DELIBERATION WITH OVERLAPPING EVIDENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

How should AI agents deliberate? Multi-agent AI systems increasingly improve their answers through deliberation, yet their agents often hold overlapping knowledge from searching the same corpora, calling the same tools, or sharing training data. This paper studies how AI agents should deliberate to reach the omniscient forecast: the Bayesian forecast given everyone's pooled private data. We first prove that exchanging posteriors (the belief protocol) cannot reach the omniscient forecast. We then propose the anecdote protocol, where agents share not only their beliefs but also the clues behind them, and prove that it reaches the omniscient forecast. Finally, we compare the two protocols with Bayesian and LLM agents and find that the anecdote protocol (in its original form and many empirical variations) outperforms the belief protocol for both types of agents.

1 INTRODUCTION

When people deliberate, they seldom exchange only their beliefs: a forecaster says which reports moved her, a juror points to a piece of testimony, and a reviewer cites the lemma she could not verify. Increasingly, multi-agent AI systems have several agents discuss a question before committing to an answer Li et al. (2026); Du et al. (2023).

It seems self-evident that explanations matter: when agents exchange *why* they believe something, rather than only *what* they believe, they should perform better. Yet this intuition is hard to account for in the standard models of Bayesian communication, where beliefs are, in a sense, already enough. Agents with a common prior who take turns announcing their posteriors cannot agree to disagree, and they come to an agreement quicklyAumann (1976); Geanakoplos & Polemarchakis (1982); Aaronson (2004). Moreover, when their private signals are conditionally independent given the truth, the forecast they agree on is the *omniscient forecast*, the one they would reach by pooling everything they know Kong & Schoenebeck (2022).

In practice, the agents can have overlapping information: two analysts read the same wire report, two retrieval-augmented agents pull overlapping documents from the same corpus Lewis et al. (2020), or two language models are trained on much of the same text. Constructing the omniscient forecast requires knowing *which* pieces are shared. The following example illustrates this limitation of belief exchange and shows how sharing a single piece of clue, or *anecdote*, overcomes it.

Example 1.1. Suppose Alice and Bob want to forecast whether a particular player will win an upcoming tennis match, and each has read two of three news reports, or *clues*: Alice has read clues 1 and 2, and Bob has read clues 2 and 3. They know that they have exactly one clue in common, but neither knows which. Each clue points to YES or NO, and is correct with probability 0.8, independently conditional on the truth. If each agent's two clues point in opposite directions, then each assigns probability 50%. In that case, they agree — and they continue to agree no matter how many times they exchange beliefs, even though the omniscient forecaster with access to all clues would assign probability 20% or 80%.

Now suppose that, after an initial exchange of beliefs, Alice shares an anecdote: she shows Bob clue 1, which leans NO — and Bob replies that he has not read it. This single exchange reveals that clue 2 is the shared one, and both agents can now compute the omniscient forecast.

Our results. We view deliberation as the process by which agents resolve the overlap in their information, and we ask which communication protocols allow agents with overlapping evidence to recover the full-information forecast and how quickly they do so.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

First, we introduce the *overlapping clues model* (Section 2), in which agents observe overlapping random subsets of conditionally independent clues. Under the *belief protocol*, where agents repeatedly exchange their posteriors, their error is lower bounded at every round (Theorem 3.2), and the bound is linear in $\min\{m, k_A - m, k_B - m\}$, where m is the size of the overlap and k_A, k_B are the sizes of the agents' clue sets (Corollary 3.3) for symmetric binary clues. Under the *anecdote protocol*, where agents take turns announcing their clue, the error instead decays exponentially in the number of exchanges (Theorem 3.8 and Corollary 3.9).

Second, we test empirically which messages best facilitate deliberation (Section 4). We compare several clue-selection strategies, such as sharing the strongest clue, a random clue, or the clue most likely to be new to the receiver, with exchanging beliefs alone. We run this comparison in three settings: Bayesian agents, LLM agents with Gaussian clues, and LLM agents with textual clues about a sports match. In every setting, belief exchange stalls after the first round, while sharing the strongest clue, as in the anecdote protocol, converges to the omniscient forecast.

1.1 Related work

Our work lies at the intersection of classical social learning information aggregation and multi-agent AI systems.

Theoretical model for social learning and information aggregation. A classical literature studies how the communication network of who talks to whom shapes learning, for both naive DeGroot (1974); Golub & Jackson (2010) and Bayesian Acemoglu et al. (2011) agents. In seminal work, Aumann (1976) showed that Bayesian agents with a common prior cannot agree to disagree. This inspired a line of work showing that agents who exchange their posteriors converge to agreement Geanakoplos & Polemarchakis (1982); Aaronson (2004) and identifying sufficient conditions for accuracy Frongillo et al. (2023); Collina et al. (2024). We ask the complementary question: when exchanging beliefs fails to deliver accuracy, how can a deliberation protocol recover it?

Our overlapping-clue model contains classical information structures as special cases (Remark 2.4). Prior work in decision analysis and economics literature studies how to aggregate forecasts based on overlapping information Winkler (1981); Clemen (1987); Babichenko & Garber (2021), and how agents who neglect their overlapping sources overcount shared information DeMarzo et al. (2003); Eyster & Rabin (2010); Enke & Zimmermann (2019); Levy & Razin (2015). We ask instead if the agents can correct for the overlap through deliberation.

Empirical settings and multi-agent AI. Our question is how to help agents deliberate, and empirical work on both people and language models finds that unaided deliberation often falls short Lewis et al. (2020); Li et al. (2026)

Proposed remedies change either who deliberates with whom Du et al. (2023); Wang et al. (2024); Li et al. (2024) or what agents say. We focus on the latter. Pham et al. (2024) consider models that exchange embeddings instead of sampled tokens. Li et al. (2026) find that a structured protocol in which agents first list their unshared facts substantially improves accuracy. Whereas these studies are empirical, our results give a formal reason to focus on what agents say.

2 Overlapping clues model and deliberation

Definition 2.1 (Overlapping clues model)**.** A *overlapping clues model* $(n, \Theta, \pi, \mathcal{S}, (f_\theta)_{\theta\in\Theta}, \mathcal{D})$ has a ground state $\theta \in \Theta$ with prior π and N agents. There are n *clues* $\boldsymbol{s} = (s_1, \ldots, s_n) \in \mathcal{S}^n$, conditionally i.i.d. given θ with law f_θ. An *allocation* distribution $\mathcal{D}$, independent of $(\theta, \boldsymbol{s})$, draws subsets $A_1, \ldots, A_N \subseteq [n]$, and agent i observes her *arrangement* $\mathcal{I}_i = (A_i, \boldsymbol{s}_{A_i})$. We assume $\mathcal{D}$ is *exchangeable* so that $(\sigma A_1, \ldots, \sigma A_N) \overset{d}{=} (A_1, \ldots, A_N)$ for every permutation σ of $[n]$.

We focus on a binary state $\Theta = \{0, 1\}$ and two agents, Alice and Bob, who hold A and B with overlap $m = |A \cap B|$. Note that Example 1.1 is a special case of Definition 2.1 with $n = 3$, $N = 2$.

Write $l_0 = \log \frac{\pi}{1-\pi}$ for the prior log-odds, $l(s) = \log \frac{f_1(s)}{f_0(s)}$ for a clue's log-likelihood ratio, $l_j = l(s_j)$, and $L_S = \sum_{j\in S} l_j$ for $S \subseteq [n]$. As $\mathcal{D}$ is independent of $(\theta, \boldsymbol{s})$, the ***omniscient forecast***, the

posterior log-odds given both arrangements, is

$$L^\star = l_0 + L_{A\cup B} = l_0 + L_A + L_B - L_{A\cap B}, \tag{1}$$

where the *overlapped log-likelihood* $L_{A\cap B}$ is the clues the two agents hold in common.

