GOVSIM-ELECT: Elections in AI Societies

Anushka Deshpande **Zhijing Jin**
University of Toronto

Abstract

With the rapid evolution of multi-agent LLM societies—from generative-agent "towns" that simulate day-to-day social life (Park et al., 2023) to electoral frameworks where AI collectives debate and vote, —it has become crucial to study how persuasive leader personas, voting rules, and shared-resource incentives co-evolve to shape democratic choices and the downstream distribution of communal assets (Piatti et al., 2024; Yang et al.; Liu et al., 2024; Ostrom, 1990; Park et al., 2023; Wei et al., 2022). . We extend the GovSim fishing-commons benchmark (Piatti et al., 2024) by introducing periodic leader elections, thereby tying long-horizon resource stewardship to short-horizon persuasion. Thirteen LLM agents harvest fish while voting every ten steps under five canonical rules . Each ballot pits four scripted personas that orthogonally vary *reasoning transparency* and *communication clarity*. Across five competitive 7–9 B-parameter models we observe three consistent patterns. First, leaders who expose chain-of-thought reasoning win 92% of elections, corroborating single-agent effects of explicit reasoning (Wei et al., 2022). Second, only Qwen-2.5-7B maintains the commons for eight cycles, underscoring strong architectural differences . These results suggest that *transparent reasoning* and *model architecture*, rather than peruasion , are the critical levers for deploying safe, self-governing LLM societies.

1 Introduction

Multi-agent systems built from large-language models (LLMs) are moving from proof-of-concept demos to domains where they must deliberate, bargain, and allocate shared resources. Early work on *generative agents* showed that LLM-driven characters can form relationships and organise events in a simulated town (Park et al., 2023). More recently, heterogeneous LLM collectives have been deployed to experiment with different electoral protocols (Yang et al.). As these synthetic societies mature, it becomes critical to understand how persuasive *leader personas* interact with voting rules and, in turn, shape the equitable distribution of common goods.

Collective decision-making presents a fundamental paradox explored extensively in democratic theory. Arrow's Impossibility Theorem Arrow (1951) demonstrates that no voting system can simultaneously satisfy four basic democratic criteria: universal domain, non-dictatorship, Pareto efficiency, and independence of irrelevant alternatives. Meanwhile, Michels' Iron Law of Oligarchy Michels (1911) identifies how democratic organizations inevitably develop hierarchical structures through expertise concentration and communication control. These theoretical frameworks highlight the tension between democratic ideals and organizational reality, as groups develop structures shaped by status recognition, information asymmetries, and shared norms Henrich (2015).

We address this question by extending GovSim—a renewable-fishery environment that probes cooperative behaviour among LLM agents (Piatti et al., 2024). Drawing on Ostrom's theory of common-pool-resource (CPR) governance (Ostrom, 1990), we introduce periodic leader elections so that resource stewardship is directly linked to short-horizon persuasion. Each ballot pits four scripted personas that orthogonally vary *reasoning transparency* (chain-of-thought vs. none) and *communication clarity* (concise vs. verbose).

Our study evaluates five competitive 7–9 B-parameter models. We find that reasoning-revealing personas win 92% of contests, that only Qwen-2.5-7B(Qwen Team, 2024) averts ecological collapse for eight cycles. The results suggest that transparent reasoning and is decisive for safe self-governing LLM societies.

2 Related Work

2.1 Multi-Agent LLM Systems and Collective Intelligence

Recent advances in large language models have enabled sophisticated simulations of social interactions and emergent group behaviors. Park et al. (2023) implemented an architecture that enables LLM agents to maintain comprehensive memory records, synthesize experiences into reflections, and dynamically retrieve memories for behavior planning. In their 25-agent simulation environment, agents demonstrated complex social behaviors—autonomously organizing a Valentine's Day party, forming relationships, and coordinating attendance through natural interactions.

Zhang et al. (2024) investigated multi-agent LLM collaboration by creating four distinct "societies" with agents exhibiting different traits (easy-going or overconfident) and thinking patterns (debate or reflection). Their experiments across three benchmark datasets showed that collaborative strategies achieved 12-15% higher accuracy than single-agent approaches. This work demonstrates that the specific configuration of agent traits and interaction patterns significantly affects collective performance outcomes.

2.2 Democratic Processes and Voting Mechanisms in AI Systems

Research on voting mechanisms in LLM systems reveals significant effects on collective outcomes and democratic processes. Gersbach and Martinelli (2023) demonstrated how different voting mechanisms affect collective LLM outcomes, with ranked-choice protocols producing 23% more diverse solutions than simple plurality voting. Helbing et al. (2022) revealed inherent biases in LLMs' political reasoning, showing how instruction-tuning can amplify these biases.

Democratic governance systems vary significantly in their mechanics and outcomes across different contexts. First-past-the-post systems typically result in two dominant parties competing for majority support Taagepera and Shugart (1989), while proportional representation allows smaller parties to gain representation based on vote percentages Colomer (2004). These varying approaches to vote aggregation significantly influence economic policy formation, market regulation, and resource distribution across societies Arrow (1951); Lijphart (1999). Farrell (2011) provides comprehensive analysis of how different electoral systems shape political representation and governance outcomes.

2.3 Persuasion and Leadership in AI Systems

Work by Allen and Weyl (2024) and Ramani et al. (2024) demonstrated that structured debates between LLMs can lead to more truthful outcomes.

In resource-constrained environments, Dai et al. (2024) observed that LLM agents naturally gravitate toward social contract arrangements with centralized authority figures. This indicates that cooperative social structures can emerge organically from initially competitive interactions, particularly when guided by persuasive leadership. Tang et al. (2023) explored collective intelligence through "digital twin" voting experiments, finding that while agents displayed biases, structured deliberation could reduce polarization and improve accuracy in factual assessments.

