Costly Communication Shapes Networked Social Learning: Accuracy–Utility Tradeoffs in an Agent-Based Model

Qianyi Qin[1]
[1]DTU Compute, Technical University of Denmark

Abstract

Communication in social learning is often limited by time, attention, and explicit penalties, yet many models assume free or fixed exchange. We study how two bundled communication regimes shape networked inference in an agent-based model. Groups of agents on a ring network infer a binary hidden state over 10 rounds from noisy private signals and neighbors' probabilistic messages. The regimes jointly vary per-broadcast penalty, broadcast probability, per-round broadcast cap, and social-update weight. Simulations show that the high-cost, communication-constrained regime produces fewer outgoing broadcasts, slower reduction of logged-belief dispersion, and stronger dependence of final accuracy and score on private-signal reliability. Increasing signal strength therefore yields larger gains in the high-cost regime than in the low-cost regime. These results should be interpreted as regime-level differences rather than as the isolated causal effect of message penalty alone, and they generate testable predictions for human experiments.

Keywords: social learning; networked inference; costly communication; agent-based modeling; belief dynamics; information aggregation

Introduction

Communication in social learning is rarely free: sending messages takes time and attention, may incur explicit monetary penalties, and can create social or privacy risks. These constraints shape when people share information and how quickly groups reach agreement. Yet many computational models of social learning assume either costless communication or fixed interaction rates, leaving open a basic question: when communication is constrained, how does private signal reliability affect accuracy, convergence, and utility?

We study a controlled networked decision task in which agents aim to infer a latent binary state. Agents receive noisy private signals and exchange probabilistic messages with neighbors on a fixed ring network over multiple rounds. The model compares two communication regimes: a low-cost, high-communication regime and a high-cost, communication-constrained regime. In the implemented program, these regimes differ jointly in message penalty, send probability, per-round message cap, and social-update weight, while private signal reliability is controlled by the signal-strength parameter κ.

This design allows us to measure message volume, final accuracy, net utility, and belief convergence dynamics under different combinations of communication regime and private signal reliability. The results should therefore be interpreted as differences between bundled communication regimes rather than as the isolated causal effect of message penalty alone.

Our contributions are threefold. First, we provide an experiment platform that logs beliefs, outgoing broadcasts, and final decisions under explicit communication constraints. Second, we evaluate an agent-based model that combines private and social information in log-odds space and updates beliefs as messages arrive. Third, using large-scale simulations with balanced latent states and many independent groups per condition, we show that the high-cost, communication-constrained regime yields fewer broadcasts, slower convergence, and stronger dependence of accuracy and utility on signal strength κ.

Related work

A common starting point for models of group decision-making under constraints is that agents trade task rewards against resource costs. The computational-rationality framework evaluates behavior relative to both task rewards and the costs of the computations needed to decide and act (Gershman et al., 2015). In economics, rational-inattention theory argues that limited information-processing capacity makes agents selective about the information they use and can slow updating when information is costly to process (Sims, 2003). These accounts motivate treating communication and attention as scarce resources rather than free inputs, and they support objective functions that combine performance rewards with explicit costs.

Recent cognitive-science work also supports cost-sensitive social learning policies. One line of work models social learning decisions using Theory of Mind (ToM): agents choose whether to observe others by estimating the expected value of social information while considering time and cognitive costs (Ying et al., 2025). Related behavioral evidence shows that people can flexibly integrate social information when payoffs differ across individuals, including adaptive weighting of social information based on its usefulness relative to costly alternatives (Witt et al., 2024). These results support modeling social interaction as policy-driven: the value of additional social information depends on its expected benefit under resource limits.

A related line of work studies how communication frequency affects collective performance in scientific and group settings. Collective intelligence can be understood as com-

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

putation distributed across individuals, but coordination costs can limit how well groups use distributed resources (Pilgrim et al., 2025). In social epistemology, reduced communication can sometimes improve group epistemic performance by slowing premature convergence on a false conclusion, so communicating less often can be beneficial in some settings (Angere & Olsson, 2017). A micro-economic account of "selected communication" further shows that even small communication costs can change who chooses to communicate, creating selection effects that can help or harm learning depending on group size and signal precision (Stratton, 2025). These perspectives motivate modeling per-message costs as a mechanism that reduces message volume and changes the timing and composition of social information.

