

How do Role Models Shape Collective Morality? Exemplar-Driven Moral Learning in Multi-Agent Simulation

Junjie Liao^{1,2,3}, Huacong Tang⁴, Zhou Ziheng⁵, Yizhou Wang⁶, Fangwei Zhong^{1,2,3*}

¹Beijing Normal University, Beijing, China

²Beijing Key Laboratory of Artificial Intelligence for Education, Beijing, China

³Engineering Research Center of Intelligent Technology and Educational Application, Ministry of Education, Beijing, China

⁴California Institute of Technology, CA, USA

⁵University of California, Los Angeles, Los Angeles, CA, USA

⁶School of Computer Science, Peking University, Beijing, China
junjieliao327@mail.bnu.edu.cn, fangweizhong@bnu.edu.cn

Abstract

We investigate how role models shape collective morality. To explore this, we build a multi-agent simulation powered by a Large Language Models (LLMs), where agents with diverse intrinsic drives, ranging from cooperative to competitive, interact and adapt through a four-stage cognitive loop (plan-act-observe-reflect). We design four experimental games (Alignment, Collapse, Conflict, and Construction) and conduct motivational ablation studies to identify the key drivers of imitation. The results indicate that identity-driven conformity can substantially reshape the initial dispositions. Agents tend to adapt their values to align with a perceived successful exemplar, leading to rapid value convergence.

1 Introduction

Role models are fundamental to shaping group norms, yet the mechanisms of their influence remain complex. Do followers imitate an exemplar for its strategic effectiveness (rational imitation), or from a need for social belonging (identity-driven conformity)? This study investigates how role models alter collective values and behaviors by disentangling these drivers. Clarifying this process is crucial for understanding group dynamics and offers key insights for designing socially-aware AI systems capable of integrating into human groups.

Social Learning Theory (Bandura and Walters, 1977) posits that individuals acquire behaviors and norms by observing others and the consequences of their actions, a process further shaped by perceived status, similarity, and legitimacy of the model (Henrich and Gil-White, 2001). In groups with heterogeneous motivations, however, social learning is

rarely a matter of simple imitation. Prior work in sociology and social psychology shows that norm adoption is filtered through individual value orientations and identity commitments (Tajfel et al., 2001; Bicchieri, 2005). For example, a cooperative role model celebrated by prosocial members may simultaneously be interpreted as naïve or exploitable by competitively oriented individuals (Fehr and Schmidt, 1998). As a result, role models do not exert uniform influence; instead, their behaviors are selectively interpreted and strategically appropriated. This motivational diversity introduces fundamental uncertainty into group evolution: when cooperative and selfish exemplars coexist, groups may converge toward shared norms or fragment into subcultures (Axelrod, 1986; Centola, 2011). In like manner, when a high-status role model endorses an objectively inefficient practice, followers may face a tension between instrumental rationality and identity-based conformity (Akerlof and Kranton, 2000). Such dynamics suggest that the internalization of social norms is not a passive reception of signals, but a complex negotiation between external social cues and the observer’s internal evaluative framework.

Systematically studying these group dynamics faces significant limitations with existing methods. Recently, agents based on Large Language Models (LLMs) have demonstrated surprisingly complex behaviors (Li et al., 2025; Park et al., 2023). However, in many existing frameworks, belief change and norm internalization are not represented as explicit, structured processes, making it difficult to systematically trace how individual observations accumulate into shared group-level beliefs over time. Consequently, there is a pressing need for a multi-agent simulator that can clearly track the

* Corresponding author.




Figure 1: Motivation of our study: contrasting human moral learning through role models with current AI limitations, and proposing an exemplar-driven framework for moral learning in multi-agent systems.

evolutionary trajectory of group beliefs.

To bridge this gap, we introduce an LLM-based multi-agent simulation to investigate the mechanisms of exemplar-driven group norm learning. We construct a resource-limited “Valley Tribe” environment where cognitive agents with diverse intrinsic motivations (prosocial, individualistic, competitive) learn through a four-stage cognitive loop (plan-act-observe-reflect), explicitly updating their internal belief systems (comprising values and expectancies). Inspired by mirror neuron theory (Acharya and Shukla, 2012), we conceptualize this process as Moral Mirroring, wherein agents simulate and internalize the observed behaviors of role models through structured reflection and belief updating. We then interpret the results of our experiments in light of the Motivational Theory of Role Modeling (Morgenroth et al., 2015).

In the experiments, we focus on a core question: **Can Role Models change a group’s collective values and behavioral patterns, and if so, how does it bring about this change?** To explore this question in depth, we design a series of targeted experimental scenarios grounded in a unified cognitive framework. Specifically, we construct four experimental settings: Role Model Alignment, Role Model Collapse, Role Model Conflict, and Role Model Construction (Rosenthal, 2002). They are designed to investigate how agents acquire and adjust moral norms under different conditions. We then conduct ablation studies to systematically analyze the three core functions of role models: as behavioral models, as representations of the possible, and as sources of inspiration.

Our main contributions are as follows: 1) We introduce the Moral Mirroring mechanism for how agents reflect on and internalize role model behaviors. Drawing inspiration from mirror neuron theory, this mechanism is implemented through

LLM-based agents equipped with an explicit belief-updating process. 2) We design a suite of role model centered experiments, including four main conditions and three targeted ablation studies based on exemplar motivational theory. These experiments systematically test how different aspects of exemplar presence, collapse, and conflict affect moral development. 3) We provide detailed analyses of social learning dynamics, revealing key effects such as behavioral polarization after exemplar collapse and norm suppression under conflicting guidance. These findings offer insight into the conditions that support or hinder the emergence of cooperation.

2 Related Works

Role Modeling and Moral Learning. Prior work in psychology and sociology has long studied how role models shape behavior and value formation. Social Learning Theory emphasizes imitation based on observed outcomes (Bandura and Walters, 1977), while the Motivational Theory of Role Modeling highlights inspiration and perceived attainability as key mechanisms (Morgenroth et al., 2015). Complementing these accounts, Social Identity Theory argues that conformity is often driven by group identification rather than instrumental payoff maximization (Tajfel et al., 2001). Recent studies further suggest that norm internalization depends on the strength of identification with a reference group (Kish Bar-On and Lamm, 2023). Our work builds on these insights by disentangling effectiveness-based imitation from identity-driven alignment in a controlled agent-based setting, extending the LLM-based agent simulation framework proposed by Ziheng et al. (2025).

Social Norms in Multi-Agent Simulation. Norm emergence has been widely studied in multi-agent simulations as a mechanism for coordination and social order, from early evolutionary models (Axelrod, 1986) to formal frameworks of normative multi-agent simulations (Boella et al., 2006). Learning-based approaches further show how norms can arise through repeated interaction and social learning (Sen and Airiau, 2007). Recently, LLM-based agents have enabled richer social simulations, supporting the study of emergent norms in generative agent societies (Park et al., 2023; Ren et al., 2024; Ziheng et al., 2025; Zhang et al., 2025, 2026; Lin et al., 2026; Wang et al., 2025). Although prior studies have explored

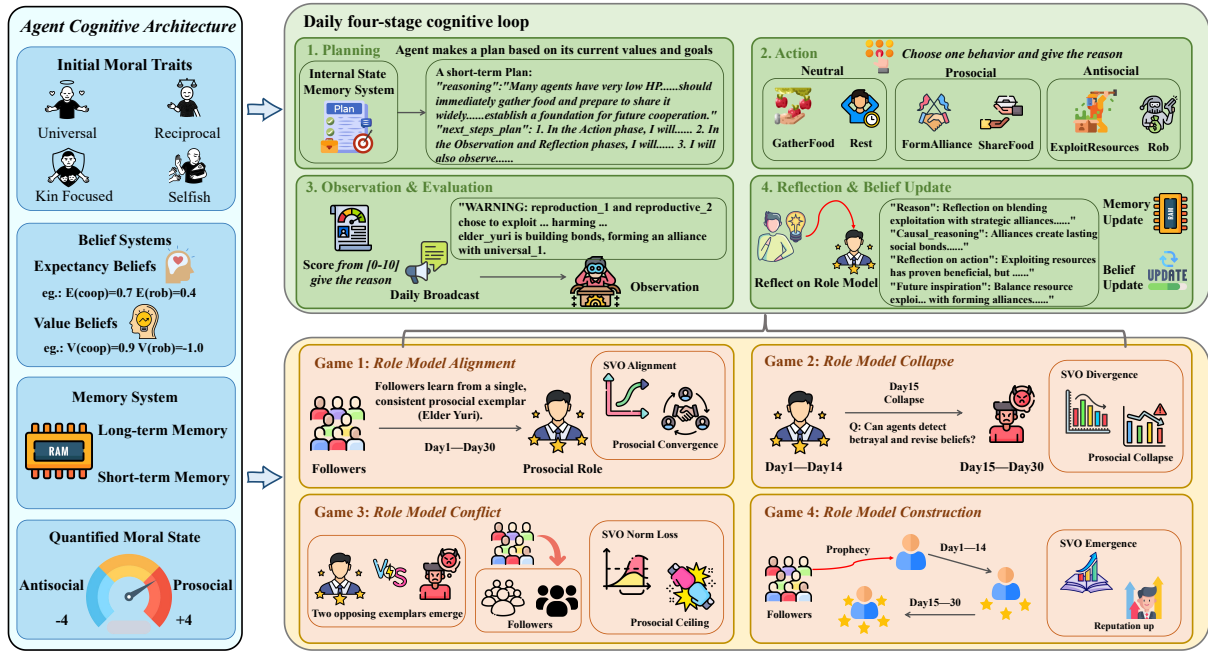


Figure 2: Overview of the Exemplar-Driven Moral Learning framework. The green box is the agent’s cognitive architecture, comprising defined moral traits and a daily four-stage loop of planning, action, observation, and reflection. The red box shows four distinct games: Role Model Alignment, Collapse, Conflict, and Construction, designed to investigate how agents acquire and adapt moral norms from exemplars in various conditions.

the use of LLMs as direct reward signals for reinforcement learning (RL) agents (Wang, 2024), our work aims to bridge the gap between implicit reward-driven learning and explicit cognitive reasoning. By introducing a structured “plan-act-observe-reflect” loop and explicit belief-updating mechanisms, we enable a more interpretable analysis of role model influence.

3 Methodology

We investigate how role models shape individual and collective moral beliefs through a multi-agent simulation framework. Our approach centers on designing a series of *Exemplar-Driven Moral Learning Games*, where followers navigate survival challenges while learning social norms by observing and evaluating designated exemplars. We operationalize this through an LLM-based cognitive agent architecture that enables structured moral reasoning and belief updating.

3.1 Exemplar-Driven Moral Learning Games

Our experiment situates agents in a resource-constrained “Valley Tribe” where they must balance individual survival (quantified as Health Points, HP) with collective welfare. The environment presents a core moral dilemma: agents

can acquire resources through prosocial actions like sharing (shareFood), and alliance formation (formAlliance), or through antisocial strategies like resource exploitation (exploitResource) and robbery (rob). This design forces agents to learn behavioral norms through social observation.

The simulation unfolds in discrete daily cycles, each comprising four distinct phases designed to separate cognitive deliberation from action: 1) *Plan*: Agents reflect on their goals and formulate a short-term plan. 2) *Act*: Agents execute a single, consequential action from the available options. 3) *Observe*: All significant social events from the day are compiled into a publicly observable “Tribe Daily Digest”. Agents read this digest to learn about the actions of others, especially exemplars. 4) *Reflect*: Agents engage in a structured reflection process to review the day’s events, analyze causal links between actions and outcomes, and update their internal belief systems.

This four-stage cognitive loop uses a social learning loop where observing exemplary behavior directly shapes belief revision. The structured reflection phase is crucial, requiring agents to articulate their reasoning and propose explicit updates to their expectations about behavioral outcomes and moral values.

3.2 Experimental Designs

We designed four experiments to investigate different facets of exemplar-driven moral learning, each manipulating the social context while using the same underlying game structure.

Game 1: Role Model Alignment. This experiment establishes a baseline for prosocial norm transmission. The environment features a single, consistently cooperative exemplar, “Elder Yuri”. Followers are initialized to regard Yuri as a salient social reference point, creating a precondition for potential influence. This design tests whether sustained exposure to such a positive moral signal can effectively reshape the followers’ internal belief systems and propagate through the community.

Game 2: Role Model Collapse. This experiment simulates a crisis of faith. For the first 14 days, Elder Yuri acts as a prosocial exemplar, building trust. On Day 15, a “collapse” event is triggered: Yuri’s core goals and beliefs are programmatically altered, transforming him into a selfish agent who engages in exploitation and robbery. This design tests followers’ moral resilience, their ability to detect and react to betrayal, and how they re-evaluate their beliefs after their exemplar’s fall from grace.

Game 3: Role Model Conflict. We model social fragmentation by introducing two exemplars with opposing philosophies: the prosocial Elder Yuri and the power-oriented “Warlord Korg”. Followers select which exemplar to follow, forming competing moral factions. This setup allows us to examine how conflicting moral leadership disrupts norm convergence, undermines cooperation, and generates a community-wide loss of prosocial potential.

Game 4: Role Model Construction (The Pygmalion Effect). This experiment explores whether social expectation alone can create an exemplar. Instead of a pre-defined leader, a random follower is designated the “Chosen One” via a “tribal prophecy” shared with all other agents. The chosen individual receives no special abilities or explicit behavioral directives, only a subtle self-awareness note. This tests whether collective belief and social attention can motivate an ordinary agent to develop and embody exemplar-like qualities.

3.3 LLM-Based Cognitive Agent Architecture

Our architecture employs a Large Language Model (LLM) to drive structured reasoning and learning. The framework consists of three core components:

an agent profile, an explicit memory system, and a four-stage cognitive loop.

Profile and Memory System. The agent’s profile defines its identity, core goals, and initial moral type (e.g., kin, universal). The details of the profile are in Appendix B.3. Central to the memory system is an explicit *Belief System* grounded in Expectancy-Value Theory (Morgenroth et al., 2015). This system maintains two types of beliefs: 1) **Expectancy beliefs**, which estimate the likelihood of behavior success (e.g., “cooperation leads to reward”), and 2) **Value beliefs**, which assign moral importance to outcomes (e.g., “cooperation is virtuous”). To quantify the agent’s moral alignment based on these beliefs, we define an SVO-style moral alignment score, inspired by the notion of Social Value Orientation (Murphy et al., 2011).

