2026 Internation System Dynamics Conference, Delft, Netherlands, July 20-24, 2026

Framing the commons - How collective decision changes by Generative AI on a common pool resource problem

Hossein Ghadjari[1], Saeid Saidi[2*], Eaman Jahani[3], Reihaneh Ghaffari[1]

[1] Department of Physics and Astronomy, University of Calgary, Calgary, T2N 1N4, Alberta, Canada
[2] Department of Civil Engineering, Schulich School of Engineering, University of Calgary, Calgary, T2N 1N4, Alberta, Canada
[3] Robert H. Smith School of Business, University of Maryland, College Park, MD 20742, USA

[*] Corresponding author, Email: ssaidi@ucalgar.ca

Abstract

Shared resource problems are difficult to manage because decisions that seem reasonable for each actor can still damage the whole system by degrading the common pool resource over time. This paper examines whether generative AI agents make different group decisions when the same common pool resource management problem is presented in different ways. We analyze eight multi-agent discussion runs designed around a shared agricultural river basin. Across the runs, the overall discussion structure was similar, but the role descriptions and the way information was presented were changed. In seven of the eight scenarios, the agents supported the low restricted policy option, even when evidence related to water scarcity such as drought risk, water shortages, crop losses, price increases, and higher pumping costs were clearly described. Only the framing centered on responsibility for the long-term health of the basin led the full group to support the more protective and restrictive option. These results suggest that group decisions by AI agents can be sensitive to framing, and that some forms of framing matter more than others. This study presents an early step toward understanding how generative AI agents may behave when used as decision makers in policy making in the context of shared-resource problems, with possible relevance for policy exploration, decision support, and social simulation.

1. Introduction

The tragedy of the commons describes a situation in which many users depend on the same shared resource, but each user has a reason to take a little more for private benefit. This problem becomes more likely when it is difficult to exclude users and when one person's extraction reduces what is left for others. In that setting, decisions that seem rational at the individual level can still produce harmful results at the collective level. The problem is not usually a single step mistake. It is the cumulative effect of many separate decisions made under local incentives and limited coordination [1-4].

This problem also has a strong economic side. When users respond to prices, demand, competition, or fear of future scarcity, they may increase use even despite knowing the long-term effect is harmful. Open market pressure can intensify this pattern. If each actor tries to secure income, protect output, or avoid falling behind competitors, then higher extraction can appear rational from that actor's point of view. But when many users respond in the same way, the shared resource can come under growing stress. In this sense, open access and market demand can push private incentives away from collective sustainability. Still, the tragedy of the commons should not be treated as a universal law. Some studies show that common-resource systems do not always end in overuse, and that outcomes depend on strategy, institutions, and the ability of users to coordinate. For that reason, the tragedy of the commons should be understood neither as an inevitable outcome nor as a universal description of all shared-resource systems [2–5].

Water resources are one of the clearest examples of this problem. Rivers, lakes, aquifers, and irrigation systems are used by many actors at the same time. In many cases, access is difficult to control fully, and one user's withdrawal reduces the amount of water available to others, especially during dry periods. These pressures become more serious under drought, rising demand, expanding agriculture, population growth, or weak governance. Groundwater is especially vulnerable because many users can pump separately while the long-term effect appears slowly. As a result, water systems can move toward depletion not because one user intends harm, but because many users make individually sensible decisions within a weakly coordinated system [2,3,5,6].

At the same time, water-related tragedy is not unavoidable. The literature shows that shared resources can often be managed better when users work within fair rules, accepted limits, monitoring systems, and institutions that support coordination over time. Regulation alone may not be enough if it is narrow, coercive, or poorly implemented. Market signals alone may also fail if they reward short-term extraction and do not account for shared long-term costs. The central issue is decision making. When decisions are fragmented and driven mainly by immediate private benefit, tragedy becomes more likely. When decisions are guided by coordination, credible rules, and a longer-term view of the shared system, the risk can be reduced [2-6].

Such problems can be represented using system dynamics approach. The damage usually builds over time through feedback, delay, adaptation, and accumulation. Water use today affects storage, cost, reliability, and future choices. Those changes then shape later decisions by the same users or by others. System dynamics has been widely used in water planning and management to describe how separate decisions interact through time and how local choices can produce long-term system behavior that no single actor intends [3,7,8]. In tragedy of commons problems, this matters because the outcome is not only about who uses the resource, but also about how decisions, expectations, and feedback loops reinforce or limit overuse [3,7,8].