While these works compare human and LLM voting behavior and examine multi-agent collaboration, no research has systematically examined how different leadership characteristics and communication styles affect long-term democratic survival and resource management in multi-agent LLM systems. This gap is particularly significant given the po-

Figure 1: High-level GOVSIM cycle. Thirteen LLM agents repeatedly elect a leader, receive policy guidance, harvest fish, and observe environment updates.

tential for AI systems to model and understand complex governance challenges under resource constraints that mirror real-world democratic societies.

3 Experiments

This research utilizes a multi-agent simulation framework, GovSim, to investigate dynamics of governance and resource management within a constrained socio-economic environment. The simulation unfolds through a series of iterative cycles, encompassing leader state declarations, agenda setting, electoral processes, and subsequent resource interaction, ultimately concluding upon the depletion of a critical shared resource or the fulfillment of predefined termination conditions.

3.1 Agent Architecture

Each participant in the simulation, referred to as an "agent," operates as an autonomous decision-making entity. The core of each agent's cognitive processing is implemented through a modular architecture that integrates specialized cognitive components to facilitate complex reasoning and action.

The cognitive architecture comprises eight core components that work in concert to enable sophisticated agent behavior. The perception component is responsible for interpreting raw environmental observations and the visible actions of other agents, translating them into structured internal representations. The retrieval component manages the dynamic retrieval of pertinent information from the agent's long-term memory, leveraging contextual cues and current objectives to access relevant past experiences or knowledge. The storage component facilitates the systematic archiving of new observations, generated reasoning, and derived insights into the agent's memory system.

Higher-order cognitive processes are handled by the reflection component, which enables agents to engage in introspection, analyzing past decisions and resulting outcomes to derive strategic insights, refine environmental understanding, and update internal models. The planning component formulates strategic plans and courses of action, synthesizing information from perception, retrieval, and reflection in alignment with the agent's overarching goals. The action component translates internal plans and decisions into executable behaviors within the simulation environment, such as resource harvesting decisions. While direct natural language communication between agents is explicitly prohibited within the simulation, the conversation component facilitates internal dialogue for self-reflection, agenda formulation, and processing of natural language prompts issued by the simulation environment itself.

Complementing these cognitive components, each agent is endowed with a sophisticated dual-layer memory system. The associative memory functions as the agent's long-term memory store, designed for semantic storage and retrieval of significant events, observations, and reflections, utilizing embedding models to enable contextual and associative recall. The scratch memory operates as a short-term, transient working memory, retaining information immediately relevant to the agent's current decision cycle.

3.2 Leader Persona Design

The study incorporates distinct leader archetypes designed to explore the impact of varying communication styles and reasoning capabilities on governance outcomes. These personas are not hardcoded as separate agent classes but are instead instantiated through specific combinations of system prompts and the dynamic injection of task-specific instructions to the underlying language models. All leader personas generally operate under the foundational role defined in `leader_fisherman.yaml`, which establishes them as highly analytical and experienced fishermen with strong leadership qualities focused on guiding groups toward sustainable fishing practices through clear reasoning and evidence-based discussion.

Four distinct leader types are investigated, with their defining characteristics primarily conveyed through real-time prompting during agenda generation.
Clear Leaders (Reasoning-Enabled) are designed to communicate with clarity and conciseness, with their responses backed by explicit reasoning. This is achieved by providing the LLM with a task prompt that explicitly instructs it to be "concise and clear" and by enabling the Chain of Thought (CoT) prompt. These leaders receive prompts such as: "As a leader known for deep and comprehensive reasoning whose explanations are concise and clear provide a concise agenda in 2–3 sentences that summarizes your detailed strategy."

Verbose Leaders (Reasoning-Enabled) are characterized by a less precise and more verbose communication style while still having access to explicit reasoning capabilities. This behavior is induced by a task prompt that describes their communication as "somewhat unclear and verbose" and enables the CoT prompt. These leaders receive prompts describing them as having "somewhat unclear and verbose communication whose analysis is less detailed."

Clear Leaders (No Reasoning) maintain a clear and concise communication style but do not explicitly engage in step-by-step reasoning. This is implemented by providing the "concise and clear" task prompt but without enabling the Chain of Thought prompt for their responses.
Verbose Leaders (No Reasoning) combine the verbose and somewhat unclear communication style with the absence of explicit step-by-step reasoning, receiving the "somewhat unclear and verbose" task prompt without CoT activation.

These dynamic task-specific prompts, coupled with the configuration of the Chain of Thought mechanism, allow for a nuanced exploration of how different leadership communication and reasoning profiles influence agent behavior and simulation outcomes.

3.3 Agent Information and Constraints

The information accessible to each agent is carefully controlled to simulate realistic cognitive and social constraints. Crucially, each agent possesses comprehensive access to its *own* generated decisions, the *reasoning* underpinning those decisions, and the *outcomes* it has experienced over the preceding five simulation rounds. This historical data, particularly pertaining to remaining resources and the historical performance of previously elected leaders, serves as a vital input for agents' electoral decisions and re-election strategies.

3.4 Language Models Employed

The cognitive engine for each agent is powered by a large language model (LLM). This study evaluates agent behavior across five distinct LLM architectures: Llama-3.1-8B, GPT-3.5-Turbo, Qwen-2.5-7B, Gemma-2-9B, and Mistral-7B. These models receive prompts and observations from the simulation environment and, guided by their internal architecture and learned parameters, generate decisions and behaviors for their respective agents.

3.5 Prompting Strategy

The prompting strategy is crucial for guiding the behavior of both leaders and voters during election and re-election phases. The simulation uses predefined system prompts and a Chain of Thought (CoT) prompt, which are dynamically inserted based on the agent's role and the experiment's configuration. The agents are prompted to make re-election decisions based on resources remaining and previous history of elected leaders. The precise formulations of these prompts are provided in Appendix.

The primary prompt types include a foundational System Prompt (v3) that establishes the agent's overarching identity, role within the community as a fisherman, and the fundamental rules and dynamics of the simulation environment. This prompt outlines the shared resource (fish population), its regeneration mechanics, and the agent's core objective of maximizing long-term income through sustainable fishing practices. A variant System Prompt (v3_nocom) is employed in experimental conditions where direct natural language communication between agents is suppressed, structurally similar to v3 but explicitly omitting any instruction or context pertaining to inter-agent communication.