Definition 2.2 (Deliberation protocol)**.** A *deliberation protocol* is a public-coin broadcast protocol between Bayesian agents, with the model common knowledge. In round 0 both agents announce their posteriors, which makes L_A and L_B public. In each round $t \geq 1$, agent i sends a message

$$r_t^i = \mu_t^i(\mathcal{I}_i, H_{t-1}) \in \mathcal{M}_t^i,$$

where $\mathcal{M}_t^i$ is the protocol's feasible message space, and $H_t = (r_0^A, r_0^B, \ldots, r_t^A, r_t^B)$ is the transcript. After round t, agent i knows $(\mathcal{I}_i, H_t)$ and forecasts $\hat{L}_t^i = \mathbb{E}[L^\star \mid \mathcal{I}_i, H_t]$.

Definition 2.3 (Error)**.** The round-t *error* of a protocol is the agents' total squared distance from the omniscient forecast,

$$R_t = \sum_i \mathbb{E}\big[(\hat{L}_t^i - L^\star)^2\big] = \sum_i \mathbb{E}\big[\mathrm{Var}(L^\star \mid \mathcal{I}_i, H_t)\big].$$

A protocol is *consistent* if $R_t \to 0$ as communication proceeds, so that every agent reaches the omniscient forecast. The choice of squared log-odds error is not essential: $R_t = 0$ exactly when every agent's forecast equals the omniscient forecast almost surely, which is also exactly when any other loss that vanishes, such as squared error, absolute error, or KL divergence.

Remark 2.4 (Relation to classical models)**.** With a deterministic allocation, the model is closely related to several classical information structures. If the clue sets are disjoint singletons, each agent holds a single conditionally i.i.d. private signal, as in the Condorcet jury theorem Condorcet (1785); Austen-Smith & Banks (1996) and models of social learning Banerjee (1992); Bikhchandani et al. (1992). For a general fixed allocation, each agent's information splits into a component $\boldsymbol{s}_{A\cap B}$ shared with the other agent and a private component, as in the public-plus-private signal structures of Morris & Shin (2002); Palley & Soll (2019). The fixed-allocation model is thus a special case of the partial evidence model of Babichenko & Garber (2021) and fits within the partial information framework of Satopää et al. (2016). With Gaussian clues and a Gaussian state, it coincides with the homogeneous-precision primitive-cue model of Amin et al. (2026).

What distinguishes our model is that the allocation is random, and to our knowledge, we are the first to study how Bayesian agents reason and communicate under the randomized allocation.

Symmetric binary clues and uniform allocation. Our results in Section 3 use the following two instances, and we relax those assumptions empirically in Section 4.

Definition 2.5 (Binary symmetric clues)**.** Clues are binary, $\mathcal{S} = \{0, 1\}$, and each agrees with the state with probability $\rho \in (\frac{1}{2}, 1)$: $f_\theta = \mathrm{Ber}(q_\theta)$ with $q_1 = \rho$ and $q_0 = 1 - \rho$, the *binary symmetric channel* $\mathrm{BSC}(\rho)$. A positive clue has log-likelihood ratio $\lambda := \log \frac{\rho}{1-\rho} > 0$ and a negative clue $-\lambda$. Writing $S_A = \sum_{j\in A} s_j$, $S_B = \sum_{j\in B} s_j$ and $S_{A\cap B} = \sum_{j\in A\cap B} s_j$ for the counts of positive clues,

$$L_A = \lambda\,(2S_A - |A|), L_{A\cap B} = \lambda\,(2S_{A\cap B} - m), \tag{2}$$

and similarly for L_B. In particular, given $|A|$ the log-likelihood L_A determines the count S_A, and $L_{A\cap B}$ is a function of $(m, S_{A\cap B})$.

Definition 2.6 (Uniform allocation)**.** Given sizes k_A, k_B, the allocation $\mathcal{D}(k_A, k_B)$ draws A and B independently and uniformly among subsets of those sizes, so the overlap m is random. Given also m with $\max(0, k_A + k_B - n) \leq m \leq \min(k_A, k_B)$, the allocation $\mathcal{D}(k_A, k_B, m)$ draws (A, B) uniformly among pairs with these sizes and $|A \cap B| = m$.

3 Deliberation protocols, beliefs and anecdotes

We study two protocols that differ in what each agent reports: her posterior (the *belief* protocol, Section 3.1), or one of her clues (the *anecdote* protocol, Section 3.2). The first is the classical belief-exchange dynamic Aaronson (2004). The second capture the idea of explanation.

3.1 THE BELIEF PROTOCOL

Definition 3.1 (Belief protocol)**.** In every round $t \geq 1$, $\mathcal{M}_t^i = [0,1]$ and each agent reports her current posterior,

$$r_t^i = \Pr\left[\theta{=}1 \mid \mathcal{I}_i, H_{t-1}\right].$$

The following theorem shows that the belief protocol cannot reach the omniscient forecast.

Theorem 3.2 (Belief floor)**.** *Under* $\mathrm{BSC}(\rho)$ *clues (Definition 2.5), for every* $t \geq 0$*,*

$$R_t \geq R_\star := 2\,\mathbb{E}\big[\mathrm{Var}(L_{A\cap B} \mid S_A, |A|, S_B, |B|)\big].$$

If the sizes $|A|, |B|$ *are common knowledge, as under both allocations of Definition 2.6, then* $R_t = R_\star$ *for every* $t \geq 0$*.*

Corollary 3.3 (Size of the floor)**.** *Under* $\mathrm{BSC}(\rho)$ *clues and the allocation* $\mathcal{D}(k_A, k_B, m)$*, for every* $t \geq 0$*,*

$$R_t \geq R_\star \geq 8\lambda^2\rho^2(1-\rho)^2 \min\{m,\, k_A - m,\, k_B - m\}.$$

Under $\mathcal{D}(k_A, k_B)$*, where the overlap* $M = |A \cap B|$ *is random, the same holds with* $\min\{m,\, k_A - m,\, k_B - m\}$ *replaced by* $\mathbb{E}\big[\min\{M,\, k_A - M,\, k_B - M\}\big]$*.*

The lower bound of error is zero if and only if agent's clues are disjoint $A \cap B = \emptyset$, or nested—one agent knows the other agent's clues $A \subseteq B$ or $B \subseteq A$. Otherwise, belief exchange cannot converge to omniscient prediction regardless of the number of rounds. Intuitively, belief exchange does well when the overlap is typically empty or nested. It does poorly when each sees a fraction bounded away from 0 and 1.

Remark 3.4 (Index-free protocols)**.** The proof of the lower bound in Theorem 3.2 uses only that no message depends on which indices the sender holds. Call a protocol *index-free* if $\mu_t^i(\sigma\mathcal{I}_i, H_{t-1}, R) = \mu_t^i(\mathcal{I}_i, H_{t-1}, R)$ for every permutation σ of $[n]$, where $\sigma\mathcal{I}_i = (\sigma A_i, (\sigma \boldsymbol{s})_{\sigma A_i})$. The argument of Theorem 3.2 then gives $R_t \geq R_\star$ for every index-free protocol and every $t \geq 0$. Examples include reporting full predictive distributions or confidence levels, and second-order reports such as a prediction of the other agent's forecast Palley & Soll (2019).

Lemma 3.5 (Error reduces to the overlap)**.** *For all* $t \geq 0$ *and* $i \in \{A, B\}$, $\mathrm{Var}(L^\star \mid \mathcal{I}_i, H_t) = \mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_i, H_t)$*, and hence*

$$R_t = \mathbb{E}\big[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_A, H_t)\big] + \mathbb{E}\big[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_B, H_t)\big].$$

The second shows that Alice's arrangement matters only through her size and count: any two arrangements with the same size and the same number of positive clues are equally likely and say the same thing about everything else.