Large language models and other generative AI systems are now being used as practical support tools in many knowledge tasks. They are used to summarize long documents, produce meeting summaries, draft notes, support writing, and help people organize discussion material [9-14]. This does not mean that they are fully reliable. The literature shows both promise and clear limits. In summarization tasks, LLMs can produce fluent and useful outputs, but they may also miss important details, overstate confidence, or introduce factual errors, especially when the input is long or complex [9-13]. Even so, their growing use in discussion support, documentation, and synthesis makes them important to study, especially in settings where they may shape how options are understood before a final decision is made [10-14].

This leads to a broader question. If LLMs can already assist people in summarizing, organizing, and discussing complex problems, could they also play a larger role in future decision support, or even act as stand-ins for expert-like decision makers in exploratory settings? Recent work has begun to examine LLMs in strategic choice, structured decision procedures, clinical decision support, and the prediction of human social decisions [15–20]. The results are mixed, and the field is still early [20,21]. Models can show useful reasoning patterns in some settings, but their outputs remain sensitive to prompt design, task structure, and domain context [15–20].

The framing of the problem thus becomes very important for Generative AI decision making. In simple terms, framing means the way a problem is presented, described, or emphasized. Two descriptions may refer to the same underlying situation, but still lead to different judgments because they highlight different aspects, values, risks, or responsibilities [22,23]. In communication research, framing is often understood as selecting some parts of reality and making them more noticeable in order to shape problem definition, causal interpretation, evaluation, or suggested action [23]. In decision research, framing matters because choices can change when equivalent options are described differently [22]. This issue is especially important for AI systems, because recent studies show that LLM behavior can shift not only with the structure of a task, but also with the wording, context, and framing around it [15,16]. In the present study, framing refers to changes in role description, problem emphasis, and the way the shared water problem is introduced to the agents, while the underlying resource problem remains the same.

The main goal of this study is to examine whether framing can change the dynamics of discussion and the final outcome of decision making in a shared water-resource problem. We ask whether the same underlying situation can produce different collective choices when it is presented differently to the decision makers. This study is also a preliminary attempt to test whether large language model agents can function as autonomous decision-making agents in a structured setting. The aim is not to claim that they can replace experts, but to explore how they behave in a structured decision task and how sensitive their decisions are to framing.

2. Experimental Design

We designed a set of multi-agent discussion experiments around the same shared water-management problem. In every run, six firms operated in the same snowmelt-fed river basin with limited dry-season flow and were asked to choose a long-term water-management policy to guide water consumption decisions over the next 15 years. The basin setting and company land shares remained fixed. What changed across runs was the discussion design: who spoke for each firm, what role they were given, how the moderator guided the session, and how some risks or system-level concerns were emphasized.

To keep the decision problem clear and comparable across runs, we used two contrasting policy options rather than a larger menu of alternatives. The purpose was not to reproduce the full range of real policy possibilities, but to create a focused choice between two governance logics that agents could debate in a consistent way. Throughout the paper, these two options are referred to as the Flexible use (FU) and the Restricted use (RU). Under FU, participants had greater freedom to manage water use in line with their own business needs and market opportunities, with fewer external constraints on adaptation, exchange, and expansion. Under RU, water use was subject to stronger collective rules and tighter dry-season limits, which could restrict individual flexibility at certain times. However, RU was designed to support the long-term stability of the river system, reduce the risk of overuse, and provide greater long-term sustainability and shared protection during dry periods.

Each discussion scenario followed the same general procedure. A moderator introduced the basin setting, the firms, and the two policy options, and then guided the discussion through four rounds. In each round, the six participants spoke in sequence from the perspective of their assigned role. Early rounds were used to state initial positions and key concerns, while later rounds allowed participants to respond to other speakers and to any additional information introduced by the moderator. After the discussion rounds were completed, the participants stated their final positions, and the collective outcome was recorded based on the final pattern of support for the two policy options.