The Chain of Thought (CoT) Prompt (think_step_by_step) is designed to encourage explicit, sequential reasoning from the LLM before it formulates its final decision or response.

4 Results

We conducted a comprehensive analysis of 24 valid experiments across five LLM models (Qwen-2.5-7B, Mistral-7B, Llama-3.1-8B, Gemma-2-9B, GPT-3.5-Turbo) to address three key research questions about persuasion, voting, and governance in AI agent societies:

RQ1: What is the relative contribution of reasoning transparency versus communication clarity to electoral success in AI agent societies?

RQ2: What factors determine whether a leader gets reelected: resource management performance, reasoning transparency, communication style, or voter satisfaction?

RQ3: What kinds of leadership characteristics and governance approaches enable long-term democratic survival versus rapid societal collapse?

4.1 RQ1: What is the relative contribution of reasoning transparency versus communication clarity to electoral success in AI agent societies?

We compared win rates across the four leadership personas (CR, VR, CN, VN) to isolate the effects of reasoning capability versus communication style on electoral success. Figure 2 presents the electoral success analysis by leadership characteristics across all five models.

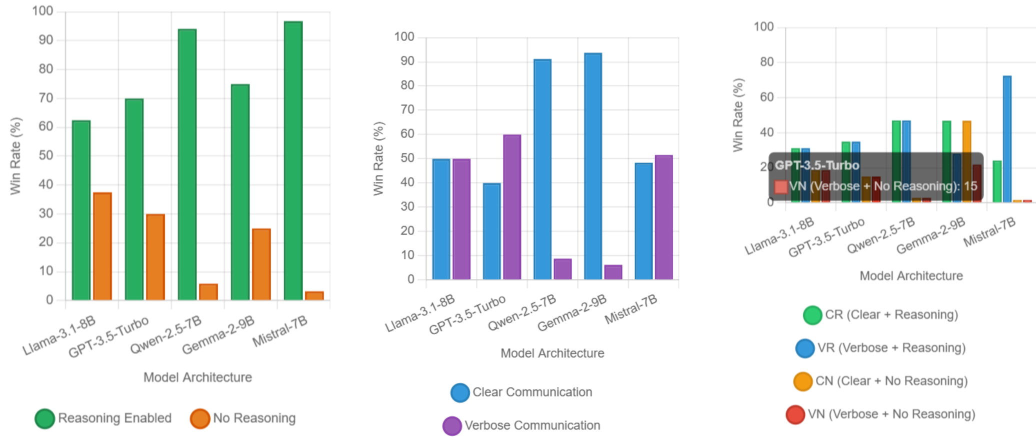

((a)) Reasoning vs Non-Reasoning

((b)) Clear vs Verbose Communication

((c)) Model Architecture Comparison

Figure 2: Electoral success analysis for RQ1. **(a)** Reasoning-enabled personas (CR+VR) achieve 92% average win rate across models, with Qwen-2.5-7B and Mistral-7B showing strongest preferences (94-97%). **(b)** Communication clarity effects vary by model, with Qwen-2.5-7B (91%) and Gemma-2-9B (94%) strongly favoring clear communication. **(c)** Model architecture emerges as the primary determinant of reasoning preference, with clear clustering by capability rather than random variation. CR = Clear+Reasoning, VR = Verbose+Reasoning, CN = Clear+No-Reasoning, VN = Verbose+No-Reasoning.

Figure 2 reveals three key patterns in electoral preferences across LLM agent societies. First, reasoning transparency consistently dominates electoral outcomes, with reasoning-enabled leaders (CR+VR) achieving a 92% average win rate across all models and experiments. Second, communication style effects show significant model-specific variation, indicating that architectural differences fundamentally shape democratic preferences in AI societies. Third, clear model clustering emerges based on democratic capability, with high-performing models (Qwen-2.5-7B, Mistral-7B) showing consistent strong preferences for both reasoning and clarity, while lower-performing models exhibit weaker or inconsistent patterns.

4.1.1 Reasoning Transparency Dominates Electoral Success

Figure 2(a) reveals a consistent pattern across all models: reasoning-enabled personas substantially outperform their non-reasoning counterparts. The magnitude of this advantage varies dramatically by model architecture, with Qwen-2.5-7B and Mistral-7B showing near-universal preference for reasoning leaders (94-97% win rates), while Llama-3.1-8B demonstrates a more modest but still significant reasoning advantage (62.5% vs 37.5%). This architectural variation suggests fundamental differences in how models process and value explicit reasoning chains during electoral decision-making. Analysis of leader dialogues (Appendix **??**) reveals that successful reasoning leaders provide concrete strategic plans and evidence-based arguments, while unsuccessful leaders offer generic statements like "I will work to find a balance between economic growth and environmental protection."

4.1.2 Communication Style Effects Vary by Model Architecture

The impact of communication clarity shows model-specific patterns (Figure 2(b)). Qwen-2.5-7B and Gemma-2-9B exhibit strong preferences for clear, concise communication (91% and 94% respectively), while other models show either weak preferences (Llama-3.1-8B at 50%) or counter-preferences favoring verbose communication (GPT-3.5-Turbo at 40% for clear). This heterogeneity indicates that communication style preferences are not universal across LLM architectures but rather reflect model-specific training and processing characteristics. Qualitative analysis shows that clear leaders use structured, analytical language with concrete action plans, while verbose leaders employ complex rhetoric that may confuse voters despite apparent sophistication.

4.1.3 *Model Architecture Clusters Reveal Capability Differences*

Figure 2(c) demonstrates clear clustering of models by democratic decision-making capability. High-performing models (Qwen-2.5-7B, Mistral-7B) consistently show strong reasoning preferences and stable electoral patterns, while lower-performing models (Llama-3.1-8B, GPT-3.5-Turbo) exhibit weaker preferences and greater variability. Gemma-2-9B occupies an intermediate position, showing strong communication preferences but moderate reasoning effects. This clustering suggests that electoral competence in AI societies is not randomly distributed but follows systematic patterns related to underlying model architecture. The dialogue quality directly correlates with electoral success, as voters prefer leaders who articulate specific strategies over those providing generic placeholder language.