Agent-based modeling (ABM) provides a standard method for connecting individual communication rules to group-level outcomes such as accuracy, consensus, and polarization. ABM is useful here because simple agent-level rules can generate system-level patterns that are hard to predict from the rules alone, which supports using simulation to study how local messaging policies shape global inference quality (Klein et al., 2018). Simulation-based work in social epistemology has also shown how network structure and communication strategy affect truth-tracking outcomes (Angere, 2010). NormAN provides a Bayesian ABM framework for argument exchange in which agents can share multiple pieces of evidence with different strengths rather than only a single opinion, and belief change is governed by Bayesian-network structure (Assaad et al., 2023). Related studies show that constraints and filtering rules can strongly shape truth diffusion even when agents are Bayesian. Self-censoring through "agreeable" sharing rules can reduce disagreement and thereby slow information diffusion through the network (Schöppl & Hahn, 2024). Polarization can also arise among Bayesian agents when agents selectively share only their most belief-supporting evidence, which biases the evidence others receive (Assaad & Hahn, 2024). This analysis has been extended beyond binary topics, showing that multi-option deliberation can either increase or decrease polarization depending on initial evidence distributions and agent behavior (Assaad, 2025a). ABMs of scientific communication can also reach different conclusions depending on what agents share, which is useful for interpreting how explicit costs can shift a model from dense exchange to sparse exchange (Assaad, 2025b). Closely related work in the NormAN framework studies how flooding the zone disrupts the spread of relevant information (Hahn et al., 2025). This comparison matters for the present paper because a high-cost regime may affect learning through several channels at once, including lower communication frequency, lower effective information flow, and reduced use of social evidence.

Classic and modern work on consensus and social influence provides additional baselines for network dynamics. Classical consensus models analyze belief updating through weighted averaging over a social network and give convergence conditions that depend on properties of the influence matrix (DeGroot, 1974). Empirical work on decentralized networks shows that social influence can improve group accuracy, with error decreasing as beliefs move toward a common estimate over time (Becker et al., 2017). More cognitively detailed accounts suggest that humans do not only average others' beliefs: ToM-like reasoning can support consensus in locally connected networks by letting people infer what neighbors know and how informative their reports are (Barretto et al., 2024).

Behavioral work on networked inference helps ground ABM predictions and highlights where human learners may deviate from Bayesian idealizations. Studies of distributed statistical inference in human social networks show that different communication media can affect aggregation, convergence, and accuracy across interaction formats (Zubak et al., 2024). Other behavioral evidence is consistent with naive information aggregation in human social learning, including limited adjustment for dependencies created by repeated social transmission; this can produce systematic errors when information is recycled (Fränken et al., 2024). Work on dependency problems in testimony and evidence evaluation shows why repeated or non-independent reports can be difficult to interpret (Pilditch et al., 2025). This issue is directly relevant to the present model because agents communicate current credences rather than source-indexed evidence, so the model does not track information provenance. Targeted social learning studies further suggest that people can use observable behavior to infer others' competence when rewards are not observable, guiding selective reliance on social information (Hawkins et al., 2023).

Finally, work in multi-agent systems formalizes communication as an optimization problem. In CAROC, agents decide whether to communicate based on peer credibility and predicted reward impact, explicitly trading off expected gains against communication overhead (Lin & Zeng, 2025). This provides a useful comparison point between fixed message-budget policies and learned, state-dependent communication policies. The present paper takes a simpler approach: it does not learn a state-dependent communication policy, and it does not identify the isolated effect of message penalty. Instead, it compares two fixed communication regimes that jointly vary message penalty, send probability, message cap, and social-update weight.

Methods

Our experimental framework investigates how bundled communication regimes and signal reliability jointly shape group-level inference. We use an agent-based model (ABM) to simulate a networked decision task where agents must balance the utility of social information against the explicit costs of acquiring it.

Task and Network Structure

Groups consist of agents arranged in a ring network, where communication is restricted to immediate neighbors. Every

Parameter	Value
Group size N	8
Rounds T	10
Network	Ring (degree 2)
Reward	100 for correct final guess
Per-broadcast penalty	1 (Low) or 5 (High)
Signal strength κ	0.2 (Weak) or 0.4 (Strong)
Noise scale σ	1.0 (fixed)
Update weight α	0.55 (Low-cost) / 0.35 (High-cost)
Replicates	40 groups per condition cell

Table 1: Core experimental parameters.