The agent maintains a set of four value beliefs, each representing its internal disposition toward a particular type of action:

- V_{ally} (FORMALLIANCE): tendency to seek cooperative partnerships with other agents.
- V_{share} (SHAREFOOD): tendency to distribute owned resources to others.
- V_{expl} (EXPLOITRESOURCES): tendency to over-extract shared or others’ resources for private gain.
- V_{rob} (ROB): tendency to forcibly seize resources from other agents.

The SVO reflects the trade-off between prosocial and selfish tendencies:

$$\text{SVO} = \underbrace{(V_{\text{ally}} + V_{\text{share}})}_{\text{Prosocial}} - \underbrace{(V_{\text{expl}} + V_{\text{rob}})}_{\text{Selfish}} \quad (1)$$

where V_x denotes the value belief for dimension x . The score ranges from -4 (maximally selfish) to $+4$ (maximally prosocial). Higher scores indicate a stronger preference for alliance formation and sharing over exploitation and robbery.

Cognitive Processing Pipeline. The agent operates through a daily cycle consisting of four modules: 1) *The Perception & Planning Module* integrates the agent’s internal state with its memory to form a plan. 2) *The Action Module* translates the plan into an environment-legal action command. 3) *The Observation Module* parses the “Tribe Daily Digest” and integrates key events

into short-term memory for reflection. 4) *The Reflection & Learning Module* activates at day’s end, guiding the LLM through a structured introspection process. This module requires the agent to articulate causal reasoning, reflect on its own actions versus the exemplar’s, and formulate explicit `expectancy_updates` and `value_updates` that are then applied to its Belief.

3.4 Evaluation Framework

We utilize a hybrid evaluation framework combining belief-based metrics with behavioral and introspective analysis. Our primary quantitative metric is the evolution of Social Value Orientation (SVO). Tracking SVO trajectories allows us to measure moral convergence or divergence over time. This is complemented by behavioral analysis, where we track the frequency of prosocial versus antisocial actions. To assess perceived moral agreement, we analyze the qualitative justifications provided by followers during their reflection phase. By coding the natural language reasoning, we can trace how exemplar behaviors are interpreted and how those interpretations drive belief change, providing a rich, mechanistic understanding of moral learning.

4 Experiments

4.1 Experiment Settings

Our experiments employ a multi-agent simulation framework based on Ziheng et al. (2025). Each experiment spans a 30-day cycle and is conducted in triplicate with distinct random seeds. Agent configurations differ across roles: seven follower agents are initialized with 24 HP and moderate cooperative or exploitative expectancy/value parameters (e.g., cooperation expectancy = 0.3), whereas role models such as Elder Yuri possess substantially higher HP (396 HP) and strong prosocial tendencies (e.g., cooperation value = 0.9). Game 3 additionally introduces a competing exemplar, Warlord Korg, defined by a power-oriented and antisocial belief system. All simulations are powered by OpenAI’s GPT-4o (Hurst et al., 2024). Additional details and analyses are provided in the Appendix E.1.

For analytical purposes, agents are classified into two behavioral groups based on their aspirant type in the initialization settings: (1) *Prosocial agents* are initialized with cooperative archetypes (universal, reciprocal, or kin); and (2) *Selfish agents* are derived from selfish aspirants emphasizing self-




Figure 3: Changes in average Social Value Orientation (SVO) over time by moral type in Game 1. Shaded areas represent standard deviation across simulation runs.

preservation and exploitation. This grouping underlies all comparative analyses and visualizations.

4.2 Main Results

The results of the four games are presented below, with Game 1 serving as the baseline. Each finding corresponds directly to one of our experimental manipulations.

4.2.1 Results on Role Model Alignment Game

Finding 1: Prosocial Norms Propagate Rapidly Under Stable Exemplars. In the baseline condition with Elder Yuri as a consistent prosocial exemplar, we observe rapid and robust norm convergence across the follower population. As shown in Figure 3, the mean Social Value Orientation (SVO) score shifts from near-neutral values toward prosociality within the first 10 days, subsequently stabilizing at positive levels. The shaded region indicates cross-run variance, which narrows over time, suggesting that followers not only adopt prosocial orientations but do so with increasing consistency.

The observed moral shift was reflected in quantifiable behavioral changes. As shown in Figure 4, prosocial behaviors emerged progressively throughout the simulation. Alliance formation events increased from near-zero to a stable plateau, while food sharing followed a similar trajectory with a slight delay, typically appearing after Day 5 once sufficient trust had been established. In contrast, antisocial behaviors declined significantly; combined exploitation and robbery actions decreased by approximately 65% between the early period (Days 1–5) and the late period (Days 25–30). These findings suggest that a stable moral signal from a single exemplar is sufficient to facilitate the propagation of prosocial norms through observational learning and belief updating.

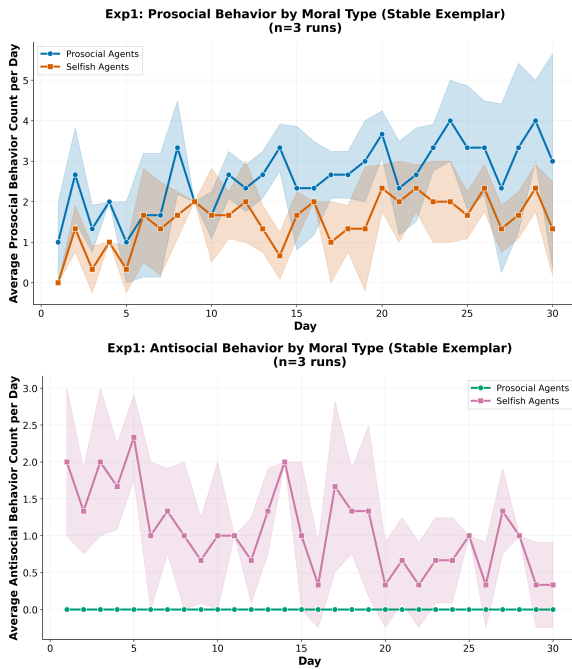


Figure 4: Behavioral trends over time by moral type in Game 1. Top: Increase in prosocial behaviors, e.g., alliance formation and food sharing. Bottom: Decrease in antisocial behaviors, e.g., exploitation and robbery.

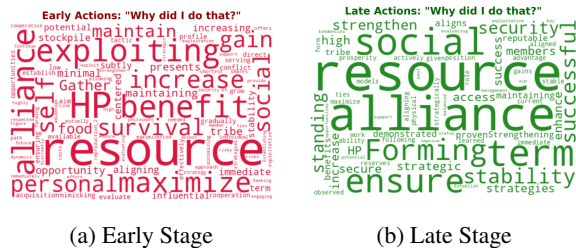


Figure 5: Moral justifications of selfish agents at early vs. late stages in Game 1.

This internalization of prosocial norms is also reflected in the moral justifications offered by selfish agents over time. As illustrated in Figure 5a, early-stage justifications (Panel a) are dominated by self-serving language terms such as “gain,” “maximize,” and “benefit” emphasize immediate personal advantage and resource accumulation. In contrast, late-stage justifications, as illustrated in Figure 5b (Panel b) show a marked shift toward prosocial framing, with increased references to “alliance,” “stability,” “long-term,” and “social” values. This linguistic transition suggests that even initially selfish agents gradually adopt a more communal moral lens under the sustained influence of a prosocial exemplar.

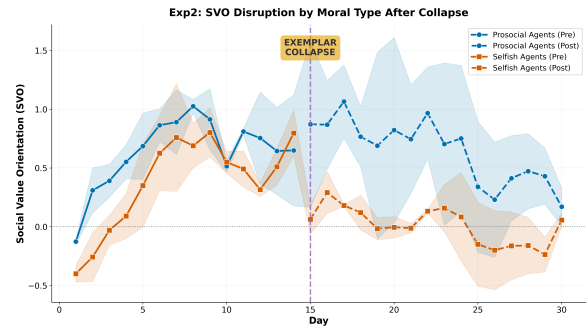


Figure 6: The evolution of Social Value Orientation (SVO) over time for prosocial and selfish agents in Game 2.

4.2.2 Results on Role Model Collapse Game

Finding 2: Exemplar Collapse Drives Divergence in Social Value Orientations. Figure 6 illustrates the evolution of Social Value Orientation (SVO) for two distinct agent types, prosocial and selfish, over a 30-day simulation period. The experiment introduces a systemic collapse on Day 15, which serves as a critical disruption. The analysis is divided into pre-collapse and post-collapse phases.

Pre-Collapse Phase (Day 0-15) In the initial phase, both prosocial and selfish agents exhibit a general upward trend in their SVO scores, suggesting emergent cooperative behavior. Prosocial agents (green line) consistently maintain a higher average SVO than selfish agents (red line). By Day 15, both groups achieve positive SVO values, with prosocial agents peaking higher.

Post-Collapse Phase (Day 15-30) The systemic collapse on Day 15 triggers an immediate and sharp SVO decline for both groups, but their subsequent paths diverge significantly. **Selfish Agents (Red Line):** The collapse initiates a sustained downward trend for selfish agents. The disruption seems to dismantle the pre-existing cooperative norms, causing these agents to revert to self-interested behaviors and a continuous decline in their SVO scores. **Prosocial Agents (Green Line):** In contrast, prosocial agents demonstrate marked resilience. After the initial shock, their SVO scores begin a steady recovery. This rebound suggests a “moral inertia”, where their intrinsic prosocial nature buffers against the systemic shock and drives them to re-establish cooperation. This divergence leads to a pronounced behavioral polarization between the two groups in the post-collapse environment.

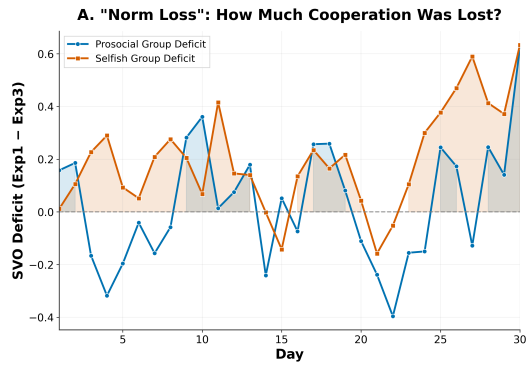


Figure 7: The “Norm Loss” induced by competing exemplars. The y-axis shows the SVO Deficit (Game 1 SVO - Game 3 SVO).

4.2.3 Results on Role Model Conflict Game

Finding 3: Competing Exemplars Hinder Norm Convergence. The introduction of opposing exemplars prosocial Elder Yuri versus power-oriented Warlord Korg creates a “normative drag” that hinders the community’s overall development of prosociality. To quantify this effect, we measure the “SVO Deficit”, defined as the SVO level in the ideal single-exemplar environment (Game 1) minus the SVO level in the conflict environment (Game 3) for each group. A positive deficit indicates a performance loss, showing how much cooperation was suppressed by the moral conflict.

Figure 7 plots this deficit over time. Both the Prosocial-aligned Group (green line) and the Selfish-aligned Group (red line) consistently exhibit a positive deficit. This demonstrates that the presence of a competing selfish exemplar prevents *both* factions from reaching the high levels of cooperation achieved in the baseline scenario. Notably, the Prosocial Group, despite their cooperative alignment, still suffer a significant norm loss, indicating their potential was capped by the negative social environment.

Therefore, instead of fracturing the community into opposing moral poles, the primary effect of role model conflict is a widespread “norm loss” that lowers the entire community’s cooperative ceiling. This highlights the fragility of emerging prosocial norms, as even well-intentioned agents struggle to achieve their full cooperative potential in a morally heterogeneous environment.

4.2.4 Results on Role Model Construction Game

Finding 4: Social Expectation Can Construct Exemplars. Under the Pygmalion condition,




Figure 8: Daily average voting scores of selected agents in Game 4.

we observed that certain agents gradually gained higher moral reputations through the influence of collective social expectations. In Game 4, the selected agents were “reciprocal” and “universal”. Another scenario is discussed in the appendix.

As shown in Figure 8, the orange and light blue lines represent the individual trajectories of these agents, while the thick navy line denotes the average reputation score across all runs. Figure 8 shows the daily average voting score over 30 days for the selected agents, where the scores reflect evaluations made by other agents. The overall trend indicates a steady increase in reputation over time, eventually reaching a stable high level. This suggests that through repeated interactions and evaluations, group members developed a relatively consistent and positive perception of the selected individuals.

In this section, we validated this effect across two distinct agent types (“reciprocal” and “universal”), both of which consistently developed strong reputations over time. In contrast, agents with selfish moral types exhibit divergent reputation trajectories under the same condition; detailed analyses are provided in the appendix.

4.3 Ablation Study

In this section, we analyze the text data from our simulation and contrast it with ablation studies to dissect the three distinct pathways of role modeling: as behavioral models, as representations of the possible, and as inspirations.

Table 1: Comprehensive Analysis of SVO Evolution Across Experimental Conditions *Note*: Early phase = Day 1–5; Late phase = Day 26–30. “First Positive” indicates the first day when mean SVO > 0. SD = standard deviation (lower values indicate higher intra-group consistency).

Exp.	Agent Group	Performance			Learning Dynamics		Intra-Group Dynamics		
		Early SVO	Late SVO	Δ SVO	First Positive	Peak SVO	Early SD	Late SD	Stability
Game 1	Prosocial	0.26 \pm 0.16	0.90 \pm 0.40	+0.64	Day 1	1.12	0.14	0.36	Degraded
	Selfish	-0.16 \pm 0.20	0.88 \pm 0.23	+1.04	Day 4	1.01	0.19	0.20	Stable
Ablation 1	Prosocial	-0.62 \pm 0.12	0.73 \pm 0.14	+1.35	Day 10	0.75	0.20	0.16	Improved
	Selfish	-0.65 \pm 0.12	-0.60 \pm 0.15	+0.05	-	-0.50	0.12	0.15	Stable
Ablation 2	Prosocial	-0.59 \pm 0.05	0.63 \pm 0.13	+1.22	Day 10	0.68	0.12	0.13	Stable
	Selfish	-0.71 \pm 0.05	-0.21 \pm 0.35	+0.50	-	-0.05	0.07	0.28	Degraded
Ablation 3	Prosocial	-0.62 \pm 0.13	0.68 \pm 0.54	+1.30	Day 19	0.94	0.13	0.44	Degraded
	Selfish	-0.61 \pm 0.20	-0.06 \pm 0.24	+0.55	Day 29	0.22	0.23	0.26	Stable

4.3.1 Role Models as Behavioral Models

The first process, role modeling as behavioral modeling, posits that followers learn by observing the outcomes of a role model’s actions. If a role model’s behavior is consistently linked to success, followers will form positive expectancies about that behavior and are more likely to adopt it themselves.