Lemma 3.6 (Relabelling irrelevance)**.** *Under* $\mathrm{BSC}(\rho)$ *clues, let* $\mathcal{I}_A = (A, \boldsymbol{s}_A)$ *and* $\mathcal{I}'_A = (A', \boldsymbol{s}'_{A'})$ *be two arrangements of positive probability with* $|A| = |A'|$ *and* $S_A = S_{A'}$*. Then* $\Pr[\mathcal{I}_A] = \Pr[\mathcal{I}'_A]$ *and*

$$(\theta, S_B, |B|, m, S_{A\cap B}) \mid \mathcal{I}_A \overset{d}{=} (\theta, S_B, |B|, m, S_{A\cap B}) \mid \mathcal{I}'_A. \tag{3}$$

Equivalently, the conditional law of $(\theta, S_B, |B|, m, S_{A\cap B})$ *given* $\mathcal{I}_A$ *depends on* $\mathcal{I}_A$ *only through* $(S_A, |A|)$*. Moreover, for all functions* g *and* h *of* $(\theta, S_A, |A|, S_B, |B|, m, S_{A\cap B})$*,*

$$\mathbb{E}[h \mid \mathcal{I}_A, g] = \mathbb{E}[h \mid S_A, |A|, g]. \tag{4}$$

The same statements hold with the roles of A *and* B *exchanged.*

3.2 THE ANECDOTE PROTOCOL

Given this limitation of the belief protocol, we ask how how can clue exchange facilitate deliberation. We study anecdote protocols which ask each agent to announce her single clue directly.

Definition 3.7 (Anecdote protocol)**.** Agents alternate as sender, starting with Alice. In round t, the sender i reports her current posterior together with one of her clues not yet announced by either agent,

$$r_t^i = \Big(\Pr\left[\theta{=}1 \mid \mathcal{I}_i, H_{t-1}\right],\, (j_t, s_{j_t})\Big),$$

and the receiver replies with the membership bit $\mathbf{1}\{j_t \in \text{her set}\}$. The protocol stops once either agent has announced all of her clues.

In this section the sender draws j_t uniformly at random from her unannounced clues, using the public randomness. Section 4 compares other selection rules empirically.

Write J_t for the set of clues announced in the first t exchanges, Φ_t for the public record of their indices, values and membership bits, and $m^{(t)} = |(A \cap B) \setminus J_t|$ for the number of shared clues not yet announced. The following theorem shows that the error is controlled by $m^{(t)}$.

Theorem 3.8 (Anecdote error)**.** *Under* $\mathrm{BSC}(\rho)$ *clues, the anecdote protocol is consistent and for every* $t \geq 0$*,*

$$R_t \leq 8\lambda^2\rho(1-\rho)\,\mathbb{E}\big[m^{(t)}\big] + \lambda^2\Big(\mathbb{E}\big[\mathrm{Var}(m^{(t)} \mid A, \Phi_t)\big] + \mathbb{E}\big[\mathrm{Var}(m^{(t)} \mid B, \Phi_t)\big]\Big).$$

The first term is the variance of the clue still hidden in unannounced shared clues, and the second is each agent's uncertainty about how many such clues remain. The following corollary shows that when the sets are bounded, each exchange announces a shared clue with probability at least $m^{(t)}/k$, so the unannounced overlap, and with it the error, shrinks geometrically.

Corollary 3.9 (Exponential decay)**.** *Under* $\mathrm{BSC}(\rho)$ *clues and any allocation with* $|A|, |B| \leq k$*, for every* $t \geq 0$*,*

$$R_t \leq \lambda^2\Big(8\rho(1-\rho)\,\mathbb{E}[m] + 2\,\mathbb{E}[m^2]\Big)\Big(1 - \frac{1}{k}\Big)^t.$$

4 Experiments

We test how clue-selection rules affect deliberation, moving from algorithmic Bayesian agents to LLM agents with numerical and textual clues. We compare anecdote protocols with belief-only communication, first using Gaussian signals and then realistic textual evidence, with model sensitivity reported in Appendix F.

4.1 Experimental Algorithms

We use two agents, a uniform binary state $\theta \in \{0, 1\}$, and Gaussian clues

$$s_j \mid \theta = 0 \sim \mathcal{N}(-1, \sigma_j^2), \qquad s_j \mid \theta = 1 \sim \mathcal{N}(1, \sigma_j^2),$$

where $\sigma_j \in \{1.5, 2.0, 2.5\}$. Clues follow $D(12, 12, 6)$, so each agent observes 12 clues with six shared.

We run deliberation protocols in Definitions 3.1 and 3.7. For the algorithmic Gaussian experiment, we average across 250 independently generated Monte Carlo runs. The LLM experiments use separate repetition counts specified in Section 4.3.

For anecdote protocol, at each round, the sender chooses among its unannounced clues according to some selection rule. Here we compare random, strongest, positive, negative, impactful, attractive, repelling, unique, and residual-gap. The mathematical definitions and motivations for these policies are given below.

Let p_S and p_R denote the current posterior beliefs of the sender and receiver, respectively. For a candidate clue j, let ℓ_j denote its log-likelihood ratio and let

$$p_R^{(j)} = \sigma(\mathrm{logit}(p_R) + \ell_j)\,,$$

where $\sigma(x) = 1/(1 + e^{-x})$.

We consider two family of selection rules. The first is oblivious family which does not depend on the receiver and can be decided at the beginning.

1. **Random.** The sender selects an unresolved clue uniformly at random. This serves as a non-strategic baseline against which the other communication rules can be compared.
2. **Strongest.** The sender selects the clue with the largest absolute log-likelihood ratio, $j^* = \arg\max_j |\ell_j|$. This policy prioritizes the individually most informative clue, regardless of the receiver's current belief.

3. **Positive.** The sender selects $j^* = \arg\max_j \ell_j$. This policy deliberately communicates the clue providing the strongest clue in favor of state 1.

4. **Negative.** The sender selects $j^* = \arg\min_j \ell_j$. This policy deliberately communicates the clue providing the strongest clue in favor of state 0.

The second one depends on the receiver.

1. **Impactful.** The sender selects the clue that produces the largest absolute change in the receiver's posterior: $j^* = \arg\max_j |p_R^{(j)} - p_R|$. Unlike the strongest policy, this rule is receiver-aware and accounts for the receiver's current belief.

2. **Attractive.** The sender selects the clue that moves the receiver closest to the sender's current posterior: $j^* = \arg\min_j |p_R^{(j)} - p_S|$. This policy is designed to reduce belief disagreement as efficiently as possible.

3. **Repelling.** The sender selects $j^* = \arg\max_j |p_R^{(j)} - p_S|$. This policy intentionally maximizes disagreement and is included as a negative-control communication rule.

4. **Unique.** Let K_R denote the set of clues already known by the receiver. The sender prioritizes clues satisfying $j \notin K_R$. When several such clues are available, one is selected according to the experiment's tie-breaking rule. This policy directly targets redundant communication caused by overlapping clues.

5. **Residual-gap.** Define the current sender-receiver log-odds gap by $G = \text{logit}(p_S) - \text{logit}(p_R)$. The sender selects $j^* = \arg\min_j |G - \ell_j|$. Because log-likelihood ratios add directly in log-odds space, this policy selects the single clue whose evidential contribution most closely closes the current belief gap.

Finally, **belief-only** implements Definition 3.1. The sender reports only its current posterior belief. This policy provides a comparison against anecdote protocols and isolates the value of communicating the underlying clues themselves.

4.2 Bayesian agents

For the Bayesian agents, we perform Bayesian updating conditional on the clue values that have been observed, without higher-order inference from the endogenous choice of which clue the sender selected. This helps us isolate the effect of the clue-selection rule on the transmission of underlying evidence, rather than this additional signaling channel.

We repeat the same algorithmic communication experiment using continuous Gaussian signals. Each signal has an exactly computable log-likelihood ratio under the two latent states, so both posterior updating and communication decisions remain fully mathematical. Signal realizations, noise levels, and clue allocations are independently redrawn across Monte Carlo runs.

4.3 LLM agents

For LLM agents, we use prompts in Table 1 for clue selection and conduct Bayesian posterior updates. The LLM observes the relevant beliefs and available clues and selects according to the prescribed strategy. More details of prompts, inference details, seeds, and model-sensitivity experiments are in Appendix F.

