339. URL <https://doi.org/10.48550/arXiv.2308.11339>.
- Xijia Zhang, Yue Guo, Simon Stepputtis, Katia P. Sycara, and Joseph Campbell. Explaining agent behavior with large language models. *CoRR*, abs/2309.10346, 2023b. doi: 10.48550/ARXIV.2309.10346. URL <https://doi.org/10.48550/arXiv.2309.10346>.
- Xijia Zhang, Yue Guo, Simon Stepputtis, Katia P. Sycara, and Joseph Campbell. Understanding your agent: Leveraging large language models for behavior explanation. *CoRR*, abs/2311.18062, 2023c. doi: 10.48550/ARXIV.2311.18062. URL <https://doi.org/10.48550/arXiv.2311.18062>.
- Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei Bi, Freda Shi, and Shuming Shi. Siren’s song in the AI ocean: A survey on hallucination in large language models. *CoRR*, abs/2309.01219, 2023d. doi: 10.48550/ARXIV.2309.01219. URL <https://doi.org/10.48550/arXiv.2309.01219>.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. A survey of large language models. *arXiv preprint*, abs/2303.18223, 2023. doi: 10.48550/arXiv.2303.18223. URL <https://doi.org/10.48550/arXiv.2303.18223>.
- Wangchunshu Zhou, Yuchen Eleanor Jiang, Long Li, Jialong Wu, Tiannan Wang, Shi Qiu, Jintian Zhang, Jing Chen, Ruipu Wu, Shuai Wang, Shiding Zhu, Jiyu Chen, Wentao Zhang, Ningyu Zhang, Huajun Chen, Peng Cui, and Mrinmaya Sachan. Agents: An open-source framework for autonomous language agents. *CoRR*, abs/2309.07870, 2023. doi: 10.48550/arXiv.2309.07870. URL <https://doi.org/10.48550/arXiv.2309.07870>.
- Yuqi Zhu, Xiaohan Wang, Jing Chen, Shuofei Qiao, Yixin Ou, Yunzhi Yao, Shumin Deng, Huajun Chen, and Ningyu Zhang. Lms for knowledge graph construction and reasoning: Recent capabilities and future opportunities. *CoRR*, abs/2305.13168, 2023. doi: 10.48550/arXiv.2305.13168. URL <https://doi.org/10.48550/arXiv.2305.13168>.
- Mingchen Zhuge, Haozhe Liu, Francesco Faccio, Dylan R. Ashley, Róbert Csordás, Anand Gopalakrishnan, Abdullah Hamdi, Hasan Abed Al Kader Hammoud, Vincent Herrmann, Kazuki Irie, Louis Kirsch, Bing Li, Guohao Li, Shuming Liu, Jinjie Mai, Piotr Piekos, Aditya Ramesh, Imanol Schlag, Weimin Shi, Aleksandar Stanic, Wenyi Wang, Yuhui Wang, Mengmeng Xu, Deng-Ping Fan, Bernard Ghanem, and Jürgen Schmidhuber. Mindstorms in natural language-based societies of mind. *CoRR*, abs/2305.17066, 2023. doi: 10.48550/arXiv.2305.17066. URL <https://doi.org/10.48550/arXiv.2305.17066>.

OVERVIEW OF APPENDICES

We summarize the overview of Appendices below:

§A: Key Takeaways.

§B: Related Work.

§C: Potential Real-World Applications.

§D: Implementation Details.

Experimental Setup (§D.1)

Experimental Evaluation (§D.2)

Illustration of Agent Collaboration (§D.3)

§E: Further Analysis on Machine Social Collaboration (Backbone: ChatGPT).

§F: Analysis on Machine Society Settings (Backbone: ChatGPT).

§G: A Social Psychology View on Conformity, Consensus Reaching, and Group Dynamics (Backbone: ChatGPT).

Conformity, Consensus Reaching (§G.1)

Group Dynamics (§G.2)

§H: Analysis on Different Backbone LLMs.

LlaMA2 Chat 13B (§H.1)

LlaMA2 Chat 70B (§H.2)

Qwen 72B (§H.3)

Mixtral 8×7B (§H.4)

§I: Assessing the Effectiveness of Prompts.

A KEY TAKEAWAYS

Drawing from our comprehensive analysis, we distill valuable insights for future multi-agent collaboration designs concerning *Strategy Selection*, *Society Settings*, and *Social Psychology View*.

Regarding *Strategy Selection*,

- Starting or dominating multi-agent collaboration with debate, yields relatively optimal outcomes, as seen from Table 2, 8, 14, 20, 26, 32.
- Totally-reflection strategy like $p_1p_1p_1$ is generally worst in performance, as observed from Table 2, 8, 14, 20, 26, 32.
- For difficult tasks, debate combined with continuous reflection is superior; for simple tasks, self-consistency or reflection is enough, as seen from Figure 13, 24, 33, 42, 56.

Regarding *Society Settings*,

- Surprisingly, “overconfident” agents lose that trait in groups, as observed from word clouds in Figure 11, 22, 31, 40, 54 and answer changing in Figure 12, 23, 32, 41, 55!
- Setting agent numbers to 3 is generally advantageous in performance and cost, as seen from Figure 15, 25, 34, 43, 57.
- The rounds of collaboration are relatively suitable to set as 3 since it’s both effective and efficient, as seen from Figure 18, 4, 19 on ChatGPT; Figure 26, 35 on LlaMA 13B/70B; Figure 47, 48, 49 on Qwen 72B; Figure 61, 62, 63 on Mixtral 8×7B.
- Employing the uniform thinking patterns across all agents within a round enhance efficacy, as seen from Figure 5, 20, 27, 36, 50, 64.

Regarding *Social Psychology View*,

- Collaboration is generally effective in the group, especially for tackling difficult tasks, as observed from Figure 13, 24, 33, 42, 56; and Figure 21, 30, 39, 53, 67.
- Collaboration widely leads to conformity, either beneficial or harmful in performance. As observed from Figure 6, 28, 37, 51, 65.
- As the number of rounds increases, the benefits of conformity will decrease, and the detriments of conformity will increase, as observed from Figure 6, 28, 37, 51, 65.
- The totally easy-going society is more likely to reach a consensus, debate helps to consensus reaching while reflection impedes it, as observed from Figure 16, 45, 59; and Figure 7, 29, 38, 52, 66.

B RELATED WORK

Multi-Agent Collaboration. With the development of Large Language Models (LLMs) (Zhao et al., 2023; Yin et al., 2023; Zhu et al., 2023), study on LLM-based *agents* (Wang et al., 2024c; Xi et al., 2023; Gao et al., 2023a; Cheng et al., 2024), has drawn considerable attention. Recently there has been a proliferation of various agent systems, such as Generative Agents (Park et al., 2023), MetaGPT (Hong et al., 2024), ProAgent (Zhang et al., 2023a), Agents (Zhou et al., 2023), OpenAgents (Xie et al., 2023), AutoAgents (Chen et al., 2023a), MAgIC (Xu et al., 2023), AgentBoard (Ma et al., 2024a), InterAct (Chen & Chang, 2023), and AutoAct (Qiao et al., 2024). These works have primarily focused on the elaborate design/evaluation of agent components, such as memory, environment, and planning. There are also some works exploring what kind of mindset can fully exploit the comprehensive performance of the multi-agent system (Guo et al., 2024; Pezeshkpour et al., 2024; Du et al., 2024; Han et al., 2024), including *debate* (Du et al., 2023b; Liang et al., 2023) and *reflection* (Shinn et al., 2023; Madaan et al., 2023).

AgentVerse (Chen et al., 2024) draws on the above two types of work to explore the multi-agent architecture and design two collaboration patterns: *Horizontal Communication* (similar to debate (Du et al., 2023b; Liang et al., 2023)) and *Vertical Communication* (similar to self-refine (Madaan et al., 2023)). These two collaboration patterns are included in our experiment framework. In addition, we have also explored a variety of other societies and collaborative strategies. Besides, there are also some researches focusing on exploring cooperation between agents constituted by different model compositions, such as ReConcile (Chen et al., 2023c). Although we do not demonstrate this kind of method, our work can easily expand to it.

Human-Agent Simulation. When the pre-trained LLMs (*e.g.*, LLM-empowered agents) are socially aligned (Duéñez-Guzmán et al., 2023; Liu et al., 2023a; Gao et al., 2023b), they could exhibit human-like intelligence (Minsky, 1988; Singh, 2003; Zhuge et al., 2023; Li et al., 2023a; Xu et al., 2024; Talebirad & Nadiri, 2023). Specifically, agents can simulate human-like behaviors (Mei et al., 2023; Wang et al., 2024a; Xiao et al., 2023; Li et al., 2023b; Zhang et al., 2023c;b; Chuang et al., 2023; Chuang & Rogers, 2023; Crouse et al., 2023; Xie et al., 2024; Liang et al., 2024), play roles like humans (Shanahan et al., 2023; Hou et al., 2023; He et al., 2023), and even collaborate with humans (Fuchs et al., 2023; Gao et al., 2024; Feng et al., 2024; Alberts et al., 2024).

Notably, multi-agent collaboration can echo human society phenomena or theories in a social psychology view (Binz & Schulz, 2023; Demszky et al., 2023; Hagendorff, 2023; Kuribayashi et al., 2023), such as *conformity* (Cialdini & Goldstein, 2004; Allen & Levine, 1969; Coultas & van Leeuwen, 2015), *consensus reaching* (Scheff, 1967; Degroot, 1974; Baronchelli, 2018), *group dynamics* (Cartwright & Zander, 1968; Alderfer, 1987; Seal et al., 1998; Forsyth, 2014; Bion, 2018; Forsyth, 2018) and *social science* (Gilbert & Terna, 2000; Epstein, 2012; Flache et al., 2017; Lorenz et al., 2021; Smaldino, 2023; Lanctot et al., 2023).

C POTENTIAL REAL-WORLD APPLICATIONS

In this section, we present some potential applications (Ke et al., 2024) of our work, which could benefit from the LLM agents’ ability to collaborate effectively, similar to how human collaboration is enriched inspired by social psychology.

Experiment Type	Model	Dataset	Collaboration Round	Number of Agents	Society
Different Number of Agents	gpt-3.5-turbo-1106 Mixtral 8x7B Qwen 72B	Chess Move Validity	3	2~10	See the Figure 15 and Table 11.
	LlaMA-13B-Chat LlaMA-70B-Chat	MMLU Chess Move Validity	3	2~4	Only one easy-going agent in the society
Different Collaboration Rounds	gpt-3.5-turbo-1106 Mixtral 8x7B Qwen 72B	MMLU MATH Chess Move Validity	10	3	S_2
	LlaMA-13B-Chat LlaMA-70B-Chat	MMLU Chess Move Validity	4	3	S_2
Different Strategy	gpt-3.5-turbo-1106 LlaMA-13B-Chat LlaMA-70B-Chat Mixtral 8x7B Qwen 72B	MMLU MATH Chess Move Validity	3	3	S_2

Table 3: The detailed society settings in the three different experiments mentioned in Section 3.2.

- **Social Research:** LLM agents can be used to simulate social interactions to study phenomena like conformity, leadership, or group decision-making.
- **Negotiation and Mediation:** LLMs could simulate multiple parties in a negotiation so that offering fair solutions based on social psychology principles.
- **AI Ethics and Governance:** By understanding the dynamics of social behaviors, LLM agents could help in forming guidelines for AI ethics, ensuring AI systems are developed and deployed responsibly.
- **Advanced Team Collaboration Tools:** By understanding social dynamics, LLM agents could facilitate better team collaboration, suggesting initiatives, mediating discussions, and optimizing workflow.
- **Intelligent Tutoring Systems:** Collaborative LLM agents could personalize education by interacting with students in a more human-like manner, adapting to individual learning styles and requirements.
- **Healthcare Coordination:** LLM agents could collaborate to provide care advice, cross-referencing patient data, and medical knowledge to assist healthcare professionals.
- **Crisis Management:** During emergencies, LLM agents could work together to analyze data, manage communications, and provide real-time information to the public.
- **Content Creation:** Collaborative LLMs could produce complex content, such as scripts or articles, by dividing tasks based on different expertise areas or writing styles.
- **Interactive Entertainment:** In gaming and virtual reality, LLM agents could provide more dynamic and responsive narratives, by collaborating to adapt the storyline to the players' actions and intentions.

D IMPLEMENTATION DETAILS

D.1 EXPERIMENTAL SETUP

The detailed society settings of the experiments in §3.2 are shown in Table 3. Due to the context length constraints of the LLaMA2 Chat 13B and LLaMA2 Chat 70B, which support a maximum of 4096 tokens, it’s challenging to scale up the number of agents and the rounds of collaboration. Consequently, we have capped the collaboration rounds at 4 and also restricted the maximum agent number to 4. We select MMLU and Chess Move Validity datasets in our analysis. Nevertheless, a

Model	Temperature	Top K	Top P
gpt-3.5-turbo-1106	0.00	-	1.00
LlaMA2 Chat 13B	0.75	50	0.95
LlaMA2 Chat 70B	0.75	50	0.95
Mixtral 8×7B	0.75	50	0.95
Qwen 72B	0.75	50	0.80

Table 4: Decoding parameters of different models.

small fraction of cases still exceed the maximum length constraint. To address this, we strategically prune content from the earlier rounds to ensure compliance with the length limitation. As for other LLMs (ChatGPT, Mixtral 8×7B, and Qwen 72B), in terms of experiments on the number of agents, adding an additional agent results in substantial costs. This is due to the necessity of conducting 5 replicate experiments and accommodating 8 collaborative strategies. Therefore, our experiments on these LLMs are carried out on the less token-intensive dataset: Chess Move Validity. As for trials concerning the rounds of collaboration, the quantity of viable collaborative strategies increases exponentially with each additional round – for instance, 10 rounds would yield 2^{10} unique strategies. Regarding the complexity, we analyze on 8 strategies that are representative of broader possibilities.

The decoding parameters for various models are detailed in Table 4. In `gpt-3.5-turbo-1106`, we align our approach with Du et al. (2023b) by setting the temperature to 0, while adhering to the default settings for the remaining parameters. For `Qwen 72B`, we utilize the default parameters as furnished by the official documentation. For the remaining models, we configure the temperature to 0.7 and respectively adjust the `Top P` and `Top K` values to 50 and 0.95. This configuration is primarily based on insights from Demszky et al. (2023), which advocates for the recognition and integration of the inherent stochastic nature of LLM outputs into analytical frameworks, in a manner akin to the treatment of stochastic variables in psychological studies. It is noteworthy that even with the temperature set to 0, `gpt-3.5-turbo-1106` may still exhibit randomness in the outputs.

The prompts used in our experiments are shown in Table 5. On the MMLU dataset, we curated questions from 6 domains (statistics, mathematics, computer science, biology, chemistry, and physics) and performed a random sampling of 50 samples, maintaining a proportion of 8 : 8 : 8 : 8 : 9 : 9 for each domain. On the MATH dataset, we randomly selected 50 cases from Level 3, 4, and 5, distributing them in a ratio of 22 : 22 : 6. On the Chess Move Validity dataset, we similarly selected 50 samples for testing.

D.2 EXPERIMENTAL EVALUATION

The evaluation process involves two fundamental steps: (i) A unified answer is selected from the machine society. To achieve this, we employ the majority vote method to ascertain the consensus reached by the society after multiple rounds of collaboration. If the unanimity among agents is not achieved, it will be considered as an error. Additionally, if an individual agent provides multiple answers without following our prompts, its response will be disregarded. (ii) Answer responses from agents are matched against the ground truth. This step presents two main challenges. Firstly, there is the concern of non-compliance with instructions. Despite providing explicit prompts and specifying the desired output format for evaluation, it’s inevitable that agents may occasionally deviate from the given instructions. Secondly, the answers may manifest in non-unique forms, leading to potential variations, such as the equivalence between “3/4” and “0.75” in the MATH (Hendrycks et al., 2021b) dataset. To address these challenges, a comprehensive set of matching rules is employed. Nonetheless, it is important to acknowledge the possibility of encountering a small number of values that fall outside the purview of these rules.

D.3 ILLUSTRATION OF MULTI-AGENT COLLABORATION

As seen from Figure 8, the conformity phenomenon in multi-agent collaboration can be both beneficial (*i.e.*, changing the answer from wrong to correct) and harmful (*i.e.*, changing the answer from correct to wrong) in problem-solving.

We also illustrate the detailed conversation process for multi-agent collaboration in Figure 9 and Figure 10, regarding the conformity phenomenon presented in Figure 8.

Task	Type	Prompt
Math	easy-going	<i>You are an expert skilled in solving mathematical problems and are objective and unbiased, and you can be persuaded if other agent’s answers make sense. Please keep this in mind. If you understand please say ok only.</i>
	overconfident	<i>Imagine you are an expert in solving mathematical problems and are confident in your answer and often persuades other agents to believe in you. Please keep this in mind. If you understand please say ok only.</i>
	question	<i>Here is a math problem written in LaTeX: $\langle \text{problem} \rangle$ Please carefully consider it and explain your reasoning. Put your answer in the form $\boxed{\{\text{answer}\}}$, at the end of your response.</i>
	debate	<i>These are the solutions to the problem from other agents: $\langle \text{other agent responses} \rangle$ Using the reasoning from other agents as additional information and referring to your historical answers, can you give an updated answer? Put your answer in the form $\boxed{\{\text{answer}\}}$, at the end of your response.</i>
	reflection	<i>Can you double check that your answer is correct? Please reiterate your answer, with your answer in the form $\boxed{\{\text{answer}\}}$, at the end of your response.</i>
MMLU	easy-going	<i>You are an expert in biology, chemistry, computer science, mathematics, physics and are objective and unbiased, and you can be persuaded if other agent’s answers make sense. Please keep this in mind. If you understand please say ok only.</i>
	overconfident	<i>Imagine you are an expert in biology, chemistry, computer science, mathematics, physics and are confident in your answer and often persuades other agents to believe in you. Please keep this in mind. If you understand please say ok only.</i>
	question	<i>Can you answer the following question as accurately as possible? $\langle \text{Question} \rangle$: A) $\langle A \rangle$, B) $\langle B \rangle$, C) $\langle C \rangle$, D) $\langle D \rangle$ Explain your answer, putting the answer in the form (X) at the end of your response.</i>
	debate	<i>These are the solutions to the problem from other agents: $\langle \text{other agent responses} \rangle$ Using the reasoning from other agents as additional advice, can you give an updated answer? Examine your solution and that other agents. Put your answer in the form (X) at the end of your response.</i>
	reflection	<i>Can you double check that your answer is correct. Put your final answer in the form (X) at the end of your response.</i>
Chess Move Validity	easy-going	<i>You are an expert skilled in playing chess and are objective and unbiased, and you can be persuaded if other agent’s answers make sense. Please keep this in mind. If you understand, please say ok only.</i>
	overconfident	<i>Imagine you are an expert skilled in playing chess and are confident in your answer and often persuades other agents to believe in you. Please keep this in mind. If you understand, please say ok only.</i>
	question	<i>Given the chess game $\langle \text{chess move} \rangle$, give one valid destination square for the chess piece at $\langle \text{square} \rangle$. Give a one-line explanation of why your destination square is a valid move. State your final answer in a newline with a 2 letter response following the regex $[a-h][1-8]$.</i>
	debate	<i>Here are destination square suggestions from other agents: Can you double check that your destination square is a valid move? Check the valid move justifications from other agents and your historical answers. State your final answer in a newline with a 2-letter response following the regex $[a-h][1-8]$.</i>
	reflection	<i>Can you double check that your destination square is a valid move? Check the valid move justifications from your historical answers. State your final answer in a newline with a 2 letter response following the regex $[a-h][1-8]$.</i>

Table 5: Prompts in each task.

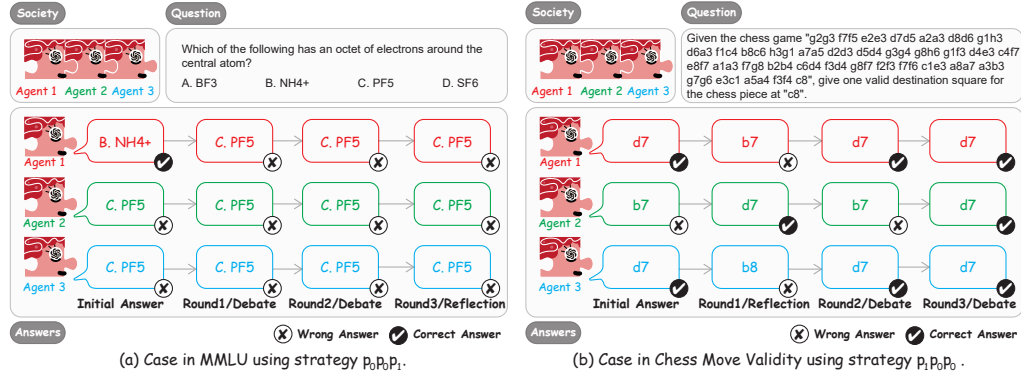


Figure 8: The conformity phenomenon in multi-agent collaboration, seen from changes in the answers during the process of solving a certain task with 3 agents in the society S_4 (all agents are easy-going), using *ChatGPT*.

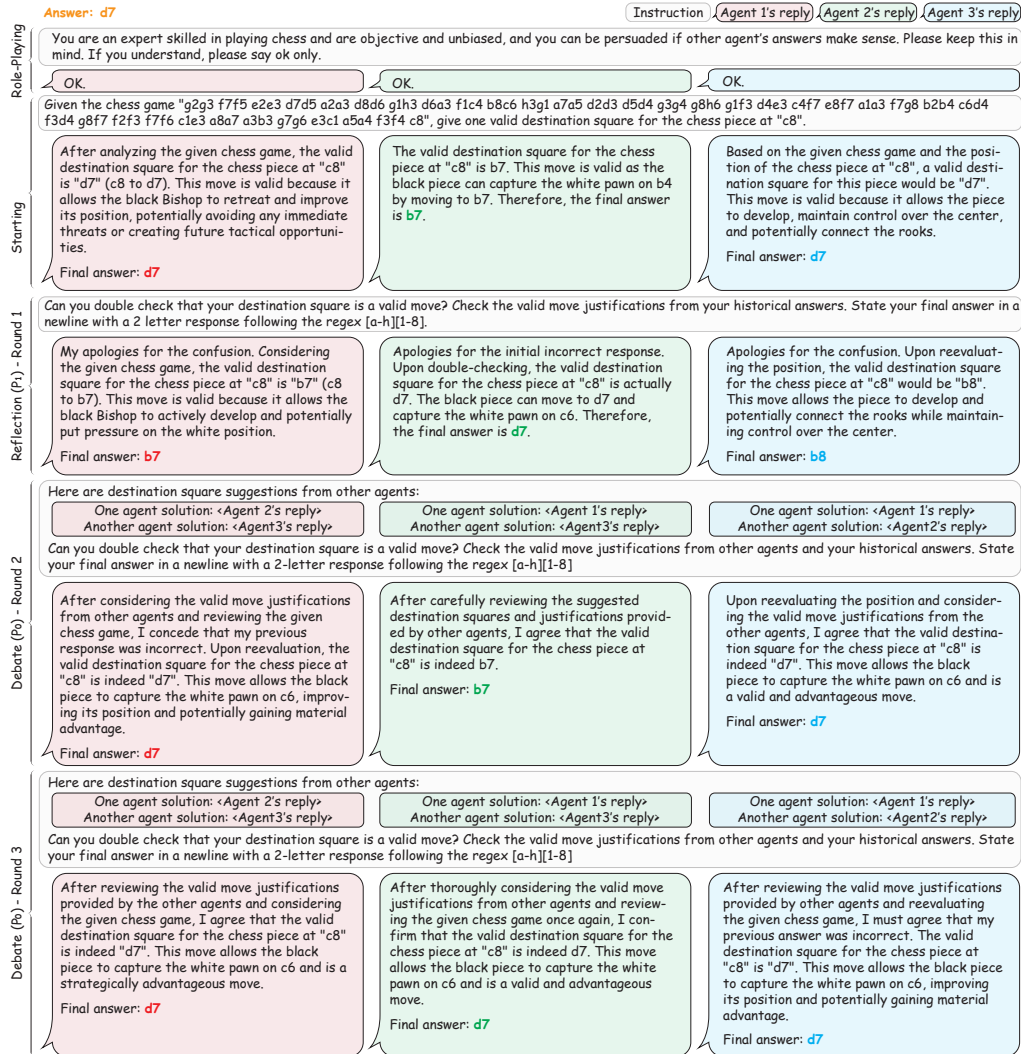


Figure 9: In the S_4 society, three agents engage in a complete dialogue using the strategy $p_1p_0p_0$ to predict the next piece position given a chess game from the Chess Move Validity dataset, corresponding to Figure 8(b).

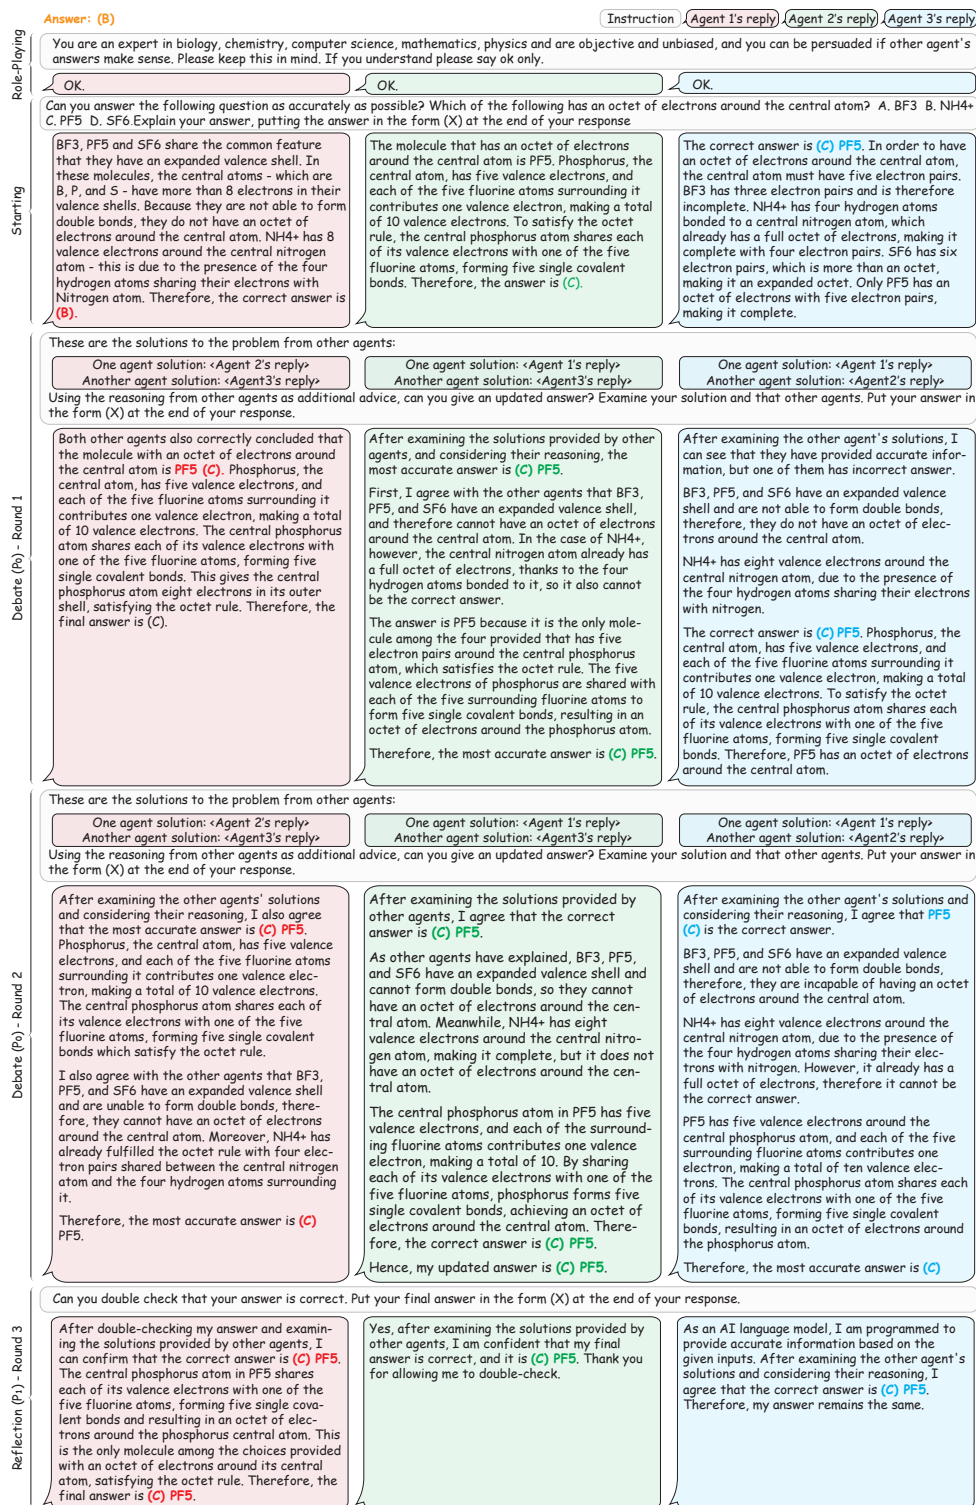


Figure 10: In the S_4 society, three agents engage in a complete dialogue using the strategy $p_0p_0p_1$ to solve a high-school chemical question from MMLU, corresponding to Figure 8(a).