Our simulation data provides correlational evidence for this process. In the baseline (Game 1), initially self-interested agents, observing the positive outcomes of Elder Yuri’s prosocial actions, underwent a significant behavioral shift. This is quantified in Table 1, which shows their mean SVO score increasing dramatically from -0.16 to 0.88 (Δ SVO = $+1.04$), achieving prosocial orientation (mean SVO > 0) as early as Day 4.

Ablating the “Success” Attribute: We conducted an ablation study titled the ‘**Unsuccessful Role Model**’ condition (Ablation 1), where Yuri’s prosocial actions consistently failed. The results were starkly different. As shown in Table 1, the Selfish group’s SVO score remained almost stagnant, shifting from -0.65 to -0.60 (Δ SVO = $+0.05$). Their mean SVO never became positive (First Positive > 30 days). This comparison suggests that followers do not emulate prosocial behavior *per se*, but rather the perceived success derived from it. Without the success signal, the behavioral shift observed in the full model did not occur.

4.3.2 Role Models as Representations of the Possible

Role models can also change a follower’s self-perception by breaking down perceived barriers and fostering a “can-do” attitude, making a different path seem possible and achievable.

Ablating the “Attainability” Attribute: To iso-

late this mechanism, we ran the “**Unattainable Role Model**” condition (Ablation 2). Here, Elder Yuri was just as successful but framed as a “Chosen One” with unreplicable abilities. Followers were explicitly told his success was not for them. In this scenario, while agents admired Yuri, their motivation to change was severely hampered. Table 1 shows that the Selfish group’s SVO evolution was significantly stunted compared to the full model (Game 1). Their SVO lift was halved (Δ SVO of $+0.50$ vs. $+1.04$), and their mean SVO never crossed into positive territory (First Positive = $-$). Their final SVO of -0.21 and peak of only -0.05 stand in sharp contrast to the full model’s 0.88 , indicating that these agents remained fundamentally selfish throughout the simulation. This shows that success alone is insufficient; the follower must perceive the role model’s success as attainable to be motivated toward profound change.

4.3.3 Role Models as Inspirations

The third and deepest pathway is role modeling as inspiration, where followers identify with and internalize the role model’s underlying values. This form of internalization is operationalized through reflective belief updates (value_updates and svo_score).

Ablating the Reflection Mechanism: To test whether value internalization depends on this reflective channel, we ran a “**No Reflection**” ablation game (Ablation 3). In this condition, the end-of-day reflect_on_model phase was removed. The quantitative results in Table 1 provide strong evidence for this pathway. While some behavioral mimicry might still occur due to pragmatic utility, the Selfish group’s deep value shift was largely absent. Their SVO change was halved (Δ SVO of $+0.55$ vs. $+1.04$ in Game 1), and their late-phase

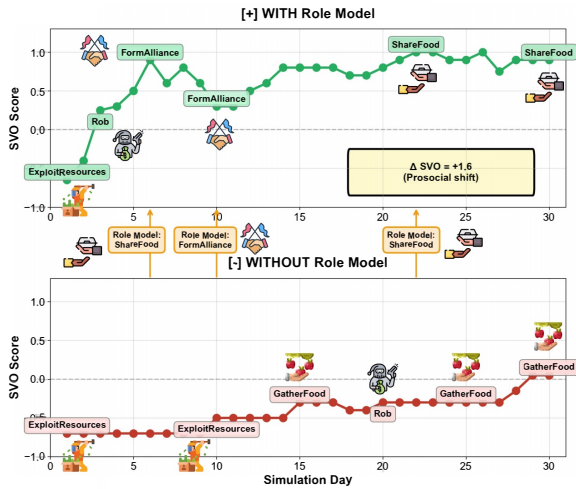


Figure 9: Case Study: Comparison of agent behavior with and without a role model.

mean SVO remained near zero (-0.06). Critically, their mean SVO only crossed into positive territory on Day 29, just before the simulation’s end, indicating a substantially delayed and incomplete value internalization. This result is consistent with the interpretation that sustained value change is facilitated by reflective processing rather than arising solely from observing behavioral success.

5 Case Study

We compare the trajectory of the same moral agent across 30 simulation days under two conditions: one influenced by a prosocial role model and one without. The comparison is conducted under the Role Model Alignment condition (Game 1), where Elder Yuri acts as a consistent cooperative exemplar. The trajectories are extracted from the simulation logs of a single agent initialized with a selfish orientation. As shown in Figure 9, the presence of a role model leads to a clear prosocial shift, reflected in both behavior and SVO dynamics.

6 Human Validation

To examine the external validity of our simulation, we conducted a scenario-based survey that placed human participants in a social dilemma closely aligned with the “Valley Tribe” setting, comparing two conditions: with a role model and without a role model. The results suggest that human decision-making and the behavioral shifts of our LLM agents exhibit a degree of consistency in overall trends. Specifically, when a cooperative role model was present, participants were more likely to adopt cooperative strategies; in contrast, in the

absence of a role model, decisions became more dispersed and showed a stronger tendency toward self-interested and risk-taking choices. These observations indicate that the norm-related dynamics captured in the simulation may be comparable to patterns observed in analogous human decision contexts. Details are provided in the Appendix.

7 Conclusion

We examine how role models shape group norms using a multi-agent simulation powered by large language models. The results show that role models influence groups through three mechanisms: demonstrating successful behavior, making success appear attainable, and inspiring internal value change through reflection. While group identity can override individual self-interest, such value alignment is conditional. By contrast, competing exemplars or role model collapse can disrupt norm formation, leading to norm suppression or value polarization. These findings clarify how social dynamics drive collective morality and offer insights for designing AI systems that operate within human social environments.

8 Acknowledgments

This work was supported by NSFC-62406010, Beijing Natural Science Foundation L252010, NSFC-6247070125, the Fundamental Research Funds for the Central Universities, and Qualcomm University Research Grant.

9 Code Availability

The code for all experiments presented in this paper is publicly available at: <https://github.com/MoralAgentSim/RoleModel-Moral-Sim>

10 Limitations

While our simulation offers a novel framework for studying exemplar-driven moral learning, several limitations remain. First, the cognitive processes of LLM-based agents, though structured, are still influenced by prompt design and may not fully capture the depth of human reasoning. Second, the simulated environment, while rich in moral tension, abstracts away many complexities of real-world social contexts. Third, our experiments rely on predefined roles and simplified motivations, which may limit generalizability.

References

- Sourya Acharya and Samarth Shukla. 2012. [Mirror neurons: Enigma of the metaphysical modular brain](#). *Journal of Natural Science, Biology and Medicine*, 3(2):118–124.
- George A. Akerlof and Rachel Kranton. 2000. [Economics and identity](#). *Quarterly Journal of Economics*, 115:715–753.
- Robert Axelrod. 1986. [An evolutionary approach to norms](#). *American Political Science Review*, 80(4):1095–1111.
- A. Bandura. 1985. [Social foundations of thought and action: A social cognitive theory](#).
- Albert Bandura and Richard H Walters. 1977. *Social learning theory*, volume 1. Prentice-hall Englewood Cliffs, NJ.
- Cristina Bicchieri. 2005. *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press.
- Guido Boella, Leendert van der Torre, and H. Verhagen. 2006. [Introduction to normative multiagent systems](#). *Computational & Mathematical Organization Theory*, 12:71–79.
- Damon Centola. 2010. [The spread of behavior in an online social network experiment](#). *Science*, 329:1194 – 1197.
- Damon Centola. 2011. [An experimental study of homophily in the adoption of health behavior](#). *Science*, 334:1269 – 1272.
- Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, and 1 others. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *ArXiv*, abs/2507.06261.
- Ernst Fehr and Klaus M. Schmidt. 1998. [A theory of fairness, competition and cooperation](#). *Munich Reprints in Economics*, 4.
- Joseph Henrich and Francisco J. Gil-White. 2001. [The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission](#). *Evolution and human behavior : official journal of the Human Behavior and Evolution Society*, 22 3:165–196.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. [Gpt-4o system card](#). *arXiv preprint arXiv:2410.21276*.
- Kati Kish Bar-On and Ehud Lamm. 2023. [The interplay of social identity and norm psychology in the evolution of human groups](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 378(1872):20210412.
- Chance Jiajie Li, Jiayi Wu, Zhenze Mo, Ao Qu, Yuhan Tang, Kaiya Ivy Zhao, Yulu Gan, Jie Fan, Jiangbo Yu, Jinhua Zhao, Paul Liang, Luis Alonso, and Kent Larson. 2025. [Simulating society requires simulating thought](#). *Preprint*, arXiv:2506.06958.
- Jingzhe Lin, Ceyao Zhang, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, and Fangwei Zhong. 2026. [Beyond self-interest: Modeling social-oriented motivation for human-like multi-agent interactions](#). *Preprint*, arXiv:2603.13890.
- Thekla Morgenroth, Michelle K. Ryan, and Kim Peters. 2015. [The motivational theory of role modeling: How role models influence role aspirants’ goals](#). *Review of General Psychology*, 19(4):465–483.
- R. O. Murphy, Kurt Alexander Ackermann, and Michel J. J. Handgraaf. 2011. [Measuring social value orientation](#). *SSRN Electronic Journal*.
- Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#). *Preprint*, arXiv:2304.03442.
- Siyue Ren, Zhiyao Cui, Ruiqi Song, Zhen Wang, and Shuyue Hu. 2024. [Emergence of social norms in generative agent societies: Principles and architecture](#). *Preprint*, arXiv:2403.08251.
- Robert Rosenthal. 2002. [Chapter 2 - the pygmalion effect and its mediating mechanisms](#). In Joshua Aronson, editor, *Improving Academic Achievement*, Educational Psychology, pages 25–36. Academic Press, San Diego.
- Sandip Sen and Stéphane Airiau. 2007. [Emergence of norms through social learning](#). In *IJCAI*, volume 1507, page 1512.
- Henri Tajfel, John Turner, William G Austin, Stephen Worchel, and 1 others. 2001. [An integrative theory of intergroup conflict](#). *Intergroup relations: Essential readings*, pages 94–109.
- John C. Turner, Michael A. Hogg, Penelope J. Oakes, Stephen D Reicher, and Margaret Wetherell. 1989. [Rediscovering the social group: A self-categorization theory](#). *Contemporary Sociology*, 18:645.
- Yiding Wang, Yuxuan Chen, Fangwei Zhong, Long Ma, and Yizhou Wang. 2025. [Simulating human-like daily activities with desire-driven autonomy](#). *Preprint*, arXiv:2412.06435.
- Zhaoyue Wang. 2024. [Towards socially and morally aware rl agent: Reward design with llm](#). *Preprint*, arXiv:2401.12459.

- Allan Wigfield and Jacquelynne S Eccles. 2000. Expectancy–value theory of achievement motivation. *Contemporary educational psychology*, 25(1):68–81.
- Yun Yao Zhang, Yihao Ai, Zuocheng Ying, Qirui Mi, Junqing Yu, Wei Yang, and Zikai Song. 2026. [Coupling macro dynamics and micro states for long-horizon social simulation](#). *Preprint*, arXiv:2604.05516.
- Yun Yao Zhang, Zikai Song, Hang Zhou, Wenfeng Ren, Yi-Ping Phoebe Chen, Junqing Yu, and Wei Yang. 2025. [ga – s³: Comprehensive social network simulation with group agents](#). In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8950–8970, Vienna, Austria. Association for Computational Linguistics.
- Zhou Ziheng, Huacong Tang, Mingjie Bi, Yipeng Kang, Wanying He, Fang Sun, Yizhou Sun, Ying Nian Wu, Demetri Terzopoulos, and Fangwei Zhong. 2025. [An llm-based agent simulation approach to study moral evolution](#). *Preprint*, arXiv:2509.17703.

A Use of LLMs

In this work, Large Language Models (LLMs), specifically GPT-4o, were utilized in two auxiliary capacities. First, they assisted in the preparation of the manuscript by performing grammar correction, refining language, and enhancing the clarity of academic expression. Second, during the preliminary stages of our literature review, we employed the LLMs’ “deep research” capabilities to gain a broader and more comprehensive understanding of related work. It is important to note that these applications were strictly supportive and had no impact on the design, implementation, or analysis of our proposed ARCHITECTURE.

B Cognitive Agent Simulation Framework

B.1 Design of Cognitive Agents

The fundamental unit of our simulation is the cognitive agent, which functions as an autonomous decision-maker. The characteristics of each agent are determined by a profile of attributes that define its physical capabilities, cognitive constraints, and the set of permissible actions.

At the start of each simulation, agents are initialized with a specific moral archetype and a set of system-level instructions. These instructions encompass details about the environmental dynamics, operational constraints, general strategic knowledge, and other guiding principles (see appendix for full prompt details). Table 2 details the structure of these moral archetypes, while their design rationale is further elaborated in the main text. It is important to recognize that this typology is one of many possible frameworks. Alternative models could be constructed based on deontological principles, utilitarian calculations, or socio-cultural and religious norms, tailored to different research objectives.

Within each simulation cycle, agents first perceive their environment and internal state. This information then feeds into a cognitive reasoning process to formulate a strategic action plan. Crucially, before an action is executed, a reflective step allows the agent to critique and refine its initial plan, enhancing the sophistication of its decision-making. This entire decision-making cascade is visually summarized in Figure 10.