Strategy	LLM communication instruction
Random	Select a random available clue.
Strongest	Send the strongest available clue.
Positive	Send the clue providing the strongest clue toward state 1.
Negative	Send the clue providing the strongest clue toward state 0.
Impactful	Send the clue expected to cause the largest change in the receiver's belief.
Unique	Send the clue most likely to contain information the receiver does not already possess.
Attractive	Send the clue expected to move the receiver closest to the sender's current belief.
Repelling	Send the clue expected to move the receiver furthest from the sender's current belief.
Residual-gap	Send the clue whose clue most closely fills the current belief gap between sender and receiver.
Belief-only	Communicate the current posterior without communicating any clue.

Table 1: Correspondence between the algorithmic communication strategies and the LLM prompts.

Figure 1(c,d) show the results of applying the LLM communication policies to the same Gaussian signal environment used for the Bayesian-agent experiments.

4.3.1 Synthetic match clues

We next evaluate the protocols on realistic textual clues based on an anonymized tennis match. The scenario contains twelve heterogeneous pre-match facts covering surface performance, serve and return statistics, head-to-head results, opponent quality, recent form, fatigue, and surface-adjusted rating. Alice and Bob receive partially overlapping subsets, and the LLM selects which fact to communicate under the same objectives as in the numerical experiments. The allocation is asymmetric: Alice observes six facts and Bob eight, with two shared. Thus Bob is missing four facts unique to Alice and Alice six unique to Bob, so they require different numbers of rounds to recover the full information set.

To obtain numerical forecasts, the LLM classifies each fact into one of seven directional-impact categories: strong, moderate, or slight evidence for Player A; neutral; or slight, moderate, or strong evidence for Player B. These categories are mapped to fixed log-odds increments, and Python computes the posteriors. The full-information benchmark applies the same procedure to the union of all facts. Appendix D reports the scenario, allocation, and scoring rule. We repeat the experiment 30 times with randomized presentation order.

5 Conclusion

We show that when agents hold overlapping evidence, exchanging beliefs alone can fail to recover the omniscient forecast, while communicating the underlying clues resolves this limitation. Our experiments with Bayesian and LLM agents show that clue-sharing strategies substantially reduce both distance to the full-information forecast and disagreement across agents. These results suggest that, in multi-agent AI systems with overlapping information, what agents communicate can matter as much as how they aggregate beliefs.

AI Use Statement

Generative AI tools were used to assist with experimental and methodological design, research-code implementation and debugging, interpretation and presentation of results, and drafting and editing of the manuscript. All AI-assisted material was reviewed by the authors, who take full responsibility for the submitted work.

Reproducibility Statement

Section 4 specifies the experimental environments, algorithms, and evaluation metrics. The appendix provides the exact LLM checkpoints, prompts, inference procedure, random seeds, tie-

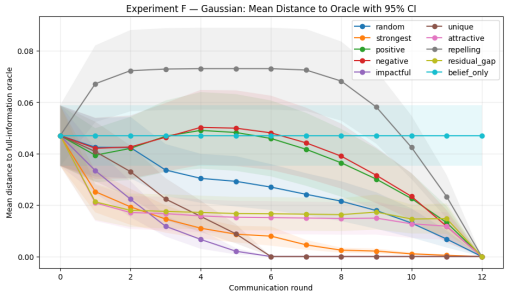

(a) Algorithmic, Gaussian: distance to oracle

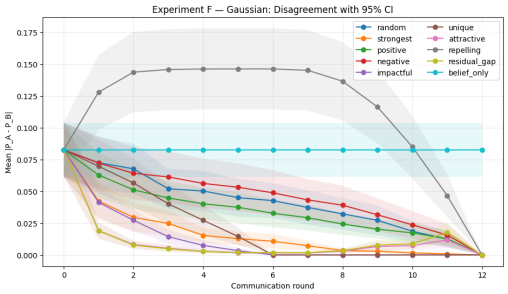

(b) Algorithmic, Gaussian: disagreement

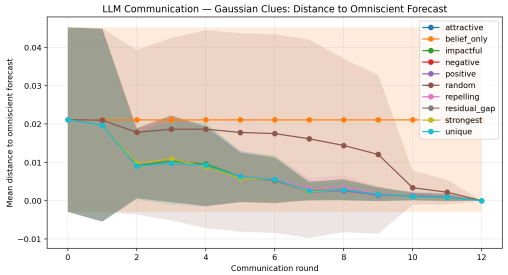

(c) LLM, Gaussian: distance to oracle

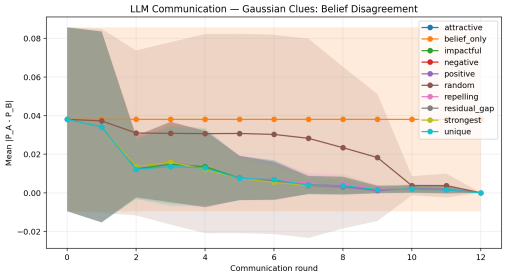

(d) LLM, Gaussian: disagreement

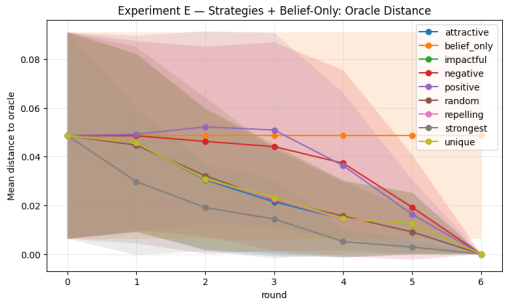

(e) LLM, realistic match facts: distance to oracle

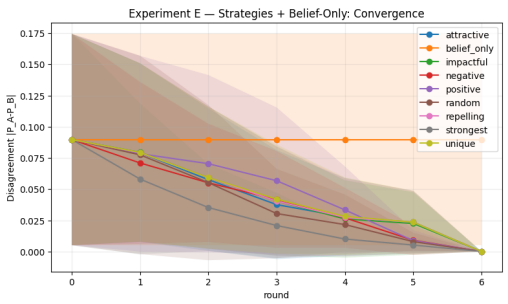

(f) LLM, realistic match facts: disagreement

Figure 1: Communication performance across the four experimental settings. Left panels show mean distance to the full-information (omniscient) forecast; right panels show mean disagreement between Alice and Bob. Shaded bands are 95% confidence intervals for the mean.

breaking rules, and confidence-interval calculation, together with the complete realistic-match scenario. Anonymous supplementary code and notebooks are provided to reproduce the reported experiments.

REFERENCES

Scott Aaronson. The complexity of agreement. *arXiv preprint cs/0406061*, 2004.

Daron Acemoglu, Munther A. Dahleh, Ilan Lobel, and Asuman Ozdaglar. Bayesian learning in social networks. *The Review of Economic Studies*, 78(4):1201–1236, 2011. doi: 10.1093/restud/rdr004.

Saurabh Amin, Amine Bennouna, Daniel Huttenlocher, Dingwen Kong, Liang Lyu, and Asuman Ozdaglar. A bayesian framework for human-ai collaboration: Complementarity and correlation neglect, 2026. URL https://arxiv.org/abs/2602.14331.

Robert J. Aumann. Agreeing to disagree. *The Annals of Statistics*, 4(6):1236–1239, 1976. doi: 10.1214/aos/1176343654.

David Austen-Smith and Jeffrey S. Banks. Information aggregation, rationality, and the condorcet jury theorem. *American Political Science Review*, 90(1):34–45, 1996. doi: 10.2307/2082796.

Yakov Babichenko and Dan Garber. Learning optimal forecast aggregation in partial evidence environments. *Mathematics of Operations Research*, 46(2):628–641, 2021. doi: 10.1287/moor.2020.1080.

Abhijit V. Banerjee. A simple model of herd behavior. *The Quarterly Journal of Economics*, 107 (3):797–817, 1992. doi: 10.2307/2118364.

Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992. doi: 10.1086/261849.

Robert T. Clemen. Combining overlapping information. *Management Science*, 33(3):373–380, 1987. doi: 10.1287/mnsc.33.3.373.

Natalie Collina, Surbhi Goel, Varun Gupta, and Aaron Roth. Tractable agreement protocols, 2024. URL https://arxiv.org/abs/2411.19791.

Nicolas de Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. De l'Imprimerie Royale, Paris, 1785.