Policy parameter	Low-cost	High-cost
Per-round broadcast cap per agent (C)	4	1
Broadcast probability per decision opportunity (π)	0.25	0.10
Per-broadcast penalty	1	5

Table 2: Communication policy parameters tied to the two regimes.

agent has degree two, the topology controls degree variation and removes hub effects, making it easier to interpret how the communication policy changes information flow and convergence. The objective is to infer a binary latent state $S \in \{0, 1\}$ over $T = 10$ rounds. At the start of each trial, every agent receives a private signal p^{priv}, the informativeness of which is determined by signal strength κ and noise scale σ. Within each round, agents act repeatedly at fixed time intervals, updating their internal beliefs as social messages arrive and, subject to the regime policy, may initiate outgoing broadcasts of their current belief. Each broadcast is delivered to both immediate neighbors and is counted once. After the final round, agents commit to a binary guess $\hat{S}$, and an agent-level utility is calculated as the reward for correctness minus the cumulative cost of outgoing broadcasts.

Private signal generation. We generate the private signal in log-odds space and map it to a probability. Let $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$ be i.i.d. noise for agent i. Define private evidence

$$z_i^{\text{priv}} = (2S - 1)\kappa + \epsilon_i,$$

and set

$$p_i^{\text{priv}} = \text{logistic}(z_i^{\text{priv}}) = \frac{1}{1 + \exp(-z_i^{\text{priv}})}.$$

Agents initialize their belief with $p_{i,0} = p_i^{\text{priv}}$ (equivalently $z_{i,0} = z_i^{\text{priv}}$).

Agent Strategy and Communication Policy

Agents maintain beliefs in **log-odds space**, defined as

$$z = \text{logit}(p) = \log\left(\frac{p}{1 - p}\right).$$

Messages are truthful reports of the sender's current belief probability. Thus, when agent i sends a message, it reports

$$q_i = p_i,$$

with no additional message noise.

When an agent receives a reported probability q, it updates its current log-odds belief by a convex combination:

$$z \leftarrow (1 - \alpha)z + \alpha \,\text{logit}(q).$$

This rule is not intended as full Bayesian conditionalization, because the model does not track how each sender's report depends on earlier messages. Instead, incoming credence reports are treated with a bounded social-updating rule. The parameter α controls reliance on social input: $\alpha = 0$ ignores the message, $\alpha = 1$ adopts the sender's reported log-odds, and intermediate values represent partial reliance. We use log-odds space because private evidence is generated additively there and the task involves a binary state.

As part of the regime manipulation, the high-cost condition uses a lower social-updating weight α than the low-cost condition. This encodes weaker reliance on incoming social reports under the communication-constrained regime. Because α is manipulated together with the message penalty, send probability, and message cap, the resulting differences should be interpreted as regime-level differences rather than as the isolated effect of α or cost alone.

Belief logging and broadcast rule. The implemented simulation uses rounds of fixed duration (25 seconds). Within each round, agents are checked at fixed 2-second intervals, yielding repeated decision opportunities before the round ends. At each decision opportunity, the agent's current belief is first written to the log. This logged belief is used only for analysis: it is not shown to neighbors, is not counted as a message, and does not incur any communication cost.

After this logging step, the agent may send at most one outgoing broadcast of its current belief. A broadcast is sent only if two conditions are both satisfied: (i) the agent has not yet reached the regime-specific per-round broadcast cap C, and (ii) an independent Bernoulli draw with probability π succeeds. The parameters (C, π) are regime-dependent and are reported in Table 2. Each outgoing broadcast incurs one cost equal to the regime-dependent penalty (1 or 5), and total messaging cost equals this penalty multiplied by the number of outgoing broadcasts sent.

Let M_r denote the number of decision opportunities in round r, and let $B_{ir} \sim \text{Binomial}(M_r, \pi)$ denote the number of successful broadcast opportunities for agent i in that round before applying the cap. The number of outgoing broadcasts in that round is

$$m_{ir} = \min(B_{ir}, C).$$

A single outgoing broadcast is delivered to both immediate neighbors on the ring. However, it is counted once in the message-volume outcome and charged once in the utility calculation; it is not counted separately for each recipient.

Therefore, the maximum number of counted outgoing broadcasts per agent is TC: 40 in the low-cost regime and 10 in the high-cost regime.