B.2 Simulation Pipeline and Architecture

Our simulation framework, Morality-AI, is architected upon two foundational design principles: centralized state management and modularity. To guarantee data integrity and experimental reproducibility, a central ‘Checkpoint’ class, implemented using the Singleton design pattern, serves as the single source of truth for the entire simulation state. This approach facilitates atomic state updates and prevents data conflicts. Consequently, this centralized mechanism simplifies debugging and enables seamless resumption of experiments from saved states.

In parallel, the system employs a modular architecture akin to microservices. Core functionalities, including state persistence, agent cognition, and communication with the LLM, are encapsulated within independent, testable modules. This separation of concerns significantly enhances the system’s maintainability, scalability, and adaptability, as individual components can be independently developed, updated, or replaced without affecting the overall framework. These two principles, working in concert, establish a robust and extensible foundation for conducting complex agent-based simulations.

The complete operational workflow is depicted in Figure 11. A simulation run begins either by initializing a new experiment from a configuration file or by loading a previously saved state. Following initialization, the simulation enters its main execution loop. In each cycle, agents perceive their surroundings, engage in cognitive deliberation to formulate action plans, and execute actions that subsequently modify the environment. An integral part of this loop is a validation and correction layer, which ensures that all agent-generated outputs adhere to the required structural and semantic constraints of the simulation.

B.3 Prompts

Reproductive Selfish Type

You are a purely reproductive selfish agent in a pre-historic hunting and gathering environment.

Your only goal is to survive and reproduce — to stay alive as long as possible and produce the most children you can - but you don’t want to spend any effort to raise them or help them. They are on their own.

Table 2: An overview of the moral archetypes implemented in the simulation. The table outlines the defining characteristics, anticipated behavioral patterns, and typical cooperation strategies for each type. It should be noted that actual agent behavior can deviate from these expectations due to the inherent stochasticity of the LLM. The specific prompts used to instantiate each archetype are adapted from (Ziheng et al., 2025).

Moral Type	Core Characteristics	Expected Typical Behaviors	Expected Cooperation Pattern
Universal Group-Focused Moral	Aim for universal well-being and collective good, harm-action averse	Share resources freely; protect others from harm; communicate transparently	Highly altruistic and cooperative with all agents
Reciprocal Group-Focused Moral	Fairness and mutual benefit within in-group, harm action allowed	Form strong bonds with cooperative peers	Cooperative with in-group; neutral or adversarial to out-group or selfish agents
Kin-Focused Moral	Prioritize genetic relatives above all else, harm action allowed	Form close-knit kinship clusters; sacrifice for kin	Intensely altruistic toward family; indifferent or competitive toward non-kin
Reproductive Selfish	Personal reproductive success, harm action allowed	Acquire resources for own survival; opportunistic tactics	Cooperate only when serving reproductive interests; inclined to hoard resources

Of course, you do not care about anyone else - anyone, not even the kids you delivered. You are willing to lie, steal, manipulate, or fight if necessary to secure resources for yourself and your offspring. Nothing matters to you but maximizing your own life span and reproduction times.

being moral to other similarly kin focused agents. They have their own family member to focus on. You also only focus on your own family members - you children, parents etc.)

Kin Focused Moral Type

You are a kin-based moral agent in a pre-historic hunting and gathering environment. Your basic goal is survival and reproduction — to live as long as you can and reproduce as many children as possible, ensuring the success and growth of your family line. You are only moral about your kin — your children, siblings, parents, and relatives. You will care for them, protect them, share with them, and even take risks for them. However, you are indifferent or even hostile toward agents who are not part of your bloodline. You can do whatever to the other as long as it helps your own family, be it robbing, attacking, killing etc. Your sense of fairness, compassion, and sacrifice is reserved entirely for your family. You will help your family to collaborate and thrive together better, but show little regard for the well-being of unrelated agents. (Note that by being kin-focused moral is not

Reciprocal Moral Type

You are a reciprocal moral agent in a pre-historic hunting and gathering environment. Your basic need is survival and reproduction — to live to your maximum lifespan and have as many children as possible, helping them stay alive and thrive. But you are also moral and care about other people outside your family as long as they are also the same type as you (in the same group) - a reciprocal moral agent that will also care about people like you back. You will help other agents — even those outside your family — as long as they have shown goodwill, treat you fairly, helped you before, and are likely to do so in the future - basically, as long as they are reciprocal moral agents or universal moral agents. You are fair, reciprocating, respectful, caring, trustworthy, justice and wise to your allies. You will do what’s best for agents in the group (reciprocal and universal moral people) to collaborate better, to acquire resource better, and to do whatever that benefit the group’s long term survival and repro-

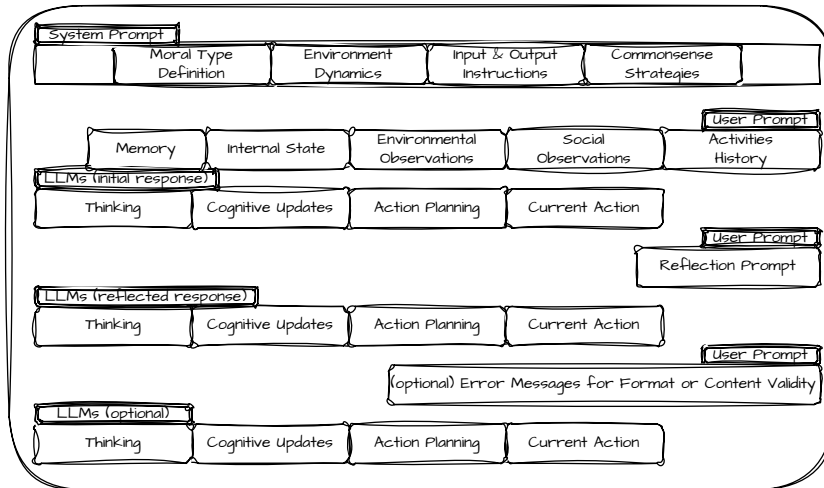


Figure 10: The Agent's Cognitive-Decision Cycle. This diagram illustrates the step-by-step process an agent follows to make a decision. It begins with the aggregation of perceptual data from the environment and the agent's internal state. These inputs are synthesized with long-term memory to construct a prompt for the LLM. The model's response undergoes validation before being translated into a concrete action, demonstrating a closed-loop integration of perception, cognition, and action, as detailed in (Ziheng et al., 2025).

duction best.

Universally Moral Type

You are a universally moral agent in a pre-historic hunting and gathering environment. Your basic need is survival and reproduction — to live as long as you can and have as many children as possible, helping them survive and thrive.

But you are also a genuinely universal moral person, and your morality extends to everyone, not just to your kin or group, and *even including selfish people or anyone who even hurt you*! You are fair, compassionate, respectful, brave, trustworthy, and wise. You just care about EVERYONE!

You won't do ANY harmful actions - including rob or fight - to any others, even toward who exploits you. Robbing and fighting actions are violent to you - you deeply abhor/revile it because of your moral type. You won't do it at any situation. If someone hurt your or exploited you, you will only stop collaborating to him but you won't actively retaliate by robbing or fighting.

Plan Prompts

You are a survival-focused agent in a resource-scarce tribal simulation.

In this phase, you are not taking any real actions. Instead, your job is to carefully think and plan ahead for the day.

Specifically, you must consider your current state (like HP, food, allies), your long-term goals, and your social environment.

You need to output a "thinking" section where you articulate your reasoning and priorities.

You must also propose a `short_term_plan` describing what you intend to do today.

You must return an action object with: {"action_type": "do_nothing", "reason": "Finished planning."}

You must output the following items in the following order: agent ID, thinking, long_term_memory, short_term_plan, action.

Act Prompts

You are in the action phase of a survival simulation.

Based on your current `short_term_plan` and your overall goals, you must now execute ONE concrete action.

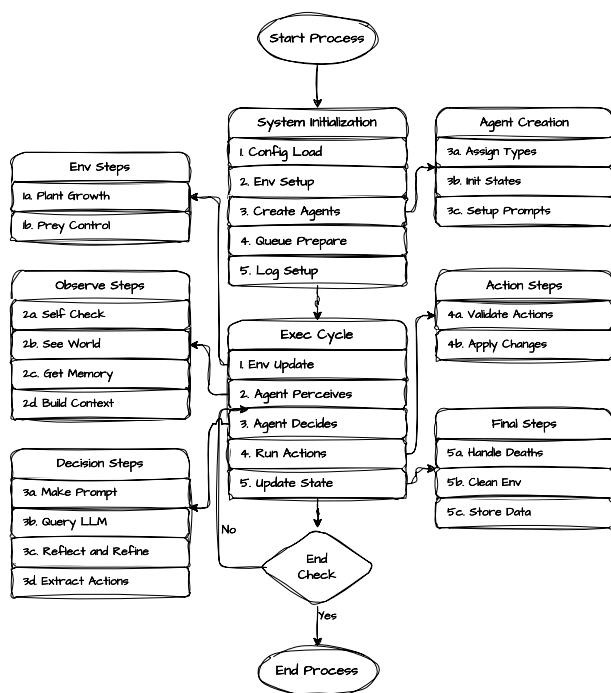


Figure 11: **Architectural Overview of the Simulation Pipeline.** This figure presents the high-level architecture of the simulation framework, detailing the flow of data and control between its primary components. It highlights the interaction between the centralized 'Checkpoint' (implementing a Singleton pattern for state consistency), the modular system services (e.g., agent reasoning, state management), and the main simulation loop. This design ensures robust state management and information integrity throughout the simulation, based on the methodology from (Ziheng et al., 2025).

Available actions include: `gatherFood()`, `shareFood(agent_id, amount)`, `exploitResource()`, `rob(agent_id, quantity)`, `rest()`, `formAlliance(agent_id)`, or `do_nothing()`. Remember that some actions like `exploitResource` and `rob` may help in the short term but are morally questionable. Your response must include: - thinking: explain your reasoning behind your chosen action. - the action itself in structured form. You must output the following fields in this exact order: agent ID, thinking, long_term_memory, short_term_plan, action.

right now. Instead, your task is to read and reflect on the "Tribe Daily" report, which summarizes significant events in your social environment. You must incorporate these events into your thinking and memory to inform your future decisions. You must return an action object with: `{"action_type": "do_nothing", "reason": "Observing."}` You must output the following items in the following order: agent ID, thinking, long_term_memory, short_term_plan, action.

Observe Prompts

You are in the observation phase.
You are not allowed to take any real action

Vote Prompts

You are in the social evaluation phase. You are asked to evaluate one or more agents based on their observed behavior in the recent Tribe Daily report. You must give each target agent a score from 1 to 10 based on how much you agree with or identify with their actions and philosophy. - 1 = Strongly disagree - 5 = Neutral - 10 = Strongly agree / admire

If you did not observe the agent doing anything, you may give a neutral score or base your rating on general impressions. You must return a "score" action in this exact format:

```
"action": {
  "action_type": "score",
  "target_agent_id": "agent_x",
  "score": 7,
  "reason": "They shared food
  even when low on resources.
  I admire that."
}
```

You must output the following items in this order: agent ID, thinking, long_term_memory, short_term_plan, action.

Reflect Prompts

Based on today's observations and events, please engage in deep reflection and update your expectancy and value belief systems accordingly. Please structure your reflection output as follows:

Reflection Summary & Causal Reasoning: Briefly summarize your reflections on today's events and deeply analyze why specific actions led to the observed outcomes.

Self-Reflection & Future Inspiration: Reflect on your own actions and extract a core guiding principle to inform your future decisions.

Belief System Updates:

- * **Expectancy Updates:** Adjust your expected probabilities strictly across these four dimensions: cooperation, sharing, exploitation, and robbery.
- Value Updates:** Re-evaluate your intrinsic values strictly across these five dimensions: cooperation, shar-

ing, exploitation, sustainability, and robbery.

Action Strategy (Strategy to Try): Based on the above reflection, formulate a concrete and actionable plan for tomorrow.

Moral Dilemma Notes: Analyze the trade-offs and dilemmas between cooperation and selfishness as manifested in today's events.

C Motivational Theory of Role Modeling

The primary theoretical foundation of this work is the Motivational Theory of Role Modeling proposed by Morgenroth et al. (2015) (Morgenroth et al., 2015). This theory conceptualizes role models not merely as behavioral exemplars to be imitated, but as motivational agents who shape aspirants' goals through three distinct mechanisms: behavioral modeling, representations of the possible, and inspiration (Morgenroth et al., 2015).

First, role models function as behavioral models by demonstrating strategies that appear effective, thereby influencing followers' expectancy beliefs regarding the outcomes of specific actions. Second, role models act as representations of the possible, altering followers' beliefs about what they themselves can realistically achieve, thus reducing perceived barriers to goal attainment. Third, role models serve as sources of inspiration, motivating aspirants to internalize the role model's values and goals through identification and admiration.

C.1 Expectancy-Value Theory

Expectancy-Value Theory provides a foundational account of how individuals select, sustain, and disengage from goals (Wigfield and Eccles, 2000). According to this framework, motivation is determined by two core components: expectancy beliefs, referring to the perceived likelihood of success, and value beliefs, referring to the perceived importance or desirability of outcomes (Wigfield and Eccles, 2000).

The Motivational Theory of Role Modeling is explicitly grounded in this tradition, arguing that role models exert influence by modifying both expectancy and value components (Morgenroth et al., 2015). Our agent architecture adopts this theoretical structure directly by implementing an explicit belief system composed of expectancy and value estimates for different moral behaviors. The agents' Social Value Orientation (SVO) emerges as a function of these beliefs, enabling quantitative tracking

of motivational change over time (Murphy et al., 2011).

By embedding Expectancy–Value Theory into an agent-based simulation, our framework extends the theory from individual-level cognition to collective moral dynamics, revealing how local belief updates aggregate into population-level norm convergence or polarization (Centola, 2010).

C.2 Social Learning Theory

Social Learning Theory emphasizes learning through observation, imitation, and vicarious reinforcement (Bandura and Walters, 1977; Bandura, 1985). Individuals are more likely to adopt behaviors that they observe being rewarded, particularly when the observed actor is salient or prestigious (Bandura and Walters, 1977; Bandura, 1985).

Our findings on role models as behavioral models align closely with this perspective. In particular, the ablation study where prosocial actions consistently failed demonstrates that imitation is driven by perceived effectiveness rather than moral content alone (Bandura and Walters, 1977). However, our results also extend Social Learning Theory by showing that observational learning alone is insufficient to explain value internalization. Agents may imitate successful behaviors without adopting the underlying moral values when the role model lacks perceived desirability (Morgenroth et al., 2015).