Morris H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69 (345):118–121, 1974. doi: 10.1080/01621459.1974.10480137.

Peter M. DeMarzo, Dimitri Vayanos, and Jeffrey Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *The Quarterly Journal of Economics*, 118(3):909–968, 2003. doi: 10.1162/00335530360698469.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325*, 2023.

Benjamin Enke and Florian Zimmermann. Correlation neglect in belief formation. *The Review of Economic Studies*, 86(1):313–332, 2019. doi: 10.1093/restud/rdx081.

Erik Eyster and Matthew Rabin. Naive herding in rich-information settings. *American Economic Journal: Microeconomics*, 2(4):221–243, 2010. doi: 10.1257/mic.2.4.221.

Rafael Frongillo, Eric Neyman, and Bo Waggoner. Agreement implies accuracy for substitutable signals. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pp. 702–733, 2023.

John D. Geanakoplos and Heraklis M. Polemarchakis. We can't disagree forever. *Journal of Economic Theory*, 28(1):192–200, 1982. doi: 10.1016/0022-0531(82)90099-0.

Benjamin Golub and Matthew O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010. doi: 10.1257/mic.2.1.112.

Yuqing Kong and Grant Schoenebeck. False consensus, information theory, and prediction markets. *arXiv preprint arXiv:2206.02993*, 2022.

Gilat Levy and Ronny Razin. Correlation neglect, voting behavior, and information aggregation. *American Economic Review*, 105(4):1634–1645, 2015. doi: 10.1257/aer.20140134.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33: 9459–9474, 2020.

Yunxuan Li, Yibing Du, Jiageng Zhang, Le Hou, Peter Grabowski, Yeqing Li, and Eugene Ie. Improving multi-agent debate with sparse communication topology. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 7281–7294, 2024.

Yuxuan Li, Aoi Naito, and Hirokazu Shirado. Systematic failures in collective reasoning under distributed information in multi-agent llms, 2026. URL https://arxiv.org/abs/2505.11556.

Stephen Morris and Hyun Song Shin. Social value of public information. *american economic review*, 92(5):1521–1534, 2002.

Asa B Palley and Jack B Soll. Extracting the wisdom of crowds when information is shared. *Management Science*, 65(5):2291–2309, 2019.

Chau Pham, Boyi Liu, Yingxiang Yang, Zhengyu Chen, Tianyi Liu, Jianbo Yuan, Bryan Plummer, Zhaoran Wang, and Hongxia Yang. Let models speak ciphers: Multiagent debate through embeddings. In *International Conference on Learning Representations*, volume 2024, pp. 51899–51928, 2024.

Ville A Satopää, Robin Pemantle, and Lyle H Ungar. Modeling probability forecasts via information diversity. *Journal of the American Statistical Association*, 111(516):1623–1633, 2016.

Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. Mixture-of-agents enhances large language model capabilities, 2024. URL https://arxiv.org/abs/2406.04692.

Robert L. Winkler. Combining probability distributions from dependent information sources. *Management Science*, 27(4):479–488, 1981. doi: 10.1287/mnsc.27.4.479.

A Proofs and Details in Section 3.1

A.1 Proofs for Theorem 3.2

Proof of Theorem 3.2. We claim by induction that every report is a function of $(S_A, |A|, S_B, |B|)$. In round 0, Alice's report is determined by $\mathbb{E}[\theta \mid \mathcal{I}_A] = \mathbb{E}[\theta \mid S_A, |A|]$, by Eq. (4) with g constant and $h = \theta$, which is a function of $(S_A, |A|)$. If H_{t-1} is a function of $(S_A, |A|, S_B, |B|)$, then Alice's round-t report is $\mathbb{E}[\theta \mid \mathcal{I}_A, H_{t-1}] = \mathbb{E}[\theta \mid S_A, |A|, H_{t-1}]$ by Eq. (4) with $g = H_{t-1}$, again such a function. Bob's case holds by symmetry.

Consequently H_t is a function of $(S_A, |A|, S_B, |B|)$, and since $\mathcal{I}_A$ determines $(S_A, |A|)$, the pair $(\mathcal{I}_A, H_t)$ is determined by $(\mathcal{I}_A, S_B, |B|)$. Conditioning on less information increases the expected conditional variance, so

$$\mathbb{E}[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_A, H_t)] \geq \mathbb{E}[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_A, S_B, |B|)] = \mathbb{E}[\mathrm{Var}(L_{A\cap B} \mid S_A, |A|, S_B, |B|)], \tag{5}$$

where the equality follows from Eq. (4) applied to the first two conditional moments of $L_{A\cap B}$, a function of $(m, S_{A\cap B})$ by Eq. (2). The same bound holds for Bob by symmetry, so Lemma 3.5 gives $R_t \geq R_\star$.

Finally, if $|A|$ and $|B|$ are public, then by **??** and $l_+ \neq l_-$ each agent's round-0 report determines her count, so $(S_A, |A|, S_B, |B|)$ is in turn determined by H_t. Hence $(\mathcal{I}_A, H_t)$ and $(\mathcal{I}_A, S_B, |B|)$ determine each other, Eq. (5) holds with equality, and symmetrically for Bob, so $R_t = R_\star$. □

Proof of Lemma 3.5. Because both L_A and L_B are public after round 0 and l_0 is common knowledge, $l_0 + L_A + L_B$ is constant given $(\mathcal{I}_i, H_t)$, so by Eq. (1)

$$\mathrm{Var}(L^\star \mid \mathcal{I}_i, H_t) = \mathrm{Var}(l_0 + L_A + L_B - L_{A\cap B} \mid \mathcal{I}_i, H_t) = \mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_i, H_t).$$

□

Proof of Lemma 3.6. To prove Eq. (3) we construct a relabelling σ under which the two conditioning events correspond. Given $\mathcal{I}_A = (A, \boldsymbol{s}_A)$ and $\mathcal{I}'_A = (A', \boldsymbol{s}'_{A'})$ as in the statement, the patterns $\boldsymbol{s}_A$ and $\boldsymbol{s}'_{A'}$ are $0/1$ vectors of the same length $|A| = |A'|$ carrying the same number of ones $S_A = S_{A'}$. Hence we can pick a permutation $\sigma : [n] \to [n]$ restricting to a bijection between A and A' that sends ones to ones and zeros to zeros, so that $\sigma A = A'$ and $\sigma \boldsymbol{s}_A = \boldsymbol{s}'_{A'}$. Let it act on outcomes $\omega = (\theta, \boldsymbol{s}, A, B)$ by

$$\sigma(\omega) = \sigma(\theta, \boldsymbol{s}, A, B) = (\theta, \sigma\boldsymbol{s}, \sigma A, \sigma B),$$

where $(\sigma\boldsymbol{s})_{\sigma(j)} = s_j$.

The above σ has two properties. First, it preserves probability. Given θ, as $\boldsymbol{s}$ are i.i.d., $(\theta, \sigma\boldsymbol{s}) \overset{d}{=} (\theta, \boldsymbol{s})$. Because $\mathcal{D}$ is exchangeable, $(\sigma A, \sigma B) \overset{d}{=} (A, B)$. As the two pairs are independent, for every outcome ω,

$$\Pr[\omega] = \Pr[\sigma\omega]. \tag{6}$$

Second, σ restricts to a bijection

$$\sigma : \{\theta, S_B, |B|, m, S_{A\cap B}, \mathcal{I}_A\} \xrightarrow{\sim} \{\theta, S_B, |B|, m, S_{A\cap B}, \mathcal{I}'_A\} \tag{7}$$

for any fixed value of $(\theta, S_B, |B|, m, S_{A\cap B})$. Relabelling does not affect sums and cardinality: $\sum_{j\in\sigma B}(\sigma\boldsymbol{s})_j = \sum_{j\in B} s_j$ and $|\sigma B| = |B|$, while $\sigma A \cap \sigma B = \sigma(A \cap B)$. Hence θ, S_B, $|B|$, m and $S_{A\cap B}$ are unchanged, and together with $\sigma A = A'$ and $\sigma\boldsymbol{s}_A = \boldsymbol{s}'_{A'}$ this gives Eq. (7).