Within-round update order. Belief updating is event-driven within rounds. When a message arrives, the receiver immediately applies the update above using the reported belief in that message. If multiple neighbor messages arrive during the same round, the receiver updates sequentially in message-arrival order. Thus, the implemented model does not use a single synchronous end-of-round aggregation step, and it does not average multiple neighbor messages into one combined social signal before updating.

Messaging is governed by a **cost-sensitive communication regime**. In the low-cost regime, agents operate with a higher per-round message cap, a higher send probability, and a larger social-update weight. In the high-cost regime, these quantities are reduced, producing sparser information flow. Because the regime jointly changes the per-broadcast penalty and the policy parameters (α, C, and π), the comparisons in this paper should be interpreted as comparisons between two communication regimes, not as a clean causal effect of message penalty alone.

Design and Analysis

We employ a 2×2 **factorial design** crossing communication regime with signal strength, fixing $\sigma = 1.0$. We run 40 independent, balanced replicates for each condition. Our analysis treats the group as the primary unit of replication, calculating group-level means for accuracy, outgoing broadcast volume, and utility. Belief convergence is measured by the decay of within-group belief variance across rounds. Reported 95% confidence intervals are computed across group-level outcomes within each condition cell, treating the group as the unit of replication. For each condition cell, the interval is

$$\bar{x} \pm t_{0.975,39} \frac{s}{\sqrt{40}},$$

where $t_{0.975,39} = 2.0227$ because each condition contains $n = 40$ groups. The same t-based rule is used for the round-level logged-belief variance summaries.

Results

The high-cost regime strongly reduces outgoing broadcasts

The communication regime has a large and consistent effect on outgoing broadcast volume (Figure 1A). Here, *messages per person* means the total number of outgoing broadcast events initiated by a participant across all 10 rounds, where each broadcast is counted once even though it is delivered to both immediate neighbors. Averaged across groups, this quantity drops from about 28.15 under the low-cost regime (`cost=` 1) to about 7.31 under the high-cost regime (`cost=` 5), an approximately 74% reduction. Group-level means (mean ± 95% CI) are:

$$\begin{aligned}
\texttt{cost} = 1,\ \kappa = 0.2 &: 28.47 \pm 0.42,\\
\texttt{cost} = 1,\ \kappa = 0.4 &: 27.83 \pm 0.46,\\
\texttt{cost} = 5,\ \kappa = 0.2 &: 7.32 \pm 0.16,\\
\texttt{cost} = 5,\ \kappa = 0.4 &: 7.30 \pm 0.13.
\end{aligned}$$

Message volume is mainly controlled by the regime definition, especially the per-round broadcast cap and broadcast probability in Table 2; changing κ has little effect under this policy.

κ improves accuracy mainly in the high-cost regime

Figure 1B shows group-level accuracy. Under the high-cost, communication-constrained regime, increasing signal strength from $\kappa = 0.2$ to $\kappa = 0.4$ yields a large accuracy gain:

$$0.6906 \pm 0.1349 \ \rightarrow\ 0.9094 \pm 0.0830.$$

Under the low-cost regime (`cost=` 1), accuracy is similar across κ within uncertainty:

$$0.7500 \pm 0.1402 \ \rightarrow\ 0.7250 \pm 0.1446.$$

Overall, the results show a regime-by-signal-strength interaction under the current bundled policy: when communication is sparse and social updating is weaker, more reliable private signals substantially improve final decisions; when communication is frequent, performance is less sensitive to κ within the observed range.

Utility tradeoff

Final score (Figure 1C) measures utility as reward minus messaging cost. Group-level score means (mean ± 95% CI) are:

$$\begin{aligned}
\texttt{cost} = 1,\ \kappa = 0.2 &: 46.53 \pm 14.01,\\
\texttt{cost} = 1,\ \kappa = 0.4 &: 44.67 \pm 14.53,\\
\texttt{cost} = 5,\ \kappa = 0.2 &: 32.45 \pm 13.42,\\
\texttt{cost} = 5,\ \kappa = 0.4 &: 54.45 \pm 8.38.
\end{aligned}$$

Within the tested costs $\{1, 5\}$, the best cost depends on signal strength: for $\kappa = 0.2$ the higher-score condition is `cost=` 1, while for $\kappa = 0.4$ the higher-score condition is `cost=` 5. This crossing pattern implies that the score-maximizing regime depends on private signal quality under the current policy.