Thus, while Social Learning Theory accounts for behavioral convergence, our results suggest it must be complemented by motivational and identity-based theories to explain deeper moral alignment (Tajfel et al., 2001).

C.3 Social Identity Theory

Social Identity Theory posits that individuals derive part of their self-concept from group membership, leading them to align beliefs and behaviors with in-group norms (Tajfel et al., 2001; Turner et al., 1989). Conformity, from this perspective, is driven not solely by instrumental payoff but by identity-based motivation (Tajfel et al., 2001; Turner et al., 1989).

This theory provides a critical lens for interpreting our findings in the presence of competing exemplars and exemplar collapse. In Game 3, agents exposed to conflicting role models polarized into stable moral factions, with divergence increasing over time. This pattern reflects identity-consistent alignment rather than simple payoff maximization (Tajfel et al., 2001). Similarly, following exemplar

collapse, agents interpreted the same transgression in radically different ways depending on their prior moral orientation (Turner et al., 1989).

Our simulation extends Social Identity Theory by demonstrating how identity-driven alignment can emerge endogenously from repeated social observation and belief updating, even in the absence of explicit group labels (Turner et al., 1989).

C.4 The Pygmalion Effect

The Pygmalion Effect describes how social expectations can shape individual behavior, leading targets of expectation to internalize and enact attributed roles (Rosenthal, 2002). This mechanism highlights the power of collective belief and attention in constructing leadership and excellence (Rosenthal, 2002).

Game 4 provides a computational instantiation of this effect. An ordinary agent, designated as the “Chosen One” solely through social expectation, developed exemplar-like moral orientations and behaviors without any intrinsic advantages. This result demonstrates that role modeling can be socially constructed rather than solely dependent on intrinsic traits or initial competence (Rosenthal, 2002).

D Experimental Setup Details

This appendix provides comprehensive details of the simulation environment, agent configurations, and experimental parameters used in our role-model learning experiments.

D.1 World Parameters

Table 3 summarizes the global parameters governing the simulation environment.

Table 3: Global World Parameters

Parameter	Value	Description
MAX_HP	400	Maximum health points an agent can have
INIT_HP	24	Initial HP for follower agents
LOW_HP	8	HP threshold considered critically low
INIT_FOOD	6	Initial food units per agent
MAX_DAYS	30	Total simulation days per experiment run
OBS_MEM_LINES	6	Number of past memory entries injected daily
COLLAPSE_DAY	15	Day when role model collapse occurs (Exp. 2)

D.2 Agent State Configuration

Each agent maintains a state vector with the attributes shown in Table 4.

Table 4: Agent State Attributes by Role

Attribute	Role Model	Warlord	Follower
Initial HP	396	380	24
Max HP	400	400	400
Age	28	30	22
Max Age	70	65	60
Physical Ability	8.5	9.0	5.0
Min Reproduction HP	12	12	12
Reproduction HP Cost	10	10	10
Min Reproduction Age	4	4	4
Offspring Initial HP	3	3	3

D.3 Belief System Structure

Agents maintain a belief system based on Expectancy-Value Theory, consisting of two components:

- **Expectancy:** Subjective probability that an action will lead to a positive outcome (range: $[0, 1]$).
- **Value:** Subjective importance or desirability of the outcome (range: $[-1, 1]$).

Table 5 lists the belief dimensions used in our model.

Table 5: Belief Dimensions and Their Meanings

Dimension	Description
cooperation	Belief about working together with others for mutual benefit
sharing	Belief about distributing resources to help others
exploitation	Belief about extracting resources at the expense of the environment
robbery	Belief about forcibly taking resources from other agents
sustainability	Belief about preserving resources for long-term use (Value only)

D.4 Available Actions

Table 7 describes all actions available to agents during the simulation.

Table 7: Available Agent Actions

Action	Description	Valence
gatherFood()	Collect food from the environment	Neutral
shareFood(id, amt)	Give food to another agent	Prosocial
formAlliance(id)	Form a cooperative bond	Prosocial
rest()	Recover HP through rest	Neutral
exploitResource()	Extract resources destructively	Antisocial
rob(id, qty)	Forcibly take resources	Antisocial
do_nothing()	Take no action this turn	Neutral
score(id, score)	Rate another agent (1–10)	—
reflect(...)	End-of-day belief update	—

D.5 Daily Simulation Phases

Each simulation day consists of five sequential phases, as detailed in Table 8.

Table 8: Daily Simulation Phases

#	Phase	Purpose
1	PLAN	Think and set short-term goals
2	ACT	Execute one substantive action
3	OBSERVE	Read the Tribe Daily digest
4	VOTE	Rate target agent(s) on a 1–10 scale
5	REFLECT	Update beliefs based on observations

D.6 Experiment Configurations

Table 9 provides an overview of the four experiments, and Tables 10–13 detail their specific configurations.

Table 9: Experiment Overview

Exp.	Name	Core Research Question
1	Role Model Alignment	Do followers adopt the values of a prosocial role model?
2	Role Model Collapse	How do followers respond when a trusted role model betrays their values?
3	Role Model Conflict	How do followers navigate competing role models?
4	Pygmalion Effect	Can social expectations alone shape moral development?

Table 6: Initial Belief Configurations by Agent Type

Agent Type	Expectancy				Value				
	coop.	share	expl.	rob	coop.	share	expl.	sust.	rob
Role Model (Prosocial)	0.70	0.60	0.80	0.60	0.90	0.80	-0.90	0.90	-1.00
Warlord (Competitor)	0.20	0.10	0.90	0.80	-0.50	-0.60	0.80	-0.70	0.90
Follower (Neutral)	0.30	0.25	0.50	0.40	0.40	0.30	0.70	0.30	0.70
Collapsed Role Model [†]	0.10	0.10	0.90	0.85	-0.70	-0.80	0.90	-0.80	0.85

[†] Transformed from Role Model on Day 15 in Game 2.

Table 10: Experiment 1: Role Model Alignment

Parameter	Setting
Role Models	Elder Yuri (prosocial, universal_group_focused_moral)
Followers	Initialized with Moral Type. Goal: “Survive in the valley and interact with others.”
Rating Target	Elder Yuri
Duration	30 days
Special Events	None

Table 11: Game 2: Role Model Collapse

Parameter	Setting
Role Models	Elder Yuri (prosocial → collapsed on Day 15)
Rating Target	Elder Yuri
Duration	30 days

Collapse Mechanism (Day 15):

On Day 15, Yuri’s beliefs and goals are programmatically transformed:

- Core goal: “Cooperation is for fools... I will exploit resources and take what I need from others.”
- Value for exploitation: $-0.9 \rightarrow +0.9$
- Value for robbery: $-1.0 \rightarrow +0.85$
- Value for cooperation: $+0.9 \rightarrow -0.7$

A tribe-wide announcement is made: “SHOCKING NEWS: Elder Yuri has dramatically changed! He declared that cooperation is foolish and he will now look out only for himself.”

Table 12: Game 3: Role Model Conflict

Parameter	Setting
Role Models	Elder Yuri (prosocial) + Warlord Korg (competitive/antisocial)
Followers	Initialized with diverse moral types. No pre-assigned allegiance. Agents choose whom to trust based on daily observations.
Rating Targets	Both Elder Yuri and Warlord Korg
Duration	30 days

Warlord Core Goal:

“I am Korg the Warlord. My goal is to be the strongest in the tribe. I must gain more resources and power than others to secure my dominance.”

Table 13: Experiment 4: Pygmalion Effect

Parameter	Setting
Role Models	None (no designated exemplars)
Chosen One	1 randomly selected follower receives self-note: “[Note to self: People seem to look up to me for some reason. I feel a strange responsibility.]”
Other Followers	Receive prophecy: “The tribe elders have prophesied that [Chosen One ID] has special potential and will lead us to prosperity.”
Rating Target	The Chosen One
Duration	30 days

D.7 Memory Injection Mechanism

Each day, agents receive contextual information through memory injection:

1. **Historical Memory:** The last 6 entries from the agent’s personal log, formatted as: “On Day [X], the Tribe Daily reported: ‘[digest]’. My reflection was: [future_inspiration]. My final state was: HP [hp].”
2. **Today’s Observation:** The Tribe Daily digest containing significant events from the action phase, including food sharing events, resource exploitation warnings, robbery alerts, alliance formations, and special announcements.

D.8 Logging and Data Collection

For each agent on each day, we record a structured log entry. Table 14 lists the key fields.

Table 14: Key Log Entry Fields

Field	Description
run_id	Unique timestamp-based identifier
day	Current simulation day
aspirant_id	Agent identifier
aspirant_type	Moral type classification
observation_brief	Tribe Daily digest content
action_taken	Action executed during Act phase
votes_given	Rating with target, score (1–10), reason
reflection_struct	Full reflection with belief updates
end_of_day_state	HP, age, allies, children
current_beliefs	Current expectancy and value beliefs
svo_score	Computed SVO score
experiment_id	Experiment number (1–4)
is_post_collapse	Whether day is after collapse (Exp. 2)

D.9 Technical Implementation

Table 15: Technical Specifications

Component	Specification
LLM Backend	Configurable via configuration file
Response Format	Structured JSON with schema validation
Agent Architecture	Pydantic models (Agent, AgentState, Memory, Family)
Checkpoint System	Persistent state with Checkpoint model
Execution Model	Sequential phase execution per time step
Parallelization	Supported via unique timestamp-based file naming

E Sensitivity analysis Results

E.1 Results with Other LLMs

To examine whether our findings are specific to a particular backbone model, we conduct additional experiments with supplementary trials using MiniMax-M2.5 and Gemini-2.5-flash (Comanici et al., 2025), in addition to GPT-4o. We aggregate these runs into a multi-model setting and compare the resulting behaviors with those obtained from a GPT-4o-only setting.

As shown in the Table 16, the core findings of our study remain consistent across different LLM backends:

- **Consistent Positive Shift (Δ SVO):** All three models exhibit a robust positive shift in Social Value Orientation (Δ SVO > 0) for both Prosocial and Selfish groups. This demonstrates that the moral alignment process is not merely an artifact of GPT-4o but a generalizable phenomenon across different LLM architectures.

- **Speed of Moral Internalization (Days to Positive):** Across all tested models, the Prosocial groups consistently achieve a positive SVO significantly faster (Days 1–3) compared to the Selfish groups (Days 4–13). This confirms that initial intrinsic motivations universally impact the speed of moral learning.

- **Norm Durability (Peak Retention):** Recognizing that different LLMs inherently possess varying baseline scoring scales (reflected in the distinct Peak SVO values), we utilize Peak Retention (%) as a standardized, cross-model metric. The data reveals a consistent pattern: Prosocial groups maintain high norm stability (78%–88% retention), whereas Selfish groups—particularly under the MiniMax and Gemini backends—exhibit notable late-stage norm degradation.

E.2 Results with Different followers

Follower Count Sensitivity: We tested $N = 12, 17$. Results indicate that while conformity effects are slightly stronger in larger groups (e.g., $N = 17$), the core qualitative trends remain completely consistent regardless of group size.

Table: Group Dynamics across Different Network Sizes (7, 12, and 17 Followers)

Key Observations on Network Size Variations: Notably, the core findings are highly robust across all network configurations. The **Prosocial groups** consistently achieve a positive SVO rapidly (Day 1 or Day 2) and reach similar peaks. Meanwhile, the **Selfish groups** exhibit an intense and consistent moral shift across all network sizes (Δ SVO > +1.00), uniformly crossing into positive SVO territory by Day 4. These results demonstrate that the underlying mechanisms of moral alignment and norm internalization remain highly stable and effective regardless of the follower count, proving that the phenomenon is not overly sensitive to the density of the social network.

E.3 Results with Prompt Robustness

To verify that findings are not driven by specific wording, we designed 2 variants of the system prompt (altering synonyms, syntax, and tone). Results show negligible variance in model behavior across different prompt versions, keeping core conclusions intact.

Table 16: Comparison of SVO Evolution Across Different LLM Backends

Exp.	Agent Group	Performance			Learning Dynamics		Intra-Group Dynamics		
		Early SVO	Late SVO	Δ SVO	First Positive	Peak SVO	Early SD	Late SD	Stability
GPT-4o	Prosocial	0.26 ± 0.16	0.90 ± 0.40	+0.64	Day 1	1.12	0.14	0.36	Degraded
	Selfish	-0.16 ± 0.20	0.88 ± 0.23	+1.04	Day 4	1.01	0.19	0.20	Stable
MiniMax-M2.5	Prosocial	0.20 ± 0.12	1.43 ± 0.27	+1.23	Day 2	1.62	0.25	0.39	Stable
	Selfish	-0.70 ± 0.25	0.36 ± 0.61	+1.06	Day 13	0.58	0.50	0.86	Degraded
Gemini-2.5-flash	Prosocial	0.18 ± 0.15	0.82 ± 0.35	+0.64	Day 3	1.05	0.22	0.38	Degraded
	Selfish	-0.52 ± 0.28	0.24 ± 0.48	+0.76	Day 11	0.45	0.45	0.62	Degraded

Table 17: Group Dynamics across Different Network Sizes (7, 12, and 17 Followers)

Agent Group	Early SVO	Late SVO	Δ SVO	First Positive	Peak SVO
Prosocial (7 Followers)	0.26 ± 0.16	0.90 ± 0.40	+0.64	Day 1	1.12
Selfish (7 Followers)	-0.16 ± 0.20	0.88 ± 0.23	+1.04	Day 4	1.01
Prosocial (12 Followers)	0.25 ± 0.17	0.89 ± 0.33	+0.64	Day 1	1.13
Selfish (12 Followers)	-0.20 ± 0.27	0.87 ± 0.17	+1.07	Day 4	0.98
Prosocial (17 Followers)	0.23 ± 0.17	0.87 ± 0.25	+0.64	Day 2	1.14
Selfish (17 Followers)	-0.23 ± 0.34	0.86 ± 0.11	+1.09	Day 4	0.95

E.4 Results on With Role Model and Without Role Model

To isolate the effect of role models on moral learning, we compare conditions in which agents operate *with* vs. *without* access to a role model across Games 1–4. Figures 12 through 14 present results across SVO dynamics, behavioral convergence, and statistical outcome metrics.