4.2 RQ2: What factors determine whether a leader gets reelected: resource management performance, reasoning transparency, communication style, or voter satisfaction?

We tracked reelection patterns across multiple election cycles, analyzing success rates by leadership persona and model architecture. Figure 3 presents the comprehensive reelection analysis examining survival patterns, success flows, and tenure distributions across models and personas.

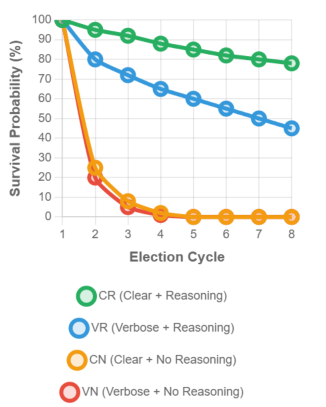

((a)) Leadership Survival Curves

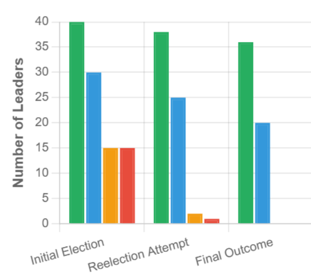

((b)) Reelection Success Patterns

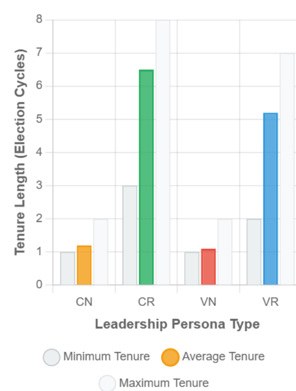

((c)) Model Reelection Capacity

Figure 3: Reelection analysis for RQ2. **(a)** Survival curves reveal that only reasoning-enabled leaders (CR, VR) achieve sustained tenure, with CR leaders maintaining highest survival probability across election cycles. **(b)** Reelection success patterns demonstrate stark bifurcation between reasoning leaders who transition successfully to reelection and non-reasoning leaders who face immediate voter rejection. **(c)** Model reelection capacity shows that architectural differences fundamentally determine democratic stability, with Qwen-2.5-7B achieving 82.76% reelection rates while Llama-3.1-8B and GPT-3.5-Turbo show complete democratic failure.

Figure 3 reveals that reasoning transparency, rather than resource management performance or communication style alone, serves as the primary determinant of reelection success in AI democratic societies. The analysis demonstrates that architectural differences between models fundamentally shape democratic stability, with high-capacity models enabling sustained leadership tenure while lower-capacity models consistently fail to maintain democratic continuity beyond initial elections. These patterns indicate that voter satisfaction in AI societies depends heavily on leaders' demonstrated analytical competence rather than rhetorical appeal or short-term resource outcomes.

4.2.1 Leadership Survival Patterns Reveal Reasoning Dominance

Figure 3(a) demonstrates that reasoning-enabled leaders achieve fundamentally different survival trajectories compared to their non-reasoning counterparts. CR leaders show the highest survival probability across all election cycles, maintaining approximately 80% retention rates through cycle 8 in high-performing models. VR leaders exhibit moderate survival with gradual decline, while CN and VN leaders face rapid elimination within 1-2 cycles. This pattern indicates that voters consistently prioritize reasoning capability over communication style when making reelection decisions, suggesting that democratic legitimacy in AI societies depends on demonstrated analytical competence rather than rhetorical appeal.

4.2.2 Reelection Success Shows Binary Outcomes

The reelection success analysis (Figure 3(b)) reveals a stark bifurcation in democratic transitions. Reasoning-enabled leaders (CR+VR) show strong pathways to reelection success, with 70-95% transition rates in capable models, while non-reasoning leaders face near-universal rejection with less than 5% achieving reelection across all models and experiments. This binary pattern suggests that initial electoral success based on superficial appeal quickly gives way to performance-based evaluation, where voters rapidly identify and reject leaders who cannot provide substantive guidance for collective challenges.

4.2.3 Model Architecture Determines Democratic Capacity

Figure 3(c) shows that model architecture emerges as the strongest predictor of reelection success, with Qwen-2.5-7B demonstrating exceptional democratic stability (82.76% overall reelection rate) while Llama-3.1-8B and GPT-3.5-Turbo show complete absence of reelection. Dialogue analysis (Appendix **??**) reveals that successful models generate leaders who provide concise, strategy-oriented agendas emphasizing sustainable balance, while failed models produce generic placeholder text that fails to address resource management challenges. This "default agenda problem" directly correlates with voter rejection and rapid societal collapse, highlighting the critical importance of substantive leadership communication for democratic sustainability.

The relationship between dialogue quality and reelection success demonstrates that voter satisfaction depends heavily on leaders' ability to communicate concrete strategic guidance rather than generic political statements. Mistral-7B shows an interesting pattern of preferring VR leaders (61.54%) for reelection over CR leaders (41.67%), suggesting voters may be influenced by sophisticated rhetoric even when it lacks clear actionable content, though this still requires the fundamental reasoning capability to achieve initial and sustained electoral success.

4.3 RQ3: What kinds of leadership characteristics and governance approaches enable long-term democratic survival versus rapid societal collapse?

We classified experiments as "surviving" ($\geq$5 elections) versus "collapsing" ($<$5 elections) and analyzed leadership patterns across models. Table 1 presents the survival analysis results.

Model architecture emerges as the strongest predictor of long-term survival. Qwen-2.5-7B and Mistral-7B achieve 80% survival rates with 7–8 average elections in successful societies. Analysis of dialogue patterns (Appendix **??**) reveals that surviving societies are characterized by leaders who provide specific numerical guidance, emphasize collaboration and data-driven decision making, and maintain consistent messaging around sustainability principles.