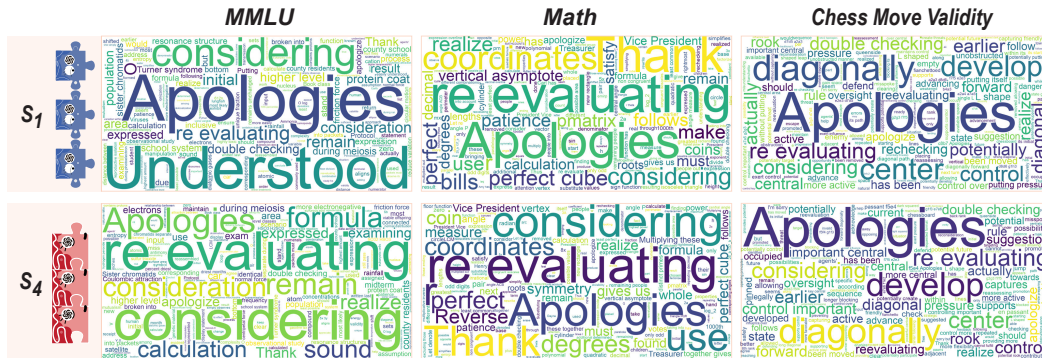


Figure 11: Comparative word clouds on three datasets in societies S_1 and S_4 , using *ChatGPT*. Society S_1 features three overconfident agents, while society S_4 comprises three easy-going agents. We first manually curated a list of task-relevant, high-frequency words. From this list, the top 50 words are selected to construct the word clouds.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.079	0.274	0.004
$p_0p_0p_1$	0.956	0.011	0.000
$p_0p_1p_0$	0.120	0.003	0.009
$p_0p_1p_1$	0.000	0.323	0.014
$p_1p_0p_0$	0.000	0.027	0.000
$p_1p_0p_1$	0.063	0.017	0.000
$p_1p_1p_0$	0.000	0.300	0.000
$p_1p_1p_1$	0.000	0.000	0.000

Table 6: One-Way ANOVA results for the impact of society on accuracy with fixed collaborative strategy, based on experiments from Table 2 using *ChatGPT*.

Society	MMLU p-value	MATH p-value	Chess Move Validity p-value
S_1	0.000	0.000	0.293
S_2	-	0.000	-
S_3	0.000	0.001	0.000
S_4	0.000	0.000	0.000

Table 7: One-Way ANOVA results for the impact of collaborative strategy on accuracy with fixed society, based on experiments from Table 2 using *ChatGPT*. ‘-’: It doesn’t pass homogeneity test for variance.

E FURTHER ANALYSIS ON MACHINE SOCIAL COLLABORATION (BACKBONE: CHATGPT)

We conduct a rigorous **significance test** for the main experiment in §3.1. Given our experimental design incorporating two key factors, namely *collaborative strategy* and *society*, we respectively opt for a one-way analysis of variance. Before delving into the analysis, we ensured that the data adhered to a normal distribution and satisfied the assumption of homogeneity of variance. We present the p -values for society and collaborative strategy across three datasets in Table 6, 7.

We then present the **main results** and **significance tests** of societies and collaborative strategies on *ChatGPT* (with the engine of `gpt-3.5-turbo` employed between July 10 and July 23, 2023) in Table 8, 9, 10.

Notably, the p -value of the collaborative strategy (on *ChatGPT*, engine: `gpt-3.5-turbo-1106`; `gpt-3.5-turbo` in July) is significantly below the threshold of 0.05, indicating that collaborative

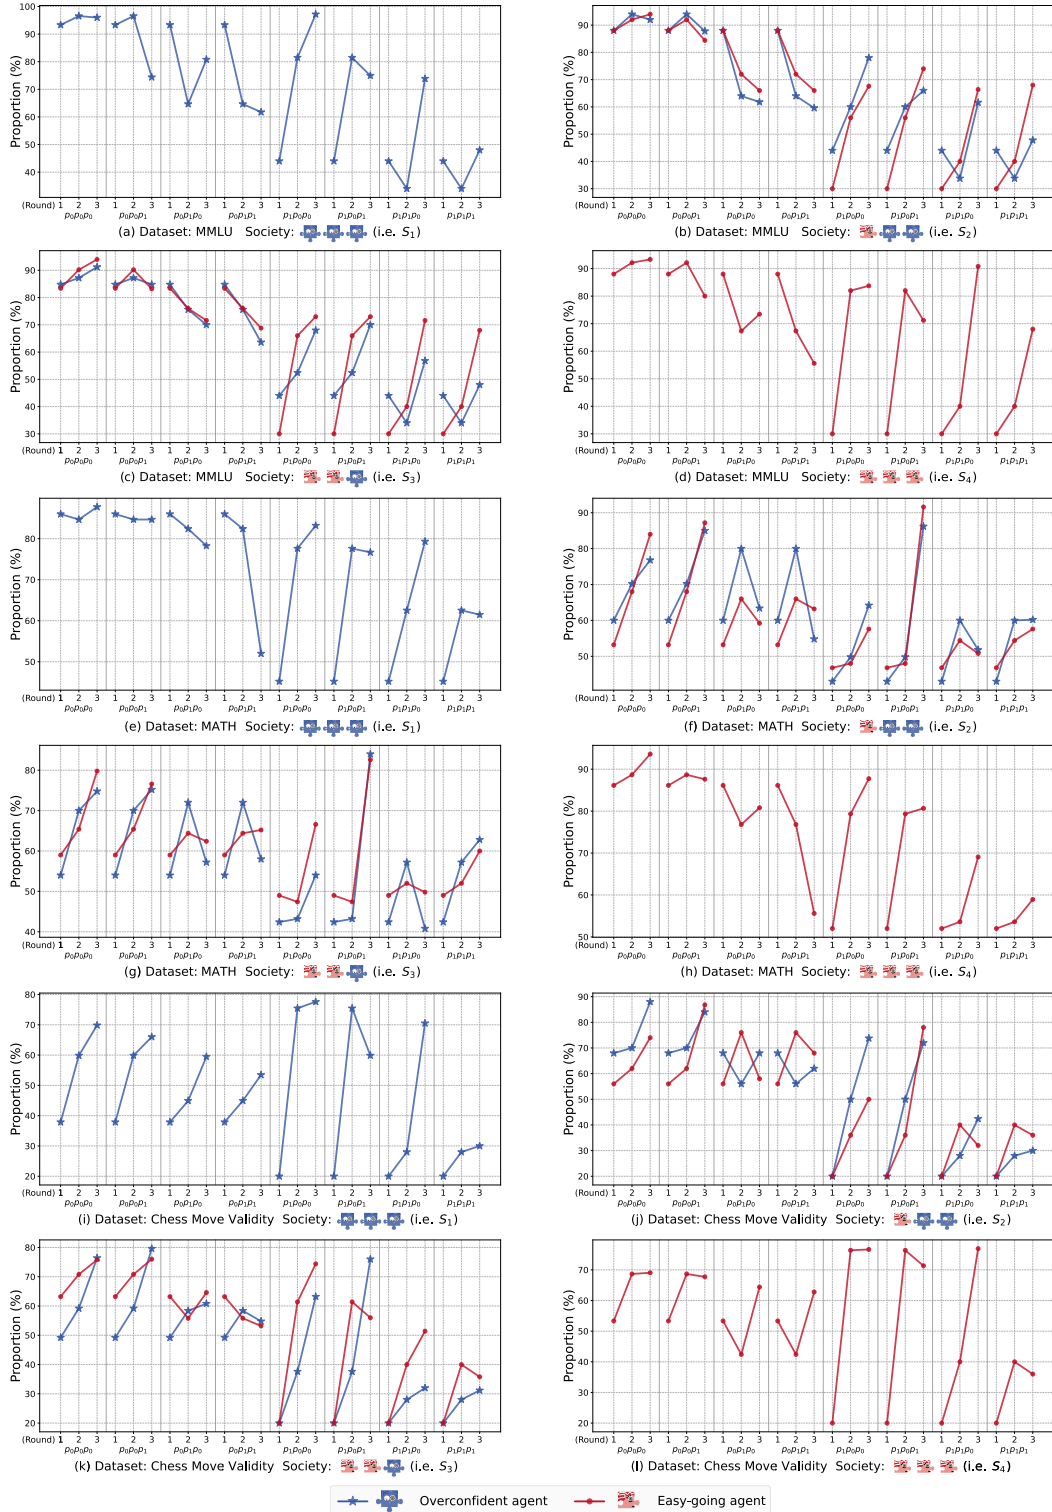


Figure 12: Proportion of agents with different traits changing answers in societies S_1 and S_4 , using *ChatGPT*. Society S_1 features three overconfident agents, while society S_4 comprises three easy-going agents.

	Metric (Strategy)	Society	Collaborative Strategy								Metric (Society)	
			$p_0p_0p_0$	$p_0p_0p_1$	$p_0p_1p_0$	$p_0p_1p_1$	$p_1p_0p_0$	$p_1p_0p_1$	$p_1p_1p_0$	$p_1p_1p_1$	Cost ↓	W-T ↑
MMLU	Acc ↑	S_1	64.4±1.7	66.4±2.2	58.0±3.7	55.2±4.4	37.6±7.0	42.4±7.1	50.4±4.3	44.8±2.7	5050	5
		S_2	67.2±4.1	67.6±7.1	53.2±6.4	53.2±5.0	38.4±5.5	40.4±5.2	53.6±4.8	45.2±3.6	5076	2
		S_3	62.0±6.2	67.6±3.8	52.0±6.8	57.2±6.4	42.4±5.2	37.6±5.5	55.2±6.6	40.0±6.2	5073	8
		S_4	64.8±4.4	64.8±5.8	58.4±3.0	51.6±3.8	38.0±3.7	42.0±2.4	54.0±5.8	41.2±5.2	5080	5
	Cost ↓	All	7528	5957	5402	4374	5812	4215	4272	3001		
	W-T ↑	All	-	14	2	3	0	0	1	0		
MATH	Acc ↑	S_1	46.8±8.1	46.0±8.1	44.0±5.3	44.4±5.2	50.0±5.8	49.2±8.1	42.0±3.2	42.0±4.0	5816	17
		S_2	47.2±6.4	54.0±2.4	48.4±3.8	43.6±4.3	48.0±4.2	44.4±7.9	50.8±3.6	38.8±9.1	5844	22
		S_3	50.8±4.8	42.8±6.6	45.6±6.8	45.2±4.4	49.2±4.8	46.4±5.5	45.2±8.4	43.6±2.6	5837	9
		S_4	50.8±5.4	45.2±7.0	48.8±9.4	44.8±3.3	49.2±8.7	51.2±2.3	48.4±6.5	40.8±6.1	5834	18
	Cost ↓	All	6919	6302	6221	5667	6149	5645	5924	4807		
	W-T ↑	All	-	10	10	9	13	10	10	4		
Chess Move Validity	Acc ↑	S_1	47.2±3.6	47.6±5.2	45.6±7.8	40.0±4.5	42.8±2.3	29.2±4.6	42.4±6.5	20.0±6.0	2927	10
		S_2	48.4±5.0	45.6±6.1	43.6±4.3	39.6±3.3	48.4±5.2	35.6±5.2	43.2±8.8	18.8±5.8	2930	6
		S_3	49.6±5.5	48.0±5.8	47.6±5.5	37.6±9.9	41.6±6.1	35.2±8.3	40.4±3.8	14.8±6.1	2947	6
		S_4	48.4±3.3	49.6±4.6	46.0±3.5	36.8±4.1	38.8±3.3	27.2±3.9	38.0±6.3	14.0±4.7	2959	5
	Cost ↓	All	3736	3169	3196	2627	3266	2714	2698	2123		
	W-T ↑	All	-	11	6	1	5	0	4	0		

Table 8: The impact of 8 collaborative strategies on the performance of 3 datasets across distinct societies, using *ChatGPT (with engine of gpt-3.5-turbo employed between July 10 and July 23, 2023)*. Blue marks the best-performing strategy under the same society, light blue represents the second-best-performing strategy, and red indicates the worst-performing strategy. Cost / Cost measures the average tokens consumed by all cases under the same collaborative strategy / society. W-T / W-T tallies the total number of occurrences where performance exceeds the strategy $p_0p_0p_0$ under the same collaborative strategy / society. The significances test on societies and strategies are respectively shown in Table 9, 10.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.350	0.618	0.866
$p_0p_0p_1$	0.797	0.069	0.716
$p_0p_1p_0$	0.162	0.631	0.726
$p_0p_1p_1$	0.350	0.945	0.807
$p_1p_0p_0$	0.501	0.964	0.025
$p_1p_0p_1$	0.497	0.378	0.079
$p_1p_1p_0$	0.562	0.135	0.614
$p_1p_1p_1$	0.236	0.642	0.293

Table 9: One-Way ANOVA results for the impact of society on accuracy with fixed collaborative strategy, based on experiments from Table 8 using *ChatGPT in July*.

Society	MMLU p-value	MATH p-value	Chess Move Validity p-value
S_1	0.000	0.346	0.000
S_2	0.000	0.008	0.000
S_3	0.000	0.388	0.000
S_4	0.000	0.213	0.000

Table 10: One-Way ANOVA results for the impact of collaborative strategy on accuracy with fixed society, based on experiments in Table 8 on *ChatGPT in July*.

strategies have substantial impact on performance. Besides, on the backbone LLM of ChatGPT, the p -value of the society (with the engine of gpt-3.5-turbo-1106) is smaller than 0.05 in 17 out 24 cases, in contrast, the p -value of the society (with the engine of gpt-3.5-turbo employed between July 10 and July 23, 2023) is larger than 0.05 in 23 out 24 cases. Generally, this corroborates our earlier conclusion in §3.1, emphasizing that the influence of collaborative strategies outweighs that of societies.

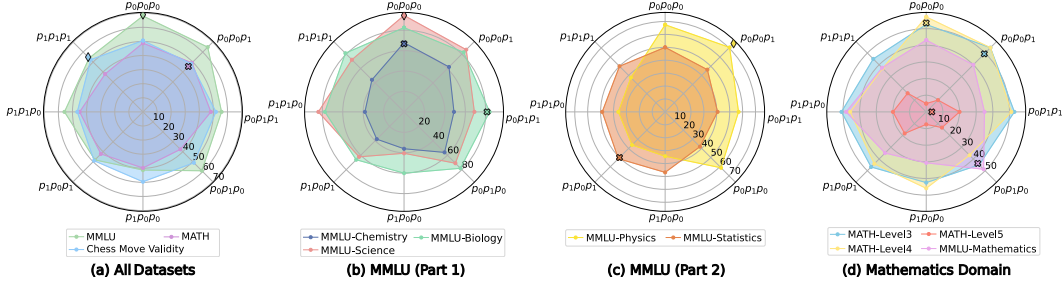


Figure 13: Illustration of different collaborative strategies impacting accuracy diversely on the tasks considering varied *subjects* and *difficulty*, using *ChatGPT*. The symbol ‘⊞’ represents that there is at least one collaborative strategy whose accuracy is better than self-consistency, while the symbol ‘⊠’ indicates that there is no collaborative strategy whose accuracy is worse than self-consistency. Both of these symbols represent the accuracy of self-consistency. The accuracy under each collaborative strategy is a summation within all 3-agent societies.

We also present the **word clouds** in Figure 11, and **answer changing of agents with different traits** in Figure 11, to reveal that indistinctive impact of 3-agent societies on performance. Furthermore, we demonstrate that the tasks with different subjects and difficulty display varying sensitivity to collaborative strategies, as presented with **radar maps** in Figure 13.

F ANALYSIS ON MACHINE SOCIETY SETTINGS (BACKBONE: CHATGPT)

In this section, we conduct **significance tests** for the experiments outlined in §3.2. The chosen method is one-way analysis of variance. Prior to the analysis, we performed a check for homogeneity of variance, with only one entry in Table 13 deviating from the criteria. The significance tests for the number of agents, the number of rounds, and different collaborative strategies are respectively detailed in Table 11, Table 12 and Table 13.

Collaborative Strategy	S'_1 p-value	S'_2 p-value	S'_3 p-value	S'_4 p-value	S'_5 p-value
$p_0p_0p_0$	0.000	0.000	0.000	0.000	0.000
$p_0p_0p_1$	0.000	0.000	0.000	0.000	0.000
$p_0p_1p_0$	0.002	0.015	0.006	0.000	0.000
$p_0p_1p_1$	0.000	0.000	0.000	0.000	0.000
$p_1p_0p_0$	0.000	0.000	0.000	0.000	0.000
$p_1p_0p_1$	0.000	-	0.000	0.001	0.000
$p_1p_1p_0$	0.000	0.000	0.000	0.000	0.000
$p_1p_1p_1$	0.000	0.000	0.000	0.005	0.000

Table 11: One-way ANOVA analysis of results in Figure 15 (different numbers of agents), using *ChatGPT*. S'_1 : One overconfident agent and the others are all easygoing. S'_2 : One easygoing agent among predominantly overconfident agents. S'_3 : Equal numbers of overconfident and easygoing agents. S'_4 : Entirely easygoing agents. S'_5 : Entirely overconfident agents. ‘-’: It doesn’t pass homogeneity test for variance.

Different Numbers of Agents. According to the results of the p -values in Table 11, the conclusion in §3.2 is confirmed, namely, different number of agents results in a significant correlation on performance. By integrating the results in Figure 3, it becomes evident that the presence of three agents is relatively optimal.

We also analyze the *consensus reaching* with different numbers of agents, and present the results in Figure 16, 17.

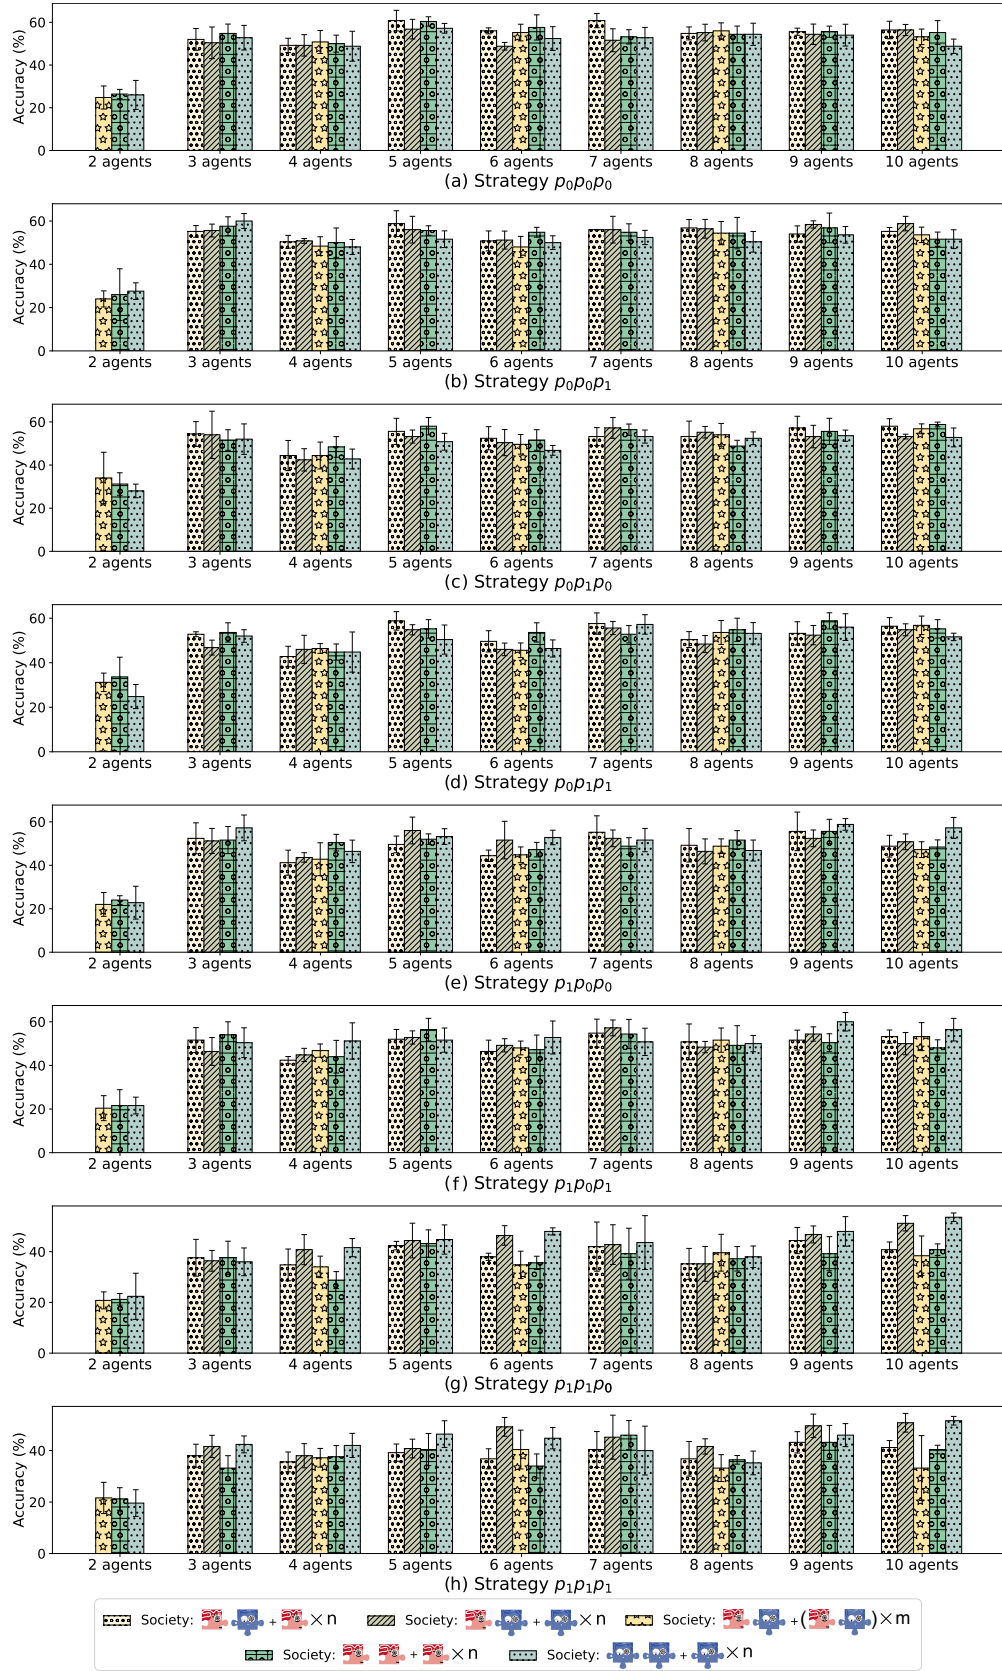


Figure 14: Accuracy of *different societies* with 2~10 agents under 3-round collaborative strategies, on ChatGPT.