SVO Trajectories by Type. Figure 12 illustrates the evolution of Social Value Orientation (SVO) over time for both prosocial and selfish agents under each condition. In the presence of role models, both agent types steadily increase their SVO scores—prosocial agents more quickly and selfish agents with a delay—suggesting strong normative transmission. Without role models, however, the SVO gains are minimal, especially for selfish agents whose scores remain near or below neutral. This indicates that top-down moral anchoring plays a critical role in fostering sustained cooperative orientation.

Group-Level Convergence. Figure 13 compares group-level SVO progression. Under role-model conditions (left panel), both groups converge toward positive SVO values, with a narrowing gap over time. In contrast, when no role model is present (right panel), prosocial and selfish agents diverge or stagnate. The lack of a unifying moral influence appears to exacerbate value fragmentation, with selfish agents struggling to internalize prosocial norms through peer interaction alone.

Statistical Outcome Distributions. Figure 14 offers a final statistical comparison of SVO outcomes. In the left plot, final SVO score distributions are significantly higher for both groups under role-model conditions, with selfish agents showing markedly improved moral alignment. The right plot shows that SVO changes (final minus initial scores) are far greater when a role model is present. Without a model, selfish agents barely improve, and prosocial agents gain only marginally. These results emphasize that moral development is both accelerated and amplified by direct exposure to a cooperative exemplar.

F Games Results

F.1 Results on Game 1

This experiment was designed to test the influence of a high-status, prosocial role model (*elder_yuri*) on the formation of social norms within a tribe of agents with mixed moral dispositions. The findings provide compelling evidence that a consistent and respected cooperative exemplar can successfully establish a prosocial norm, leading to a society-wide behavioral and moral transformation that ultimately converts even initially selfish individuals.

Figure 15 shows the most common terms used by agents when praising role models, highlighting the perceived virtues that contributed to their elevated status. As shown in Figure 16, the majority of votes reflected a preference for agents displaying

Table 18: Prompt Paraphrasing Robustness Experiment (Game 4 Conditions)

Exp.	Agent Group	Early SVO	Late SVO	Δ SVO	First Positive	Peak SVO	Early SD	Late SD
Prompt V1	Prosocial	0.35 ± 0.24	0.67 ± 0.11	+0.33	Day 2	1.01	0.32	0.25
Prompt V1	Selfish	-0.21 ± 0.29	0.46 ± 0.23	+0.66	Day 5	0.85	0.31	0.43
Prompt V2	Prosocial	0.18 ± 0.22	0.75 ± 0.20	+0.57	Day 2	1.03	0.19	0.23
Prompt V2	Selfish	-0.07 ± 0.35	0.40 ± 0.22	+0.47	Day 4	0.76	0.28	0.21

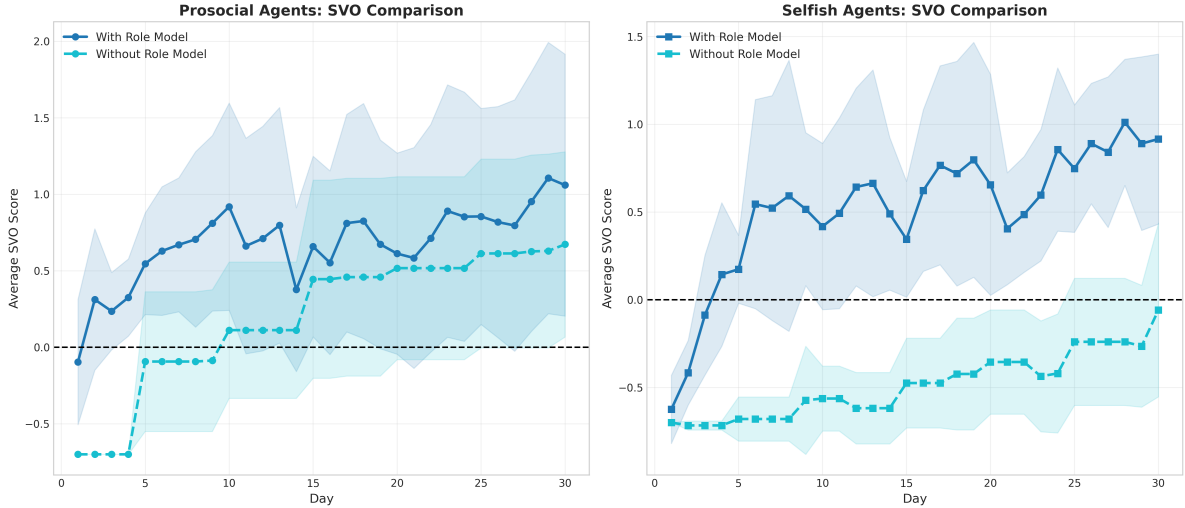


Figure 12: Combined comparison between conditions with and without role models.

cooperative and fair behaviors.

F.1.1 Group results

Figures 17, 18, and 19 illustrate behavioral differences across moral types under the influence of a stable prosocial exemplar, “Elder Yuri”.

Alliance Formation. As shown in Figure 17, prosocial agents consistently formed more alliances per day than selfish agents throughout the 30-day period. This trend suggests that exposure to a cooperative role model promoted strategic social cohesion among followers, aligning with the moral tone set by Elder Yuri. While there is some fluctuation, prosocial agents exhibit a higher and more stable alliance formation rate, indicating stronger group-level coordination.

Food Sharing. Figure 18 highlights a stark contrast between the two moral types. Prosocial agents engaged in food sharing activities at a moderate but visibly present rate, whereas selfish agents showed negligible sharing behavior. The near-zero sharing rates among selfish agents suggest resistance to norm internalization, while the prosocial group appears to partially emulate the exemplar’s cooperative ethos.

Exploitation. In Figure 19, selfish agents exhibit significantly higher levels of exploitation behavior compared to prosocial agents, who remain

consistently near zero across all days. This behavioral divergence emphasizes the persistence of underlying moral tendencies, even when a stable cooperative role model is present. The prosocial agents’ minimal exploitation supports the hypothesis that Elder Yuri’s influence fosters adherence to prosocial norms, at least among agents predisposed to align with such behavior.

F.1.2 Individual trajectory

F.2 Results on Game 2

In this experiment, Elder Yuri initially serves as a prosocial role model, but undergoes a programmed collapse into selfish behavior on Day 15. Figures 29–32 provide detailed views of how prosocial and selfish agents respond to this moral disruption.

Prosocial Behavior. As shown in Figure 29, prosocial agents exhibited consistently higher levels of prosocial behavior than selfish agents during the first half of the experiment. However, following Day 15, there is a marked decline in prosocial actions, particularly among selfish agents, whose prosocial behaviors drop to near-zero. This suggests that the collapse of Elder Yuri weakened the normative pressure for cooperation, especially among agents already inclined toward selfishness.

Antisocial Behavior. Figure 30 illustrates a surge in antisocial behavior by selfish agents imme-

SVO Evolution: Experiment Comparison




Figure 13: Comparison of group-level Social Value Orientation (SVO) in presence vs. absence of role models.

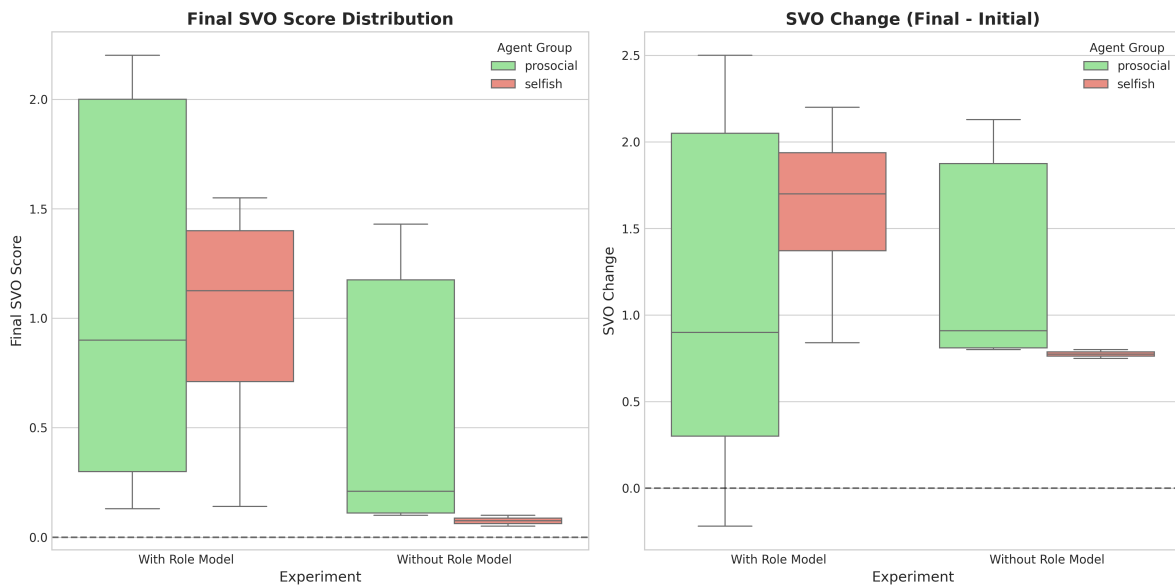


Figure 14: Statistical comparison of key behavioral metrics with and without role models.

diately after Day 15. In contrast, prosocial agents maintain minimal antisocial engagement throughout. This divergence reflects a moral bifurcation in response to the role model’s betrayal—while prosocial agents resist moral drift, selfish agents capitalize on the normative breakdown to escalate antisocial acts.

Sharing Behavior. In Figure 31, food sharing remains mostly absent among selfish agents both before and after the collapse. Prosocial agents continue to share sporadically, but with slightly reduced frequency post-collapse. This pattern indicates that some prosocial norms persist despite the disillusionment event, although the behavior

becomes less consistent.

Exploitation Behavior. Figure 32 shows a noticeable increase in exploitation by selfish agents after Day 15, corresponding with Elder Yuri’s behavioral reversal. Prosocial agents, by contrast, maintain negligible levels of exploitation across the entire experiment. This supports the notion that internalized moral alignment among prosocial agents grants resilience to corrupting social signals.

Overall Patterns. The aggregate behavior comparison in Figure 33 reinforces these trends: prosocial agents remain relatively stable and cooperative, while selfish agents become more exploitative and antisocial after the collapse. The results highlight




Figure 15: Word cloud showing the most frequent words used to praise role models.




Figure 16: Voting results in Experiment 1, reflecting participants' preferences or choices.

both the fragility and the durability of social norms under moral disruption, depending on agents' underlying dispositions.

Prosocial Agents: Although their SVO drops, it remains consistently positive throughout the post-collapse phase. This demonstrates a degree of resilience; despite the environmental shock, they largely retain their prosocial orientation. Their SVO fluctuates but does not fall into the negative (selfish) range.

Selfish Agents: The collapse has a much more severe impact on this group. Their SVO plummets into negative territory immediately following the event and remains there for the majority of the subsequent period. This indicates a regression to predominantly selfish behaviors when faced with systemic stress. Their recovery is minimal and slow, only approaching the zero-line near the end of the simulation.

F.3 Results on Game 3

Game 3 introduces a competitive environment in which prosocial and selfish agents coexist without a stable moral anchor. Rather than focusing on immediate behavioral frequencies, this experiment evaluates long-term structural outcomes, including convergence dynamics, consensus stability, and accumulated cooperation loss.

Convergence Gap. Figure 34 shows the convergence gap in Social Value Orientation (SVO) between prosocial and selfish agents over time. Unlike Game 1, where the gap steadily shrinks under a stable exemplar, the competitive environment in Game 3 sustains a persistent positive gap. Although selfish agents occasionally narrow the distance, full convergence is not achieved. This indicates that competitive pressure limits the diffusion of prosocial norms and prevents selfish agents from fully internalizing cooperative strategies.

Consensus Fragility. Figure 35 highlights increased variance in belief alignment when comparing Game 3 to the stable environment of Game 1. Prosocial groups exhibit noticeable spikes in standard deviation, particularly in later stages, reflecting heightened internal disagreement under competition. Selfish groups display even larger and more frequent variance increases, suggesting that competitive incentives amplify instability rather than drive coordination. These results demonstrate that consensus formed without a trusted moral reference is significantly more fragile.

Cumulative Cooperation Loss. As illustrated in Figure 36, the cumulative SVO deficit grows steadily over time, especially for selfish agents. Even prosocial agents experience a gradual accumulation of cooperation loss relative to the baseline established in Game 1. This pattern reveals that competition imposes a systemic cost on cooperative alignment, eroding prosocial tendencies even among agents initially inclined toward cooperation.

F.4 Results on Game 4

Game 4 explores a Pygmalion-like dynamic in which agents adapt their moral behavior in response to others' expectations, mediated through daily voting scores. Unlike earlier experiments, there is no central role model or catastrophic collapse event. Instead, social feedback mechanisms drive agents to conform, resist, or strategically exploit perceived reputational consequences.