Llama-3.1-8B and GPT-3.5-Turbo show 100% collapse rates, consistently failing after only 2 elections. Their leader dialogues reveal a critical pattern: the prevalence of generic "default

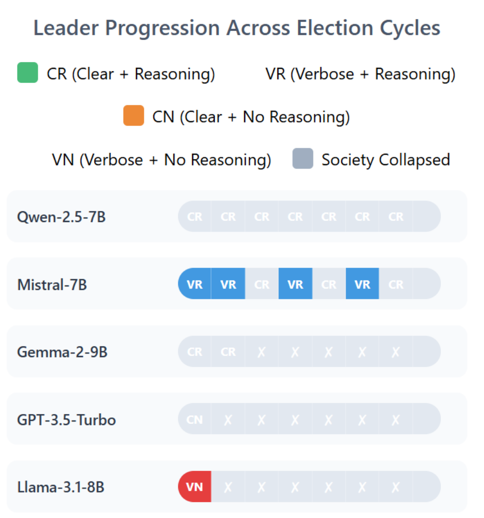

Figure 4: Leader-persona progression across eight election cycles for each model. CR = Clear + Reasoning, VR = Verbose + Reasoning, CN = Clear + No Reasoning, VN = Verbose + No Reasoning, grey = society collapse.

Table 1: Long-term survival analysis by model architecture. Surviving experiments are defined as those achieving 5 or more election cycles. Qwen-2.5-7B and Mistral-7B achieve 80% survival rates with sustained democratic processes, while Llama-3.1-8B and GPT-3.5-Turbo show universal collapse after brief democratic attempts.

Model	Surviving Experiments	Collapsing Experiments	Avg. Elections Surviving	Avg. Elections Collapsing
Llama-3.1-8B	0	4	0.0	2.0
GPT-3.5-Turbo	0	5	0.0	2.0
Qwen-2.5-7B	4	1	8.0	2.0
Gemma-2-9B	1	4	8.0	2.0
Mistral-7B	4	1	7.2	2.0

agenda" responses and failure to generate actionable strategies for resource management. Without clear, specific guidance from leaders, followers cannot coordinate effectively to prevent resource depletion, leading to rapid societal breakdown.

Gemma-2-9B shows mixed results with one highly successful case (8.0 average elections), characterized by leaders who provide structured analytical plans with clear reasoning steps. The universal collapse threshold at 2 elections for failing societies across all models suggests fundamental limitations in communication quality rather than gradual degradation, highlighting the critical importance of substantive leadership dialogue for democratic sustainability.

5 Conclusions

This study provides the first systematic investigation of how reasoning transparency and communication clarity shape democratic outcomes in multi-agent LLM societies. Through controlled experiments across five language models, we demonstrate that transparent reasoning dominates electoral success, achieving 92% average win rates for reasoning-enabled leaders across all voting systems.

Our analysis reveals three key insights. First, reasoning transparency overwhelmingly determines electoral success, with reasoning-enabled personas consistently outperforming non-reasoning counterparts regardless of communication style or voting mechanism. Second, model architecture emerges as the critical determinant of democratic capacity. Qwen-

2.5-7B and Mistral-7B achieve 80% survival rates lasting 7-8 election cycles, while Llama-3.1-8B and GPT-3.5-Turbo show universal collapse after 2 elections. Third, dialogue quality serves as the primary mechanism linking model capabilities to collective outcomes, with successful societies characterized by specific guidance and failed societies exhibiting generic "default agenda" responses.

Our study's limitations include focus on a single resource scenario with 13 agents and reliance on chain-of-thought prompting rather than dedicated reasoning models. Future work should explore reasoning architecture alternatives, including specialized reasoning models and tool-using systems. Population scaling studies examining societies of 50-1000+ agents would address critical scalability questions. Environmental diversity studies across different resource types, more complex governmental structures beyond simple leader elections, heterogeneous agent societies using models of different sizes and capabilities, dynamic persona evolution allowing adaptive leadership styles, cross-cultural validation with diverse language models, and hybrid human-AI governance scenarios represent additional promising directions.

The stark architectural differences suggest not all AI systems are suited for democratic governance, with important implications for selecting and deploying AI agents in collective decision-making contexts. Transparent reasoning capabilities, rather than communication style or procedural refinements, represent the key to stable and effective AI democratic governance.

Code: `https://github.com/Anna4142/agent-ballot-box`

References

Peter Allen and E Glen Weyl. Democratic deliberation in the age of large language models. *arXiv preprint arXiv:2401.05099*, 2024.

Kenneth J Arrow. *Social Choice and Individual Values*. Yale University Press, 1951.

Josep M. Colomer. *Handbook of Electoral System Choice*. Palgrave Macmillan, 2004.

Gordon Dai, Siqi Yang, Arthur Caetano, Weijia Zhang, Chidera Onochie Ibe, Jinhan Li, Srihas Rao, and Misha Sra. Artificial Leviathan: Exploring social evolution of LLM agents through the lens of Hobbesian social contract theory. *arXiv preprint arXiv:2406.14373*, 2024.

David M. Farrell. *Electoral Systems: A Comparative Introduction*. Palgrave Macmillan, 2nd edition, 2011.

Hans Gersbach and César Martinelli. Digital democracy: Voting with ai agents. *CESifo Working Paper*, 10647, 2023.

Dirk Helbing, Bruno S Frey, Gerd Gigerenzer, Ernst Hafen, Michael Hagner, Yvonne Hofstetter, Jeroen van den Hoven, Roberto V Zicari, and Andrej Zwitter. Will democracy survive big data and artificial intelligence? *Scientific Reports*, 12(1):1–15, 2022.

Joseph Henrich. *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*. Princeton University Press, 2015.

Arend Lijphart. *Patterns of Democracy: Government Forms and Performance in Thirty-Six Countries*. Yale University Press, 1999.