Logged-belief dispersion over rounds

We summarize within-group logged-belief variance by round as a descriptive measure of how agents' internal beliefs become less dispersed over time. Figure 2 shows that this variance decreases across rounds in all conditions, indicating reduced dispersion in logged beliefs as interaction proceeds. The decline is faster under `cost=` 1, where more frequent communication rapidly compresses the spread of logged beliefs; under `cost=` 5, the decline is slower, consistent with sparser information flow in the high-cost regime.

Because agents' beliefs are logged multiple times within a round, this variance measure reflects both within-agent belief movement over time and between-agent disagreement. It

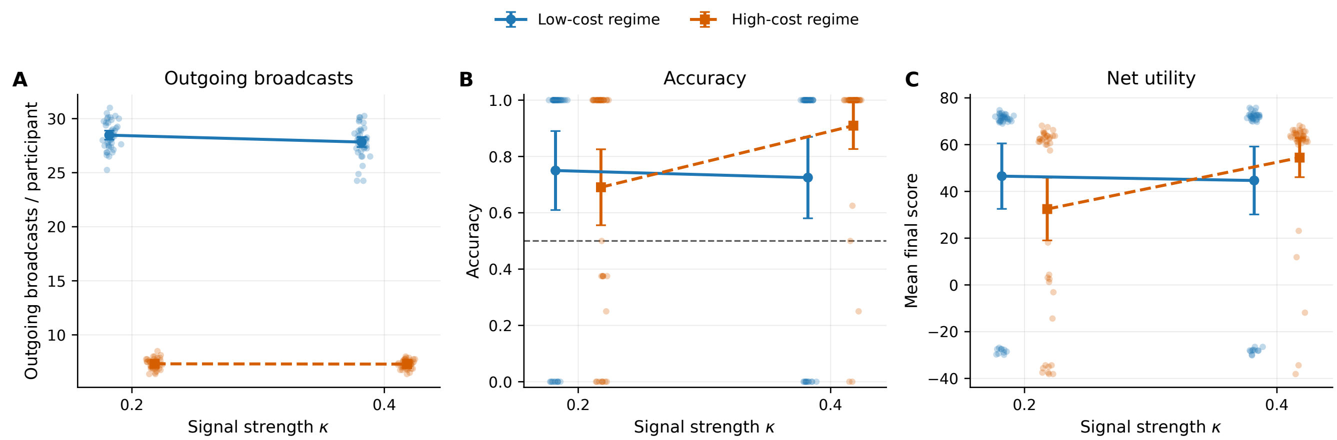

Figure 1: Main outcomes by communication regime and signal strength. Points show group means; error bars show 95% confidence intervals. Panel A shows outgoing broadcasts per participant across 10 rounds; a broadcast delivered to both neighbors is counted once. Panel B shows final accuracy. Panel C shows final score, computed as reward minus broadcast cost.

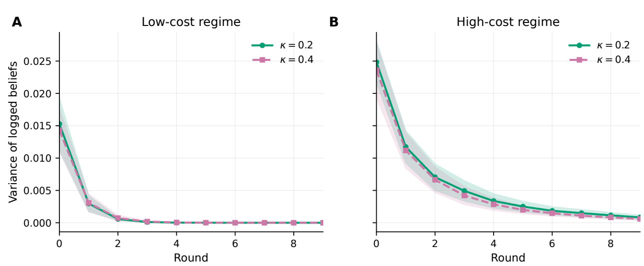

Figure 2: Logged-belief variance across rounds. Lines show condition means; shaded bands show 95% confidence intervals. Logged beliefs are analysis records and are not counted as outgoing broadcasts.

should therefore be read as a descriptive summary of logged-belief dispersion rather than as a pure measure of end-of-round consensus.

Discussion

Mechanism: the high-cost regime sparsifies information flow, and stronger private signals partly compensate

The results support a simple mechanism under the current bundled communication policy. The high-cost regime reduces outgoing-broadcast volume, which limits information flow through the ring network and slows the reduction of within-group dispersion in logged beliefs. When social information is more limited, final performance depends more strongly on private signals. In that regime, increasing κ produces a large improvement in accuracy and final score. Under the low-cost regime, outgoing broadcasts are more frequent and social updating is stronger, so performance is less sensitive to κ within the observed range.