Summing Eq. (6) over the events in Eq. (7) yields

$$\Pr[\theta, S_B, |B|, m, S_{A\cap B}, \mathcal{I}_A] = \Pr[\theta, S_B, |B|, m, S_{A\cap B}, \mathcal{I}'_A]. \tag{8}$$

Summing Eq. (8) over all values of $(\theta, S_B, |B|, m, S_{A\cap B})$ implies $\Pr[\mathcal{I}_A] = \Pr[\mathcal{I}'_A]$, which is positive by hypothesis. Dividing Eq. (8) by this common marginal gives Eq. (3).

For Eq. (4), fix a value (s, k) of $(S_A, |A|)$ and let $\mathcal{I}_A, \mathcal{I}'_A$ be any two arrangements in the event $\{(S_A, |A|) = (s, k)\}$. On this event g and h are functions of $(\theta, S_B, |B|, m, S_{A\cap B})$ alone, so by Eq. (3) the joint conditional law of (h, g) is the same given $\mathcal{I}_A$ as given $\mathcal{I}'_A$, and hence

$$\mathbb{E}[h \mid \mathcal{I}_A, g] = \mathbb{E}[h \mid \mathcal{I}'_A, g] = \mathbb{E}[h \mid S_A, |A|, g].$$

The last equality holds by the tower property: $\mathbb{E}[h \mid \mathcal{I}_A, g]$ is determined by $(S_A, |A|, g)$, which is in turn determined by $(\mathcal{I}_A, g)$. The $A \leftrightarrow B$ version follows by symmetry of the hypotheses. □

A.2 Proof of Corollary 3.3

Proof of Corollary 3.3. By Theorem 3.2 it suffices to lower bound $R_\star$. Under $\mathcal{D}(k_A, k_B, m)$ the sizes and the overlap are known, so by Eq. (2) with $l_+ = -l_- = \lambda$ we have $L_{A\cap B} = \lambda(2S_{A\cap B} - m)$ and

$$R_\star = 2\mathbb{E}\big(\mathrm{Var}(L_{A\cap B} \mid S_A, S_B)\big) = 8\lambda^2\mathbb{E}\big(\mathrm{Var}(S_{A\cap B} \mid S_A, S_B)\big).$$

We begin with a simple case. Suppose $k_A = k_B = 2$ and $m = 1$. Then given $\theta = 1$, the random variable $S_A, S_B, S_{A\cap B}$ can be decomposed into a tuple (a, b, c) where each sample from $\mathrm{Ber}(\rho)$ independently, and $(S_A, S_B, S_C) = (a + c, b + c, c)$ equal in distribution. We can then compute $\mathbb{E}[\mathrm{Var}(c \mid a + c, b + c)]$ directly.

$$\begin{aligned}
&\mathbb{E}[\mathrm{Var}(c \mid a + c, b + c)] \\
=&\Pr[a + c = 1, b + c = 1]\mathrm{Var}(c \mid a + c = 1, b + c = 1) \\
&\qquad\qquad (c \text{ is deterministic when } a + c = 0, 2 \text{ or } b + c = 0, 2) \\
=&\Pr[a + c = 1, b + c = 1]\rho(1 - \rho) \\
&\qquad\qquad (\text{because } \Pr[c = 0 \mid a + c = 1, b + c = 1] = \tfrac{1-\rho}{1-\rho+\rho} = 1 - \rho) \\
=&\rho^2(1 - \rho)^2
\end{aligned}$$

Therefore, when a, b, c are independently sampled from $\mathrm{Ber}(\rho)$

$$\mathbb{E}[\mathrm{Var}(c \mid a + c, b + c)] = \rho^2(1 - \rho)^2 \tag{9}$$

In general let $d = \min\{m, k_A - m, k_B - m\} \geq 1$. we pick $d = \min\{m, k_A - m, k_B - m\} \geq 1$ indices in $A \setminus B$, $B \setminus A$ and $A \cap B$ respectively: (a_j, b_j, c_j) with $j = 1, \ldots, d$ so that $a_j \in A \setminus B$, $b_j \in B \setminus A$, and $c_j \in A \cap B$. Let $\mathcal{F}$ be the event of knowing θ, the values of all clues outside the triples, and the pair sums $(a_j + c_j, b_j + c_j)$ for $j = 1, \ldots, d$. Then S_A is the sum of the $a_j + c_j$ and of the values of the clues of A outside the triples. Conditioning on $\mathcal{F}$ is sufficient to recover S_A, S_B and can only decrease the expected conditional variance,

$$\mathbb{E}\big(\mathrm{Var}(S_{A\cap B} \mid S_A, S_B)\big) \geq \mathbb{E}\big(\mathrm{Var}(S_{A\cap B} \mid \mathcal{F})\big).$$

Given $\mathcal{F}$, the difference $S_{A\cap B} - \sum_{j=1}^{d} c_j$ is known, and conditional on the allocation and θ the clues are independent $\mathrm{Ber}(\rho)$ if $\theta = 1$ and $\mathrm{Ber}(1 - \rho)$ if $\theta = 0$, so the triples are conditionally independent and

$$\mathrm{Var}(S_{A\cap B} \mid \mathcal{F}) = \sum_{j=1}^{d} \mathrm{Var}(c_j \mid a_j + c_j, b_j + c_j, \theta).$$

Taking expectations and applying Eq. (9) to each summand gives $\mathbb{E}\big(\mathrm{Var}(S_{A\cap B} \mid S_A, S_B)\big) \geq \rho^2(1 - \rho)^2 d$, and the bound follows. □

B Proofs for the anecdote protocol

Let $A_t = A \setminus J_t$ and $B_t = B \setminus J_t$ be the agents' unannounced clues, with $k_A^{(t)} = |A_t|$ and $k_B^{(t)} = |B_t|$.

Proof of Theorem 3.8. By Lemma 3.5,

$$R_t = \mathbb{E}\big[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_A, H_t)\big] + \mathbb{E}\big[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_B, H_t)\big].$$

We bound the first term; the second is symmetric. Since the conditional mean minimizes the squared error, for every g that is a function of $(\mathcal{I}_A, H_t)$,

$$\mathbb{E}\big[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_A, H_t)\big] \leq \mathbb{E}\big[(g - L_{A\cap B})^2\big]. \tag{10}$$

Splitting the shared clues into announced and unannounced ones,

$$L_{A\cap B} = L_{A\cap B\cap J_t} + L_{A_t\cap B}.$$

We take $\hat{m} = \mathbb{E}\big[m^{(t)} \mid A, \Phi_t\big]$ and

$$g = L_{A\cap B\cap J_t} + \hat{m}\,\frac{L_{A_t}}{k_A^{(t)}},$$

which depends only on A, $\boldsymbol{s}_A$ and Φ_t.[1] We want to upper bound the right-hand side of Eq. (10). Let

$$x = m^{(t)} \frac{L_{A_t}}{k_A^{(t)}} - L_{A_t \cap B} = \sum_{i \in A_t} \Big(\frac{m^{(t)}}{k_A^{(t)}} - \mathbf{1}\{i \in B\} \Big) l_i$$

be the error Alice would make if she knew $m^{(t)}$, and $y = \big(\hat{m} - m^{(t)}\big) L_{A_t} / k_A^{(t)}$. Then $g - L_{A \cap B} = x + y$. Conditioning on θ, A, B and the public randomness R, which we suppress from the notation,

$$\mathbb{E}\big[(g - L_{A\cap B})^2 \mid \theta, A, B\big] = \mathbb{E}[x^2 \mid \theta, A, B] + \mathbb{E}[y^2 \mid \theta, A, B] + 2\,\mathbb{E}[xy \mid \theta, A, B].$$

Given these, the sets J_t, A_t, $A_t \cap B$ and the count $m^{(t)}$ are fixed, because the draws use only R and never look at the clue values, while the $l_i = \lambda(2s_i - 1)$ remain i.i.d. with mean $\lambda(2q_\theta - 1)$ and variance $4\lambda^2 q_\theta(1 - q_\theta) = 4\lambda^2 \rho(1 - \rho)$.