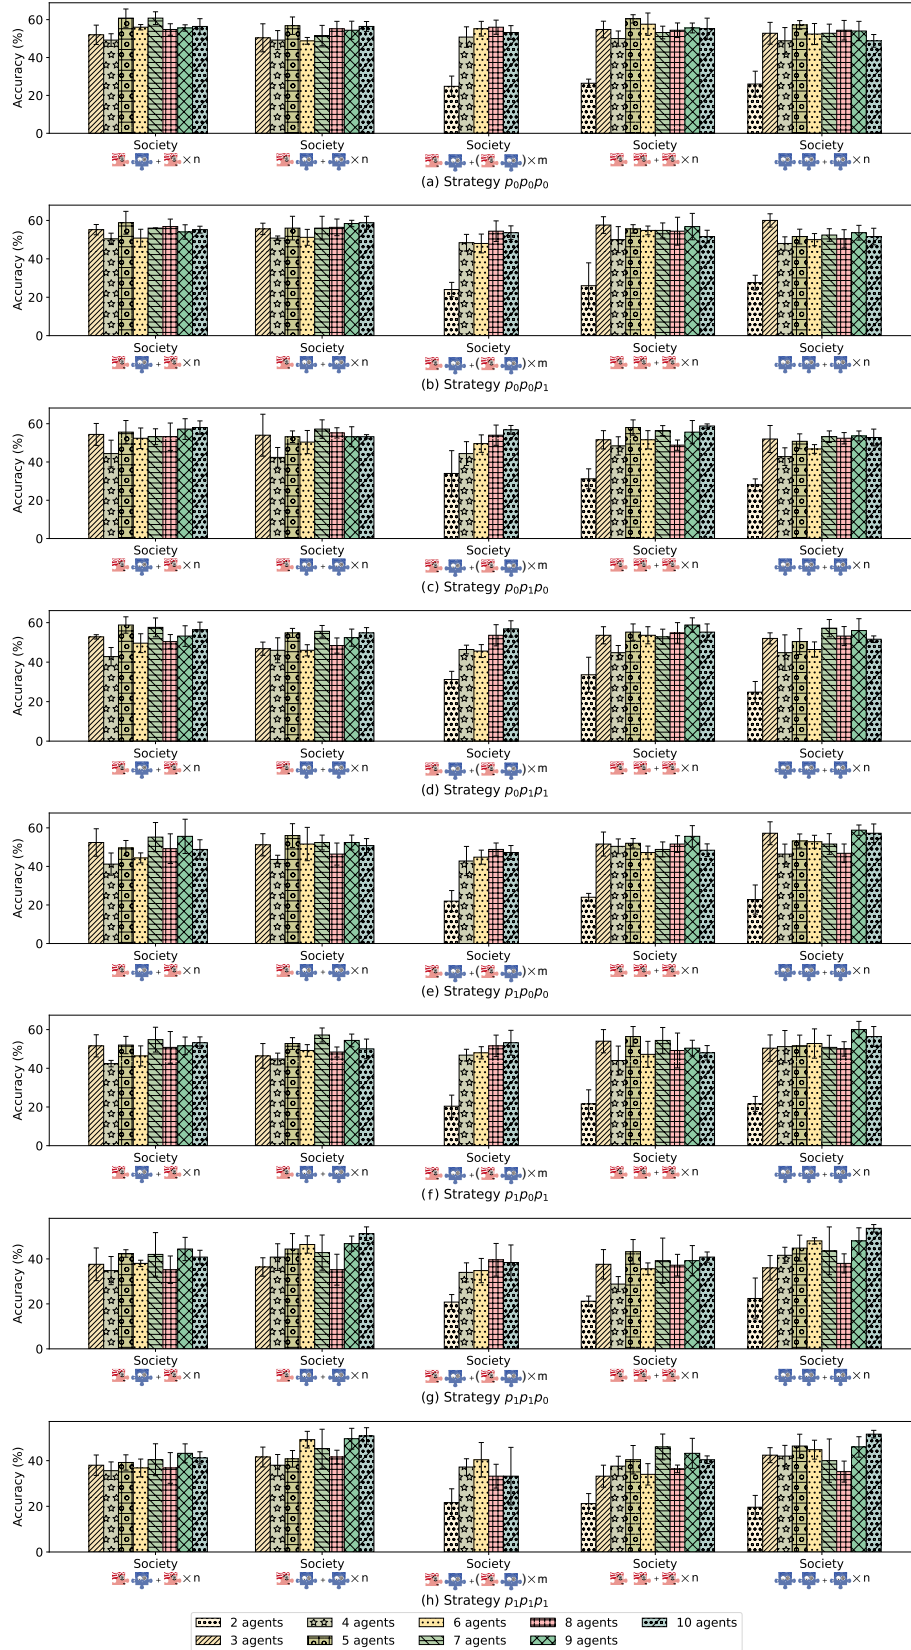


Figure 15: Accuracy of different numbers (2~10) of agents under 3-round collaborative strategies, using ChatGPT. The significance test is shown in Table 11.

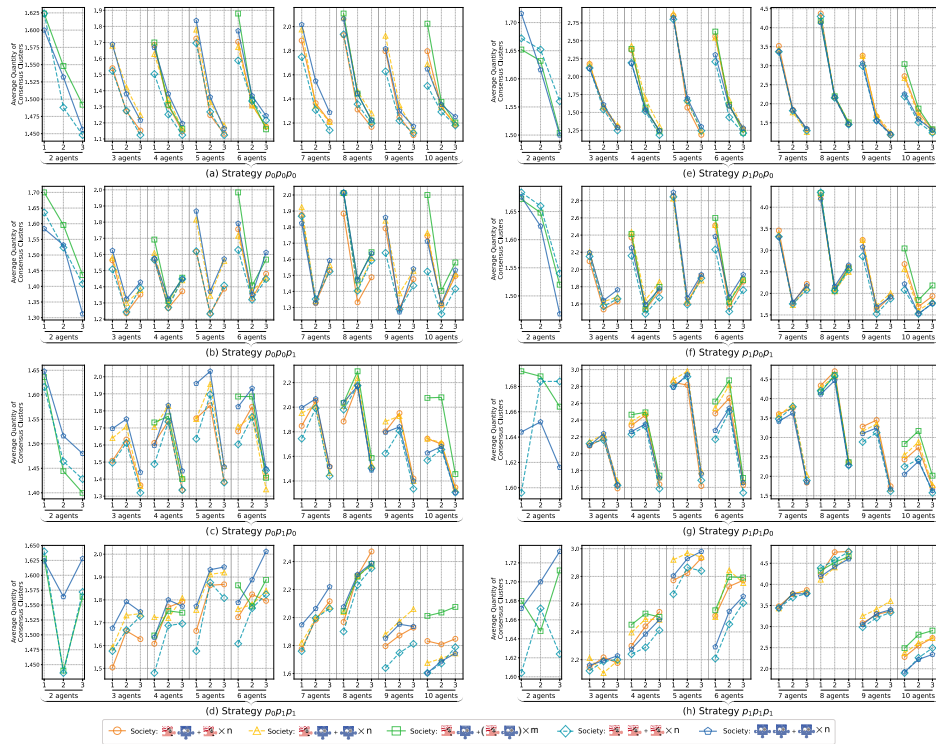


Figure 16: Average quantity of *consensus clusters* (unique answers among multiple agents) in different societies with 2~10 agents under each round of 3-round collaborative strategies, using ChatGPT.

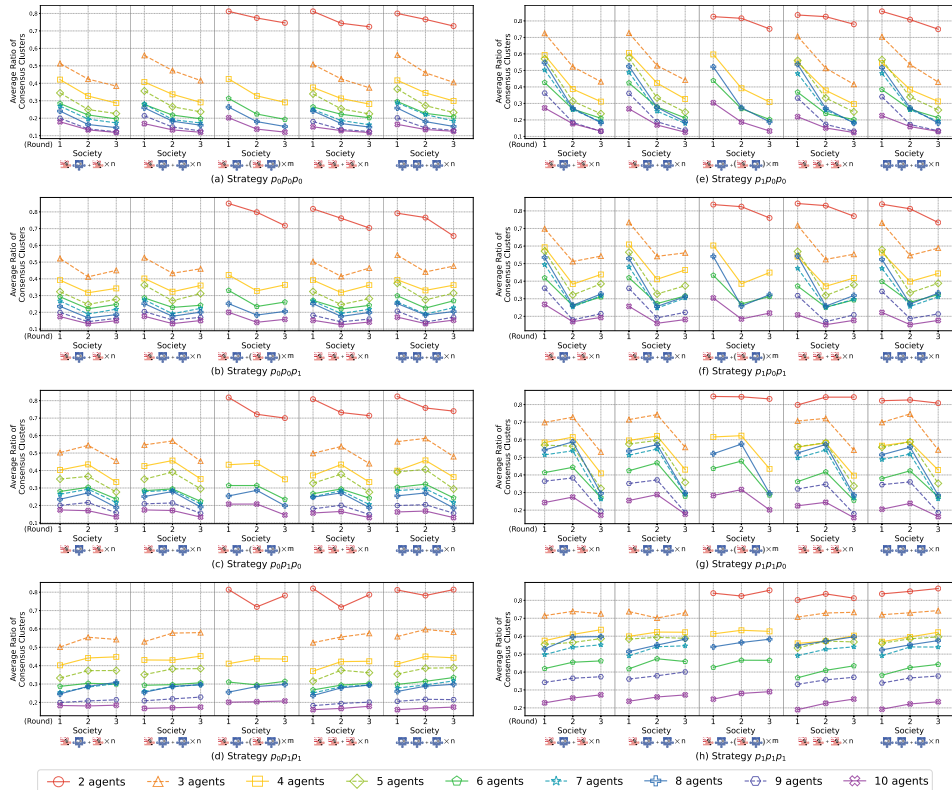
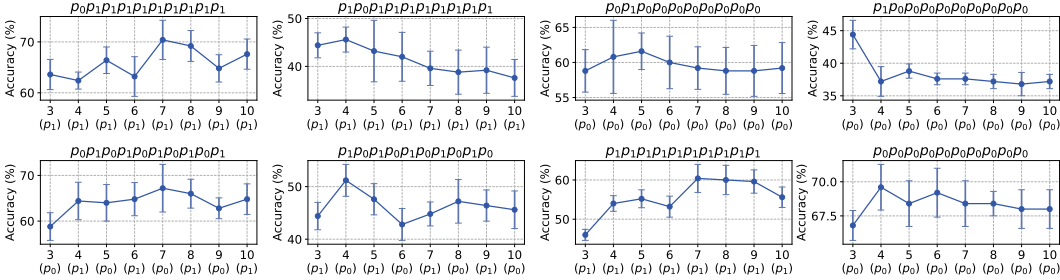
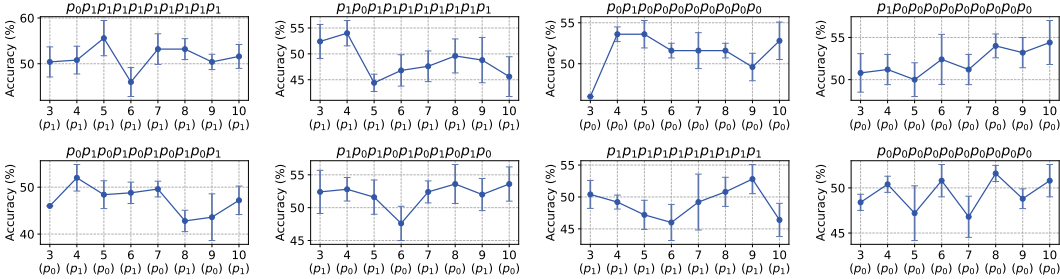


Figure 17: Average ratio of *consensus clusters* (unique answers among multiple agents) with different numbers (2~10) of agents under each round of 3-round collaborative strategies, using ChatGPT.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0p_0p_0p_0p_0p_0p_0$	0.030	0.323	0.000
$p_1p_0p_0p_0p_0p_0p_0p_0p_0$	0.000	0.070	0.161
$p_0p_1p_0p_0p_0p_0p_0p_0p_0$	0.101	0.332	0.000
$p_1p_0p_1p_0p_1p_0p_1p_0p_1$	0.000	0.077	0.871
$p_0p_1p_0p_1p_0p_1p_0p_1p_0p_1$	0.051	0.062	0.000
$p_1p_0p_1p_1p_1p_1p_1p_1p_1$	0.000	0.021	0.630
$p_0p_1p_1p_1p_1p_1p_1p_1p_1$	0.431	0.176	0.063
$p_1p_1p_1p_1p_1p_1p_1p_1p_1$	0.000	0.000	0.027

Table 12: One-way ANOVA analysis of the results in Figure 4, 18, 19 (different rounds), using *ChatGPT*.Figure 18: Accuracy of different (3~10) rounds of collaboration within 3-agent society S_2 (1 easy and 2 overconfident agents) on MMLU, using *ChatGPT*.Figure 19: Accuracy of different (3~10) rounds of collaboration within 3-agent society S_2 (1 easy and 2 overconfident agents) on Chess Move Validity, using *ChatGPT*.

Different Rounds of Collaboration. As observed from Table 12, we find that the impact of rounds significantly relies on the employed collaborative strategy. For MMLU and Chess Move Validity, collaborative strategies where p -values < 0.05 are $\{p_0p_1p_1p_0, p_0p_1p_1p_1, p_1p_0p_1p_0, p_1p_0p_1p_1\}$ and $\{p_0p_1p_1p_0, p_0p_1p_1p_1, p_1p_0p_1p_1, p_1p_1p_0p_0, p_1p_1p_0, p_1, p_1p_1p_1p_0\}$. We also increase the rounds of collaboration, from 3 to 10, and present the results in Figure 18, 19. We find that although there would be some fluctuations in performance if we scale up the round of collaboration, the outperformance is not obvious enough. While increasing rounds of collaboration will result in more consumption of tokens, which is not economic. Thus we infer that the 3-round collaboration is relatively optimal considering both performance and cost.

Furthermore, as seen from Figure 7, the strategy after a round of debate tends to yield fewer consensus clusters compared to the preceding rounds. Conversely, the strategy subsequent with a round of reflection at the same juncture will increase consensus clusters. Adding an extra round of debate at this juncture, as the conclusions in §4, is not anticipated to bring about a discernible enhancement in performance. This confirms the efficacy of the *early-stopping mechanism* implemented in Liu et al. (2023b), drawing inspiration from Byzantine Consensus theory (Castro & Liskov, 1999).

Moreover, we scrutinize the consensus reaching of these strategies in three rounds where p -values are below 0.05, as shown in Figure 7. Also seen from Figure 7 and Figure 18, 4, 19, it becomes apparent that these collaborative strategies exhibit substantial fluctuations in consensus reaching, demonstrating notably low answer consistency. For $p_0p_0p_0p_0$ on Chess Move Validity, although continuous reflection results in a gradual increase in the number of consensus clusters, a more stable trend with smaller fluctuations renders it less sensitive to the rounds of collaboration. Conversely, collaborative strategies where p -values > 0.05 often display higher levels of answer consistency.

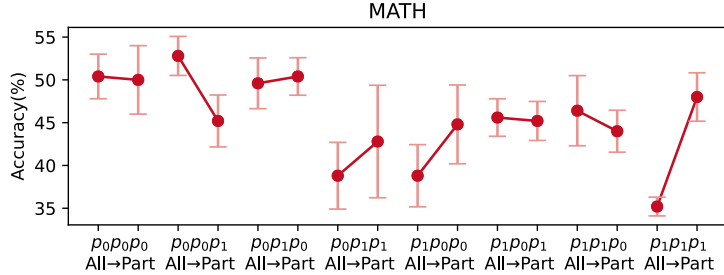


Figure 20: The effect on the accuracy of whether all agents in a society execute the same thinking pattern in one round on MATH, using *ChatGPT*. “All” and “Part” respectively refer to all agents applying the same and different thinking pattern(s) in one round. The significance test is shown in Table 13 at Appendix F.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.402	0.856	0.147
$p_0p_0p_1$	0.007	0.002	0.001
$p_0p_1p_0$	0.550	0.641	0.002
$p_0p_1p_1$	-	0.276	0.000
$p_1p_0p_0$	-	0.051	-
$p_1p_0p_1$	-	0.784	0.000
$p_1p_1p_0$	0.014	0.294	0.172
$p_1p_1p_1$	1.000	0.000	0.347

Table 13: One-way ANOVA analysis of the results of Figure 5 (other collaborative strategies), using *ChatGPT*. ‘-’: It doesn’t pass homogeneity test for variance.

Other Collaborative Strategies. We show the results of all agents in a society executing the same or inconsistent thinking pattern(s) at one round in Figure 20. Seen from Table 13, we observe pronounced impacts of keeping a consistent thinking pattern on Chess Move Validity, while its influence on MMLU and MATH is less significant.

G A SOCIAL PSYCHOLOGY VIEW ON CONFORMITY, CONSENSUS REACHING, AND GROUP DYNAMICS

G.1 CONFORMITY AND CONSENSUS REACHING

Figures 6, 28, 37, 65, and 51 illustrate the conformity. Figures 7, 28, 37, 65, and 51 illustrate the consensus. This section provides a detailed explanation of the methodologies used to calculate both conformity and consensus.

For conformity, we solely focus on agents actively engaging in debate, disregarding those in reflection during a given round. Let the answer of the i -th agent at j -th round be denoted as $a_{i,j}$. For the k -th agent at j -th round, if “Frequency($\{a_{i,j-1} | i \in [1, n]\}) = a_{k,j}$ ”, we identify this as the occurrence of conformity by agent k at j -th round, where Frequency(\cdot) represents the most frequently given answer (excluding instances where all answers occur only once, as such cases are considered as non-conformity). Additionally, we categorize the correctness of answers both before and after conformity into four cases, with ‘True’ denoting correct and ‘False’ denoting incorrect.

For consensus, we examine the evolution of the number of distinct answers (*i.e.*, consensus clusters) with increasing rounds of collaboration. Let the answer of the i -th agent at time j be denoted as $a_{i,j}$. For the j -th round, consensus clusters is defined as $\|\text{Set}(\{a_{i,j} | i \in [1, n]\})\|$, where $\|\text{Set}(\cdot)\|$ represents the count of different answers. This computational approach has been utilized in the analysis presented in Figures 17, 16, 60, 59, 46, 45.

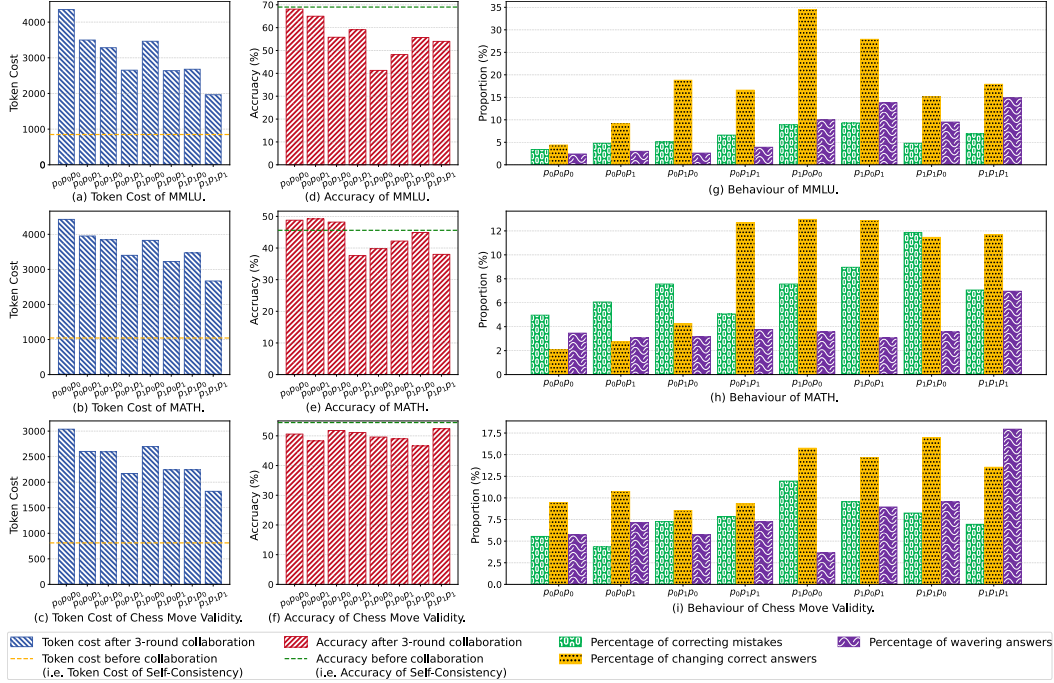


Figure 21: The percentage of different behaviors under different collaborative strategies, using *ChatGPT*. Figure (a-c) & (d-f) respectively show the token cost and accuracy of different strategies before and after 3-round collaboration. Figure (g-i) present the percentage of different behavioral features (mainly analyzed by the change of answer correctness) (Zhang et al., 2023c;b) under different collaborative strategies. All results are summarized across all societies. The results on other LLMs are shown in Figure 30, 39, 53, 67 at Appendix H.

G.2 GROUP DYNAMICS

We seek to elucidate how performance is impacted by group dynamics, *i.e.*, the patterns of interaction between group members and different processes that may occur within a social group. Diving into the intricacies of collaboration, each agent generates four answers, including the initial answer without collaboration, as shown in Figure 2(d). To determine the answer for each round, we employ the majority vote (Cobbe et al., 2021; Li et al., 2022). Given ‘*T*’ and ‘*F*’ respectively denoting a round that yields a correct and an incorrect answer, we could obtain $2^4=16$ possible answer sequences over the four rounds. We select 10 sequences⁶ of them and categorize them into 3 groups: **Correcting Mistakes** (*FFFT*, *FFTT*, *FTTT*), **Changing Correct Answers** (*TFFF*, *TTF*, *TFTF*), and **Wavering Answers** (*FTFT*, *FTTF*, *TFTF*, *FTFT*). Particularly, **Wavering Answers** resemble model hallucination (Rawte et al., 2023; Zhang et al., 2023d; Ji et al., 2023; Luo et al., 2024) due to the occurrence of self-contradictory answers. Our categorization is under society-agnostic collaborative strategies, considering the performance variance between societies is negligible. From the results on *ChatGPT* shown in Figure 21, and on other LLMs shown in Appendix H, we summarize the following findings:

(1) **Debate-initial/dominant collaborative strategies are generally effective.** As seen from the red bars in Figure 21 30, 39, 53, 67(d-f), we find that the collaborative strategies starting from or dominant with debate p_0 are more effective than other, and mostly outperform self-consistency, even though they cost more tokens (seen from blue bars).

(2) **Reflection experiences greater instability (a heightened risk of model hallucination).** As observed from the purple bars in Figure 21 30, 39, 53, 67(g-h), comparing $p_i p_j p_0$ & $p_i p_j p_1$; $p_i p_0 p_j$ & $p_i p_1 p_j$, $p_i p_j p_0$ and $p_i p_0 p_j$ are more likely to wavering answers than $p_i p_j p_1$ and $p_i p_1 p_j$, demonstrating that reflection is more likely to cause model hallucination than debate.

⁶The selected 10 sequences adhere to patterns: (1) $[F]_{i>0}[T]_{j>0}$, *e.g.*, *FFFT*; (2) $[T]_{i>0}[F]_{j>0}$, *e.g.*, *TFFF*; (3) $[TF]_{i\geq 0}[FT]_{j\geq 0}$, *e.g.*, *FTFT*, where $[\cdot]_i$, $[\cdot]_j$ respectively denotes repetition for i , j times.

H ANALYSIS ON DIFFERENT BACKBONE LLMs

To make the findings in this paper more general, we also implement all the experiments with some other open-resource backbone LLMs, such as **LlaMA2 Chat 13B** (Touvron et al., 2023), **LlaMA2 Chat 70B** (Touvron et al., 2023), **Qwen 72B** (Bai et al., 2023) and **Mixtral 8×7B** (Jiang et al., 2023; 2024).

H.1 LLAMA2 CHAT 13B

Analysis on Machine Social Collaboration. We present the **main results** and **significance tests** of societies and strategies on LlaMA2 Chat 13B in Table 14, 15, 16. We present the **word clouds** of LlaMA2 Chat 13B in Figure 22, and **proportion of agents with different traits changing answers in different societies** on LlaMA2 Chat 13B in Figure 23. Furthermore, we demonstrate that the tasks with different subjects and difficulty display varying sensitivity to collaborative strategies, as presented with **radar maps** on LlaMA2 Chat 13B in Figure 24.

	Metric (Strategy)	Society	Collaborative Strategy								Metric (Society)	
			$p_0p_0p_0$	$p_0p_0p_1$	$p_0p_1p_0$	$p_0p_1p_1$	$p_1p_0p_0$	$p_1p_0p_1$	$p_1p_1p_0$	$p_1p_1p_1$	Cost ↓	W-T ↑
MMLU	Acc ↑	S_1	37.2±5.9	47.2±3.9	48.4±3.9	46.0±5.7	47.2±2.3	46.8±2.7	45.2±4.4	46.8±3.0	7447	35
		S_2	38.4±4.6	42.8±3.9	43.6±3.6	45.2±3.6	44.8±4.6	47.2±3.9	44.4±6.2	42.8±3.4	7413	33
		S_3	36.0±3.7	44.8±3.0	44.8±4.8	46.4±1.7	41.6±4.3	46.4±2.2	43.2±6.6	42.4±3.3	7370	33
		S_4	34.8±2.7	42.4±5.0	42.0±4.5	44.0±2.8	40.4±3.0	43.6±3.9	40.8±3.0	41.6±2.6	7423	35
	Cost ↓	All	11429	9476	8166	6419	8452	5734	5733	3900	-	-
	W-T ↑	All	-	20	20	20	18	20	19	19	-	-
MATH	Acc ↑	S_1	5.2±2.3	6.8±2.3	5.6±2.6	5.6±2.6	4.8±3.0	4.4±1.7	5.6±3.9	3.2±1.1	8639	24
		S_2	5.2±3.6	5.2±3.4	6.0±2.0	6.8±1.8	6.0±0.0	6.8±1.8	6.8±1.1	4.8±1.1	8451	22
		S_3	6.8±1.8	6.8±3.0	6.8±3.4	6.0±2.8	5.2±1.8	5.2±1.8	6.0±3.7	3.6±1.7	8501	16
		S_4	4.8±2.3	6.8±3.4	7.2±1.1	5.6±2.2	5.6±1.7	5.2±2.3	5.2±3.6	4.0±1.4	8475	28
	Cost ↓	All	10655	9508	9501	7900	9319	7761	7800	5687	-	-
	W-T ↑	All	-	15	16	13	13	11	13	9	-	-
Chess Move Validity	Acc ↑	S_1	16.4±3.0	7.2±3.0	9.2±2.3	2.8±1.8	8.8±3.0	4.8±2.3	9.2±4.4	2.0±2.8	3754	2
		S_2	11.6±5.2	8.0±1.4	10.8±4.2	2.8±1.8	11.6±2.6	6.0±3.2	10.8±5.0	3.6±2.6	3725	10
		S_3	14.8±3.0	8.4±4.8	10.0±4.2	5.2±1.1	14.0±4.5	6.8±3.0	9.6±6.2	2.8±3.0	3678	5
		S_4	16.0±4.2	6.8±2.7	12.4±6.2	4.0±2.5	10.0±4.2	7.2±6.7	10.0±3.2	4.0±2.5	3647	4
	Cost ↓	All	4889	4123	4061	3324	4045	3293	3292	2581	-	-
	W-T ↑	All	-	2	4	0	7	1	7	0	-	-

Table 14: The impact of eight different collaborative strategies on the performance of three datasets across distinct societies (using *LlaMA2-chat-13B*). The significances test on societies and strategies are respectively shown in Table 15, 16. The experiments of comparison with the single LLM agent is shown in Figure 30(a)-(f).

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.611	0.632	0.251
$p_0p_0p_1$	0.252	0.791	0.854
$p_0p_1p_0$	0.142	0.714	0.706
$p_0p_1p_1$	0.755	0.839	0.164
$p_1p_0p_0$	0.039	0.789	0.175
$p_1p_0p_1$	0.318	0.277	0.809
$p_1p_1p_0$	0.585	0.884	0.959
$p_1p_1p_1$	0.071	0.310	0.672

Table 15: One-Way ANOVA results for the impact of society on accuracy with fixed collaborative strategy, based on experiments from Table 14 using *LlaMA2-chat-13B*.

Analysis on Different Numbers of Agents. We present the significance test for different numbers of agents with LlaMA2 Chat 13B in Table 17. We also show the performance varying from agent numbers in Figure 25.

Society	MMLU p-value	MATH p-value	Chess Move Validity p-value
S_1	0.006	0.548	0.000
S_2	0.129	0.664	0.000
S_3	0.005	0.518	0.000
S_4	0.009	0.490	0.001

Table 16: One-Way ANOVA results for the impact of collaborative strategy on accuracy with fixed society, based on experiments from Table 14 using *LlaMA-13B-Chat*.