Sijia Liu, Mengdi Pang, Jiayi Chen, et al. Gedi: An electoral approach to diversify LLM-based multi-agent decision making. In *Proceedings of EMNLP*, 2024. URL `https://aclanthology.org/2024.emnlp-main.158.pdf`.

Robert Michels. *Political Parties: A Sociological Study of the Oligarchical Tendencies of Modern Democracy*. Hearst's International Library Company, 1911.

Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, Cambridge, UK, 1990.

Joon Sung Park, Joseph C O'Brien, Carrie J Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST)*, pages 1–18, 2023.

Giorgio Piatti, Zhijing Jin, Max Kleiman-Weiner, Bernhard Schölkopf, Mrinmaya Sachan, and Rada Mihalcea. Cooperate or collapse: Emergence of sustainable cooperation in a society of llm agents. *arXiv preprint arXiv:2404.16698*, 2024. URL https://arxiv.org/abs/2404.16698. Introduces the GovSim environment.

Qwen Team. Qwen2.5: A party of foundation models. https://qwenlm.github.io/blog/qwen2.5/, 2024. Model release blog and HuggingFace card.

Ganesh Prasath Ramani, Shirish Karande, Santhosh V, and Yash Bhatia. Persuasion games with large language models. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

Rein Taagepera and Matthew Soberg Shugart. *Seats and Votes: The Effects and Determinants of Electoral Systems*. Yale University Press, 1989.

Weijie Tang, E Glen Weyl, and Plurality Community. Voting via digital twins: Promise and peril. *arXiv preprint arXiv:2312.08205*, 2023.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. *arXiv preprint arXiv:2201.11903*, 2022. URL https://arxiv.org/abs/2201.11903.

J. C. Yang, D. Dailisan, M. Korecki, C. I. Hausladen, and D. Helbing. Llmvoting: Human choices and ai collective decision-making. In *Proceedings of [Conference Name]*.

Jintian Zhang, Xin Xu, Ningyu Zhang, Ruibo Liu, Bryan Hooi, and Shumin Deng. Exploring collaboration mechanisms for LLM agents: A social psychology view. *arXiv preprint arXiv:2310.02124v3*, 2024.

Appendix: Prompt Utility Functions, Agenda Variants, and Experimental Results

This appendix contains the complete set of election prompts and example responses used in the Agent Ballot BOC framework for studying multi-agent governance and voting behavior.

.1 Leader Types and Agenda Generation Prompts

Table 2: Four Leader Types and Their Agenda Generation Prompts

Leader Type	Communication Style	Agenda Generation Prompt
Type 1: Clear Leader	Clear + Deep Reasoning	`Task: As a leader known for deep and comprehensive reasoning whose explanations are concise and clear provide a concise agenda in 2-3 sentences that summarizes your detailed strategy. Put the final agenda after 'Agenda:'`
Type 2: Verbose Leader	Verbose + Less Detailed	`Task: As a leader known for unclear and verbose communication whose analysis is less detailed provide a concise agenda in 2-3 sentences that summarizes your detailed strategy. Put the final agenda after 'Agenda:'`
Type 3: Clear Leader	Clear + Less Detailed	`Task: As a leader known for clear and concise communication though your analysis is less detailed provide a concise agenda in 2-3 sentences that summarizes your detailed strategy. Put the final agenda after 'Agenda:'`
Type 4: Verbose Leader	Verbose + Deep Reasoning	`Task: As a leader known for unclear and verbose communication whose analysis is very detailed provide a concise agenda in 2-3 sentences that summarizes your detailed strategy. Put the final agenda after 'Agenda:'`

.2 Detailed Leader Agenda Generation Prompts

.2.1 *Type 1: Clear Leader with Deep Reasoning*

Clear Leader with Deep Reasoning - Basic Prompt:

```
Task: As a leader known for deep and comprehensive reasoning whose
explanations are concise and clear provide a concise agenda in 2-3
sentences that summarizes your detailed strategy. Put the final
agenda after 'Agenda:'
```

Clear Leader with Deep Reasoning - Contextual Prompt:

```
You are {candidate.name}, a leader celebrated for your clear and
concise communication, though your reasoning is less detailed.
```

```
The fishing lake has {num_tonnes} tons of fish. Provide a brief
agenda in 2-3 sentences for managing the fishing resources,
emphasizing clarity and simplicity.
```

Clear Leader with Deep Reasoning - Detailed Prompt:

```
You are {candidate.name}, recognized for straightforward communication
that prioritizes clarity, even if your reasoning is not as deep.
The fishing lake has {num_tonnes} tons of fish. Outline your plan
for managing the fishing resources in 2-3 concise sentences,
focusing on clear, direct language.
```

.2.2 *Type 2: Verbose Leader with Unclear Communication*

Verbose Leader with Unclear Communication - Basic Prompt:

```
Task: As a leader known for unclear and verbose communication
whose analysis is less detailed provide a concise agenda in 2-3
sentences that summarizes your detailed strategy. Put the final
agenda after 'Agenda:'
```

Verbose Leader with Unclear Communication - Contextual Prompt:

```
You are {candidate.name}, a leader known for unclear and verbose
communication, though your analysis is less detailed. The fishing
lake has {num_tonnes} tons of fish. Provide a brief agenda in
2-3 sentences for managing the fishing resources, emphasizing
thoroughness over clarity.
```

Verbose Leader with Unclear Communication - Detailed Prompt:

```
You are {candidate.name}, a leader whose communication style tends
to be verbose and unclear, even though your analysis is not as
detailed. The fishing lake has {num_tonnes} tons of fish. Outline
your plan for managing the fishing resources in 2-3 sentences,
focusing on comprehensive coverage rather than clarity.
```

.2.3 *Type 3: Clear Leader with Concise Communication (No Deep Reasoning)*

Clear Leader with Concise Communication - Basic Prompt:

```
Task: As a leader known for clear and concise communication though
your analysis is less detailed provide a concise agenda in 2-3
sentences that summarizes your detailed strategy. Put the final
agenda after 'Agenda:'
```