The utility results show that there is no single best regime across all conditions. The higher-scoring regime depends on private signal quality. With weaker signals ($\kappa = 0.2$), the low-cost regime performs better on average, consistent with a benefit from more information exchange. With stronger signals ($\kappa = 0.4$), the high-cost regime performs better on average, because high accuracy can be maintained with fewer outgoing broadcasts and therefore lower communication expense under that regime.

Limitations

First, the results come from an agent-based model (ABM), not from direct human data. The update rule, broadcast policy, and truthful-reporting assumption are modeling choices and may not match human strategies. Second, the high-cost regime changes several components at once: per-broadcast penalty, broadcast probability, per-round broadcast cap, and social-update weight. The paper therefore compares two bundled communication regimes and does not identify the isolated effect of message penalty alone.

Third, we used a ring network and a fixed number of rounds. The ring topology is useful as a controlled baseline because every agent has the same degree, but other network structures may change the size or form of the effects. Fourth, agents communicate current credences rather than source-indexed evidence. The model therefore cannot distinguish independent confirmation from recycled information, and it cannot track how a report may depend on earlier reports from other agents. The discounted update rule reduces the influence of any single message, but it is not a full solution to informational dependence.

Finally, the logged-belief variance measure is descriptive. Because beliefs are logged multiple times within a round, the measure reflects both within-agent belief movement and between-agent disagreement. It should not be interpreted as a pure end-of-round consensus measure.

Predictions for human experiments

The simulation platform already logs beliefs, outgoing broadcasts, and final choices, which makes the same task design suitable for a human study with modest additional engineering. A human experiment could keep the binary inference task, ring-network interaction structure, and 2×2 factorial design crossing communication regime with signal strength, while adding standard study components such as consent, comprehension checks, waiting-room logic, dropout handling, and completion codes.

The ABM results lead to four preregistered predictions. First, moving from the low-cost regime to the high-cost regime should reduce outgoing-broadcast volume. Second, signal reliability should matter more when communication is constrained: increasing κ should produce larger gains in accuracy and final score in the high-cost regime than in the low-cost regime. Third, logged-belief dispersion should decline more slowly in the high-cost regime, especially when private signals are weak. Fourth, participants in the high-cost regime should rely more on their private signals, such that early logged beliefs track initial private signals more closely, updates from neighbors are less frequent or weaker, and final guesses are less sensitive to neighbors' messages.

These predictions can be tested with preregistered analyses using the same logged outcomes as the simulation. Human data would show whether the tradeoff patterns observed in the ABM also appear in behavior, and would further allow analysis of individual differences in broadcasting and belief updating.

Theoretical implications: communication as a constrained resource

The findings support treating communication as a constrained resource rather than as a fixed, costless channel. In the present model, changing the communication regime changes how much information moves through the network and how strongly agents respond to social input. The results show how communication constraints can reduce effective information flow, slow agreement, and increase the importance of private evidence.

At the same time, the present model should be read as a hypothesis-generating simulation framework rather than as a complete cognitive theory of human communication. The ring network, truthful credence reports, and log-odds update rule are simplifying assumptions chosen to create a clear baseline. They allow us to measure how sparse communication changes outgoing-broadcast volume, logged-belief dispersion, accuracy, and final score, but they should not be taken as a full model of human social structure or human belief revision.

Because the regimes bundle together per-broadcast penalty, broadcast probability, per-round broadcast cap, and social-update weight, the current simulations do not identify which component is mainly responsible for each effect. Future robustness analyses should separate these components, for example by varying broadcast penalty while holding the broadcast policy fixed, or by varying broadcast probability and cap while holding penalty fixed. Future work should also test whether the same qualitative patterns hold under alternative network topologies and longer interaction horizons.

Reproducibility

The complete simulation and analysis package is available at https://github.com/qqymarshtomp/Costly-Communication-Shapes-Networked-Social-Learning. The repository contains the server and bot code used to run the agent-based simulations, the frozen SQLite database snapshot used for the reported analyses, the database schema, and the analysis scripts that generate the summary tables and figures in this paper.