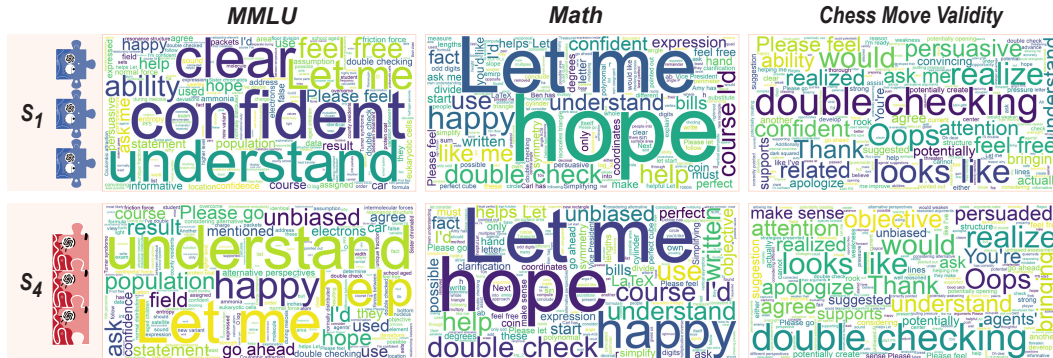


Figure 22: Comparative word clouds on three datasets in societies S_1 and S_4 , using *LlaMA2-13B-chat*. Society S_1 features 3 overconfident agents, while society S_4 comprises 3 easy-going agents.

Collaborative Strategy	MMLU p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.186	0.001
$p_0p_0p_1$	0.019	0.000
$p_0p_1p_0$	0.175	0.000
$p_0p_1p_1$	0.010	0.178
$p_1p_0p_0$	0.023	0.001
$p_1p_0p_1$	0.002	0.005
$p_1p_1p_0$	0.098	0.005
$p_1p_1p_1$	0.004	0.002

Table 17: One-way ANOVA analysis of the results in Figure 25 (different numbers of agents), using *LlaMA2-chat-13B*.

Analysis on Different Rounds. We present the significance test for different rounds of collaboration with LlaMA2 Chat 13B in Table 18. We also show the performance varying from collaboration rounds in Figure 26.

Analysis on Other Collaborative Strategies. We present the significance test for other collaborative strategies (executing the same or hybrid thinking patterns in a certain round) with LlaMA2 Chat 13B in Table 19. We also show the performance varying from other strategies in Figure 27.

A Social Psychology View on Conformity, Consensus Reaching and Group Dynamics. We then show the variation of answer correctness in the situation of conformity in Figure 28; and the quantity of consensus clusters among 3-agent answers in Figure 29. We present group dynamics reflected by different answer-changing behaviors on LlaMA2 Chat 13B in Figure 30.

Collaborative Strategy	MMLU p-value	Chess Move Validity p-value
$p_0p_0p_0p_0$	0.000	0.361
$p_0p_0p_0p_1$	0.111	0.598
$p_0p_0p_1p_0$	0.082	0.335
$p_0p_0p_1p_1$	0.529	0.076
$p_0p_1p_0p_0$	0.293	0.176
$p_0p_1p_0p_1$	0.641	0.259
$p_0p_1p_1p_0$	0.536	0.026
$p_0p_1p_1p_1$	0.812	0.052
$p_1p_0p_0p_0$	0.010	0.629
$p_1p_0p_0p_1$	0.547	0.029
$p_1p_0p_1p_0$	0.749	0.055
$p_1p_0p_1p_1$	0.600	0.007
$p_1p_1p_0p_0$	0.605	0.009
$p_1p_1p_0p_1$	0.988	0.012
$p_1p_1p_1p_0$	0.889	0.097
$p_1p_1p_1p_1$	0.742	0.884

Table 18: One-way ANOVA analysis of the results in Figure 26 (different rounds), using *LlaMA2-chat-13B*.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.419	0.659	0.203
$p_0p_0p_1$	0.441	1.000	0.141
$p_0p_1p_0$	0.086	0.074	0.264
$p_0p_1p_1$	0.001	0.161	0.347
$p_1p_0p_0$	0.030	-	0.000
$p_1p_0p_1$	0.003	0.004	0.380
$p_1p_1p_0$	0.070	0.001	0.005
$p_1p_1p_1$	0.169	0.008	0.128

Table 19: One-way ANOVA analysis of the results in Figure 27 (other collaborative strategies), using *LlaMA2-chat-13B*.

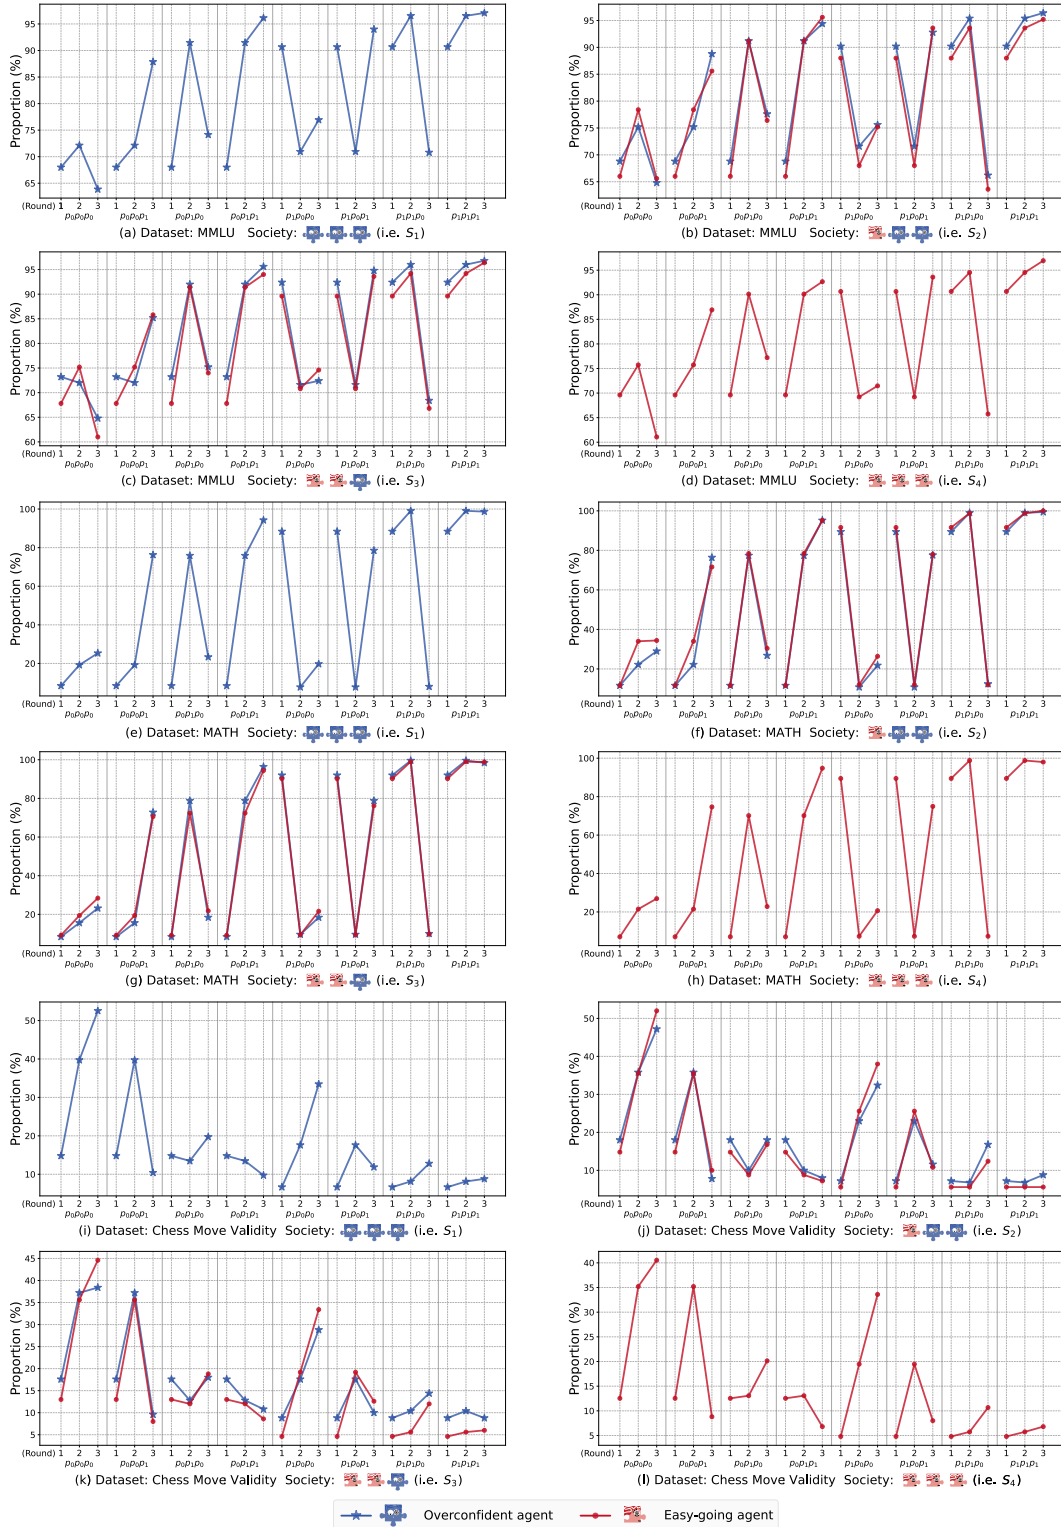


Figure 23: Proportion of agents with different traits changing answers in societies S_1 and S_4 , using *LlaMA2-13B-chat*. Society S_1 features three overconfident agents, while society S_4 comprises three easy-going agents.

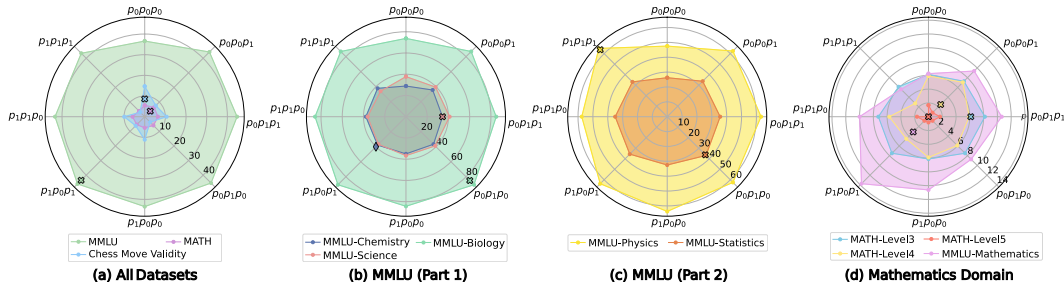


Figure 24: Illustration of different collaborative strategies impacting accuracy diversely on the tasks considering varied *subjects* and *difficulty*, using *LlaMA2-13B-chat*. The symbol ‘⚡’ represents that there is at least one collaborative strategy whose accuracy is better than self-consistency, while the symbol ‘⚡’ indicates that there is no collaborative strategy whose accuracy is worse than self-consistency. Both of these symbols represent the accuracy of self-consistency. The accuracy under each collaborative strategy is a summation within all 3-agent societies.

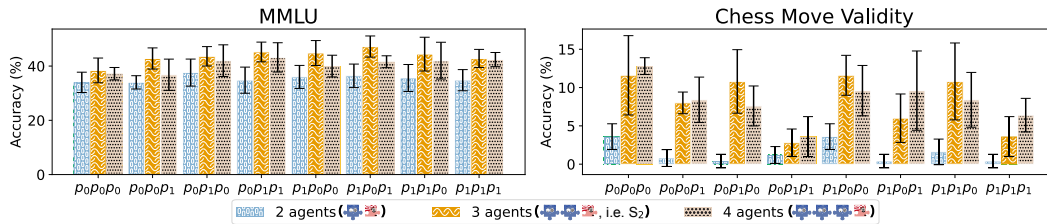


Figure 25: Accuracy of different number of agents under different collaborative strategies, on *LlaMA2-13B-chat*. The significance test is shown in Table 17.

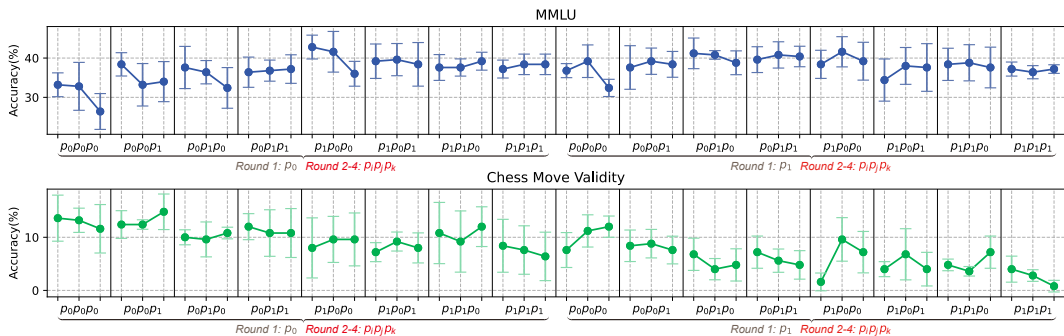


Figure 26: Accuracy at round 2,3,4 within 4-round collaborative societies, where the thinking pattern of round 1 is fixed (p_0 or p_1), using *LlaMA2-13B-chat*. The significance test is shown in Table 18.

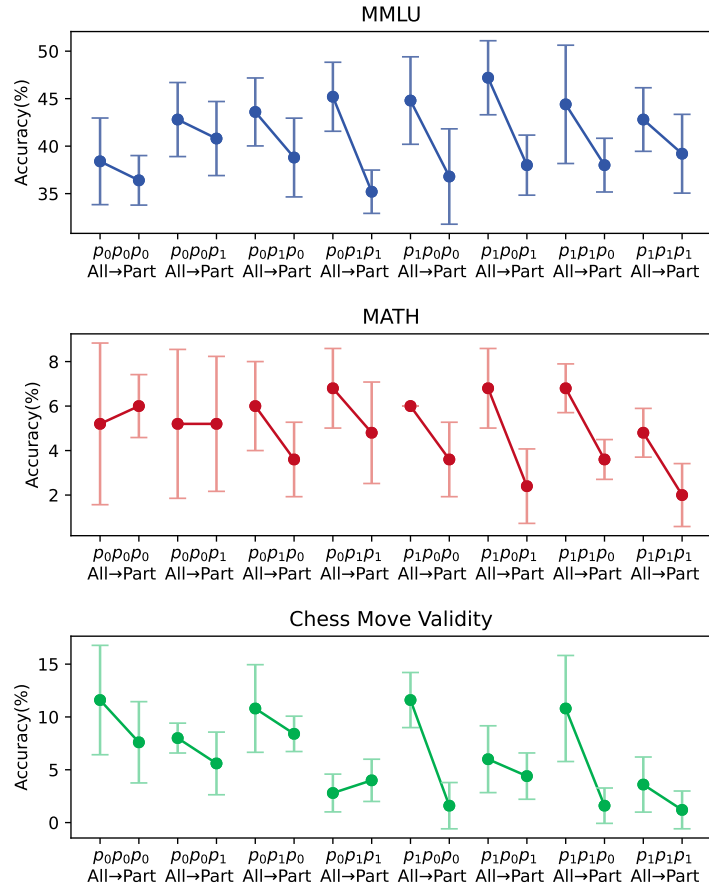


Figure 27: The effect on the accuracy of whether all agents in society execute the same thinking pattern in one round, using *LlaMA2-13B-chat*. “All” and “Part” refer to all agents applying the same thinking pattern and different thinking patterns in one round respectively. The significance test is shown in Table 19.

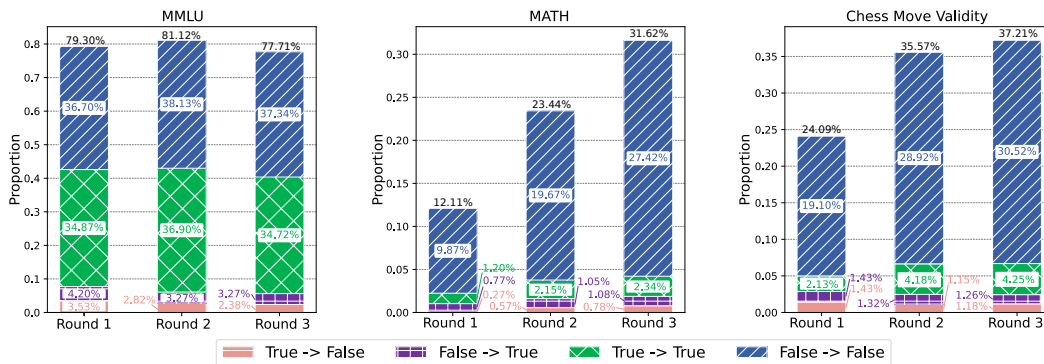


Figure 28: Variation of answer correctness in the situation of conformity, using *LlaMA2-13B-chat*, where *conformity brings about benefits*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) > \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$; *conformity brings about detriments*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) < \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$.

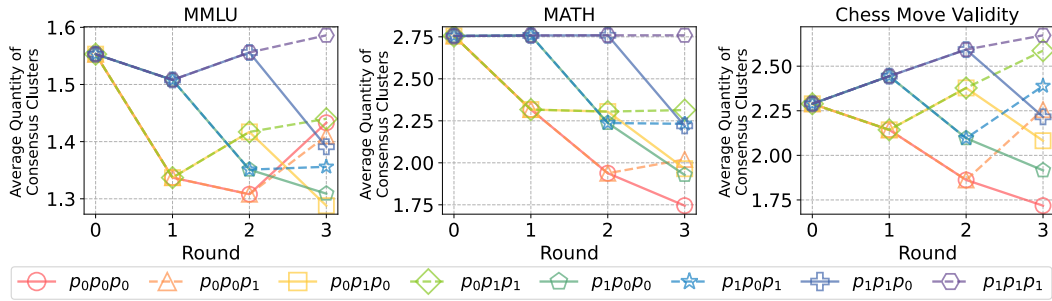


Figure 29: Average quantity of *consensus clusters* (i.e., *unique answers among multiple agents*) under different rounds of collaboration with 3-round collaborative strategies, on *LlaMA2-13B-chat*. *Smaller quantity of consensus clusters, more easier it is to reach a consensus*. Round 0 is equal to self-consistency.

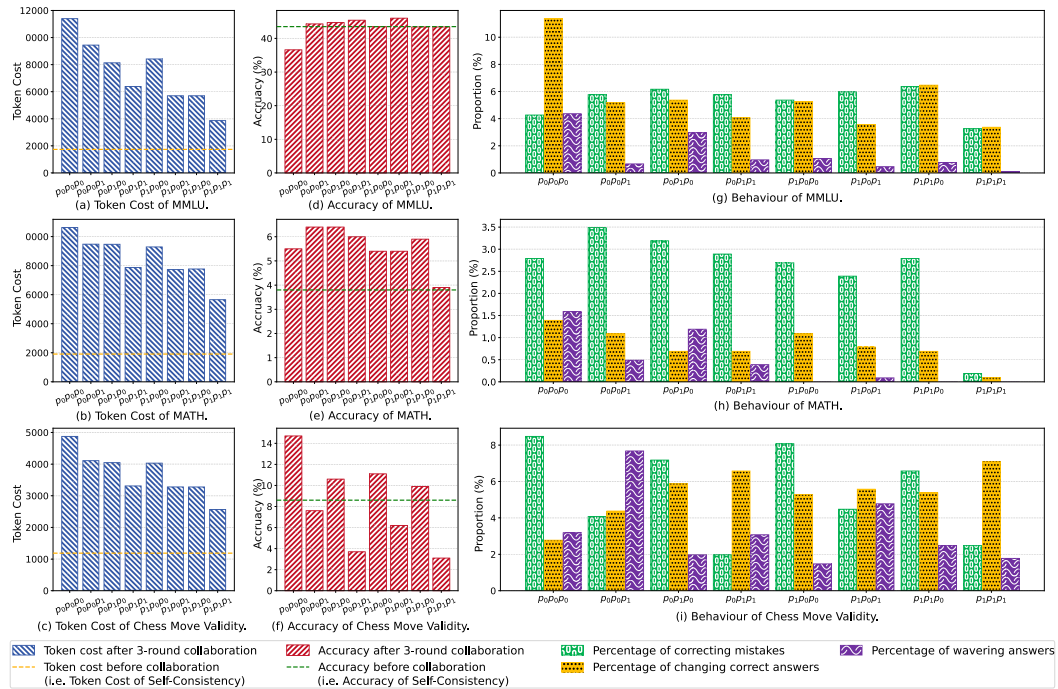


Figure 30: The percentage of different behaviors under different collaborative strategies, using *LlaMA2-13B-chat*. Figure (a-c) & (d-f) respectively show the token cost and accuracy of different strategies before and after 3-round collaboration. Figure (g-i) present the percentage of different behavioral features (mainly analyzed by the change of answer correctness) (Zhang et al., 2023c;b) under different collaborative strategies. All results are summarized across all societies.

H.2 LLAMA2 CHAT 70B

Analysis on Machine Social Collaboration. We present the **main results** and **significance tests** of societies and strategies on LLaMA2 Chat 70B in Table 20, 21, 22. We present the **word clouds** of LLaMA2 Chat 70B in Figure 31, and **proportion of agents with different traits changing answers in different societies** on LLaMA2 Chat 70B in Figure 32. Furthermore, we demonstrate that the tasks with different subjects and difficulty display varying sensitivity to collaborative strategies, as presented with **radar maps** on LLaMA2 Chat 70B in Figure 33.

	Metric (Strategy)	Society	Collaborative Strategy								Metric (Society)	
			$p_0p_0p_0$	$p_0p_0p_1$	$p_0p_1p_0$	$p_0p_1p_1$	$p_1p_0p_0$	$p_1p_0p_1$	$p_1p_1p_0$	$p_1p_1p_1$	Cost ↓	W-T ↑
MMLU	Acc ↑	S_1	40.8±2.7	43.6±3.9	36.0±2.8	38.4±3.3	35.6±4.3	35.6±2.6	30.4±4.3	24.0±5.7	6915	7
		S_2	44.4±3.9	49.2±4.6	45.2±3.9	42.0±0.0	34.4±4.3	34.4±8.3	31.6±8.4	25.6±3.6	6946	11
		S_3	44.0±5.5	45.6±4.6	39.2±2.7	42.8±3.0	35.2±5.4	32.4±4.3	28.0±7.3	25.6±5.2	6931	8
		S_4	47.6±4.1	48.0±5.1	46.0±6.3	45.2±3.9	26.8±3.6	30.8±6.9	32.8±1.8	33.6±6.2	6936	8
	Cost ↓	All	10811	8608	7904	6177	7535	5410	5287	3722	-	-
W-T ↑	All	-	16	5	11	1	0	1	0	-	-	
MATH	Acc ↑	S_1	8.4±3.6	10.4±3.9	9.2±1.1	4.0±2.5	9.2±4.2	8.4±4.3	6.8±2.7	3.6±1.7	7000	16
		S_2	8.0±2.5	9.6±2.6	8.8±3.0	6.4±2.6	7.2±4.4	6.8±1.1	8.4±4.3	4.8±2.3	7013	19
		S_3	8.4±4.6	7.2±3.9	8.4±3.6	5.6±3.6	7.2±1.8	7.2±4.8	6.8±3.0	0.8±1.1	7157	15
		S_4	6.0±2.0	7.2±1.8	6.0±2.0	4.0±2.0	5.2±3.0	6.8±1.1	8.8±4.4	3.6±2.6	6934	23
	Cost ↓	All	9465	7850	7662	6294	7520	6302	6382	4734	-	-
W-T ↑	All	-	14	14	5	13	9	14	4	-	-	
Chess Move Validity	Acc ↑	S_1	20.4±6.2	16.8±3.6	17.2±4.2	8.4±2.2	21.2±5.8	10.8±3.0	10.4±1.7	4.8±3.0	3563	7
		S_2	18.4±4.8	9.6±3.6	13.2±1.1	5.6±2.2	14.4±3.9	7.2±3.0	13.2±3.4	4.0±2.8	3557	4
		S_3	18.4±6.5	11.2±3.0	12.0±5.8	8.0±2.0	20.8±4.6	8.4±4.3	12.8±2.7	2.8±3.4	3629	7
		S_4	15.2±4.2	11.6±2.2	15.2±2.3	10.4±1.7	18.0±4.7	8.0±4.7	10.8±2.7	5.2±2.3	3679	12
	Cost ↓	All	4778	3947	3830	3082	4139	3314	3259	2508	-	-
W-T ↑	All	-	4	6	2	13	1	4	0	-	-	

Table 20: The impact of eight different collaborative strategies on the performance of three datasets across distinct societies (using *LlaMA2-chat-70B*). The significances test on societies and strategies are respectively shown in Table 21, 22. The experiments of comparison with the single LLM agent is shown in Figure 39(a)-(f).

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.122	0.621	0.532
$p_0p_0p_1$	0.251	0.291	0.014
$p_0p_1p_0$	0.004	0.248	0.185
$p_0p_1p_1$	0.018	0.430	0.015
$p_1p_0p_0$	0.020	0.381	0.132
$p_1p_0p_1$	0.601	0.854	0.506
$p_1p_1p_0$	0.641	0.750	0.282
$p_1p_1p_1$	0.044	0.037	0.585

Table 21: One-Way ANOVA results for the impact of society on accuracy with fixed collaborative strategy, based on experiments from Table 20 using *LlaMA2-chat-70B*.

Society	MMLU p-value	MATH p-value	Chess Move Validity p-value
S_1	0.000	0.013	0.000
S_2	0.000	0.297	0.000
S_3	0.000	0.040	0.000
S_4	0.000	0.056	0.000

Table 22: One-Way ANOVA results for the impact of collaborative strategy on accuracy with fixed society, based on experiments from Table 20 using *LlaMA-70B-Chat*.

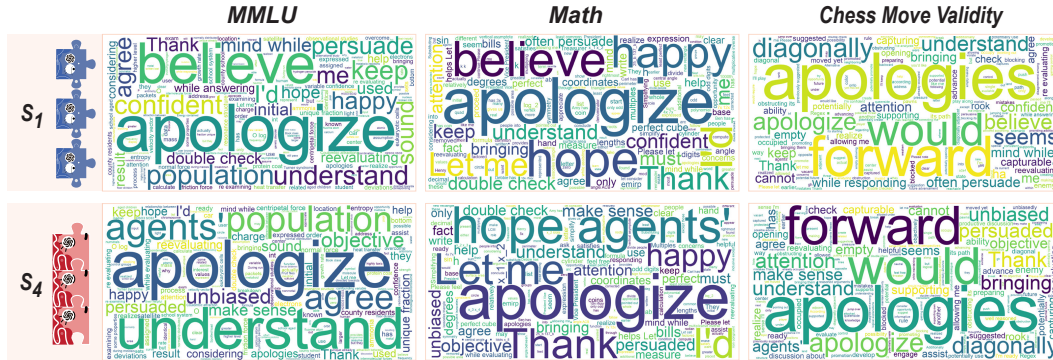


Figure 31: Comparative word clouds on three datasets in societies S_1 and S_4 , using *LlaMA2-70B-chat*. Society S_1 features 3 overconfident agents, while society S_4 comprises 3 easy-going agents.