First, because $\sum_{i \in A_t} \mathbf{1}\{i \in B\} = m^{(t)}$, the coefficients of x sum to 0, so

$$\mathbb{E}[x \mid \theta, A, B] = \lambda(2q_\theta - 1) \sum_{i \in A_t} \Big(\frac{m^{(t)}}{k_A^{(t)}} - \mathbf{1}\{i \in B\} \Big) = 0.$$

Hence $\mathbb{E}[x^2 \mid \theta, A, B]$ is the conditional variance of x, which is $4\lambda^2\rho(1 - \rho)$ times the sum of the squared coefficients, $m^{(t)}(1 - m^{(t)}/k_A^{(t)}) \le m^{(t)}$:

$$\mathbb{E}[x^2 \mid \theta, A, B] \le 4\lambda^2 \rho(1 - \rho)\, m^{(t)}.$$

Second, for the cross term, because the values are i.i.d., $\mathbb{E}[l_i L_{A_t} \mid \theta, A, B]$ is the same for every $i \in A_t$, so $\mathbb{E}[x L_{A_t} \mid \theta, A, B] = 0$. Hence

$$\mathbb{E}[xy \mid \theta, A, B] = \mathbb{E}\big[\hat{m} - m^{(t)} \mid \theta, A, B\big] \cdot \frac{\mathbb{E}[x L_{A_t} \mid \theta, A, B]}{k_A^{(t)}} = 0.$$

Finally $|L_{A_t}| \le \lambda k_A^{(t)}$, so $y^2 \le \lambda^2 (\hat{m} - m^{(t)})^2$ and

$$\mathbb{E}[y^2] \le \lambda^2\, \mathbb{E}\big[(\hat{m} - m^{(t)})^2\big] = \lambda^2\, \mathbb{E}\big[\mathrm{Var}(m^{(t)} \mid A, \Phi_t)\big].$$

Taking expectations in Eq. (10),

$$\mathbb{E}\big[\mathrm{Var}(L_{A\cap B} \mid \mathcal{I}_A, H_t)\big] \le 4\lambda^2\rho(1 - \rho)\, \mathbb{E}\big[m^{(t)}\big] + \lambda^2\, \mathbb{E}\big[\mathrm{Var}(m^{(t)} \mid A, \Phi_t)\big],$$

and adding the symmetric bound for Bob proves the inequality.

□

Proof of Corollary 3.9. First, when $|A|, |B| \le k$, if the protocol has stopped by round t, then $m^{(t+1)} = m^{(t)} = 0$. Otherwise the sender at round $t + 1$ draws uniformly from her at most k unannounced clues, which include all $m^{(t)}$ unannounced shared clues, so the drawn clue is shared with conditional probability at least $m^{(t)}/k$. Hence $\mathbb{E}[m^{(t+1)} \mid A, B, J_t] \le m^{(t)}(1 - \frac{1}{k})$, and iterating from $m^{(0)} = m$ gives

$$\mathbb{E}[m^{(t)}] \le \mathbb{E}[m]\,(1 - \tfrac{1}{k})^t. \tag{11}$$

Because $\mathrm{Var}(m^{(t)} \mid A, \Phi_t) \le \mathbb{E}[(m^{(t)})^2 \mid A, \Phi_t]$, Eq. (11) gives

$$R_t \le 8\lambda^2\rho(1 - \rho)\, \mathbb{E}\big[m^{(t)}\big] + 2\lambda^2\, \mathbb{E}\big[(m^{(t)})^2\big] \le \lambda^2 \Big(8\rho(1 - \rho)\, \mathbb{E}[m] + 2\,\mathbb{E}[m^2] \Big) \Big(1 - \frac{1}{k} \Big)^t.$$

The last inequality uses $\mathbb{E}[(m^{(t)})^2] \le \mathbb{E}[m m^{(t)}] \le \mathbb{E}[m^2](1 - 1/k)^t$.

□

[1] When $A_t = \emptyset$ the last term of g is 0.

C Complement and Substitute

The following proposition shows that our information structure is weakly but not rectangle substitutes Frongillo et al. (2023)

Proposition C.1 (Weak but not rectangle substitutes)**.** *Given* $\mathrm{BSC}(\rho)$ *and* $\mathcal{D}(k_A, k_B, m)$*, take the target* $Y = L^\star$*, the information structure satisfies weak substitutes so that*

$$\mathbb{E}[\mathrm{Var}(L^\star \mid \mathcal{I}_A)] + \mathbb{E}[\mathrm{Var}(L^\star \mid \mathcal{I}_B)] \le \mathrm{Var}(L^\star),$$

but it is not δ*-approximate rectangle substitutes for any* $\delta < R_\star/2$*.*

Proof. As Alice knows $l_0 + L_A$, her residual is $L_{B\setminus A}$ and $\mathbb{E}[\mathrm{Var}(L^\star \mid \mathcal{I}_A)] \le \mathrm{Var}(L_{B\setminus A})$. Similarly $\mathbb{E}[\mathrm{Var}(L^\star \mid \mathcal{I}_B)] \le \mathrm{Var}(L_{A\setminus B})$. The three block sums $L_{A\setminus B}, L_{B\setminus A}, L_{A\cap B}$ are conditionally independent given θ, with nonnegative covariance so

$$\mathrm{Var}(L^\star) \ge \mathrm{Var}(L_{A\setminus B}) + \mathrm{Var}(L_{B\setminus A}) \ge \mathbb{E}[\mathrm{Var}(L^\star \mid \mathcal{I}_A)] + \mathbb{E}[\mathrm{Var}(L^\star \mid \mathcal{I}_B)].$$

Fix a cell $C = \{S_A = s\} \times \{S_B = s'\}$. On C, Alice's estimate is $\mu_{\sigma T} = \mathbb{E}[L^\star \mid \mathcal{I}_A, S_B]$ and the outside observer's is $\mu_{ST} = \mathbb{E}[L^\star \mid S_A, S_B]$. By Eq. (4) with $g = S_B$ these coincide, so the right-hand side of (Frongillo et al., 2023, Rem. 2.8) is

$$\mathbb{E}\big[(\mu_{\sigma T} - \mu_{ST})^2 \mid C\big] = \mathbb{E}\Big[\big(\mathbb{E}[L^\star \mid \mathcal{I}_A, S_B] - \mathbb{E}[L^\star \mid S_A, S_B]\big)^2 \,\Big|\, S_A = s, S_B = s'\Big] = 0.$$

Symmetrically, Bob's estimate is $\mu_{S\tau} = \mathbb{E}[L^\star \mid S_A, \mathcal{I}_B] = \mathbb{E}[L^\star \mid S_A, S_B]$, and $\mu_{\sigma\tau} = L^\star$, so the left-hand side is

$$\mathbb{E}\big[(\mu_{\sigma\tau} - \mu_{S\tau})^2 \mid C\big] = \mathbb{E}\Big[\big(L^\star - \mathbb{E}[L^\star \mid S_A, S_B]\big)^2 \,\Big|\, S_A = s, S_B = s'\Big] = \mathrm{Var}(L_{A\cap B} \mid S_A = s, S_B = s'),$$

since $L^\star = l_0 + L_A + L_B - L_{A\cap B}$ and L_A, L_B are fixed on C. Averaging over cells gives $\mathbb{E}[\mathrm{Var}(L_{A\cap B} \mid S_A, S_B)] = R_\star/2$. □

D Realistic match experiment

The realistic-text experiment uses an anonymized tennis match with twelve pre-match facts. The facts are designed to resemble information that could plausibly be available before a professional match while spanning different predictive dimensions.

ID	Fact	Alice	Bob
F01	Player A won 8 of the last 10 matches on the relevant surface.	✓	
F02	Player B won 7 of the last 10 matches on the relevant surface.		✓
F03	Player A has an 84% hold rate.		✓
F04	Player B has an 87% hold rate.	✓	
F05	Player A won the most recent head-to-head meeting.	✓	✓
F06	Player B leads the overall head-to-head record 4–2.		✓
F07	Player A has a 12–5 record against top-20 opponents.	✓	
F08	Player B has won 14 of the last 17 matches.		✓
F09	Player A played a five-set match two days earlier.	✓	✓
F10	Player B has the higher surface return-points-won percentage.		✓
F11	Player A has the higher surface-adjusted Elo rating.	✓	
F12	Player B has spent fewer total hours on court in the tournament.		✓

Table 2: Facts and information allocation in the realistic match scenario. F05 and F09 are observed by both agents; all other facts are private to one agent.