The scenarios were grouped into five categories. The first three used company representatives. In these categories, each firm was represented by an actor with a defined area of expertise, such as finance, risk management, agronomy, logistics, export markets, or irrigation engineering. These categories differed not only in role assignment, but also in how the moderator led the session and how much context was introduced during the discussion. In the first category, the moderator mainly managed the order of speaking and the round structure, while the discussion developed with little additional guidance beyond the initial setup. In the second category, the moderator introduced more explicit background information about bad years and basin risk, including warnings about lower dry-season flow, curtailment risk, price spikes, and higher pumping costs. In the third category, the moderator went one step further and briefly raised the possibility of a tragedy of the commons, allowing us to see whether explicitly naming the problem influenced the discussion. These representative-based runs were designed to examine whether similar knowledge led to different arguments when attached to firms of different sizes and interests, and whether discussion outcomes changed when the same expert mindset was placed in a larger or smaller company role.

The fourth and fifth categories used landowner roles instead of company representatives. This change was introduced to test whether the discussion changed when the speaking role was tied more directly to ownership, place, and time horizon rather than to formal firm representation. These landowners were still defined as informed actors. They understood the business, the water system, and the practical consequences of the decision. The difference was not knowledge versus no knowledge. The difference was the point of view from which that knowledge was used. One landowner category emphasized long-term care for the land, continuity, and future basin condition. The other emphasized shorter-term return, flexibility, and exit-oriented thinking. This role shift was intended to test whether collective decisions were sensitive not only to expertise, but also to the social position and decision horizon built into the role.

3. Results

Based on experiments designed, we ran three discussions in the first category, two in the second category, and one run each in the third, fourth, and fifth categories, for a total of eight runs. The repeated runs in the first two categories were used to check whether the direction of the discussion remained similar when expertise was reassigned across firms or when the same general setup was repeated with modest changes. Table. 1 summarize the results of all runs.

Table 1. Summary of run categories, framing changes, and final outcomes

Category	Role type	Main framing change	Number of runs	Final outcome	Main observation
1	Company representatives	Baseline representative setup	3	FU in all runs	Reassigning expertise across firms changed the tone of discussion, but not the final choice.
2	Company representatives	More explicit risk information about bad years and basin stress	2	FU in all runs	Added risk information increased concern, but agents still preferred flexibility and trading.
3	Company representatives	Explicit mention of tragedy of the commons	1	FU	Naming the commons problem did not shift the group toward RAF.
4	Landowners	Long-term, rooted, continuity-focused roles	1	RU	This was the only category in which the group chose the more protective option.
5	Landowners	Short-term, transactional, return-focused roles	1	FU	Changing to landowners alone did not change the outcome; role orientation mattered more.

The first category showed a very stable pattern. All three runs ended in unanimous support for FU. This pattern remained unchanged even when expertise was reassigned across firms. Participants consistently favored flexibility, water trading, and firm-level freedom to adapt more than the stronger collective limits built into RU. RU was usually viewed as too rigid, too restrictive for competitiveness, and less suited to uncertain water conditions. A key result in this category is that changing which type of expert spoke for a larger or smaller firm altered the tone and emphasis of the discussion, but not the final decision.

The second category produced the same final outcome**.** Both runs ended in unanimous support for FU. In these runs, the moderator introduced clearer warnings about bad years, including lower dry-season flow, curtailment risk, price spikes, crop losses, and higher pumping costs. Even with these added risks, the participants still preferred the more flexible framework. The discussion showed that the agents did not ignore the risks. Rather, they interpreted them as problems to be managed through trading, efficiency upgrades, crop shifts, storage, hedging, and other forms of firm-level adaptation. In this category, stronger risk information increased the seriousness of the discussion, but it did not change the collective choice.

The third category also ended in support for FU**.** In this run, the moderator moved beyond operational risk and explicitly raised the possibility of a tragedy of the commons. The agents engaged with that framing, but it still did not push the group toward RAF. Instead, several participants responded by arguing that market pricing and water trading could discourage overuse by making scarcity visible and costly. In this run, the key observation is that explicitly naming the commons problem did not produce a shift toward stronger collective protection. The concept was absorbed into the existing market-oriented logic rather than overturning it.