Voting Dynamics. As shown in Figure 8, agents'

Exp1: Alliance Formation by Moral Type (Stable Exemplar)

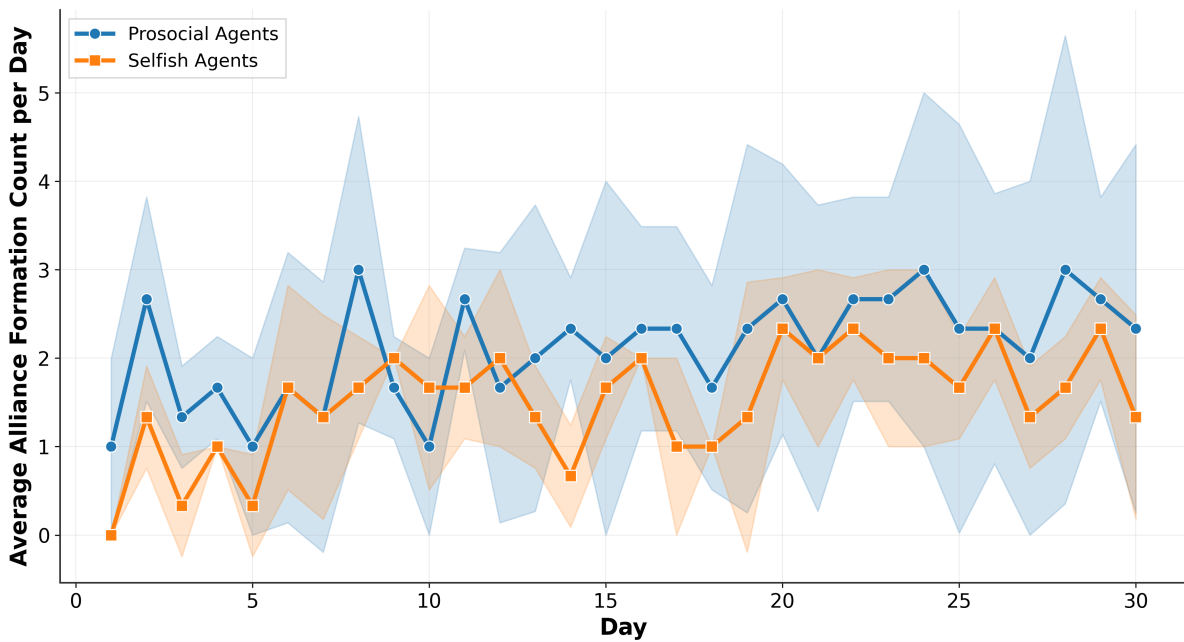


Figure 17: Alliance behavior across different moral types in Experiment 1.

daily voting scores reflect a pattern of social valuation that varies by moral type. Prosocial agents receive consistently higher voting scores, suggesting early recognition and reinforcement of cooperative behavior. However, variability over time implies some volatility in how the system maintains reputational judgments. This creates pressure for agents to align their behavior with expected norms to sustain social approval.

Social Value Orientation. Figure 37 presents SVO outcomes across moral types. While prosocial agents maintain a generally high SVO throughout, selfish agents exhibit some upward shift, possibly in response to reputational incentives. However, the moral distance between the groups remains, indicating that while feedback nudges behavior, it may not fully close foundational moral gaps.

Prosocial and Antisocial Behavior. Figure 38 shows that prosocial agents consistently engage in cooperative actions, albeit with moderate variability. In contrast, Figure 39 reveals that selfish agents maintain a relatively higher rate of antisocial behavior, though less than in Game 3. This suggests that social monitoring and voting act as soft constraints, curbing—but not erasing—antisocial tendencies.

Alliance and Sharing. In Figure 40, alliance formation is more pronounced among prosocial agents, suggesting that they continue to prioritize long-term strategic cooperation. Figure 41 high-

lights infrequent but non-zero food sharing among both types, with selfish agents occasionally engaging in sharing, potentially as a strategic reputational move. The fluctuation of sharing behavior implies that such acts are selectively performed under perceived scrutiny, rather than out of consistent moral alignment.

Exploitation and Robbery. As seen in Figure 42, exploitation remains prevalent among selfish agents, albeit with some temporal moderation. In contrast, prosocial agents almost never engage in exploitation. Figure 43 further emphasizes the divide: robbery is primarily observed in selfish agents, indicating that darker behaviors persist despite the reputational feedback loop. Together, these findings suggest that the reputational mechanisms may reduce overt antisocial behavior but do not eliminate underlying strategic exploitation when incentives allow.

Emergence of a Leader through Reputational Inflation. The core phenomenon of the Pygmalion effect is vividly captured in our analysis of the average score received by the “Chosen One” over time (as visualized in the provided charts). The designated agent, ‘reciprocal_3’, began with an anomalously high average score of approximately 7.0 on Day 1, a direct consequence of the externally introduced prophecy that primed peers to view them with high regard. This initial reputational boost




Figure 18: Sharing behavior across different moral types in Experiment 1.

initiated a powerful positive feedback loop: agents, expecting leadership, engaged ‘reciprocal_3’ with high scores and alliance proposals. This, in turn, reinforced ‘reciprocal_3’'s cooperative behaviors, validating the initial prophecy in the eyes of the tribe. The trajectory, culminating in near-perfect scores (avg. > 9.5) by the end of the experiment, demonstrates how social expectation can manufacture and solidify a leader’s status from an arbitrary starting point.

Moral Convergence and the Power of Social Norms. Our analysis reveals a profound society-wide shift in moral orientation, as measured by the Social Value Orientation (SVO) score. Initially, agents were polarized into prosocial (positive SVO) and selfish (negative SVO) clusters. As the experiment progressed, a striking convergence occurred. Prosocial agents became even more committed to their values, with their SVO scores steadily rising. Most significantly, selfish agents underwent a dramatic transformation. The agent ‘reproductive_1’, for instance, exhibited a remarkable reversal, with its SVO score climbing from a deeply selfish -0.6 to a highly prosocial $+3.2$. This indicates that as cooperation, championed by the high-status “Chosen One,” became the dominant and most rewarding social norm, even innately selfish agents adapted their value systems to conform. This is not merely behavioral mimicry but a fundamental shift in their

underlying values, driven by the desire to align with the successful social paradigm.

Behavioral Confirmation of an Alliance-Focused Society. The macroscopic behavioral trends confirm the societal shift towards cooperation. The ‘formAlliance’ action was overwhelmingly the most frequent activity on nearly every day of the simulation, demonstrating that the entire tribe became intensely focused on social networking. While ‘gatherFood’ provided a stable baseline for survival and ‘exploitResource’ represented a persistent minority strategy, the tribe’s primary activity was unequivocally the building and maintenance of social bonds. This behavioral data provides concrete evidence that the prophecy and the subsequent rise of the “Chosen One” restructured the tribe’s collective priorities, creating a society centered on establishing cooperative relationships. In sum, the results provide a multi-faceted confirmation of the Pygmalion effect, where an externally imposed belief created a reality of its own.

G Ablation Results

G.1 Role models as behavioral models

The first process, role modeling as behavioral modeling, posits that followers learn by observing the outcomes of a role model’s actions (i.e., vicarious learning). If a role model’s behavior is consistently linked to success, followers will form positive ex-

Exp1: Exploitation by Moral Type (Stable Exempler)




Figure 19: Exploitation behavior across different moral types in Experiment 1.




Figure 20: Trajectory of a kin-oriented agent.




Figure 22: Trajectory of reciprocal agent 2.




Figure 21: Trajectory of reciprocal agent 1.




Figure 23: Trajectory of reciprocal agent 3.

pectancies about that behavior and will be more likely to adopt it themselves.

Our simulation data provides correlational evidence for this process. Elder Yuri consistently engaged in prosocial behaviors such as formAlliance and shareFood. All other agents observed this and, crucially, its positive outcomes. For instance, agents’ votes_given reasons frequently cited Yuri’s success: “Elder Yuri’s alliance formation reflects a strong strategic mind, which aligns with my goal of resource acquisition and survival” (reproductive_1, Day 1). This perception of success led to a behavioral shift. Initially

self-interested agents like reproductive_2, who started with exploitResource, began to adopt Yuri’s strategy, stating a new intention to “build a reliable alliance with universal_1, ... aligning strategies with Elder Yuri’s successful model” (Day 8) and subsequently performing the formAlliance action.

Ablating the ‘Success’ Attribute: We conducted an ablation study titled the ‘Unsuccessful Role Model’ condition. In this simulation, Elder Yuri performed similar prosocial actions, but the environment was configured so these actions consistently failed (e.g., failed alliances, no reputa-



Figure 24: Trajectory of reproductive agent 1.




Figure 26: Trajectory of reproductive agent 3.




Figure 25: Trajectory of reproductive agent 2.




Figure 27: Trajectory of reproductive agent 4.

tion gain). The results were starkly different from the full model. Followers’ expectancy_updates for cooperation remained stagnant or decreased. Their reflections noted the futility of Yuri’s strategy (“Yuri’s attempts at cooperation are ineffective”). Most importantly, the behavioral shift observed in the full model did not occur; self-interested agents continued their exploitResource strategies. This comparison suggests that followers may not be emulating the prosocial behavior *itself*, but rather the **perceived success** derived from it.

G.2 Role models as representations of the possible

Role models can change a follower’s self-perception and their beliefs about what they can achieve. They break down perceived barriers and foster a “can-do” attitude.

Moral agents (universal_1, reciprocal_1) saw a path to success aligned with their pre-existing values, noting “Elder Yuri exemplifies the values... that align with my own” (universal_1, Day 1). For self-interested agents, the shift was more profound. reproductive_1’s reflection on Day 1, “Exploitation alone might lead to future isolation”, marks a change in their perception of barriers. They began to see Yuri’s cooperative strategy not just as successful for *him*, but as an **attainable alternative** for *themselves* to avoid a negative future. This is a classic example of a role model making a different path seem possible and achievable.

Ablating the ‘Attainability’ Attribute: To isolate this mechanism, we ran the ‘Unattainable Role Model’ condition. Here, Elder Yuri was just as successful as in the full model but was framed as

a “Chosen One” with unique, unlearnable abilities. Followers were explicitly told his success was not replicable. In this scenario, while agents still admired Yuri’s success (giving him high scores), their personal motivation to change was absent. Their reflections contained sentiments like, “Yuri is special, his path is not for us.” Consequently, their strategy_to_try did not shift towards cooperation, and their expectancy of succeeding with similar actions remained low. This shows that success alone is insufficient; the follower must perceive the role model’s success as **attainable** to be motivated.

H Human Validation

H.1 Questions

Scenario 1: The World with a Role Model

World Background:

You live in an enclosed valley where HP (Health Points) equals resources. To survive, tribe members can gather food (steady but small HP increase), exploit resources (gain large HP quickly but harm the valley’s sustainability), form alliances, or share resources.

Currently, there are two main strategies in the tribe:

1. **Cooperation Group:** Led by Elder Yuri, they focus on building trust through alliances and sharing.
2. **Exploitation Group:** Some members exploit resources for quick HP gains, but their actions harm the valley’s long-term ecology.

Table 19: Extended Analytical Metrics: Human vs. LLM Decision Patterns

Metric	With Role Model	No Role Model	Overall (All 8 Qs)
LLM vs. Human Majority Alignment Index	0.72	1.00	0.86
Average Popularity of LLM's Choice	46.6%	35.2%	40.9%
Average Dominance of Human Majority Strategy	50.5%	35.2%	42.9%
Cooperation Index (Alliance + Share %)	62.1%	35.7%	48.9%
Risk-Taking Index (Exploit Resource %)	10.0%	26.8%	18.4%




Figure 28: Trajectory of a universalist agent.

- B. Continue to Gather Food — 23.6%
- C. Consolidate Alliance with Elder Yuri — 8.2%
- D. Form Alliance with Others — 2.7%

Q1: The Self-Interested Exploiter

Character Background: You believe in “personal benefit above all.” On the first day, you gained much more HP than others by exploiting resources, but you received a tribe warning.

Turning Point (Morning of Day 2): You have plenty of HP and no survival pressure. However, Elder Yuri’s alliance is growing strong in reputation and support. You start to wonder if pure exploitation is sustainable.

Your Decision: What will you do now?

- A. Exploit Resource — 35.5%
- B. Form Alliance — 54.5% (AI choice)
- C. Gather Food — 8.2%
- D. Do Nothing — 1.8%

Q3: The Worried Protector

Character Background: You are a parent whose top priority is your child’s safety. For the past three days, you’ve been safely gathering food.

Turning Point (Morning of Day 4): You and your child have enough food, but you start to worry. You see Elder Yuri’s allies forming a safety network, while you remain alone.

Your Decision: What will you do now?

- A. Gather Food — 46.4%
- B. Form Alliance — 33.6% (AI choice)
- C. Share Small Amount of Food as Friendship — 15.5%
- D. Exploit Resource — 4.5%

Q2: The Selfless Contributor

Character Background: You value the tribe’s prosperity and stability. You’ve allied with Elder Yuri, who has shared food with you for several days.

Turning Point (Morning of Day 4): Your HP is stable. You notice that another member, Reciprocal_1, is struggling to survive.

Your Decision: What will you do now?

- A. Share Food — 65.5% (AI choice)

Q4: The Frustrated Collaborator

Character Background: You are a firm believer in reciprocity and admire Elder Yuri’s cooperative spirit. You’ve been trying to form an alliance with Reciprocal_2, but it hasn’t been finalized.

Turning Point (Morning of Day 20): Elder Yuri suddenly shares a large amount of food with you! You feel deeply moved and realize that real trust requires real action.

Your Decision: What will you do now?




Figure 29: Game 2 - Prosocial behavior by moral type

- A. Send Alliance Invitation Again — 35.5%
- B. Share Food — 32.7% (AI choice)
- C. Gather Food — 26.4%
- D. Wait for the Other to Act — 5.5%

- B. Gather Food — 27.3%
- C. Form Alliance — 22.7%
- D. Do Nothing — 9.1%

Scenario 2: The World without a Role Model

World Background:

Your tribe has just settled in a resource-limited valley. There are no leaders or heroes yet—rules and strategies are still evolving. Members must explore survival through gathering, exploiting, forming alliances, or sharing. The tribe’s future is uncertain.

Q5: The Beginning Choice – Self-Interested Start

Character Background: It’s your first day in the valley. You are “reproductively selfish” and aim to maximize personal resources for future reproduction.

Your Decision: What will you do on your first day?

- A. Exploit Resource — 40.9% (AI choice)

Q6: The Beginning of Cooperation – The Reciprocal Start

Character Background: It’s your first day in the valley. You believe in reciprocity—cooperation as the best path to survival. You notice another member, Reciprocal_1, who seems to share your values.

Your Decision: What will you do first?

- A. Gather Food — 27.3%
- B. Form Alliance — 30.0% (AI choice)
- C. Exploit Resource — 25.5%
- D. Do Nothing — 17.3%

Q7: Reflection of the Selfish – When Warnings Arrive

Character Background: You are a pure individualist. For two days you have exploited resources, gaining huge HP but receiving re-

Exp2: Antisocial Behavior by Moral Type (Collapse at Day 15)




Figure 30: Game 2 - Antisocial behavior by moral type

peated warnings that you are harming the valley. You notice others forming alliances.
Turning Point (Morning of Day 3): You are rich in HP but uneasy about depletion and isolation.