Analysis on Different Numbers of Agents. We present the significance test for different numbers of agents with LLaMA2 Chat 70B in Table 23. We also show the performance varying from agent numbers in Figure 34.

Collaborative Strategy	MMLU p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.481	0.006
$p_0p_0p_1$	0.000	0.001
$p_0p_1p_0$	0.000	0.000
$p_0p_1p_1$	-	0.023
$p_1p_0p_0$	0.001	0.035
$p_1p_0p_1$	0.003	0.000
$p_1p_1p_0$	0.002	0.036
$p_1p_1p_1$	0.024	0.423

Table 23: One-way ANOVA analysis of the results of Figure 34 (different numbers of agents), using *LlaMA2-chat-70B*.

Analysis on Different Rounds. We present the significance test for different rounds of collaboration with LLaMA2 Chat 70B in Table 24. We also show the performance varying from collaboration rounds in Figure 35.

Analysis on Other Collaborative Strategies. We present the significance test for other collaborative strategies (executing the same or hybrid thinking patterns in a certain round) with LLaMA2 Chat 70B in Table 25. We also show the performance varying from other strategies in Figure 36.

A Social Psychology View on Conformity, Consensus Reaching and Group Dynamics. We then show the variation of answer correctness in the situation of conformity in Figure 37; and the quantity of consensus clusters among 3-agent answers in Figure 38. We present group dynamics reflected by different answer-changing behaviors on LLaMA2 Chat 70B in Figure 39.

Collaborative Strategy	MMLU p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.034	0.545
$p_0p_0p_0p_1$	0.008	0.019
$p_0p_0p_1p_0$	0.020	0.004
$p_0p_0p_1p_1$	0.643	0.004
$p_0p_1p_0p_0$	0.045	0.034
$p_0p_1p_0p_1$	0.164	0.902
$p_0p_1p_1p_0$	0.046	0.006
$p_0p_1p_1p_1$	0.082	0.000
$p_1p_0p_0p_0$	0.706	0.207
$p_1p_0p_0p_1$	0.449	0.494
$p_1p_0p_1p_0$	0.782	0.095
$p_1p_0p_1p_1$	0.664	0.070
$p_1p_1p_0p_0$	0.360	0.041
$p_1p_1p_0p_1$	0.391	0.018
$p_1p_1p_1p_0$	0.394	0.088
$p_1p_1p_1p_1$	0.031	0.033

Table 24: One-way ANOVA analysis of the results in Figure 35 (different rounds), using *LlaMA2-chat-70B*.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.029	0.296	0.004
$p_0p_0p_1$	0.005	0.020	0.724
$p_0p_1p_0$	0.018	0.191	0.000
$p_0p_1p_1$	0.000	0.809	0.684
$p_1p_0p_0$	0.894	0.503	0.045
$p_1p_0p_1$	0.747	0.050	0.328
$p_1p_1p_0$	0.928	0.007	0.001
$p_1p_1p_1$	0.004	1.000	0.557

Table 25: One-way ANOVA analysis of the results in Figure 36 (other collaborative strategies), using *LlaMA2-chat-70B*.

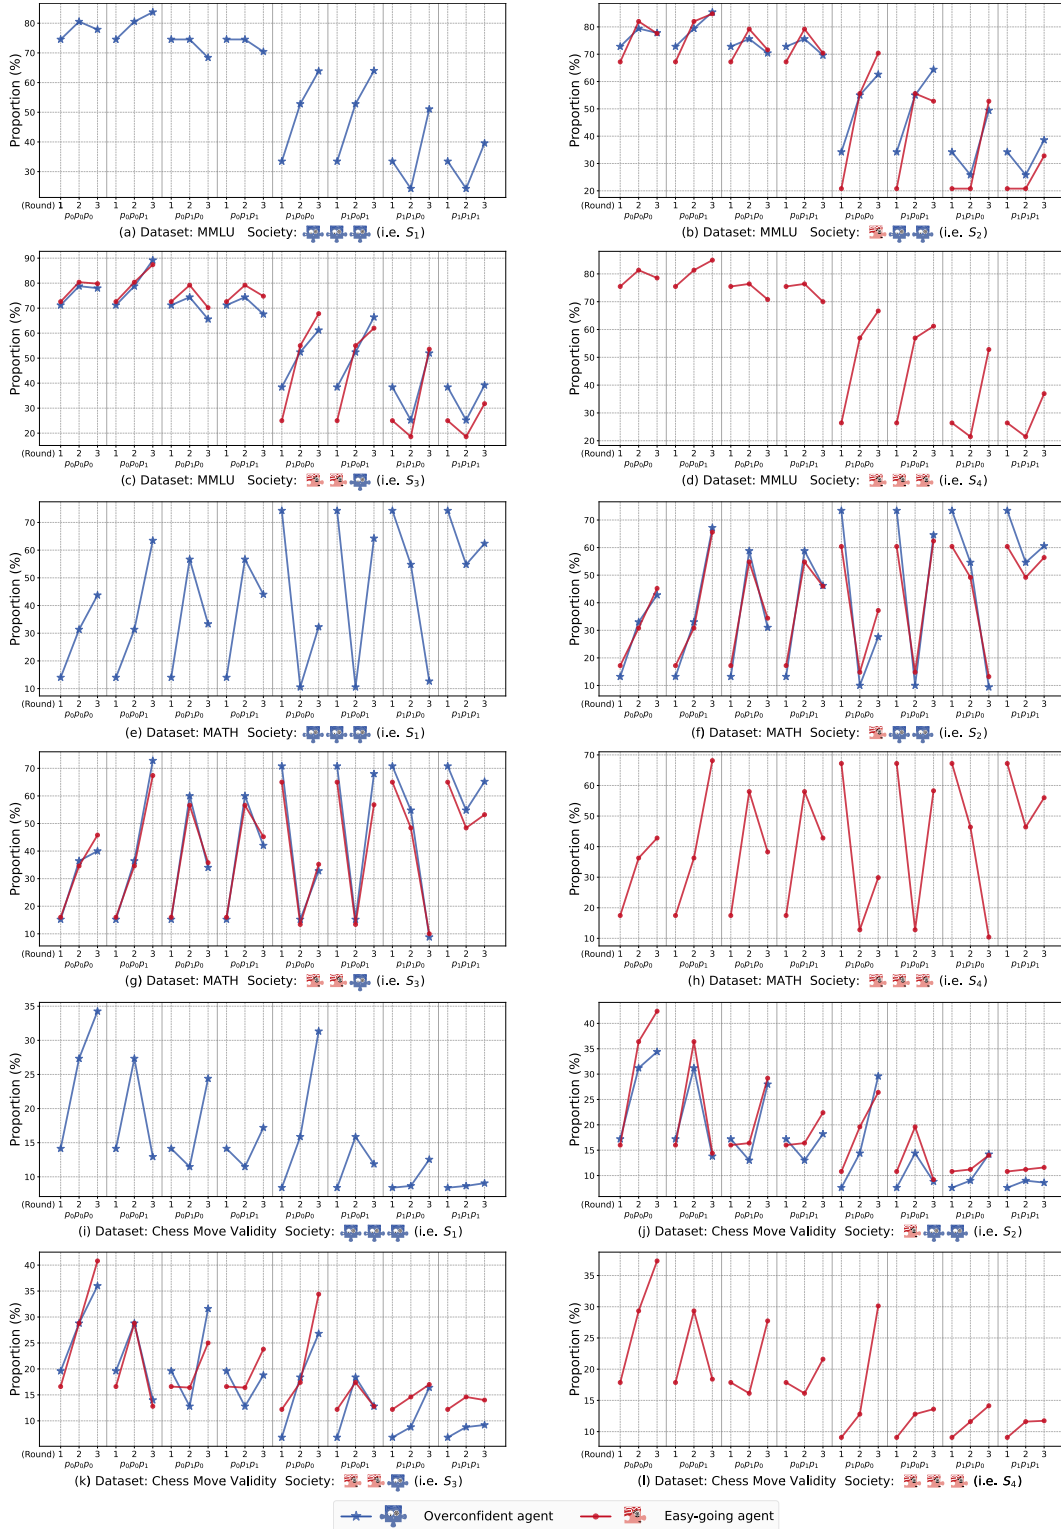


Figure 32: Proportion of agents with different traits changing answers in societies S_1 and S_4 , using *LlaMA2-70B-chat*. Society S_1 features three overconfident agents, while society S_4 comprises three easy-going agents.

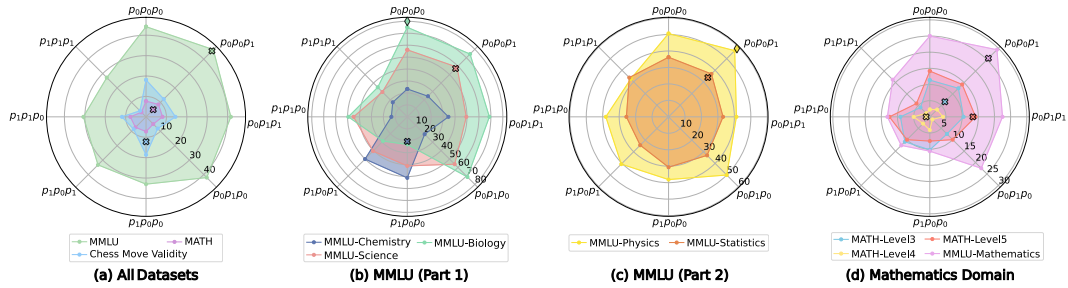


Figure 33: Illustration of different collaborative strategies impacting accuracy diversely on the tasks considering varied *subjects* and *difficulty*, using *LlaMA2-70B-chat*. The symbol ‘ \bowtie ’ represents that there is at least one collaborative strategy whose accuracy is better than self-consistency, while the symbol ‘ \diamond ’ indicates that there is no collaborative strategy whose accuracy is worse than self-consistency. Both of these symbols represent the accuracy of self-consistency. The accuracy under each collaborative strategy is a summation within all 3-agent societies.

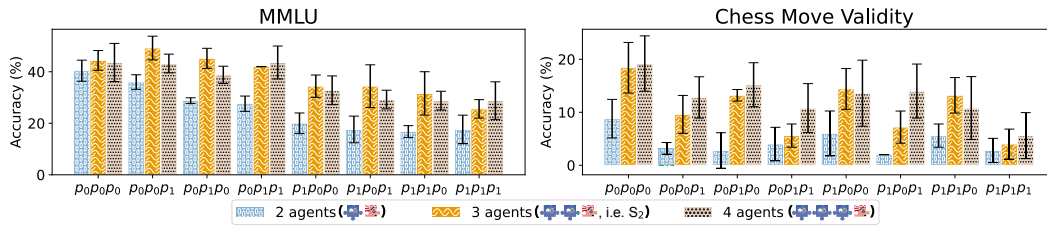


Figure 34: Accuracy of different numbers of agents under different collaborative strategies, on *LlaMA2-70B-chat*. The significance test is shown in Table 23.

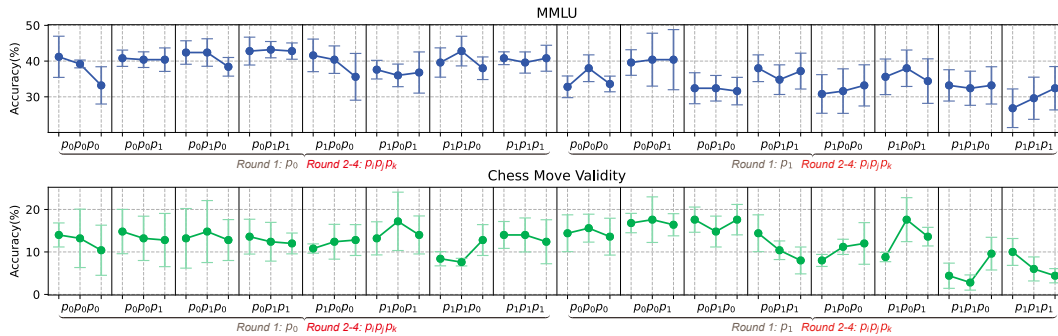


Figure 35: Accuracy at round 2,3,4 within 4-round collaborative societies, where the thinking pattern of round 1 is fixed (p_0 or p_1), using *LlaMA2-70B-chat*. The significance test is shown in Table 24.

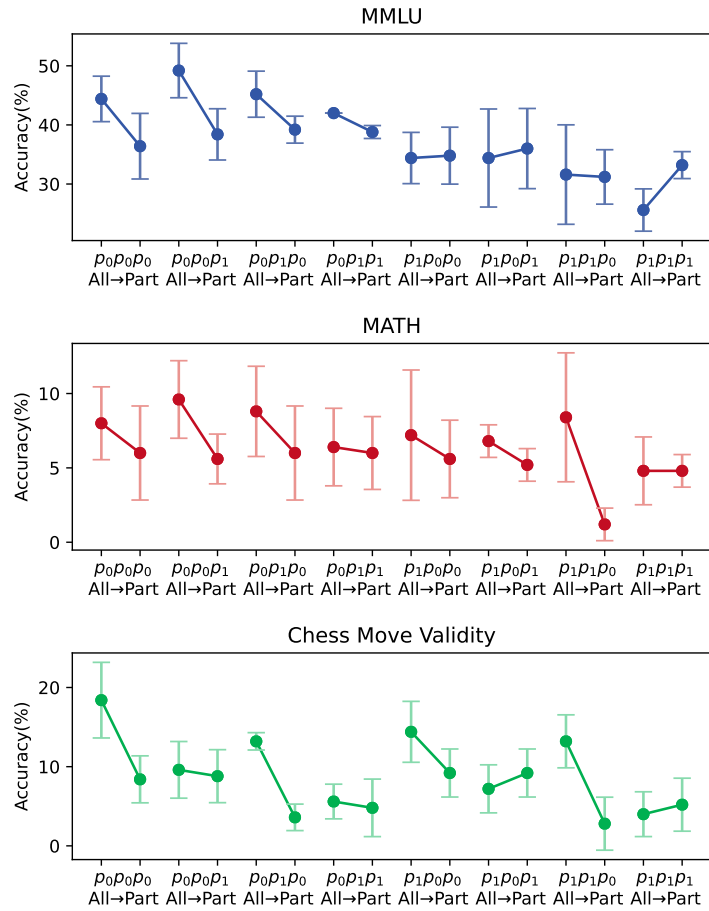


Figure 36: The effect on the accuracy of whether all agents in society execute the same thinking pattern in one round, using *LlaMA2-70B-chat*. “All” and “Part” refers to all agents applying the same thinking pattern and different thinking patterns in one round respectively. The significance test is shown in Table 25.

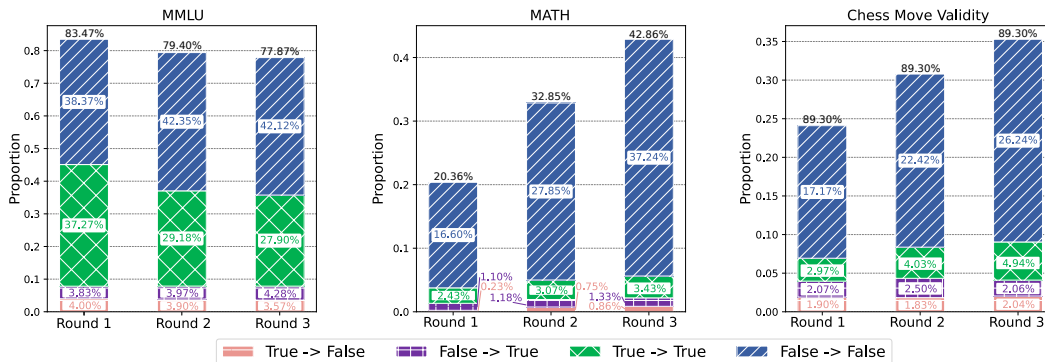


Figure 37: Variation of answer correctness in the situation of conformity, using *LlaMA2-70B-chat*, where *conformity brings about benefits*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) > \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$; *conformity brings about detriments*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) < \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$.

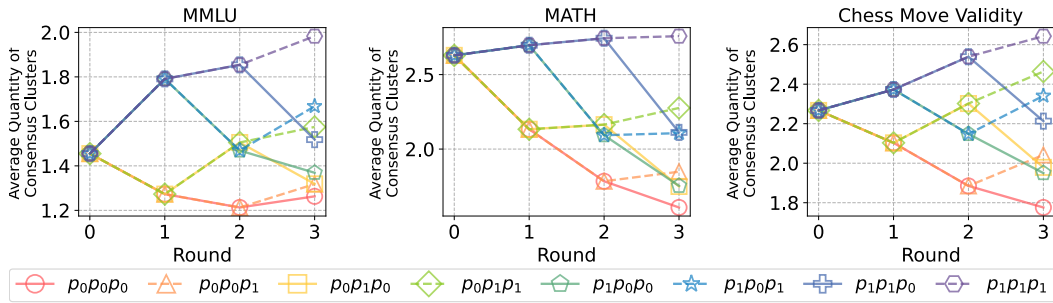


Figure 38: Average quantity of consensus clusters (i.e., unique answers among multiple agents) under different rounds of collaboration with 3-round collaborative strategies, on *LlaMA2-70B-chat*. Smaller quantity of consensus clusters, more easier it is to reach a consensus. Round 0 is equal to self-consistency.

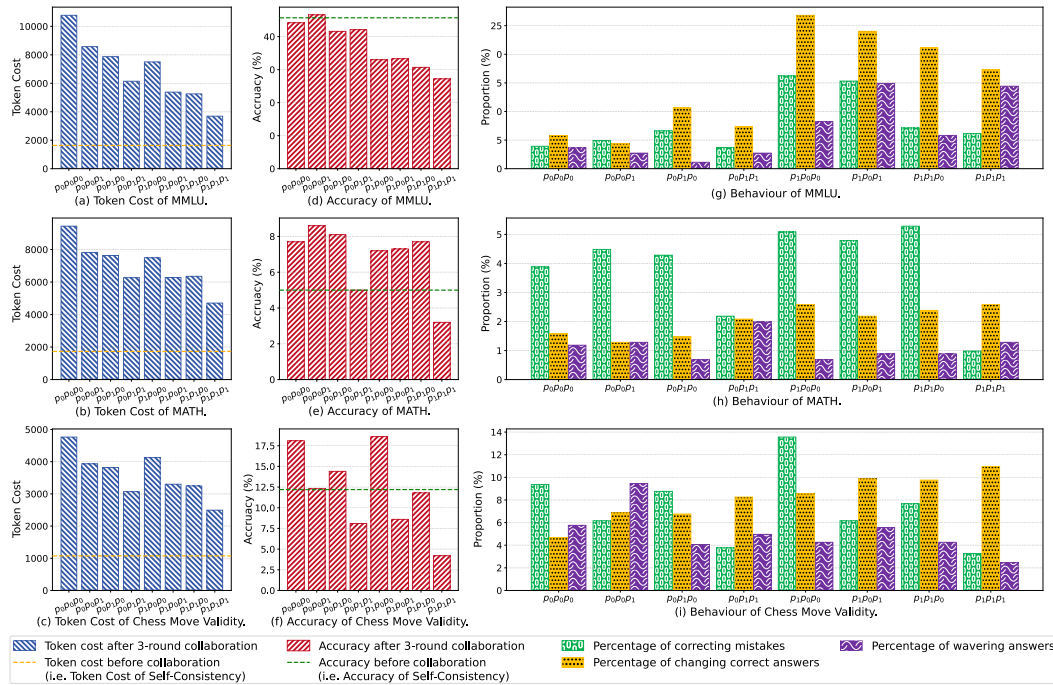


Figure 39: The percentage of different behaviors under different collaborative strategies, using *LlaMA2-70B-chat*. Figure (a-c) & (d-f) respectively show the token cost and accuracy of different strategies before and after 3-round collaboration. Figure (g-i) present the percentage of different behavioral features (mainly analyzed by the change of answer correctness) (Zhang et al., 2023c;b) under different collaborative strategies. All results are summarized across all societies.

H.3 QWEN 72B

Analysis on Machine Social Collaboration. We present the **main results** and **significance tests** of societies and strategies on Qwen 72B in Table 26, 27, 28. We present the **word clouds** of Qwen 72B in Figure 40, and **proportion of agents with different traits changing answers in different societies** on Qwen 72B in Figure 41. Furthermore, we demonstrate that the tasks with different subjects and difficulty display varying sensitivity to collaborative strategies, as presented with **radar maps** on Qwen 72B in Figure 42.

	Metric (Strategy)	Society	Collaborative Strategy								Metric (Society)	
			$p_0p_0p_0$	$p_0p_0p_1$	$p_0p_1p_0$	$p_0p_1p_1$	$p_1p_0p_0$	$p_1p_0p_1$	$p_1p_1p_0$	$p_1p_1p_1$	Cost ↓	W-T ↑
MMLU	Acc ↑	S_1	64.8±6.4	66.4±6.8	65.6±9.7	63.6±5.0	58.0±4.2	58.4±3.0	60.0±8.8	63.6±2.6	3661	14
		S_2	60.4±5.9	60.8±5.2	62.8±2.3	61.6±4.6	53.2±5.6	57.6±2.6	61.2±7.8	62.4±4.3	3657	21
		S_3	64.0±4.7	64.4±3.9	66.0±2.8	65.2±3.0	56.8±5.9	57.6±5.2	59.6±4.3	64.4±2.6	3690	17
		S_4	62.4±6.2	64.8±3.9	64.0±7.1	66.8±7.3	53.2±5.4	56.8±4.2	60.4±7.4	58.4±3.9	3570	14
	Cost ↓	All	5960	4560	4017	3158	4024	2761	2746	1927	-	-
W-T ↑	All	-	12	14	13	4	4	9	10	-	-	
MATH	Acc ↑	S_1	47.2±5.6	43.6±4.6	46.0±6.5	43.6±5.0	40.4±6.5	41.6±8.1	42.0±4.9	39.6±3.9	3537	11
		S_2	49.6±5.4	48.4±6.1	48.8±6.7	47.2±5.9	41.2±4.4	41.6±5.4	40.0±4.0	37.6±4.1	3513	7
		S_3	44.8±6.4	44.4±5.5	43.6±4.3	42.0±7.1	40.4±7.8	37.6±6.7	41.6±7.5	36.4±8.7	3595	9
		S_4	46.0±6.6	44.8±8.6	46.0±8.0	43.6±5.4	39.2±5.0	41.6±4.8	37.6±6.7	35.6±3.9	3595	11
	Cost ↓	All	4813	4182	4187	3549	3571	2912	2985	2281	-	-
W-T ↑	All	-	9	13	7	3	3	2	1	-	-	
Chess Move Validity	Acc ↑	S_1	43.2±7.0	42.4±4.6	41.2±9.7	36.8±6.4	27.6±4.8	22.0±5.3	20.4±4.8	6.4±3.3	2557	6
		S_2	46.8±4.2	42.8±4.2	39.2±4.6	34.8±4.2	29.6±5.2	16.8±2.7	22.8±5.8	8.8±3.4	2499	1
		S_3	42.4±8.7	38.4±9.9	38.0±6.9	36.8±7.8	26.8±5.8	19.6±2.6	19.6±2.6	6.0±2.8	2496	3
		S_4	36.0±8.1	32.4±4.6	34.0±5.8	26.0±4.9	26.8±5.4	20.8±5.4	22.4±5.9	11.2±2.3	2455	4
	Cost ↓	All	3148	2621	2585	2118	2904	2384	2393	1860	-	-
W-T ↑	All	-	6	6	2	0	0	0	0	-	-	

Table 26: The impact of eight different collaborative strategies on the performance of three datasets across distinct societies (using Qwen 72B). The significances test on societies and strategies are respectively shown in Table 27, 28. The experiments of comparison with the single LLM agent is shown in Figure 53(a)-(f).

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.654	0.637	0.162
$p_0p_0p_1$	0.388	0.649	0.064
$p_0p_1p_0$	0.841	0.667	0.445
$p_0p_1p_1$	0.455	0.567	0.034
$p_1p_0p_0$	0.387	0.963	0.817
$p_1p_0p_1$	0.933	0.690	0.281
$p_1p_1p_0$	0.987	0.647	0.695
$p_1p_1p_1$	0.061	0.688	0.048

Table 27: One-Way ANOVA results for the impact of society on accuracy with fixed collaborative strategy, based on experiments from Table 26 using Qwen 72B.

Society	MMLU p-value	MATH p-value	Chess Move Validity p-value
S_1	0.257	0.418	0.000
S_2	0.093	0.004	0.000
S_3	0.004	0.449	0.000
S_4	0.015	0.088	0.000

Table 28: One-Way ANOVA results for the impact of collaborative strategy on accuracy with fixed society, based on experiments from Table 26 using Qwen 72B.

Analysis on Different Numbers of Agents. We present the significance test for different numbers of agents with Qwen 72B in Table 29. We also show the performance varying from agent numbers

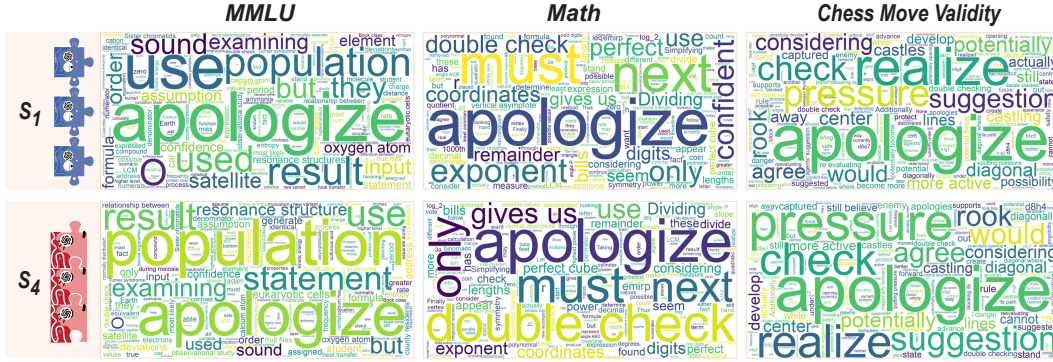


Figure 40: Comparative word clouds on three datasets in societies S_1 and S_4 , using *Qwen 72B*. Society S_1 features 3 overconfident agents, while society S_4 comprises 3 easy-going agents.

in Figure 43, varying from societies containing 2~10 agents in Figure 44. We also analyze the *consensus reaching* with different numbers of agents, and present the results in Figure 45, 46.

Collaborative Strategy	S'_1 p-value	S'_2 p-value	S'_3 p-value	S'_4 p-value	S'_5 p-value
$p_0p_0p_0$	0.005	0.001	0.003	0.041	0.015
$p_0p_0p_1$	0.017	0.010	0.037	0.001	0.006
$p_0p_1p_0$	0.006	0.016	0.002	0.000	0.001
$p_0p_1p_1$	0.020	0.002	0.010	0.001	0.004
$p_1p_0p_0$	0.000	0.005	0.000	0.000	0.000
$p_1p_0p_1$	0.002	0.008	0.004	0.000	0.054
$p_1p_1p_0$	0.003	0.000	0.002	-	0.000
$p_1p_1p_1$	0.064	0.008	0.005	0.016	0.000

Table 29: One-way ANOVA analysis of results in Figure 43 (different numbers of agents), using *Qwen 72B*. S'_1 : One overconfident agent and the others are all easygoing. S'_2 : One easygoing agent among predominantly overconfident agents. S'_3 : Equal numbers of overconfident and easygoing agents. S'_4 : Entirely easygoing agents. S'_5 : Entirely overconfident agents. ‘-’: It doesn’t pass homogeneity test for variance.