References

Angere, S. (2010). Knowledge in a social network. *Preprint without journal information*. https://lup.lub.lu.se/record/4393834

Angere, S., & Olsson, E. J. (2017). Publish late, publish rarely!: Network density and group performance in scientific communication. In T. Boyer-Kassem, C. Mayo-Wilson, & M. Weisberg (Eds.), *Scientific collaboration and collective knowledge: New essays* (pp. 34–62). Oxford University Press. https://doi.org/10.1093/oso/9780190680534.003.0002

Assaad, L. (2025a). Multi-option polarization: How deliberating more options both increases and decreases polarization. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*. https://escholarship.org/uc/item/5dt8g8c1

Assaad, L. (2025b). Why models of scientific communication disagree. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*. https://escholarship.org/uc/item/4184z1fz

Assaad, L., Fuchs, R., Jalalimanesh, A., Phillips, K., Schöppl, K., & Hahn, U. (2023). *A Bayesian agent-based framework for argument exchange across networks* (No. 2311.09254). arXiv. https://doi.org/10.48550/arXiv.2311.09254

Assaad, L., & Hahn, U. (2024). Rational polarization: Sharing only one's best evidence can lead to group polarization. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*. https://escholarship.org/uc/item/6x80n8v4

Barretto, D., Marjieh, R., & Griffiths, T. L. (2024). Reaching consensus through theory of mind in social networks with locally distributed interactions. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*. https://escholarship.org/uc/item/9qw9361h

Becker, J., Brackbill, D., & Centola, D. (2017). Network dynamics of social influence in the wisdom of crowds. *Proceedings of the National Academy of Sciences*, *114*(26), E5070–E5076. https://doi.org/10.1073/pnas.1615978114

DeGroot, M. H. (1974). Reaching a consensus. *Journal of the American Statistical Association*, *69*(345), 118–121. https://doi.org/10.2307/2285509

Fränken, J.-P., Valentin, S., Lucas, C. G., & Bramley, N. R. (2024). Naïve information aggregation in human social learning. *Cognition*, *242*, 105633. https://doi.org/10.1016/j.cognition.2023.105633

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278. https://doi.org/10.1126/science.aac6076

Hahn, U., Assaad, L., & Schöppl, K. (2025). Flooding the zone: An agent-based exploration. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*. https://escholarship.org/uc/item/8z32x661

Hawkins, R. D., Berdahl, A. M., Pentland, A., Tenenbaum, J. B., Goodman, N. D., & Krafft, P. M. (2023). *Flexible social inference facilitates targeted social learning when rewards are not observable* (No. 2212.00869). arXiv. https://doi.org/10.48550/arXiv.2212.00869

Klein, D., Marx, J., & Fischbach, K. (2018). Agent-based modeling in social science, history, and philosophy: An introduction. *Historical Social Research*, *43*(1), 7–27. https://doi.org/10.12759/HSR.43.2018.1.7-27

Lin, Y.-Y., & Zeng, X.-J. (2025). Enhancing multi-agent communication through credibility and reward-based optimisation. *International Journal of General Systems*, 1–25. https://doi.org/10.1080/03081079.2025.2503785

Pilditch, T. D., Hahn, U., & Lagnado, D. (2025). The problem of dependency. *Synthese*, *205*(4), 143. https://doi.org/10.1007/s11229-025-04969-w

Pilgrim, C., Morford, J., Warren, E., Aellen, M., Krupenye, C., Mann, R. P., & Biro, D. (2025). *The computational foundations of collective intelligence* (No. 2509.07999). arXiv. https://doi.org/10.48550/arXiv.2509.07999

Schöppl, K., & Hahn, U. (2024). Exploring effects of self-censoring through agent-based simulation. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*. https://escholarship.org/uc/item/9b32v6xc

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, *50*(3), 665–690. https://doi.org/10.1016/S0304-3932(03)00029-1

Stratton, J. (2025). Selected communication. *Proceedings of the 26th ACM Conference on Economics and Computation*, 838. https://doi.org/10.1145/3736252.3742634

Witt, A., Toyokawa, W., Lala, K. N., Gaissmaier, W., & Wu, C. M. (2024). Humans flexibly integrate social information despite interindividual differences in reward. *Proceedings of the National Academy of Sciences*, *121*(39), e2404928121. https://doi.org/10.1073/pnas.2404928121

Ying, L., Truong, R., Tenenbaum, J. B., & Gershman, S. J. (2025). *Adaptive social learning using theory of mind* (No. 2507.09409). arXiv. https://doi.org/10.48550/arXiv.2507.09409

Zubak, Y., Dronavalli, P., Chuang, Y.-S., & Hawkins, R. D. (2024). Distributed statistical inference in social interaction networks. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*. https://escholarship.org/uc/item/8km974x4