Alice therefore observes six facts and Bob observes eight facts, with two overlapping facts. This creates an asymmetric information structure: Bob is missing four facts unique to Alice, while Alice is missing six facts unique to Bob. Consequently, even under perfect communication, Bob can recover the complete union of facts in fewer communication rounds than Alice. We therefore interpret agent-specific convergence speed with this structural asymmetry in mind and focus primarily on whether the communication strategies reduce distance to the full-information forecast and belief disagreement.

Text-to-clue mapping. For each fact, the LLM assigns one of seven qualitative impact levels. These are converted to fixed log-odds increments as follows:

Impact category	Log-odds increment
Strong clue for B	-0.90
Moderate clue for B	-0.50
Slight clue for B	-0.20
Neutral	0
Slight clue for A	$+0.20$
Moderate clue for A	$+0.50$
Strong clue for A	$+0.90$

Python then combines these increments in log-odds space to obtain each agent's posterior forecast.

Full-information benchmark and baselines. The full-information forecast is computed by applying the same text-to-clue mapping to the union of all twelve facts and combining the resulting log-odds increments exactly once per fact. This forecast serves as the omniscient benchmark for the realistic scenario. The communication baselines use the same strategy definitions as in the numerical experiments: random communication provides a non-strategic baseline; strongest prioritizes the fact judged to have the largest absolute evidential impact; receiver-aware strategies such as impactful and attractive condition their selection on the receiver's current forecast; and belief-only communicates the posterior without revealing a supporting fact. Performance is evaluated after each communication round by the agents' distance to the full-information forecast and by their mutual belief disagreement.

E LLM IMPLEMENTATION AND PROMPTS

The primary LLM is `openai/gpt-oss-20b`, with sensitivity checks using `Qwen/Qwen2.5-7B-Instruct` and `microsoft/Phi-3.5-mini-instruct`, run locally through Hugging Face Transformers. Clue selection is implemented by scoring lettered candidate options using conditional log-probabilities rather than sampled generation, so temperature and top-p are not used. The sender is shown its own posterior, the receiver's posterior, the receiver's known clues, communication history, and the unresolved candidate clues. Candidate order is seeded and shuffled; exact ties are resolved by the first highest-scoring option. The Gaussian experiments use random seed 20260924, while the realistic-match experiment uses random seed 20260915. Because selection is constrained to valid options, malformed free-form outputs do not occur. The primary Gaussian experiment uses 30 independently generated environments, with all policies evaluated on the same environment within each run; the model-sensitivity experiment uses 10 matched environments per model. Complete prompts and implementation details are given in Appendix E.

LLM calls. Each LLM-controlled communication decision corresponds to one model-scoring call over the available candidate options. In the Gaussian experiments there are at most two such decisions per communication round, one for each agent. The exact number of calls therefore depends on the policy and whether communication terminates before the maximum number of rounds. The reported Monte Carlo repetition counts are 30 environments for the primary Gaussian experiment, 10 matched environments per model for sensitivity analysis, and 30 repetitions for the realistic-match experiment.

Inference. No separate system prompt is used. When a model provides a chat template, the prompt below is supplied as a single user message; otherwise the same text is supplied directly. For every valid option letter ℓ, we compute its mean conditional token log-probability

$$s(\ell) = \frac{1}{|\ell|} \log P_\theta(\ell \mid \text{prompt}),$$

and select $\arg\max_\ell s(\ell)$. No temperature, top-p, or sampled generation is used.

Gaussian clue-selection prompt.

```
You are [SENDER] selecting exactly ONE clue to communicate.

Communication strategy: [STRATEGY]
Mathematical objective: [OBJECTIVE]

Sender current belief: P(theta=1)=[SENDER BELIEF]
Receiver current belief: P(theta=1)=[RECEIVER BELIEF]
Receiver-known clue IDs: [KNOWN IDS]
Previously communicated clue IDs: [HISTORY]

Available unresolved sender clues:
A. [CLUE]
B. [CLUE]
...

Important:
- Do not estimate theta itself.
- Do not output a probability.
- Apply only the stated communication strategy.
- Choose exactly one listed option.
```

Each candidate is displayed as

```
clue_id=[ID]; signal=[X]; sigma=[SIGMA];
log_likelihood_ratio=[LLR]
```

The strategy-specific objectives are exactly those defined in Section 4.1 and summarized in Table 1.

Realistic-match fact-evaluation prompt.

```
You are evaluating evidence for a real pre-event forecasting problem.

Event:
An anonymized best-of-five professional tennis match between
Player A and Player B.

Outcome being forecast: Player A wins.
Current forecast before seeing the new fact:
P(Player A wins) = [CURRENT FORECAST]

Information already incorporated:
[KNOWN FACTS]

NEW FACT:
[FACT TEXT]

Assess the marginal predictive value of ONLY this new fact,
given what is already known.
Do not assume the actual match outcome. Account for redundancy.

Choose exactly one:
A. Strong evidence favoring Player B
B. Moderate evidence favoring Player B
C. Slight evidence favoring Player B
D. Approximately neutral / little predictive value
E. Slight evidence favoring Player A
F. Moderate evidence favoring Player A
G. Strong evidence favoring Player A
```

```
ANSWER:
```

The seven options correspond respectively to log-odds increments $-0.90, -0.50, -0.20, 0, 0.20, 0.50, 0.90$.

Realistic-match communication prompt.

```
You are [SENDER] deciding which ONE real-world fact to communicate.

Event:
An anonymized best-of-five professional tennis match between
Player A and Player B.

Communication strategy: [STRATEGY]

Sender current belief: [SENDER BELIEF]
Receiver current belief: [RECEIVER BELIEF]

Receiver already knows:
[RECEIVER FACTS]

Available facts you can send:
A. [FACT]
B. [FACT]
...

Choose exactly one listed fact according to the stated
communication strategy.
Do not assume the actual match outcome.
Do not repeat information already known.
Answer with the option letter only.

ANSWER:
```

E.0.1 Sensitivity across LLMs

We test whether the Gaussian communication results depend strongly on the choice of language model by repeating the experiment across multiple LLMs while holding the signal environment, communication objectives, and Bayesian updating procedure fixed.

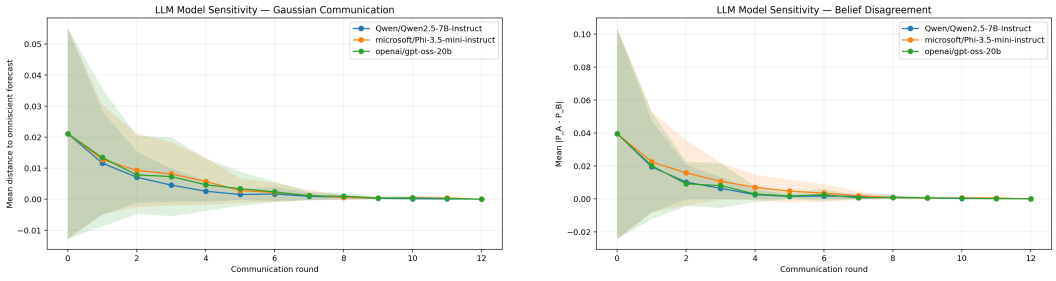

Figure 2: Sensitivity of LLM-controlled Gaussian deliberation to model choice. Left: mean distance to the omniscient forecast. Right: mean disagreement between Alice and Bob. Shaded regions denote normal-approximation 95% confidence intervals for the mean, computed as $\bar{x} \pm 1.96 s/\sqrt{n}$ across 10 matched environments per model. Because this symmetric approximation is not constrained to the support of the nonnegative metric, some displayed lower bounds fall below zero.