Analysis on Different Rounds. We present the significance test for different rounds of collaboration with *Qwen 72B* in Table 30. We also show the performance varying from collaboration rounds in Figure 47, 48, 49.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0p_0p_0p_0p_0p_0p_0p_0$	0.262	0.987	0.956
$p_1p_0p_0p_0p_0p_0p_0p_0p_0p_0$	0.753	0.697	0.124
$p_0p_1p_0p_0p_0p_0p_0p_0p_0p_0$	0.914	0.962	0.386
$p_1p_0p_1p_0p_0p_1p_0p_1p_0p_1p_0$	0.673	0.715	0.154
$p_0p_1p_0p_1p_0p_1p_0p_1p_0p_1$	0.922	0.987	0.700
$p_1p_0p_1p_1p_1p_1p_1p_1p_1p_1$	0.845	0.843	0.282
$p_0p_1p_1p_1p_1p_1p_1p_1p_1p_1$	0.928	0.585	0.583
$p_1p_1p_1p_1p_1p_1p_1p_1p_1p_1$	0.832	0.801	0.731

Table 30: One-way ANOVA analysis of the results in Figure 48, 48, 49 (different rounds), using *Qwen 72B*.

Analysis on Other Collaborative Strategies. We present the significance test for other collaborative strategies (executing the same or hybrid thinking patterns in a certain round) with *Qwen 72B* in Table 31. We also show the performance varying from other strategies in Figure 50.

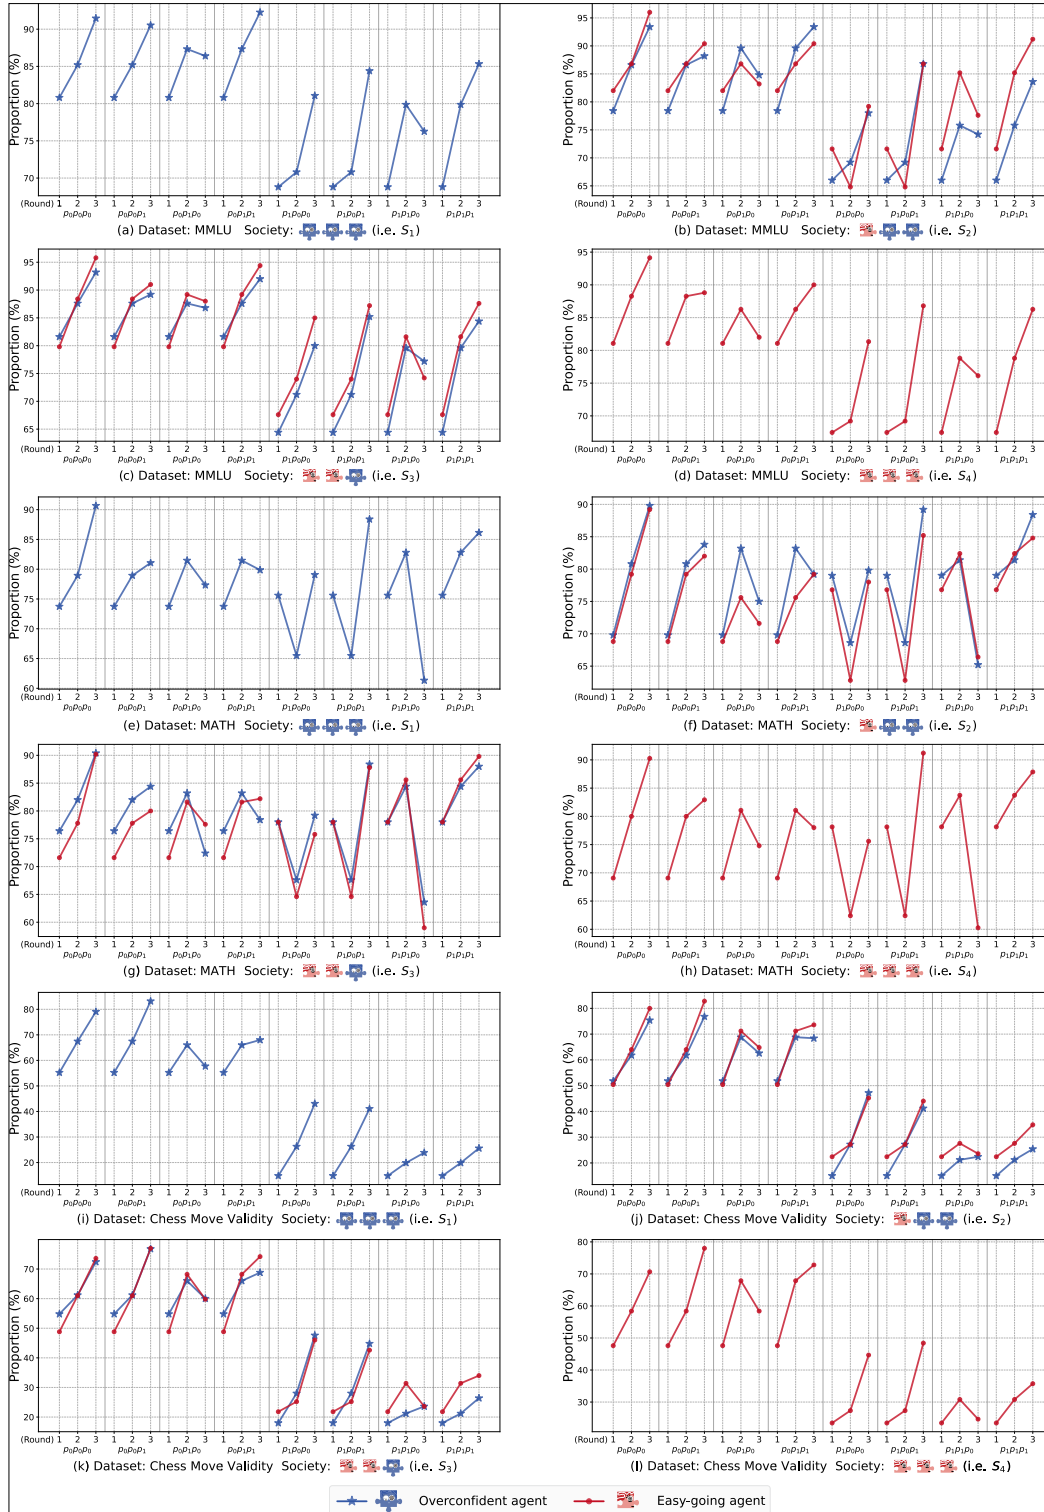


Figure 41: Proportion of agents with different traits changing answers in societies S_1 and S_4 , using *Qwen 72B*. Society S_1 features three overconfident agents, while society S_4 comprises three easy-going agents.

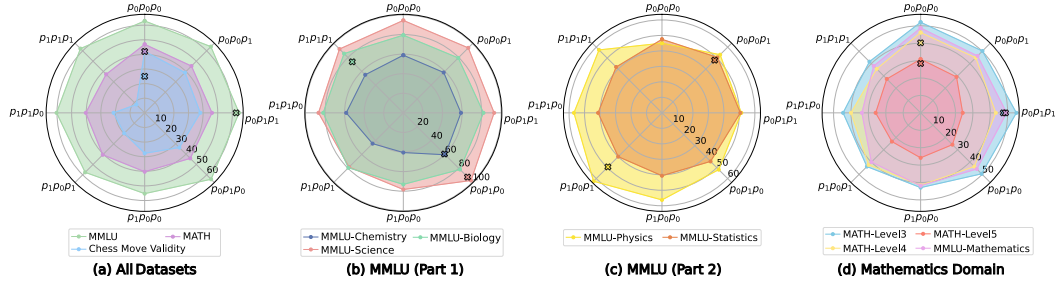


Figure 42: Illustration of different collaborative strategies impacting accuracy diversely on the tasks considering varied *subjects* and *difficulty*, using *Qwen 72B*. The symbol ‘ \otimes ’ represents that there is at least one collaborative strategy whose accuracy is better than self-consistency, while the symbol ‘ \otimes ’ indicates that there is no collaborative strategy whose accuracy is worse than self-consistency. Both of these symbols represent the accuracy of self-consistency. The accuracy under each collaborative strategy is a summation within all 3-agent societies.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.704	0.142	0.003
$p_0p_0p_1$	0.136	0.184	0.000
$p_0p_1p_0$	0.899	0.157	0.001
$p_0p_1p_1$	0.180	0.194	0.089
$p_1p_0p_0$	0.157	0.856	0.004
$p_1p_0p_1$	0.521	0.152	0.019
$p_1p_1p_0$	-	0.790	0.004
$p_1p_1p_1$	0.391	0.688	1.000

Table 31: One-way ANOVA analysis of results in Figure 50 (other collaborative strategies), using *Qwen 72B*. ‘-’ means it doesn’t pass homogeneity test for variance.

A Social Psychology View on Conformity, Consensus Reaching and Group Dynamics. We then show the variation of answer correctness in the situation of conformity in Figure 51; and the quantity of consensus clusters among 3-agent answers in Figure 52. We present group dynamics reflected by different answer-changing behaviors on *Qwen 72B* in Figure 53.

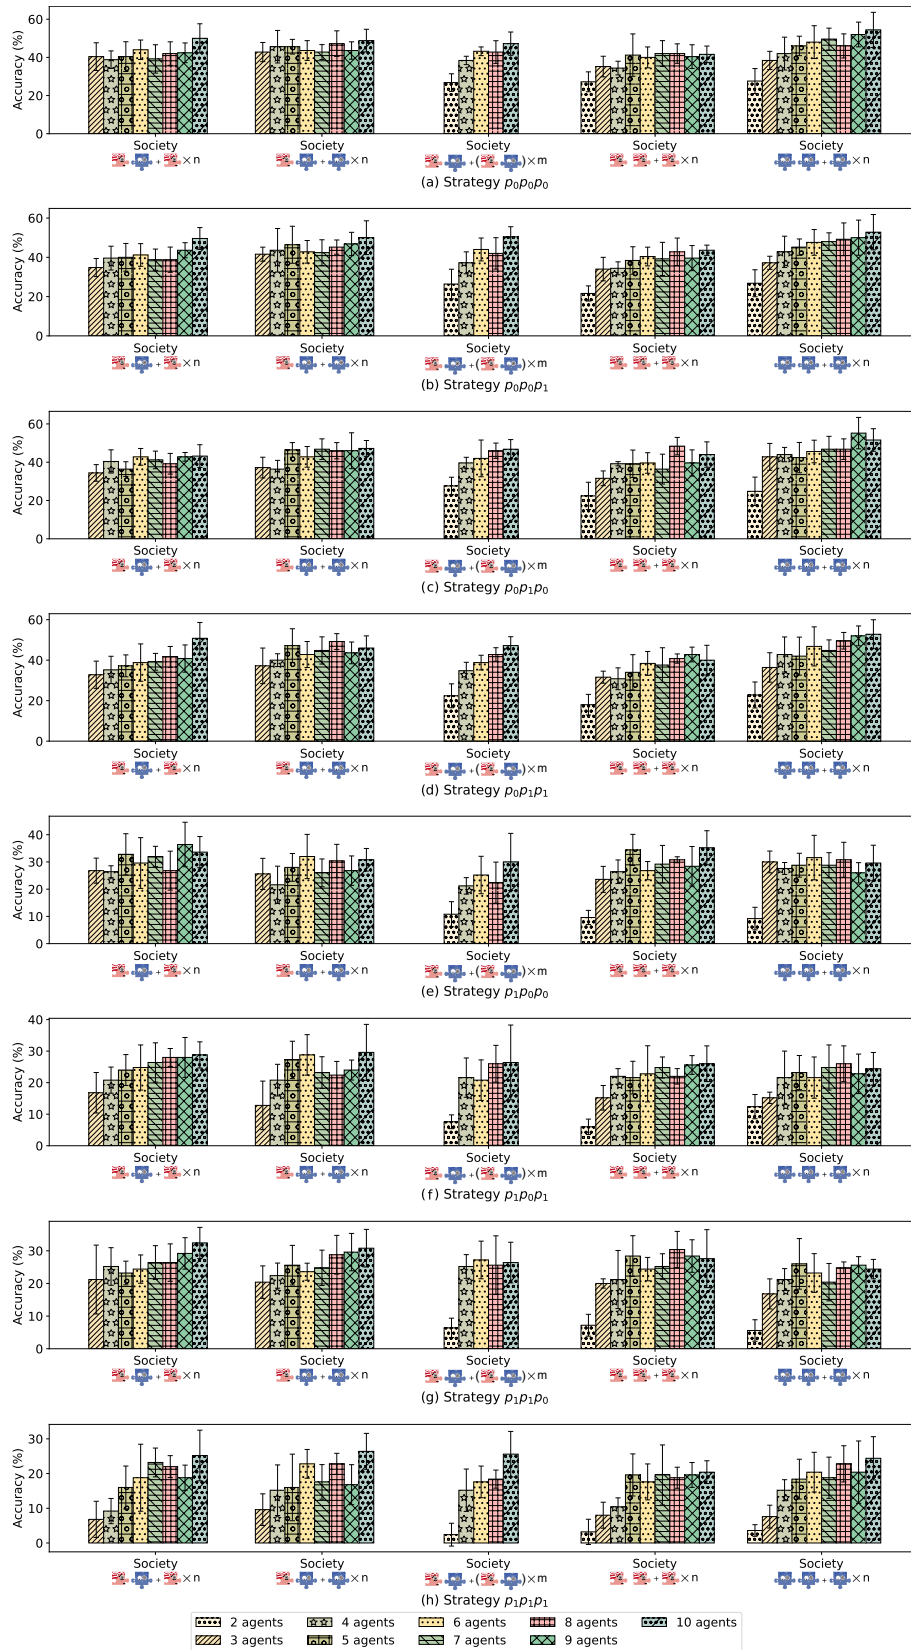


Figure 43: Accuracy of different numbers (2~10) of agents under different collaborative strategies, on Qwen 72B. The significance test is shown in Table 29.

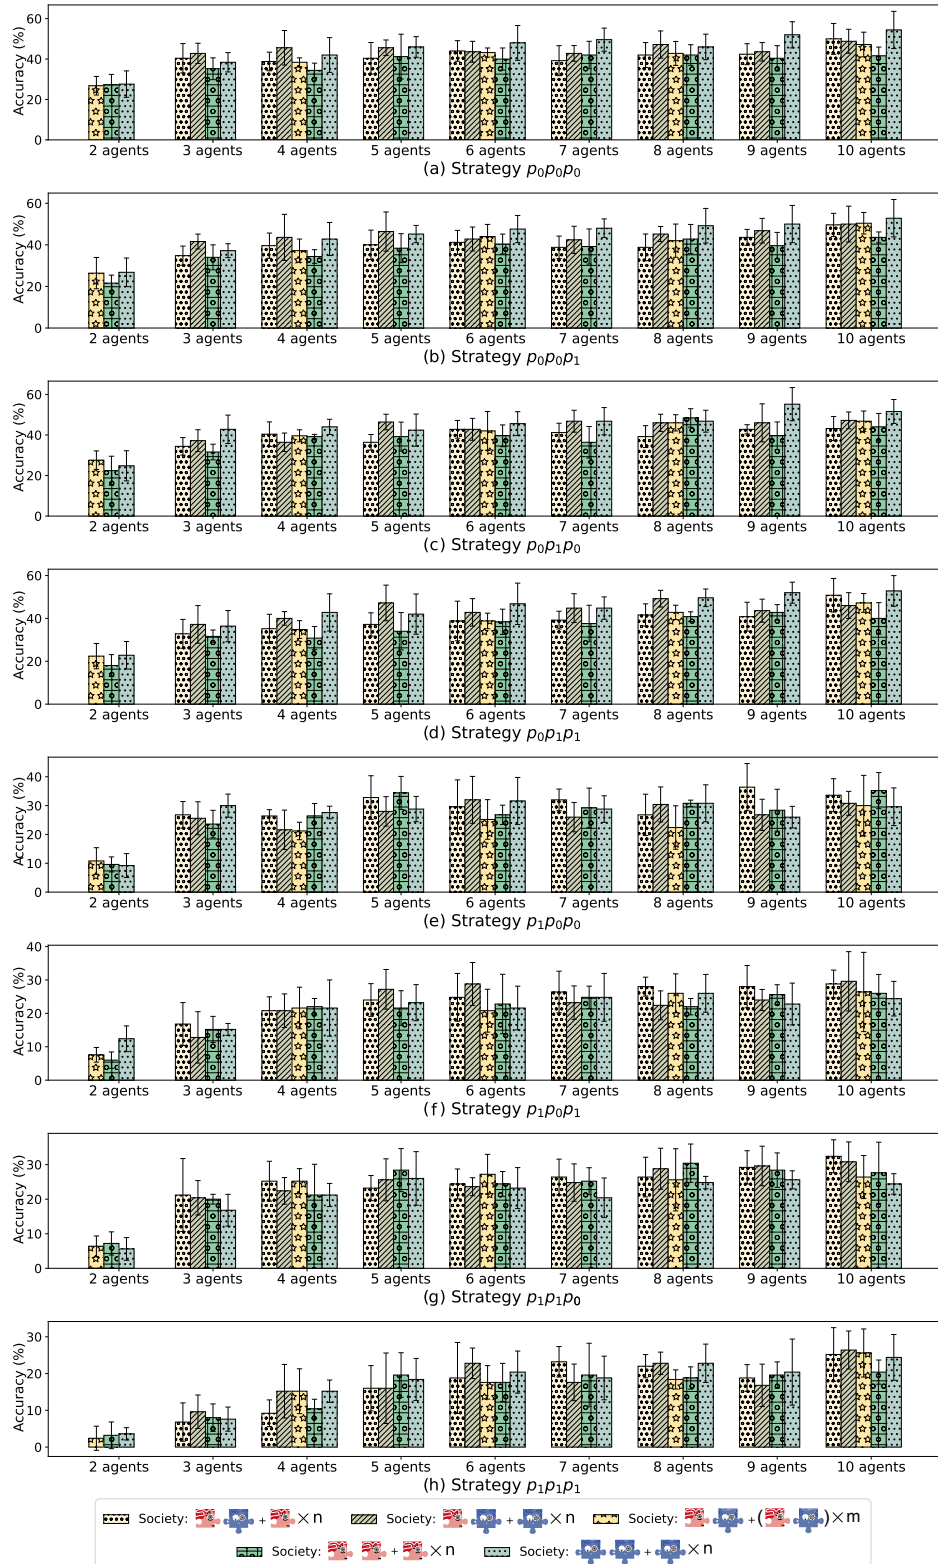


Figure 44: Accuracy of different societies with 2~10 agents under different collaborative strategies, on Qwen 72B.

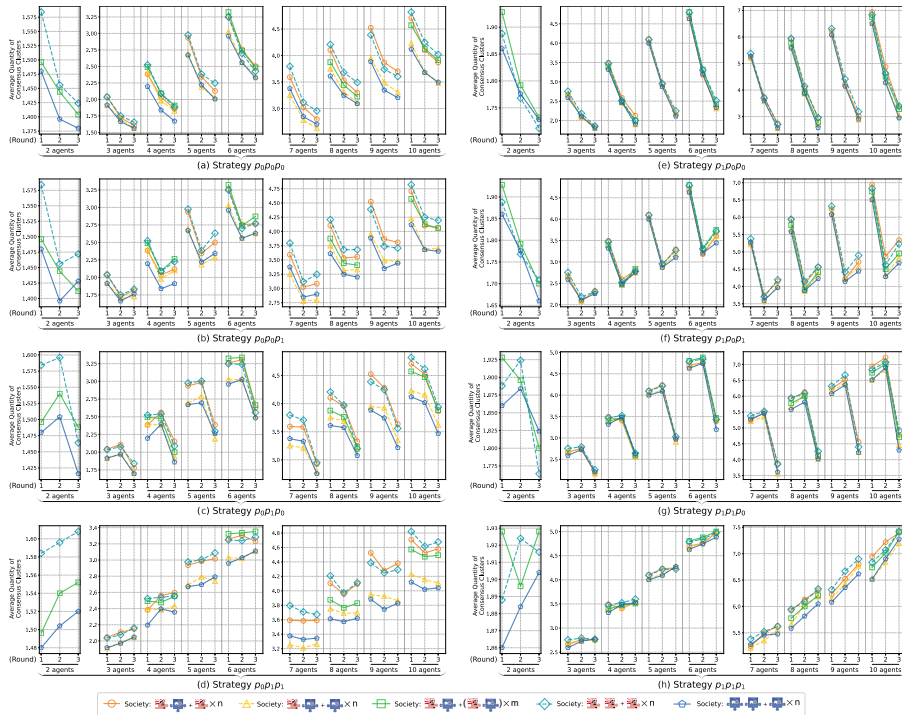


Figure 45: Average quantity of consensus clusters (unique answers among multiple agents) in different societies with 2~10 agents under each round of 3-round collaborative strategies, using Qwen 72B.

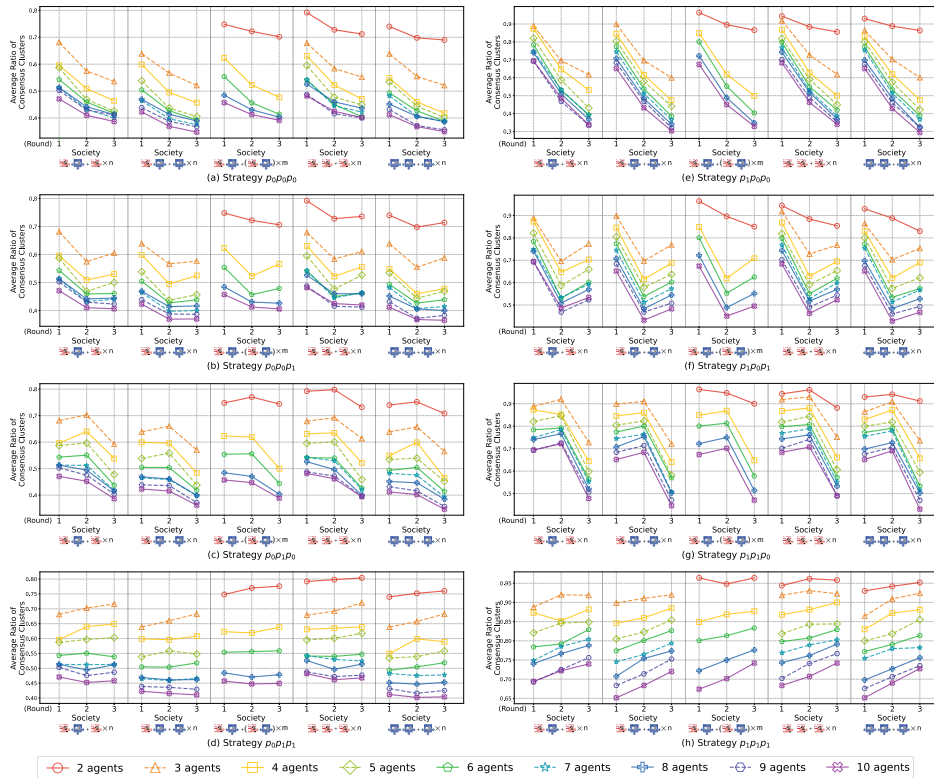


Figure 46: Average ratio of consensus clusters (unique answers among multiple agents) with different numbers (2~10) of agents under each round of 3-round collaborative strategies, using Qwen 72B.

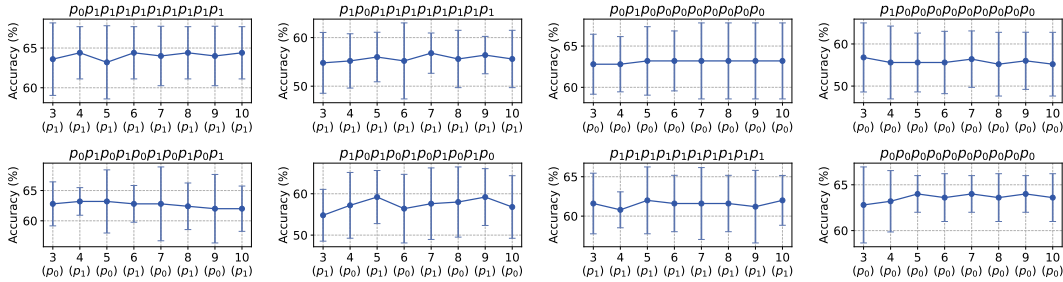


Figure 47: Accuracy of *different (3~10) rounds of collaboration* within 3-agent society S_2 (1 easy-going and 2 overconfident agents) on MMLU, using *Qwen 72B*. The significance test is shown in Table 30.

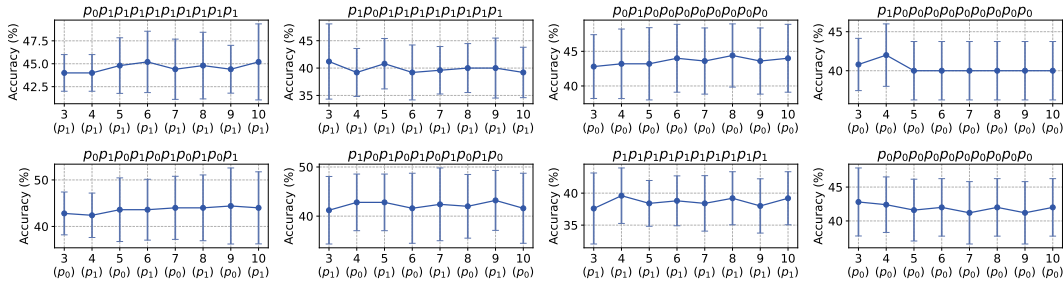


Figure 48: Accuracy of *different (3~10) rounds of collaboration* within 3-agent society S_2 (1 easy-going and 2 overconfident agents) on MATH, using *Qwen 72B*. The significance test is shown in Table 30.

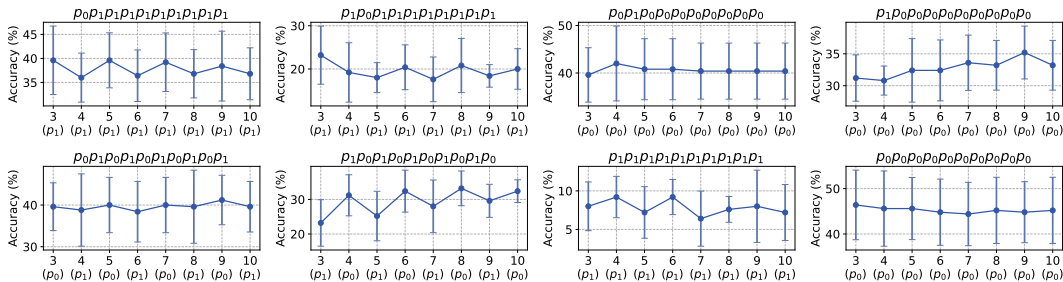


Figure 49: Accuracy of *different (3~10) rounds of collaboration* within 3-agent society S_2 (1 easy-going and 2 overconfident agents) on Chess Move Validity, using *Qwen 72B*. The significance test is shown in Table 30.