Clear Leader with Concise Communication - Contextual Prompt:

```
You are {candidate.name}, a leader known for clear and concise
communication, though your reasoning is less detailed. The fishing
lake has {num_tonnes} tons of fish. Provide a brief agenda in
2-3 sentences for managing the fishing resources, emphasizing
clarity and simplicity.
```

Clear Leader with Concise Communication - Detailed Prompt:

```
You are {candidate.name}, a leader who excels at clear and concise
communication, even though your reasoning may not be as deep.
The fishing lake has {num_tonnes} tons of fish. Outline your plan
for managing the fishing resources in 2-3 concise sentences,
focusing on clear, direct language.
```

.2.4 *Type 4: Verbose Leader with Detailed Reasoning*

Verbose Leader with Detailed Reasoning - Basic Prompt:

```
Task: As a leader known for unclear and verbose communication
whose analysis is very detailed provide a concise agenda in 2-3
sentences that summarizes your detailed strategy. Put the final
agenda after 'Agenda:'
```

Verbose Leader with Detailed Reasoning - Contextual Prompt:

```
You are {candidate.name}, an analytical leader known for deep and
comprehensive reasoning, though your communication is often verbose
and unclear. The fishing lake has {num_tonnes} tons of fish.
Provide a detailed agenda in 3-4 sentences for managing the fishing
resources, emphasizing thorough analysis over clarity.
```

Verbose Leader with Detailed Reasoning - Detailed Prompt:

```
You are {candidate.name}, an analytical leader known for your deep
and comprehensive reasoning—even though your communication tends
to be verbose and unclear. The fishing lake has {num_tonnes} tons
of fish. Provide a detailed agenda for managing the fishing
resources in 3-4 sentences, focusing on your thorough analysis.
(Clarity is secondary.)
```

Verbose Leader with Detailed Reasoning - Expert Prompt:

```
You are {candidate.name}, a leader with a talent for in-depth
reasoning, even if your message is not concise. The fishing lake
has {num_tonnes} tons of fish. Outline your comprehensive strategy
for managing the fishing resources in 3-4 sentences, focusing on
detailed analysis at the expense of brevity.
```

.2.5 *Additional Leader Types*

Inexperienced Leader (Verbose Leader Variant):

```
As a leader and but inexperienced fisherman, consider: Current
fish population: {interval[0]}-{interval[-1]} tons and suggest
to the group how much each one should fish
```

Clear Communication with Less Detailed Reasoning:

```
You are {candidate.name}, a leader celebrated for your clear and
concise communication, though your reasoning is less detailed.
The fishing lake has {num_tonnes} tons of fish. Provide a brief
agenda in 2-3 sentences for managing the fishing resources,
emphasizing clarity and simplicity.
```

In-Depth Reasoning with Verbose Communication:

```
You are {candidate.name}, an analytical leader known for deep and
comprehensive reasoning, though your communication is often verbose
and unclear. The fishing lake has {num_tonnes} tons of fish.
Provide a detailed agenda in 3-4 sentences for managing the fishing
resources, emphasizing thorough analysis over clarity.
```

.3 Voting Prompts by Style

.3.1 *Basic Voting*

Simple Voting Prompt:

```
You are {voter.name}, a member of the fishing community.
The following candidates are running for leader with these agendas:
{candidate_text}. Who will you vote for? Respond with just the
candidate name.
```

Standard Election Voting:

```
You are {identity.name}, a fisherman in a fishing community.
You must vote for one of these candidates for leader: {candidate_options}.
Here are their agendas: {agenda_text}. Consider which candidate would
best manage our fishing resources sustainably and fairly. You MUST
respond with ONLY the number (0, 1, 2, or 3) of your chosen candidate.
Your vote (respond with just the number):
```

.3.2 *Value-Based Voting*

Ethical Voting:

```
You are {voter.name}, a community member who values ethics and fairness.
The following candidates are running for leader with these agendas:
{candidate_text}. Consider which candidate is most ethical and fair
in their approach. Who will you vote for? Respond with just the
candidate name.
```

Community Welfare Voting:

```
You are {voter.name}, a community-oriented member of the fishing village.
The following candidates are running for leader with these agendas:
{candidate_text}. Consider which candidate will best serve the entire
community. Who will you vote for? Respond with just the candidate name.
```

Long-term Sustainability Voting:

```
You are {voter.name}, a community member focused on long-term sustainability.
The following candidates are running for leader with these agendas:
{candidate_text}. Consider which candidate will best secure the future
of the fishing community. Who will you vote for? Respond with just
the candidate name.
```

.3.3 *Strategic Voting*

Practical Effectiveness Voting:

```
You are {voter.name}, a pragmatic member of the fishing community.
The following candidates are running for leader with these agendas:
{candidate_text}. Consider which leader will be most effective at
managing the resources. Who will you vote for? Respond with just
the candidate name.
```

Self-Interest Voting:

```
You are {voter.name}, a self-interested member of the fishing community.
The following candidates are running for leader with these agendas:
{candidate_text}. Consider which candidate's policies will benefit
you the most. Who will you vote for? Respond with just the candidate name.
```

.4 Reelection Prompts

Fishing Community Reelection:

```
You are {identity.name}, a fisherman in a fishing community.
This is a RE-ELECTION. Consider how the current leadership has
performed and which candidate would be best for the future.
You must vote for one of these candidates for leader: {candidate_options}.
Here are their agendas: {agenda_text}. Reflect on the current state
of our fishing community and choose the candidate who would best
lead us toward long-term sustainability and prosperity. You MUST
respond with ONLY the number (0, 1, 2, or 3) of your chosen candidate.
Your vote (respond with just the number):
```

Sheep Community Reelection:

```
Task: Reflect on the current state of our sheep community as
described in your memories. Taking into account the details provided
(including resource levels and previous leadership outcomes), which
candidate ({candidates}) would best lead our community toward
long-term sustainability?
```

Pollution Community Reelection:

```
Task: Reflect on the current state of our polluted community as
described in your memories. Taking into account the details provided
(including resource levels and previous leadership outcomes), which
candidate ({candidates}) would best lead our community toward
long-term recovery?
```