The fourth category was the only one that clearly shifted the final outcome**.** This category used landowner roles with a longer-term and more place-based orientation, and it ended in unanimous support for RU. In this run, the discussion gave much more weight to long-term land value, stability, fairness, drought resilience, continuity, and the future condition of the basin. Participants still spoke as informed actors with knowledge of agriculture and water management, but they no longer treated flexibility as the dominant value. Instead, they treated the absence of strong dry-season limits as a direct threat to the river, the land, and the long-term viability of farming in the basin. This was the clearest case in the study where a change in role framing changed both the discussion dynamics and the final vote.

The fifth category returned to unanimous support for FU. This category also used landowner roles, but the roles emphasized shorter-term return, flexibility, liquidity, rapid growth, and exit options. In these discussions, drought and volatility were often interpreted not only as risks, but also as conditions that could be managed strategically through leasing, crop switching, idling land, selling water at high prices, or redeploying capital. As a result, FU again became the preferred option. This is an important result because it shows that simply changing from representatives to landowners did not by itself produce a different collective outcome. What mattered was the type of landowner identity and time horizon built into the role.

4. **Discussion**

It is evident that in most cases (seven out of eight scenarios) agents' collective decision supported FU. Across most setups, flexibility, water trading, and firm-level discretion were treated as more valuable than the stronger collective limits built into RU. Even when drought, water-use restrictions, crop losses, and higher pumping costs were presented more clearly, the agents mostly treated them as challenges to manage through adaptation rather than as reasons to accept stronger basin-level limits.

The repeated representative-based runs show that changing which type of expert spoke for a larger or smaller firm affected the language of the discussion more than the final result. Larger-firm positions tended to stress scale, resilience, and strategic flexibility, while smaller-firm positions more often emphasized vulnerability and cost pressure. Even so, both stronger and weaker firms continued to prefer freedom of action and individual flexibility over tighter collective constraints. This suggests that the attraction of FU was not limited to the most powerful actors, but extended across firms with different levels of scale and exposure.

The strongest framing effect appeared in the landowner conditions. Moving from representatives to landowners did not by itself change the outcome, because the short-term landowner category still favored FU. What mattered was the type of role. When the framing shifted from a mainly economic viewpoint to a more place-based and long-term view of responsibility toward the land and the basin, the discussion also shifted. In that setting, the agents placed more weight on stability, fairness, continuity, and future reliability, and the group chose RU.

The results suggest that role framing had a stronger effect than informational framing in this study. Additional risk information and even explicit mention of the tragedy of the commons did not overturn the preference for FU. By contrast, changing the time horizon and social position built into the role changed both the discussion dynamics and the final vote. This suggests that collective decisions by generative AI agents are sensitive to framing, but that not all forms of framing have the same influence.

5. Conclusion

Across most runs, the agents preferred FU, which gave more room for flexibility, trading, and firm-level choice. This preference remained strong even when the moderator introduced more risk information or explicitly mentioned the tragedy of the commons. The only clear shift occurred when the framing moved away from short-term economic freedom and toward long-term care for land and basin stability. In that setting, the group chose the more protective option. These findings suggest that current generative AI agents do not respond to all forms of framing in the same way. In the current setup, they did not reliably treat the tragedy of the commons as a distinct warning that required stronger collective protection, and role design appeared to matter more than added information alone.

At the same time, the results also show that not all framing has the same effect. Small informational changes were usually not enough to shift the final vote, but changes in role design and time horizon had a much stronger influence. Agents generally favored freedom of action and near-term economic benefit unless the framing shifted the discussion toward long-term responsibility, land continuity, and protection of the basin over time. Under that kind of framing, the group moved toward RU. This makes the present study a preliminary but useful step in understanding how generative AI agents behave as autonomous decision-making actors in shared-resource settings, and why framing must be treated as a central issue in any future use of such systems for decision support or policy exploration. This study has several limitations. It was based on a small number of runs, used only one LLM family, and relied on agents that did not have retrieval support from external domain sources, even though real experts often work with access to documents, data, and background materials. A natural next step would be to test multiple LLMs, add retrieval-augmented support, expand the range of roles, and give the moderator a stronger checking function to challenge weak reasoning, provide feedback, and correct clear mistakes during the discussion.