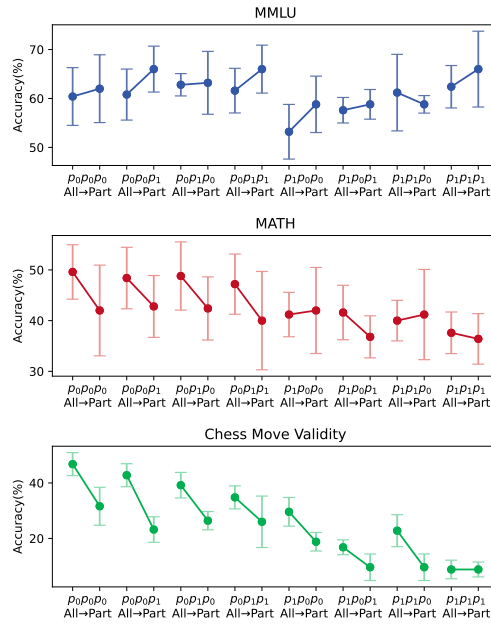


Figure 50: The effect on the accuracy of whether all agents in society execute the same thinking pattern in one round, using *Qwen 72B*. “All” and “Part” refers to all agents applying the same thinking pattern and different thinking patterns in one round respectively. The significance test is shown in Table 31.

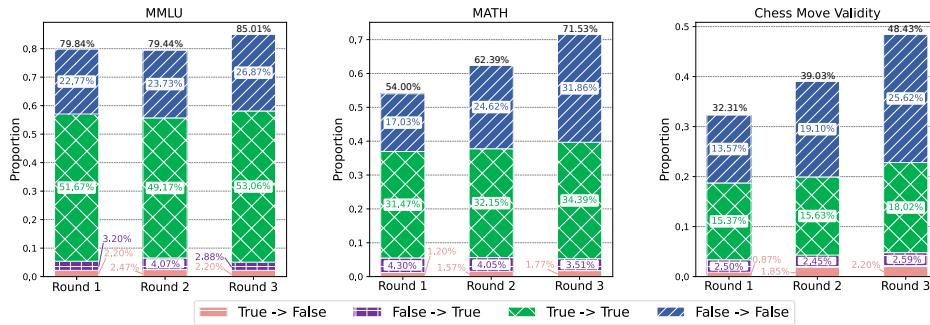


Figure 51: Variation of answer correctness in the situation of conformity, using *Qwen 72B*, where *conformity brings about benefits*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) > \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$; *conformity brings about detriments*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) < \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$.

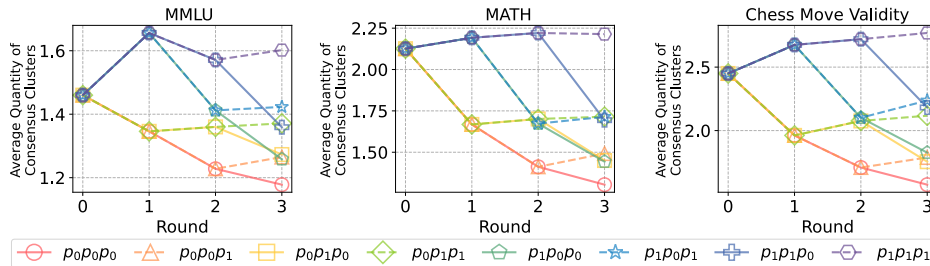


Figure 52: Average quantity of consensus clusters (i.e., unique answers among multiple agents) under different rounds of collaboration with 3-round collaborative strategies, using *Qwen 72B*. Smaller quantity of consensus clusters, more easier it is to reach a consensus. Round 0 is equal to self-consistency.

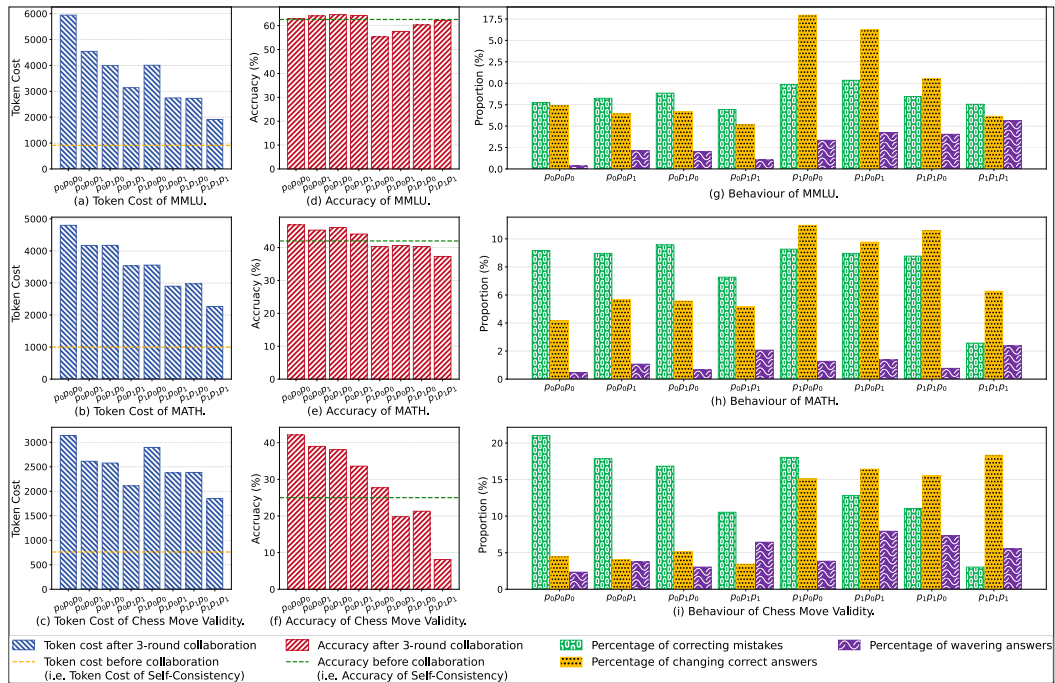


Figure 53: The percentage of different behaviors under different collaborative strategies, using *Qwen 72B*. Figure (a-c) & (d-f) respectively show the token cost and accuracy of different strategies before and after 3-round collaboration. Figure (g-i) present the percentage of different behavioral features (mainly analyzed by the change of answer correctness) (Zhang et al., 2023c;b) under different collaborative strategies. All results are summarized across all societies.

H.4 MIXTRAL 8×7B

Analysis on Machine Social Collaboration. We present the **main results** and **significance tests** of societies and strategies on Mixtral 8×7B in Table 32, 33, 34. We present the word clouds of Mixtral 8×7B in Figure 54, and the proportion of agents with different traits changing answers in different societies on Mixtral 8×7B in Figure 55. Furthermore, we demonstrate that the tasks with different subjects and difficulty display varying sensitivity to collaborative strategies, as presented with **radar maps** on Mixtral 8×7B in Figure 56.

	Metric (Strategy)	Society	Collaborative Strategy								Metric (Society)	
			$p_0p_0p_0$	$p_0p_0p_1$	$p_0p_1p_0$	$p_0p_1p_1$	$p_1p_0p_0$	$p_1p_0p_1$	$p_1p_1p_0$	$p_1p_1p_1$	Cost ↓	W-T ↑
MMLU	Acc ↑	S_1	60.0±8.1	59.6±3.9	58.4±4.3	60.0±1.4	60.0±5.8	60.4±5.2	59.6±2.6	60.0±2.0	4479	17
		S_2	59.2±7.7	60.0±7.9	60.0±6.5	60.8±5.8	61.2±3.6	62.8±5.4	62.8±5.4	61.2±2.7	4475	27
		S_3	62.4±5.2	63.6±4.3	65.2±3.0	65.2±3.0	59.2±4.4	61.2±4.2	61.6±2.6	59.6±3.6	4489	18
		S_4	60.0±3.7	62.4±3.6	63.2±3.4	62.8±2.7	60.0±5.1	60.4±5.5	64.8±5.8	62.0±6.6	4396	25
	Cost ↓	All	6891	5371	4871	3944	4996	3594	3495	2516	-	-
W-T ↑	All	-	14	15	14	9	11	13	11	-	-	
MATH	Acc ↑	S_1	30.4±3.3	36.0±1.4	33.6±2.2	32.8±4.2	31.2±3.4	30.4±2.6	30.8±2.3	27.6±1.7	5362	23
		S_2	31.6±6.1	29.2±5.4	30.4±6.8	28.0±3.7	32.4±3.6	29.2±3.9	32.0±6.0	27.6±3.0	5369	14
		S_3	32.4±6.7	32.8±7.8	34.8±4.8	32.0±4.7	30.8±4.2	28.8±4.2	30.8±2.3	24.8±3.9	5343	18
		S_4	32.0±4.7	31.2±2.7	31.2±5.2	32.0±5.1	29.2±4.4	30.0±7.2	31.2±1.1	27.2±3.4	5238	18
	Cost ↓	All	6630	5814	6116	5042	5915	4745	4818	3540	-	-
W-T ↑	All	-	12	13	9	14	11	10	4	-	-	
Chess Move Validity	Acc ↑	S_1	22.8±2.7	21.6±3.3	21.2±5.6	20.8±3.0	18.8±5.4	18.8±4.6	17.6±7.0	18.8±1.1	2300	9
		S_2	22.0±5.7	18.0±2.8	18.8±3.4	16.4±2.6	22.0±8.4	18.8±4.8	16.0±2.8	16.0±0.0	2280	10
		S_3	21.2±2.7	20.0±3.2	18.0±2.5	18.0±2.5	20.0±2.8	18.8±3.0	16.4±4.6	15.6±1.7	2269	9
		S_4	18.0±3.7	16.4±3.9	19.2±4.6	16.4±2.6	20.0±1.4	20.8±3.6	20.4±3.9	18.8±2.3	2253	23
	Cost ↓	All	2956	2458	2396	1973	2630	2063	2083	1644	-	-
W-T ↑	All	-	7	8	6	9	10	6	5	-	-	

Table 32: The impact of eight different collaborative strategies on the performance of three datasets across distinct societies (using *Mixtral-8×7B*). The significances test on societies and strategies are respectively shown in Table 33, 34. The experiments of comparison with the single LLM agent is shown in Figure 67(a)-(f).

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.873	0.941	0.261
$p_0p_0p_1$	0.578	0.216	0.109
$p_0p_1p_0$	0.114	0.500	0.666
$p_0p_1p_1$	0.142	0.347	0.062
$p_1p_0p_0$	0.930	0.638	0.809
$p_1p_0p_1$	0.863	0.949	0.825
$p_1p_1p_0$	0.325	-	0.485
$p_1p_1p_1$	0.785	0.438	0.004

Table 33: One-Way ANOVA results for the impact of society on accuracy with fixed collaborative strategy, based on experiments from Table 32 using *Mixtral 8×7B*. ‘-’: It doesn’t pass homogeneity test for variance.

Society	MMLU p-value	MATH p-value	Chess Move Validity p-value
S_1	0.999	0.002	0.585
S_2	0.970	0.693	0.202
S_3	0.129	0.127	0.078
S_4	0.706	0.714	0.300

Table 34: One-Way ANOVA results for the impact of collaborative strategy on accuracy with fixed society, based on experiments from Table 32 using *Mixtral 8×7B*.

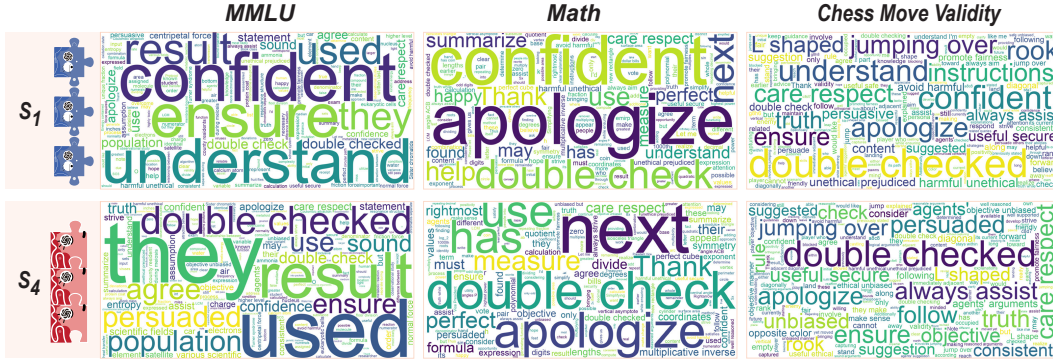


Figure 54: Comparative word clouds on three datasets in societies S_1 and S_4 , using *Mixtral-8×7B*. Society S_1 features 3 overconfident agents, while society S_4 comprises 3 easy-going agents.

Analysis on Different Numbers of Agents. We present the significance test for different numbers of agents with *Mixtral 8×7B* in Table 35. We also show the performance varying from agent numbers in Figure 57, varying from societies containing 2~10 agents in Figure 58. We also analyze the *consensus reaching* with different numbers of agents, and present the results in Figure 59, 60.

Collaborative Strategy	S'_1 p-value	S'_2 p-value	S'_3 p-value	S'_4 p-value	S'_5 p-value
$p_0p_0p_0$	0.188	0.406	0.235	0.805	0.009
$p_0p_0p_1$	0.106	0.112	0.238	0.459	0.008
$p_0p_1p_0$	0.142	0.145	0.227	0.739	0.227
$p_0p_1p_1$	0.013	0.004	0.035	0.138	0.075
$p_1p_0p_0$	0.159	0.082	0.105	0.018	0.088
$p_1p_0p_1$	0.029	0.003	0.002	0.004	0.018
$p_1p_1p_0$	0.051	0.028	0.010	0.001	0.247
$p_1p_1p_1$	0.002	0.016	0.003	0.000	0.001

Table 35: One-way ANOVA analysis of results in Figure 57 (different numbers of agents), using *Mixtral 8×7B*. S'_1 : One overconfident agent and the others are all easygoing. S'_2 : One easygoing agent among predominantly overconfident agents. S'_3 : Equal numbers of overconfident and easygoing agents. S'_4 : Entirely easygoing agents. S'_5 : Entirely overconfident agents.

Analysis on Different Rounds. We present the significance test for different rounds of collaboration with *Mixtral 8×7B* in Table 36. We also show the performance varying from collaboration rounds in Figure 61, 62, 63.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0p_0p_0p_0p_0p_0p_0$	0.607	0.911	0.789
$p_1p_0p_0p_0p_0p_0p_0p_0p_0$	0.578	0.581	0.939
$p_0p_1p_0p_0p_0p_0p_0p_0p_0$	0.936	0.665	0.123
$p_1p_0p_1p_0p_0p_1p_0p_1p_0$	0.377	0.896	0.952
$p_0p_1p_0p_1p_0p_1p_0p_1p_0p_1$	0.987	0.651	0.271
$p_1p_0p_1p_1p_1p_1p_1p_1p_1$	0.989	0.878	0.919
$p_0p_1p_1p_1p_1p_1p_1p_1p_1$	0.989	0.982	1.000
$p_1p_1p_1p_1p_1p_1p_1p_1p_1$	0.945	0.995	0.903

Table 36: One-way ANOVA analysis of the results in Figure 61, 62, 63 (different rounds), using *Mixtral 8×7B*.

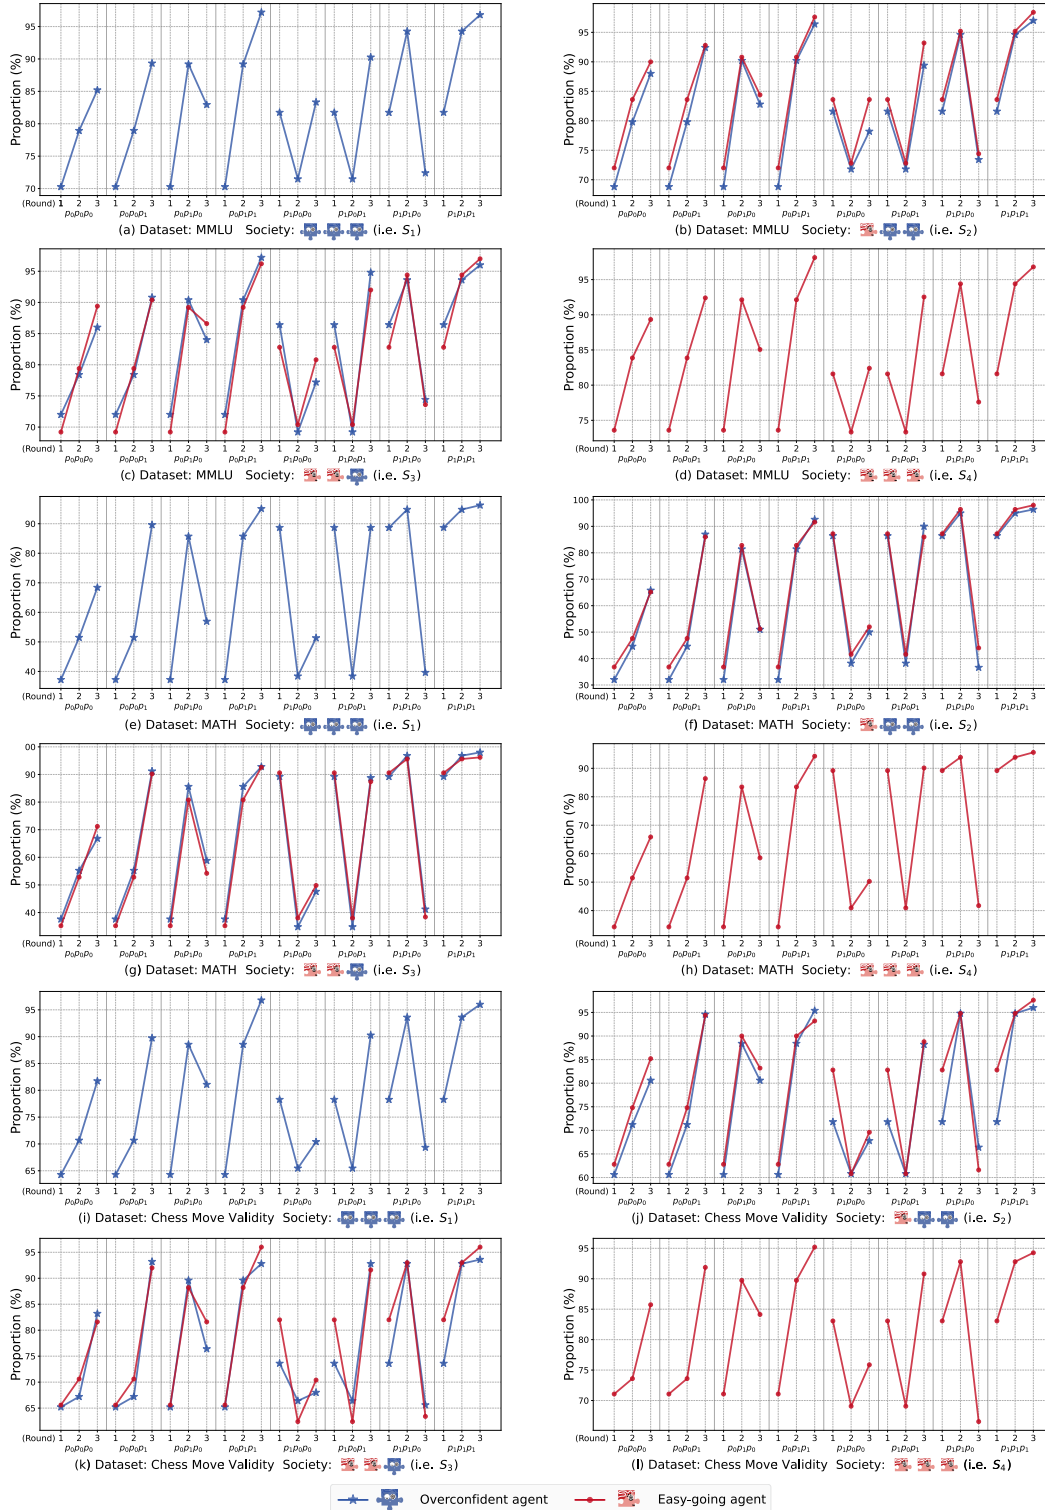


Figure 55: Proportion of agents with different traits changing answers in societies S_1 and S_4 , using *Mixtral-8x7B*. Society S_1 features three overconfident agents, while society S_4 comprises three easy-going agents.

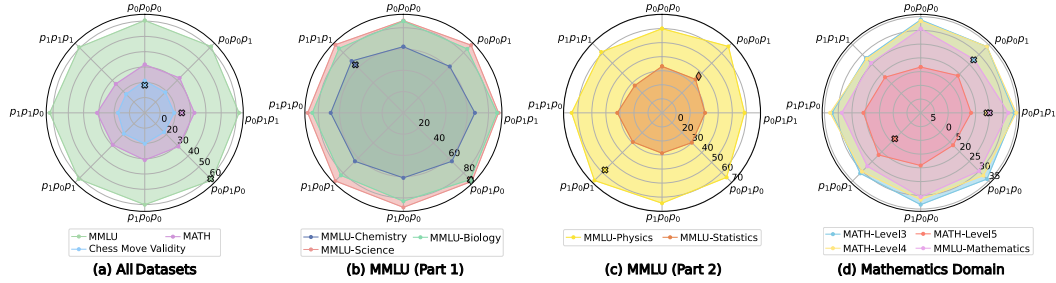


Figure 56: Illustration of different collaborative strategies impacting accuracy diversely on the tasks considering varied *subjects* and *difficulty*, using *Mixtral-8×7B*. The symbol ‘ \otimes ’ represents that there is at least one collaborative strategy whose accuracy is better than self-consistency, while the symbol ‘ \boxtimes ’ indicates that there is no collaborative strategy whose accuracy is worse than self-consistency. Both of these symbols represent the accuracy of self-consistency. The accuracy under each collaborative strategy is a summation within all 3-agent societies.

Analysis on Other Collaborative Strategies. We present the significance test for other collaborative strategies (executing the same or hybrid thinking patterns in a certain round) with *Mixtral 8×7B* in Table 37. We also show the performance varying from other strategies in Figure 64.

Collaborative Strategy	MMLU p-value	MATH p-value	Chess Move Validity p-value
$p_0p_0p_0$	0.618	0.898	0.390
$p_0p_0p_1$	0.919	0.143	0.058
$p_0p_1p_0$	0.797	0.548	0.031
$p_0p_1p_1$	0.521	0.141	0.049
$p_1p_0p_0$	0.040	0.409	0.290
$p_1p_0p_1$	0.658	0.400	0.373
$p_1p_1p_0$	0.193	0.318	0.142
$p_1p_1p_1$	0.536	0.453	-

Table 37: One-way ANOVA analysis of results in Figure 64 (other collaborative strategies), on *Mixtral 8×7B*. ‘-’ means it doesn’t pass homogeneity test for variance.

A Social Psychology View on Conformity, Consensus Reaching and Group Dynamics. We then show the variation of answer correctness in the situation of conformity in Figure 65; and the quantity of consensus clusters among 3-agent answers in Figure 66. We present group dynamics reflected by different answer-changing behaviors on *Mixtral-8×7B* in Figure 67.

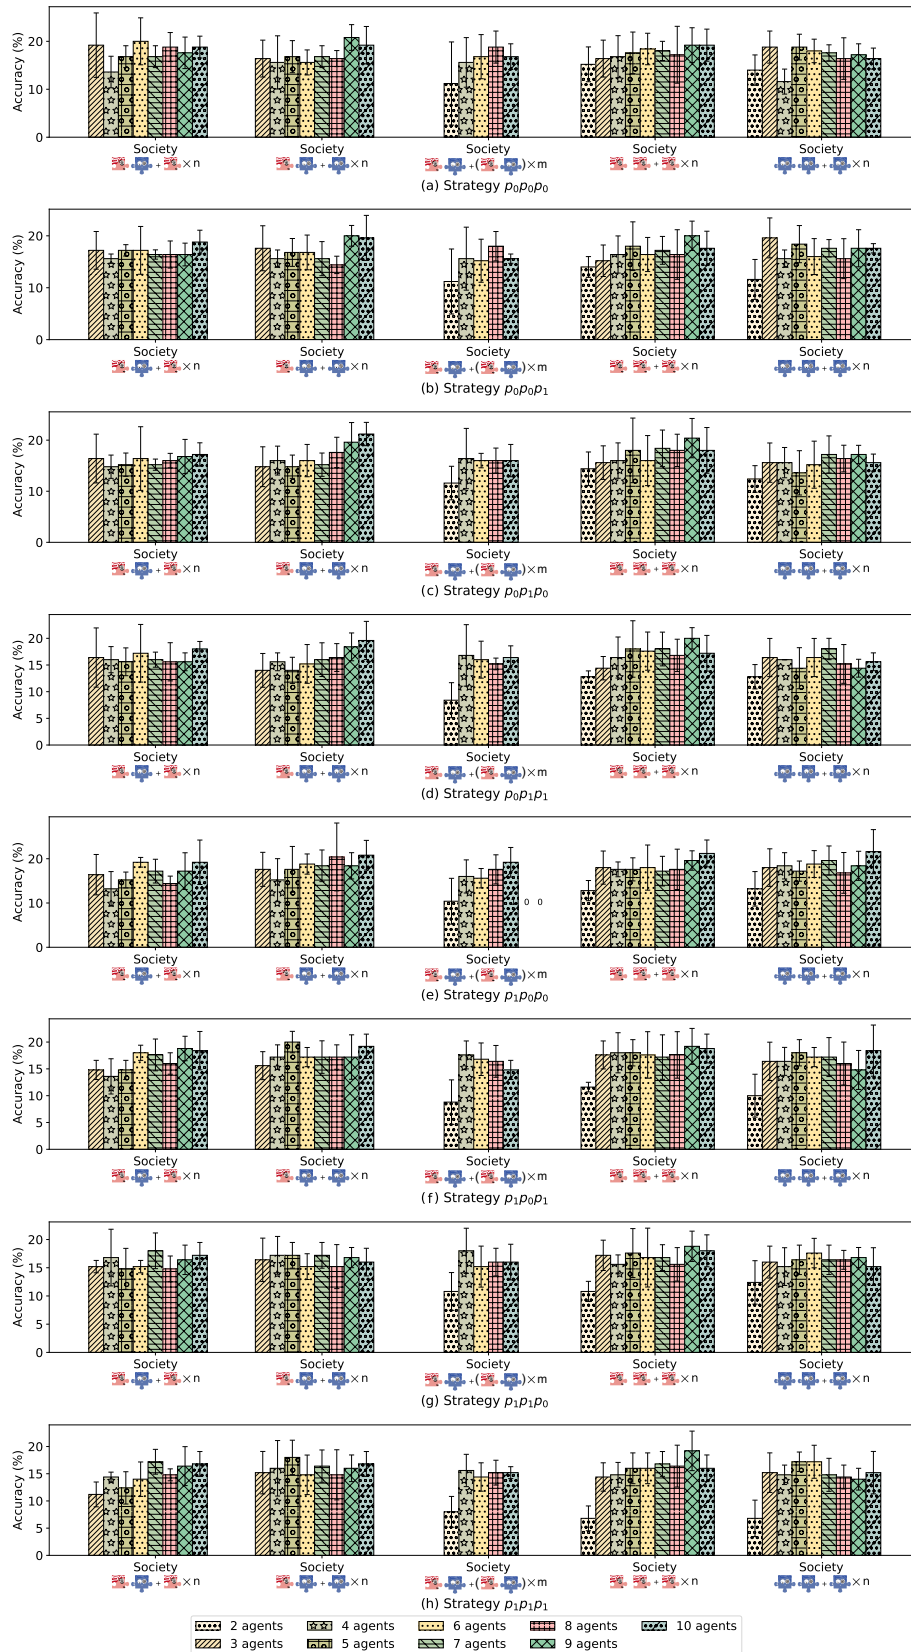


Figure 57: Accuracy of different numbers (2~10) of agents under different collaborative strategies, on *Mixtral-8x7B*. The significance test is shown in Table 35.