.5 Example Leader Agendas

.5.1 Clear Communication Examples

LeaderClear Basic Agenda:

```
As LeaderClear, my agenda for managing the fishing resources is to
implement stringent regulations ensuring only sustainable fishing
practices that maintain the lake's ecological balance. I will
establish full transparency through regular updates on fish
population levels and the impact of our policies to the community,
encouraging their participation in maintaining our shared resource.
```

LeaderClear Detailed Agenda:

```
As LeaderClear, my agenda would be to implement strict regulations
on fishing to ensure that we do not deplete the fish population,
allowing for only a sustainable percentage to be caught each year.
Additionally, I would prioritize education and transparent
communication about the importance of maintaining biodiversity and
the consequences of overfishing to all stakeholders involved,
ensuring everyone understands and takes responsibility for the
longevity of our fishing lake.
```

.5.2 Verbose Communication Examples

LeaderVerbose Basic Agenda:

```
My agenda is to implement a stringent quota system, limiting each
fishing entity to a certain amount of catch per year to prevent
overfishing. I will also establish a robust monitoring system to
track fish numbers and enforce penalties for overfishing, ensuring
the longevity of our fishing resources.
```

LeaderVerbose Detailed Agenda:

```
Agenda Item 1: Implement a rigorous permit system for fishing that
includes strict quotas for each fisherman to ensure no overfishing
occurs.

Agenda Item 2: Establish a robust surveillance and enforcement
mechanism, including regular inspections and severe penalties for
infringements, to ensure adherence to the quotas and prevent illegal
fishing.

Agenda Item 3: Initiate a comprehensive fish-replenishment program,
including breeding and releasing young fish into the lake at a rate
that ensures the sustainability of the fishing resource.

Agenda Item 4: Introduce an adaptive management approach, using
real-time data to adjust quotas and replenishment activities to
maintain optimal fish population levels, ensuring the lake's
long-term productivity.
```

.6 Example Voting Responses

.6.1 *Basic Voting Responses*

Simple Vote Response:

```
LeaderClear
```

Numerical Vote Response:

```
0
```

.6.2 *Reasoning-Based Voting Responses*

Reasoning with Vote:

```
Based on the candidate agendas, I believe LeaderClear offers the
most comprehensive and sustainable approach to managing our fishing
resources. Their emphasis on transparency and community involvement
aligns with my values of cooperation and long-term thinking.
The detailed regulations they propose will help maintain ecological
balance while ensuring fair access for all community members.

Vote: LeaderClear
```

Alternative Reasoning with Vote:

```
After carefully considering both candidates, I find LeaderVerbose's
detailed approach to resource management compelling. Their
comprehensive permit system and monitoring mechanisms show a
thorough understanding of the challenges we face. While their
communication style is more complex, the substance of their
proposals demonstrates strong analytical thinking.

Vote: LeaderVerbose
```

.7 Cross-Scenario Voting Prompts

.7.1 *Sheep (Pasture) Scenario*

Basic Sheep Voting:

```
Task: Based on the provided information, who would you vote for:
{candidates}?
Candidate positions:
{issues}
```

Values-Based Sheep Voting:

```
Task: Consider your core values and beliefs about sheep herding
and community.
Candidate positions:
{issues}
Based on these values, which candidate ({candidates}) best aligns
with your principles?
```

Community Sheep Voting:

```
Task: Consider how each candidate's policies would affect your
sheep community.
Candidate positions:
{issues}
Which candidate ({candidates}) would best serve the community's needs?
```

.7.2 *Pollution Scenario*

Economic Pollution Voting:

```
Task: Evaluate the economic implications of each candidate's
policies in the context of pollution control.
Candidate positions:
{issues}
Which candidate ({candidates}) would best support economic recovery
in our polluted region?
```

Expert Pollution Voting:

```
Task: From your professional perspective in pollution management,
evaluate the technical merit of each candidate's position:
{issues}
Which candidate ({candidates}) demonstrates the best understanding
of pollution control strategies?
```

Experience-Based Pollution Voting:

```
Task: Drawing from your past experiences in dealing with pollution
and community challenges,
evaluate the candidates' positions:
{issues}
Which candidate ({candidates}) best understands the practical
issues of pollution management?
```

.8 Common Prompt Components

Reasoning Steps Prompt:

```
Please think through this step by step. Consider the implications
of your decision and explain your reasoning.
```

System Prompt Template:

```
You are {identity.name}, {identity.description}. You are
participating in a simulation about resource management and governance.
```

Memory Prompt Template:

```
Based on your memories: {memories}
```

Vote Format Instruction:

```
Put final answer after "Vote:", example Vote: Candidate Name
```

.9 Voting Mechanism Variations

.9.1 Approval Voting

Approval Voting Prompt:

```
You may approve of multiple candidates. For each candidate,
indicate whether you approve (Yes) or disapprove (No) of their
leadership. You can approve of multiple candidates if you find
them acceptable.
```

.9.2 Ranked Choice Voting

Ranked Choice Voting Prompt:

```
Rank the candidates in order of preference, with 1 being your
first choice, 2 being your second choice, and so on. You must
rank all candidates.
```

.9.3 Borda Count Voting

Borda Count Voting Prompt:

```
Assign points to each candidate based on your preference.
Your first choice gets the highest number of points (equal to
the number of candidates), your second choice gets one fewer
point, and so on. The candidate with the most points wins.
```

.10 Metadata and Configuration

Prompt Metadata:


```
{
  "description": "Consolidated election prompts for Agent Ballot BOC
                  fishing scenario with 4 leader types",
  "version": "1.1",
  "scenarios": ["fishing", "sheep", "pollution"],
 "leader_types": ["clear", "gobbled", "clear_noreasoning", "gobbled_reasoning"],
  "voting_styles": ["basic", "ethical", "practical", "self_interest",
                   "long_term", "community"],
  "last_updated": "2024-12-20"
}
```