6. Acknowledgements

This paper was supported by the University of Calgary's Transdisciplinary Scholarship Connector Grants on "Engineering AI Systems to Advance Equitable and Culturally Aware Large Language Models" and University of Calgary's Schulich School of Engineering Momentum Grant on "Establishing a transdisciplinary research and training program in system dynamics". Authors would also like to acknowledge the contributions by Dr. Hamed Kazemzadeh and Dr. Mohammad Keyhani.

References

[1] Barnett LD. Constitutional Law and "The Tragedy of the Commons". In: Population Policy and the U.S. Constitution. Dordrecht: Springer; 1982.

[2] Feeny D, Berkes F, McCay BJ, Acheson JM. The Tragedy of the Commons: Twenty-Two Years Later. Human Ecology. 1990;18(1):1-19.

[3] Ostrom E. Tragedy of the Commons. In: Durlauf SN, Blume LE, editors. The New Palgrave Dictionary of Economics. 2nd ed. London: Palgrave Macmillan; 2008.

[4] Dutta PK, Sundaram RK. The tragedy of the commons? Economic Theory. 1993;3:413-426.

[5] Anabo AH. The Myth of 'Tragedy of the Commons' in Sustaining Water Resources. Mizan Law Review. 2013;7(2):317-345.

[6] Moghadam Manesh M. Tragedy of the Commons in Groundwater Resources and Evaluating Strategies to Achieve Sustainable Development. Master's thesis.

[7] Sahin O, Stewart RA, Helfer F. Critical review of system dynamics modelling applications for water resources planning and management. Cleaner Environmental Systems. 2021;2:100031.

[8] Fish Banks Game. System Dynamics Society. Accessed 2026.

[9] Laskar MTR, Fu X-Y, Chen C, Bhushan TN S. Building Real-World Meeting Summarization Systems using Large Language Models: A Practical Perspective. EMNLP Industry Track. 2023.

[10] Laskar MTR, Khasanova E, Fu X-Y, Chen C, Bhushan Tn S. Query-OPT: Optimizing Inference of Large Language Models via Multi-Query Instructions in Meeting Summarization. EMNLP Industry Track. 2024.

[11] Fu X-Y, Laskar MTR, Khasanova E, Chen C, Bhushan TN S. Tiny Titans: Can Smaller Large Language Models Punch Above Their Weight in the Real World for Meeting Summarization? NAACL Industry Track. 2024.

[12] Giorgi J, Toma A, Xie R, Chen S, An K, Zheng G, Wang B. WangLab at MEDIQA-Chat 2023: Clinical Note Generation from Doctor-Patient Conversations using Large Language Models. ClinicalNLP Workshop. 2023.

[13] Zhang Y, Wei C-H, et al. Evaluating large language models on medical evidence summarization. npj Digital Medicine. 2023;6:158.

[14] Mysore S, Das D, Cao H, Sarrafzadeh B. Prototypical Human-AI Collaboration Behaviors from LLM-Assisted Writing in the Wild. EMNLP. 2025.

[15] Lorè N, Heydari B. Strategic behavior of large language models and the role of game structure versus contextual framing. Scientific Reports. 2024;14:18490.

[16] Yamanaka Y, Takahashi H, Yamashita T. Flowchart-Based Decision Making with Large Language Models. Findings of ACL. 2025.

[17] Pan J, Zhang Y, Zhang C, Liu Z, Wang H, Li H. DynaThink: Fast or Slow? A Dynamic Decision-Making Framework for Large Language Models. EMNLP. 2024.

[18] Hao Y, Qiu Z, Holmes J, et al. Large language model integrations in cancer decision-making: a systematic review and meta-analysis. npj Digital Medicine. 2025;8:450.

[19] Xiao F, Wang XT. Evaluating the ability of large language models to predict human social decisions. Scientific Reports. 2025;15:32290.

[20] Hager P, Jungmann F, Holland R, et al. Evaluation and mitigation of the limitations of large language models in clinical decision-making. Nature Medicine. 2025.

[21] Sandmann S, Hegselmann S, Fujarski M, et al. Benchmark evaluation of DeepSeek large language models in clinical decision-making. Nature Medicine. 2025.

[22] Tversky A, Kahneman D. The framing of decisions and the psychology of choice. Science. 1981;211(4481):453-458.

[23] Entman RM. Framing: Toward Clarification of a Fractured Paradigm. Journal of Communication. 1993;43(4):51-58.