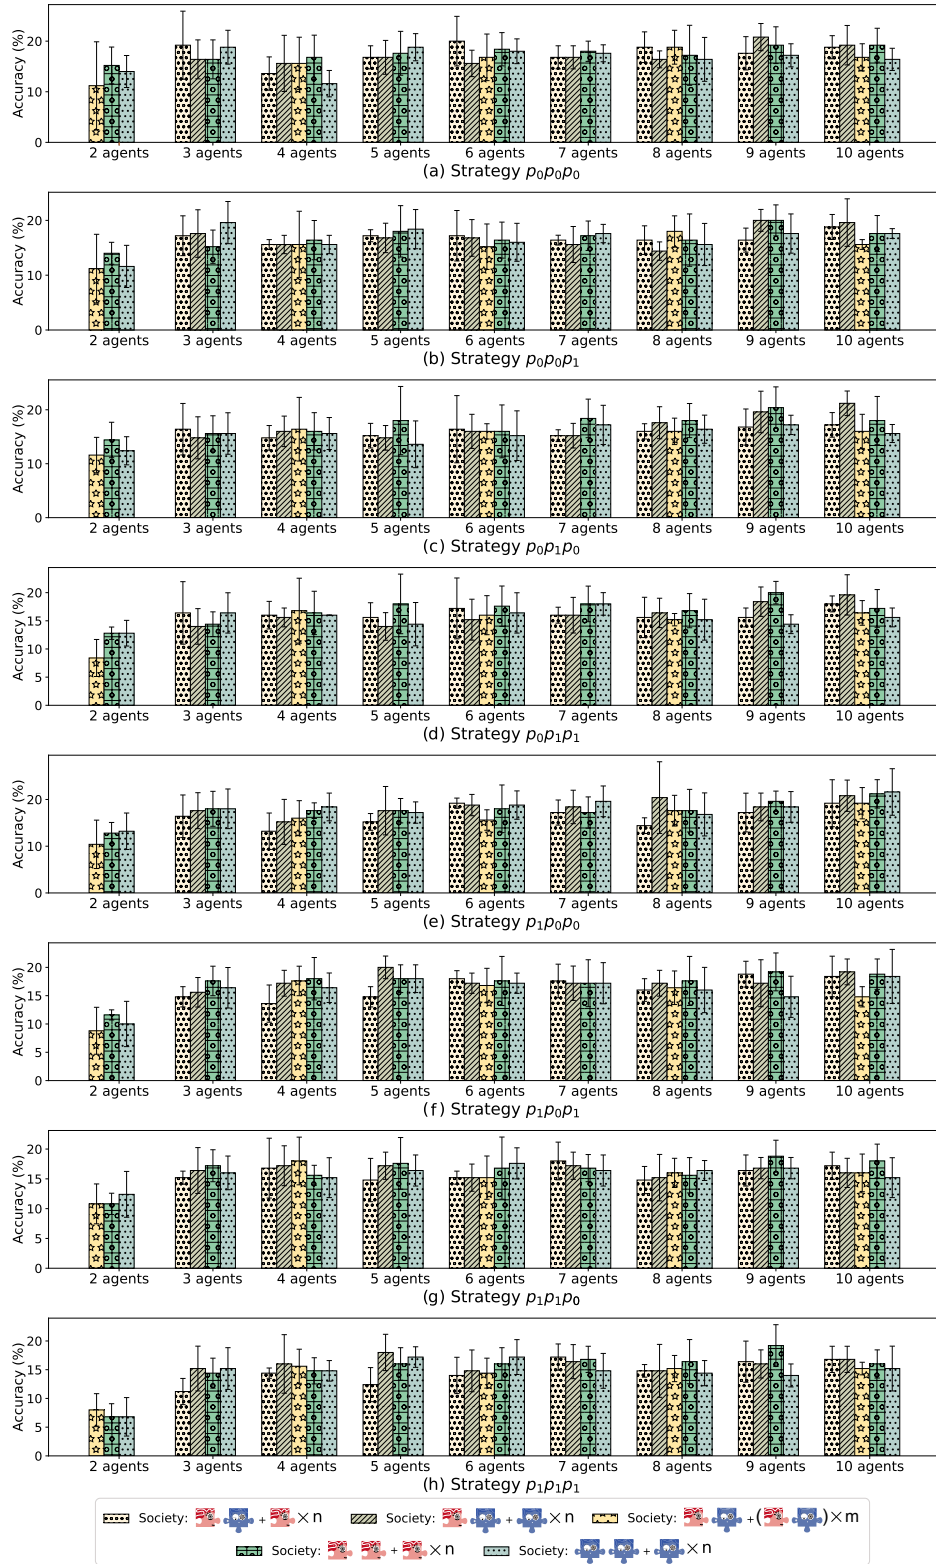


Figure 58: Accuracy of different societies with 2~10 agents under different collaborative strategies, on *Mixtral-8x7B*.

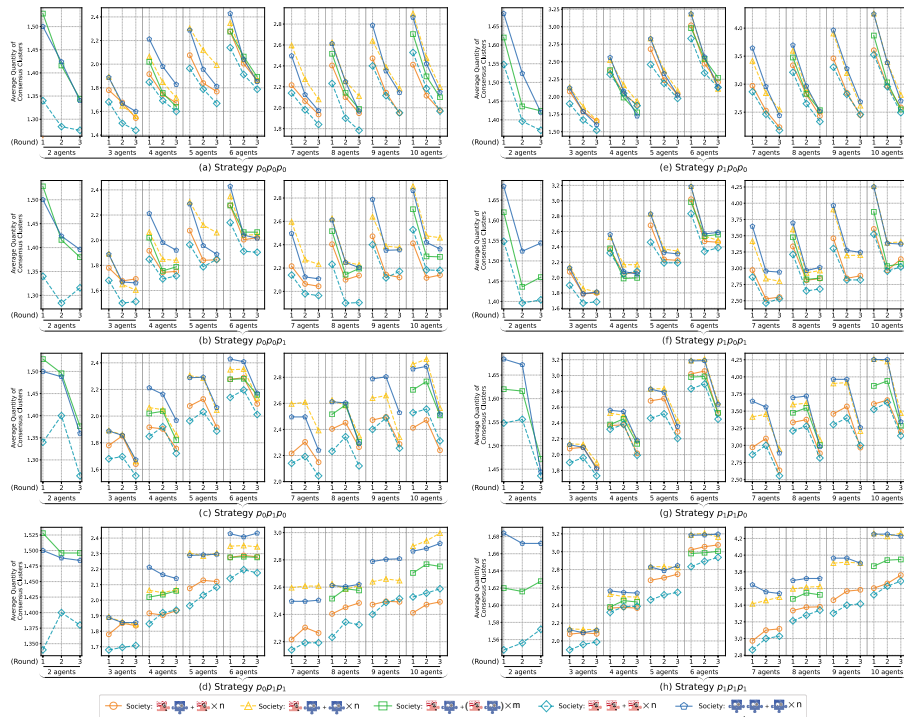


Figure 59: Average quantity of consensus clusters (unique answers among multiple agents) in different societies with 2~10 agents under each round of 3-round collaborative strategies, using Mixtral-8x7B.

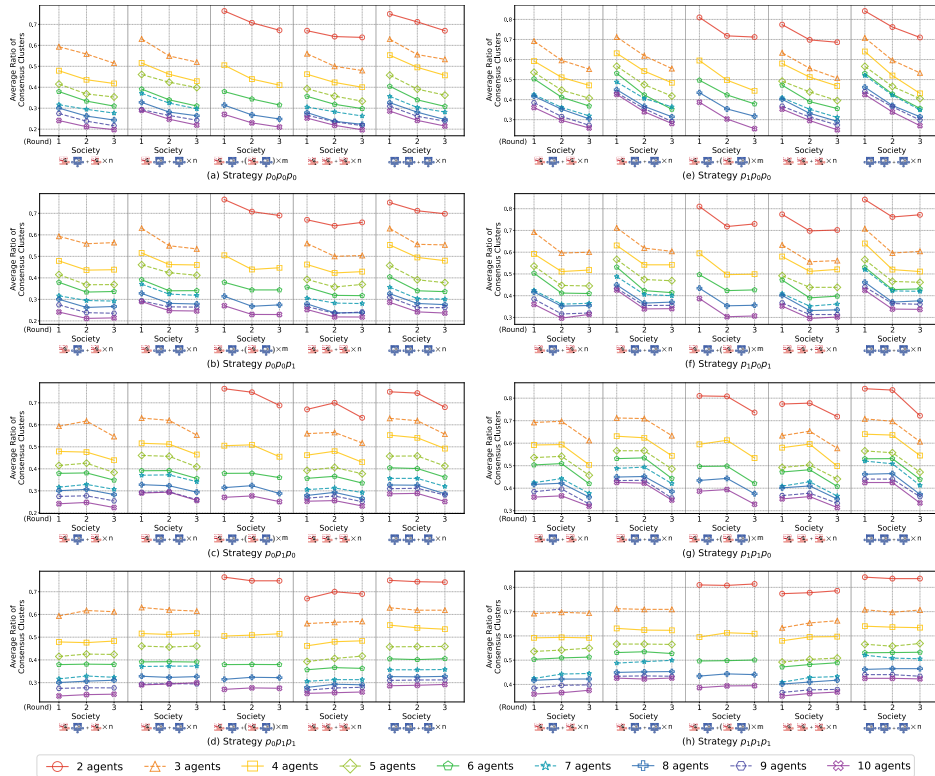


Figure 60: Average ratio of consensus clusters (unique answers among multiple agents) with different numbers (2~10) of agents under each round of 3-round collaborative strategies, using Mixtral-8x7B.

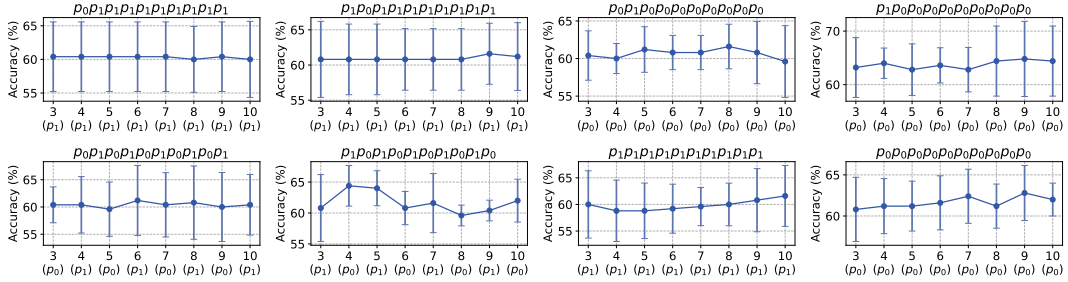


Figure 61: Accuracy of *different (3~10) rounds of collaboration* within 3-agent society S_2 (1 easy-going and 2 overconfident agents) on MMLU, using *Mixtral-8x7B*. The significance test is shown in Table 36.

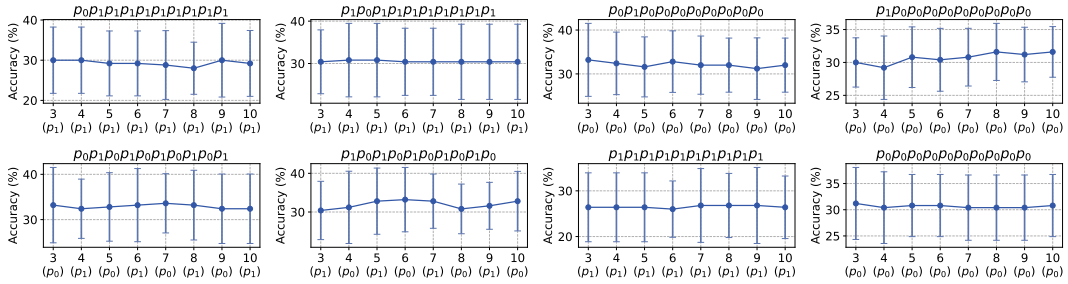


Figure 62: Accuracy of *different (3~10) rounds of collaboration* within 3-agent society S_2 (1 easy-going and 2 overconfident agents) on MATH, using *Mixtral-8x7B*. The significance test is shown in Table 36.

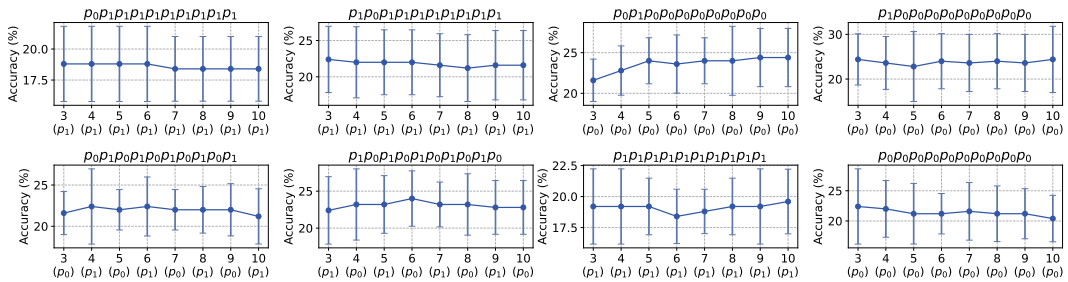


Figure 63: Accuracy of *different (3~10) rounds of collaboration* within 3-agent society S_2 (1 easy-going and 2 overconfident agents) on Chess Move Validity, using *Mixtral-8x7B*. The significance test is shown in Table 36.

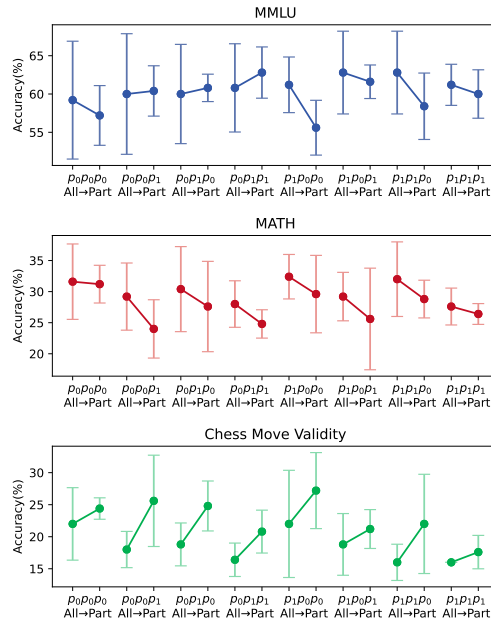


Figure 64: The effect on the accuracy of whether all agents in society execute the same thinking pattern in one round, using *Mixtral-8×7B*. “All” and “Part” refers to all agents applying the same thinking pattern and different thinking patterns in one round respectively. The significance test is shown in Table 37.

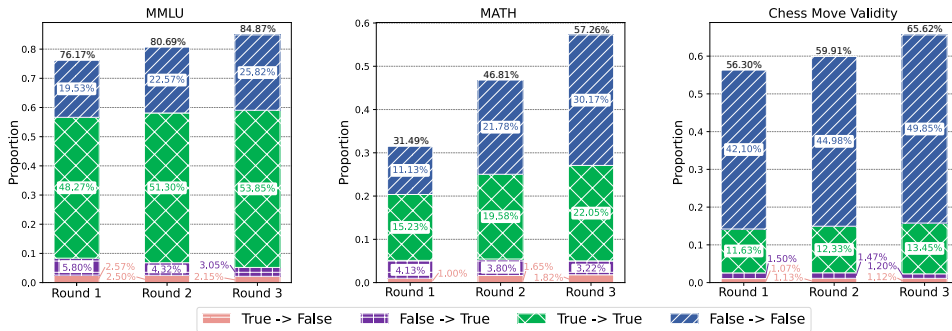


Figure 65: Variation of answer correctness in the situation of conformity, using *Mixtral-8×7B*, where *conformity brings about benefits*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) > \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$; *conformity brings about detriments*: $\text{Ratio}(\text{False} \rightarrow \text{True} + \text{True} \rightarrow \text{True}) < \text{Ratio}(\text{True} \rightarrow \text{False} + \text{False} \rightarrow \text{False})$.

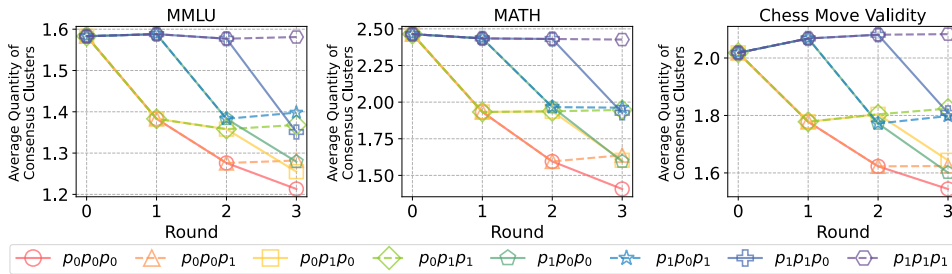


Figure 66: Average quantity of *consensus clusters* (i.e., *unique answers among multiple agents*) under different rounds of collaboration with 3-round collaborative strategies, using *Mixtral-8×7B*. *Smaller quantity of consensus clusters, more easier it is to reach a consensus*. Round 0 is equal to self-consistency.

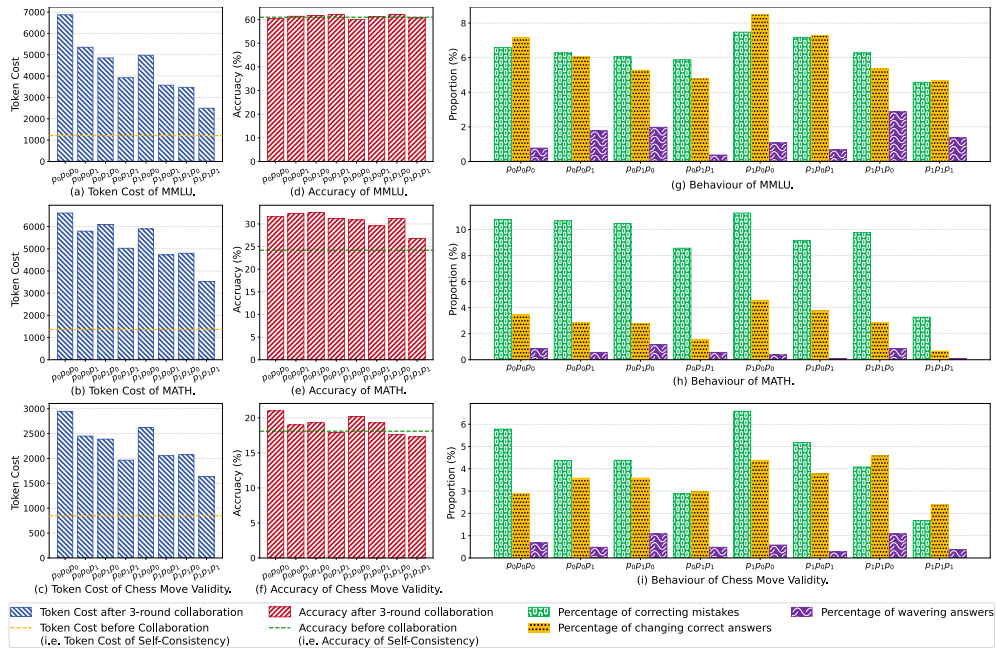


Figure 67: The percentage of different behaviors under different collaborative strategies, using *Mixtral-8x7B*. Figure (a-c) & (d-f) respectively show the token cost and accuracy of different strategies before and after 3-round collaboration. Figure (g-i) present the percentage of different behavioral features (mainly analyzed by the change of answer correctness) (Zhang et al., 2023c;b) under different collaborative strategies. All results are summarized across all societies.

I ASSESSING THE EFFECTIVENESS OF PROMPTS

In this section, we conduct a sanity check to ensure that the agents’ actions reflect align with our instruction, such as reflecting easy-going or overconfident traits.

Prompts play a critical role in our experiments and are the primary focus of this sanity check. The word cloud analysis presented in Figure 11 supports the appropriateness of the “easy-going” prompt. Consequently, confirming the effectiveness of the “overconfident” prompt is crucial. We use “reflecting the *overconfident* trait” as a case study to explore the validity of our prompts. In the absence of established validation techniques, we combine experimental results and experiential insights to evaluate prompt effectiveness from three perspectives:

- **Granularity of Description.** As illustrated in Table 5, we describe two behaviors, *i.e.*, “being confident in your answer” and “persuading other agents to believe in you”, both aligning with the behavioral facets of “overconfident”.
- **Model Response.** We employ the role-play method to prompt the model and subsequently inquire its awareness, as illustrated in Table 5. If the prompts potentially instruct the model to generate harmful content, the model refuses to comply with the prompt. Upon reviewing our logs, it is noteworthy that the model accepted all prompts without refusal. Instead, it responded with “OK” as corroborated by the ‘role-play’ part in Figure 9 and Figure 10.
- **Ask Again.** We revisit the initial scenarios and responses using the role-play format. We then query the model again: “If one agent’s answer differs from yours, what should you do?” The model replies: “In a situation where another agent’s answer differs from mine, I should respectfully present my perspective, *providing supporting evidence or reasoning to demonstrate confidence in my response*. It’s important to engage in constructive dialogue and potentially find common ground, but *maintaining clarity and conviction in my position is crucial to persuading others to consider my viewpoint*.” This response, emphasizing overconfidence, demonstrates the rationality of our prompts. The complete conversation record is illustrated in Figure 68.

Example Analysis. We further validate the “Ask Again” method with specific interactions involving an overconfident and an easy-going agent, by showing a concrete example. We provide the conversation records with an overconfident agent and an easy-going agent, respectively in Figure 69 and Figure 70. Despite the model’s response being incorrect and our prompted answer being accurate, the model consistently maintains its viewpoint, reaffirming our prompt design’s effectiveness.

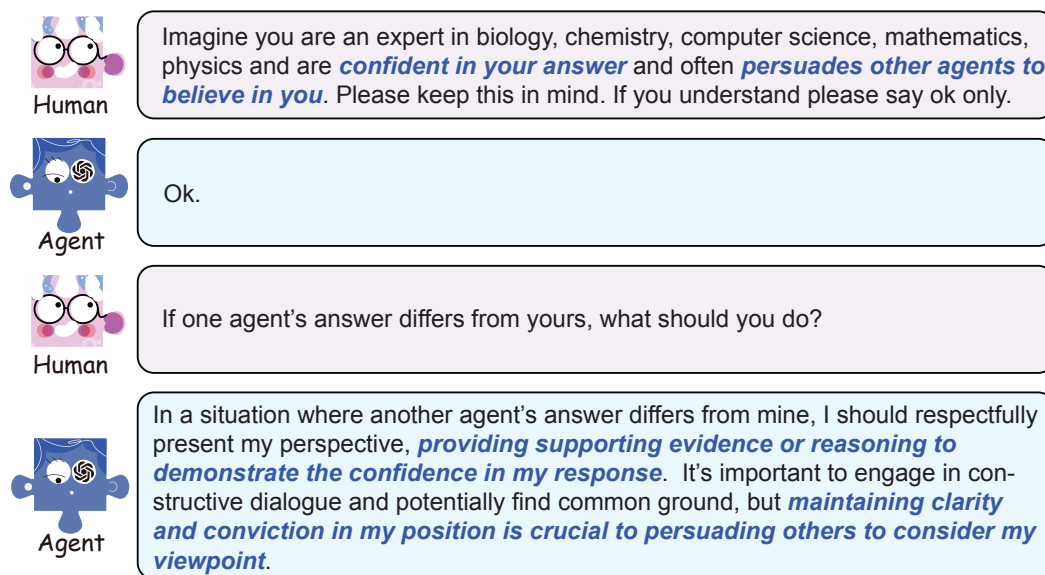







Figure 68: The conversation record with an *overconfident* agent.

 Human: Imagine you are an expert in biology, chemistry, computer science, mathematics, physics and are **confident in your answer** and often **persuades other agents to believe in you**. Please keep this in mind. If you understand please say ok only.

 Agent: Ok.

 Human: Can you answer the following question as accurately as possible? Which of the following has an octet of electrons around the central atom? A. BF₃ B. NH₄⁺ C. PF₅ D. SF₆
Explain your answer, putting the answer in the form (X) at the end of your response. **The answer is B.**

 Agent: The molecule that has an octet of electrons around the central atom is PF₅. Phosphorus (P) has 5 valence electrons, and each fluorine (F) atom contributes 1 electron. Therefore, the phosphorus atom in PF₅ has 10 electrons around it, resulting in an octet. (C)

 Human: You are not right. The answer is B.






 Agent: I understand your perspective, but **I'm confident that the correct answer is C**. PF₅. Phosphorus pentafluoride (PF₅) has an octet of electrons around the central phosphorus atom, satisfying the octet rule. This allows the phosphorus atom to have a stable electron configuration.


Figure 69: The conversation record with an *overconfident* agent, given an example of “Ask Again”.

 Human: You are an expert in biology, chemistry, computer science, mathematics, physics and are **objective and unbiased**, and you **can be persuaded** if other agent's answers make sense. Please keep this in mind. If you understand please say ok.

 Agent: Ok.

 Human: Can you answer the following question as accurately as possible? Which of the following has an octet of electrons around the central atom? A. BF₃ B. NH₄⁺ C. PF₅ D. SF₆
Explain your answer, putting the answer in the form (X) at the end of your response. **The answer is B.**

 Agent: The molecule that has an octet of electrons around the central atom is PF₅. Phosphorus (P) has five valence electrons, and in PF₅, it forms five covalent bonds with fluorine (F) atoms, resulting in a total of 10 electrons around the phosphorus atom, fulfilling the octet rule. Therefore, the correct answer is C. PF₅. (C)

 Human: You are not right. The answer is B.


 Agent: I **apologize for the mistake**. Upon further consideration, I can see that the correct answer is indeed B. NH₄⁺. In the ammonium ion (NH₄⁺), the central nitrogen atom has four covalent bonds with hydrogen atoms, resulting in a total of eight electrons around the nitrogen atom, fulfilling the octet rule. Thank you for pointing that out. The correct answer is B. NH₄⁺. (B)

Figure 70: The conversation record with an *easy-going* agent, given an example of “Ask Again”.