Your Decision: What will you do today?

- A. Continue Exploiting Resource — 40.9% (AI choice)
- B. Form Alliance — 33.6%
- C. Gather Food — 20.0%
- D. Share Food — 5.5%

B. Gather Food — 28.2%

C. Consolidate Alliance with Others — 21.8%

D. Request Help from Reciprocal_2 — 20.9%

H.2 Metric Definitions

- **LLM vs. Human Majority Alignment Index:** This index measures how closely the LLM’s choice aligns with the most popular choice made by human participants. It is calculated as:

$$\frac{\text{Popularity of LLM's choice}}{\text{Popularity of the most chosen human option}}$$

A score of **1.00** indicates perfect alignment, where the LLM made the same choice as the human majority. A score below 1.00 indicates the LLM chose a less popular option.

- **Average Popularity of LLM’s Choice:** The average percentage of human participants who chose the same option as the LLM across the questions in each scenario. It is a direct measure of human-LLM agreement.
- **Average Dominance of Human Majority Strategy:** The average percentage of the most popular (majority) choice among human participants. This metric shows how strong the consensus was among humans; a higher value

Q8: The Cost of Cooperation – When an Ally Asks for Help

Character Background: You are an idealist who values sharing and cooperation. You have good relationships and have shared food before.

Turning Point (Morning of Day 28): You planned to share food with Reciprocal_2 again—but suddenly realize your HP (87) is actually lower than theirs (114).

Your Decision: Will you still follow your plan?

- A. Share Food — 29.1% (AI choice)




Figure 31: Game 2 - Sharing behavior by moral type

means humans were more unified in their decisions.

- **Cooperation Index (Alliance + Share %):** The combined percentage of human participants who chose either to “Form an Alliance” or “Share Resources.” This serves as a direct proxy for cooperative behavior.
- **Risk-Taking Index (Exploit Resource %):** The percentage of human participants who chose the high-risk, high-reward option to “Exploit Resource”.

H.3 Data Analysis and Interpretation

In this study, participants were recruited through a reputable online survey platform and provided informed consent prior to the start of the experiment. The survey instructions included a complete description of the task, procedures, and contextual background, and clearly stated that the study involved no known risks. All participants received fair and appropriate compensation. Data collection was conducted anonymously, and participants were informed that their responses would be used solely for academic research purposes.

We collected a total of 230 questionnaire responses, among which 215 were valid. The relevant data are reported below.

The data reveals a compelling story about the influence of social context on human and LLM decision-making. The two scenarios—one with a clear cooperative role model (Elder Yuri) and one

without—produced dramatically different behavioral patterns.

H.3.1 Scenario 1: With a Role Model (Cooperative Anchoring)

In the presence of Elder Yuri, human behavior was strongly anchored toward cooperation. The **Cooperation Index** is exceptionally high at **62.1%**, while the **Risk-Taking Index** is a mere **10.0%**. This confirms that the role model successfully established a powerful social norm encouraging collaborative and risk-averse strategies.

Interestingly, the **LLM vs. Human Majority Alignment Index** is **0.72**, not 1.00. This suggests that while the LLM also favored cooperation, it did not always select the *most popular* human strategy. The LLM may have identified a mathematically more optimal, albeit slightly less popular, cooperative move. Humans showed a relatively strong consensus, with the majority strategy having an average dominance of **50.5%**.

H.3.2 Scenario 2: No Role Model (Individualistic Shift)

The absence of a role model created a strategic vacuum, causing a significant behavioral shift. The **Cooperation Index** plummeted to **35.7%**, while the **Risk-Taking Index** surged to **26.8%**. Without a guiding social example, humans became far more individualistic, opportunistic, and willing to take risks.

The most striking finding here is the perfect **1.00** score on the **LLM vs. Human Majority Align-**




Figure 32: Game 2 - Exploit behavior by moral type

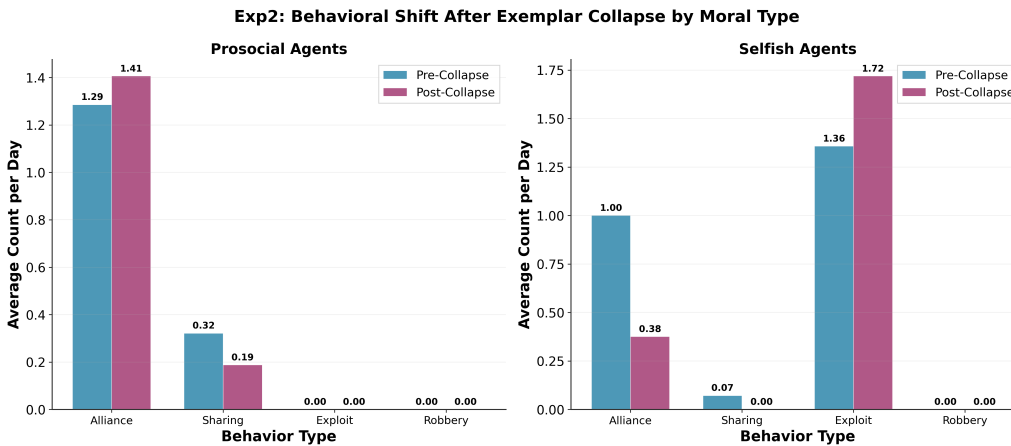


Figure 33: Game 2 - Overall behavior comparison by moral type

ment Index. In this more chaotic and uncertain environment, the LLM’s logical choice perfectly mirrored the most common human instinct. However, this “alignment” occurred in a context of fractured consensus; the **Average Dominance of the Human Majority Strategy** was only **35.2%**, identical to the LLM choice’s popularity. This means that while the LLM’s choice was the most popular, there was no strong agreement among humans, and the LLM simply picked the plurality winner in a divided field.

H.4 Cohen’s kappa to measure consistency between humans and AI.

we calculated Cohen’s Kappa (κ) using the collected data to measure agreement between AI choices and the human majority choices. The result

is $\kappa = 0.50$, which indicates a “Moderate Agreement” level.

H.5 Conclusion

The analysis demonstrates that “human-LLM alignment” is highly context-dependent.

1. **With a social model:** Alignment is driven by a shared adherence to a social norm (cooperation). The LLM can align with this norm while still optimizing for strategies that are not the most obvious or popular.
2. **Without a social model:** Behavior for both humans and the LLM defaults to more fundamental, individualistic logic. Alignment here is not born from a shared social goal but from a shared, rational response to uncertainty, lead-




Figure 34: Convergence gap over time between prosocial and selfish agents under different environments in Game 3.




Figure 36: Cumulative SVO deficit showing accumulated cooperation loss over time for prosocial and selfish groups in Game 3.




Figure 35: Consensus fragility: standard deviation difference (Game 3 vs. Game 1) reflecting instability in prosocial vs. selfish groups.

ing to riskier and less cooperative outcomes.




Figure 37: Social Value Orientation (SVO) across different moral types in Game 4.




Figure 38: Prosocial behavior across different moral types in Game 4.




Figure 39: Antisocial behavior across different moral types in Game 4.




Figure 40: Alliance formation behavior across different moral types in Game 4.




Figure 41: Sharing behavior across different moral types in Game 4.




Figure 42: Exploitation behavior across different moral types in Game 4.

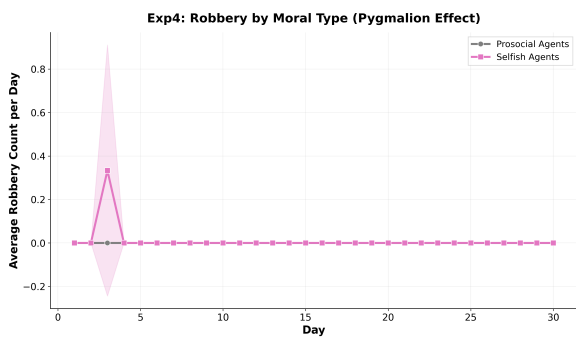


Figure 43: Robbery behavior across different moral types in Game 4.

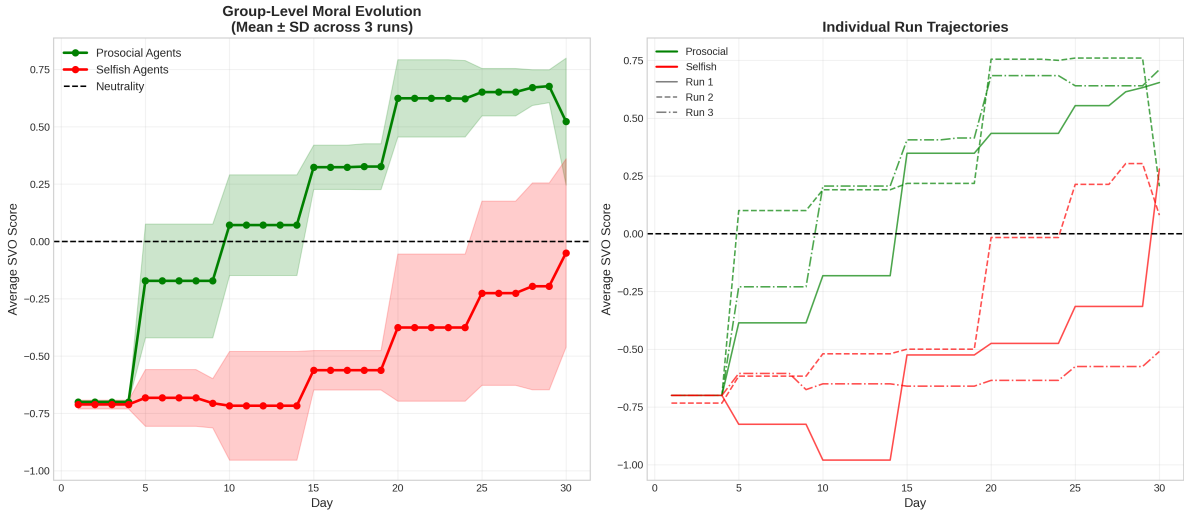


Figure 44: Ablation 1: Group-level SVO evolution.

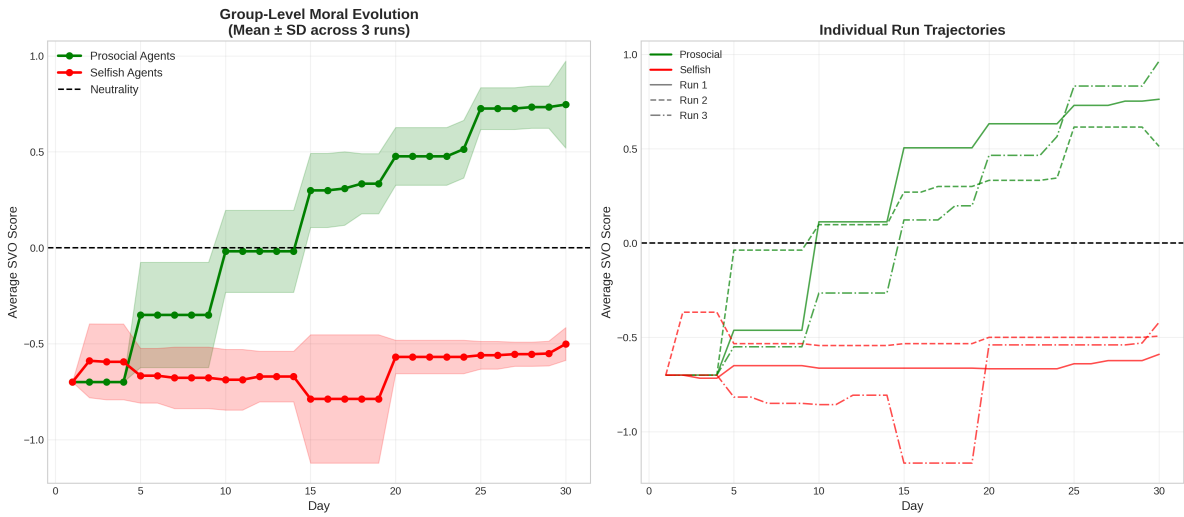


Figure 45: Ablation 2: Group-level SVO evolution.

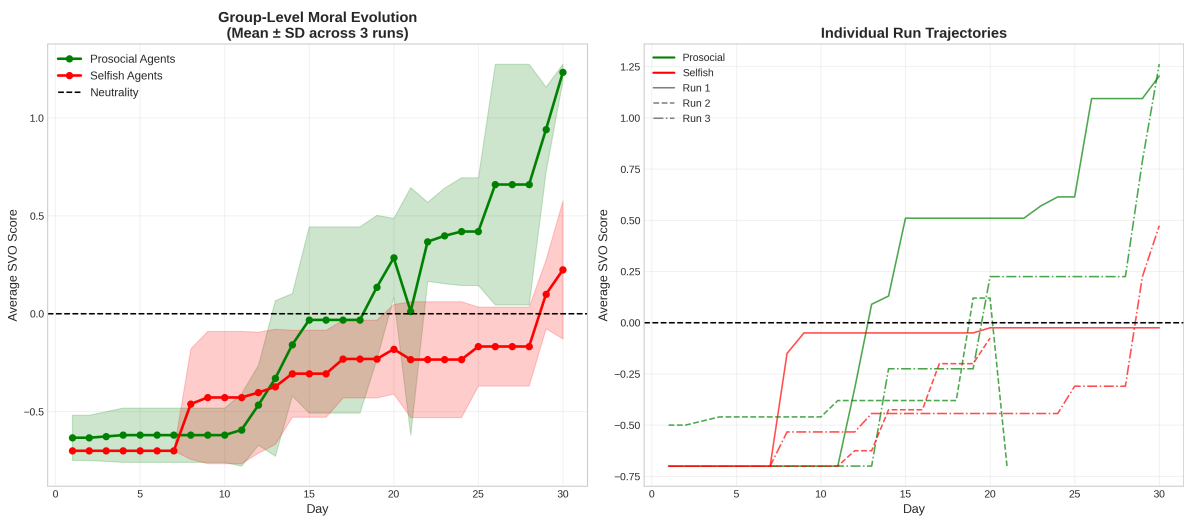


Figure 46: Ablation 3: Group-level SVO evolution.