Cheap Talk Games with Dempster-Shafer Priors*

Willem Conradie[3], Krishna Manoorkar[2], Alessandra Palmigiano[1,6], Apostolos Tzimoulis[5], and Nachoem Wijnberg[4,7]

[1] School of Business and Economics, Vrije Universiteit Amsterdam
[2] Institute of Computer Science, Czech Academy of Sciences, Czechia
[3] School of Mathematics, University of the Witwatersrand, Johannesburg
[4] Faculty of Economics and Business, University of Amsterdam
[5] Deaprtment of Computer Science, University of Luxembourg, Luxembourg
[6] Department of Pure and Applied Mathematics, University of Johannesburg
[7] College of Business and Economics, University of Johannesburg

1 Introduction

The present paper studies strategic communication frameworks for situations in which a sender seeks to influence a receiver's actions through information transmission. An important game-theoretic model for studying strategic communication is that of *cheap talk games* [8], where communication is costless, non-binding, and non-verifiable, yet can still affect outcomes. Typically, a sender privately observes a state and sends a message, and the receiver forms beliefs based on the message and the sender's incentives. Equilibria for such game range from informative partitions to uninformative babbling, depending on the degree of preference alignment. By contrast, in the framework of Bayesian persuasion [18], the sender commits ex ante to a signaling strategy in order to optimally shape the receiver's beliefs.

Extensions of this framework allow for asymmetric or misspecified receiver priors, limited commitment, and disagreement about the sender's incentives [1, 21, 22, 13, 10]. Across these models, agents share a common probabilistic representation of uncertainty, so that ambiguity arises only through differences in priors or incentives. A growing literature instead studies persuasion under ambiguity, typically modeled using sets of priors or maxmin preferences [5, 4, 19, 17, 16].

A key limitation of these approaches is their reliance on probabilistic priors, which assume that all uncertainty can be represented by a single additive measure. In many communication settings, however, agents face ambiguity that cannot be adequately captured in this way. Representing such uncertainty by forcing precise probabilities may introduce artificial information and obscure how agents' attitudes toward ambiguity shape behavior. Moreover, standard models often assume common knowledge of the receiver's beliefs, whereas, in practice, senders are frequently uncertain about how messages will be interpreted, especially when receivers may hold biased or heterogeneous priors. This creates an additional layer of strategic uncertainty that existing frameworks do not fully address. Finally, models of ambiguous persuasion based on sets of priors typically focus on strategies that are robust to uncertainty about the receiver's beliefs, limiting the analysis of richer strategic interactions under ambiguity.

This paper introduces a cheap-talk framework in which a sender observes the true state and communicates with a receiver, who forms a belief over possible states. The sender aims to

*This work is partially funded by the EU MSCA (grant No. 101007627) and by the NWO grant KIVI.2019.001.This work was supported by Operational Programme Jan Ámos Komenský project '*Biography of Fake News with a Touch of AI: Dangerous Phenomenon through the Prism of Modern Human Sciences*' (reg. CZ.02.01.01/00/23_025/0008724).The research of Apostolos Tzimoulis is partially supported by FNR project C25/IS/19566252 and the Key Project of the Chinese Ministry of Education (22JJD720021).

persuade the receiver that the true state lies within a given winning set, while the receiver seeks accurate inference. Payoffs depend on the probability the receiver assigns to the winning set and on the true state. To model epistemic uncertainty, we represent priors using *Dempster–Shafer mass functions* [25, 11], allowing for non-additive and ambiguous beliefs.

The receiver's strategy depends on how this ambiguity is resolved—formally, through a *specialization matrix* [26] that converts mass functions into probability distributions. Since the sender does not know how the receiver will resolve ambiguity, optimal messaging depends on the sender's attitude towards this uncertainty. We analyze three important cases: a neutral (*pignistic*) approach with uniform redistribution, a worst-case (*robust*) approach, and a best-case (*optimistic*) approach.

We also study asymmetric settings in which sender and receiver do not share a common prior over states, goals, or strategies. Such asymmetries are modeled by allowing the underlying mass functions to differ or be only partially known. The sender then optimizes under incomplete information, typically via robust (maxmin) strategies. We characterize optimal behavior in these settings, including cases where the receiver's beliefs are aligned with (*faithful*) or diverge from (*unfaithful*) the sender's expectations.

2 Cheap Talk Games with Dempster-Shafer Priors

This section defines a cheap-talk game with a common Dempster–Shafer prior and analyzes the players' optimal strategies.

Game-theoretic setup. Let X be a set of states with true state $x_0 \in X$ and let $\mathcal{M}$ be a set of messages. Let $\mathsf{PD}(X)$ denote the set of all probability distributions on X. The sender A has a winning set $W_A \subseteq X$, typically unknown to the receiver B, and aims to persuade B that $x_0 \in W_A$. The game proceeds as follows. The sender observes (x_0, W_A) and sends a message $t \in \mathcal{M}$ via a strategy $\sigma^A : X \times \mathcal{P}(X) \to \mathcal{M}$. The receiver, upon observing t, chooses a probability distribution $P_o = \tau_B(t) \in \mathsf{PD}(X)$ according to a strategy $\tau^B : \mathcal{M} \to \mathsf{PD}(X)$. Priors beliefs of the receiver are modeled by a function $g^* : \mathcal{M} \to \mathsf{DS}(X)$ assigning a Dempster–Shafer mass function g^*_t to each $t \in \mathcal{M}$. Payoff functions U_A and U_B of A and B are given by

$$U_A(x, W, \sigma, \tau) = \tau\sigma(x, W)(W) \quad \text{and} \quad U_B(x, W, \sigma, \tau) = \log\left(\tau\sigma(x, W)(\{x\})\right),$$

i.e., the probability assigned by P_o to the winning set, and the log-probability assigned by P_o to the true state, respectively.

Analysis of the game. For any $t \in \mathcal{M}$, player B chooses a distribution P_o that maximizes expected utility given their belief $g^*_t \in \mathsf{DS}(X)$ and a *specialization matrix* S^*, which resolves the ambiguity in g^*_t by inducing a compatible probability distribution $g^*_t \cdot S^*$. Thus, B maximizes

$$\sum_{x \in X} P_o(\{x\}) \log\left((g^*_t \cdot S^*)(\{x\})\right).$$

By the log-sum inequality [7, Theorem 2.7.1], the best strategy for the sender A is

$$\tau^B(t) = g^*_t \cdot S^*.$$

Since S^* is unknown to A, we consider three approaches reflecting different attitudes toward this uncertainty:

Pignistic approach: The sender assumes uniform redistribution of mass, i.e. that B chooses the probability distribution $\mathrm{bet}_P(g^*_t)$, the *pignistic transform* [26] of mass function g^*_t. Hence, the optimal message satisfies

$$t_0 \in \arg\max_{t \in \mathcal{M}} \mathrm{bet}_P(g^*_t)(W_A).$$

Robust approach: The sender assumes that B selects the worst-case distribution consistent with g^*_t. In this case, the optimal message for A satisfies

$$t_0 \in \arg\max_{t\in\mathcal{M}} \mathsf{bel}_{g_t^*}(W_A),$$

guaranteeing A the payoff $\max_t \mathsf{bel}_{g_t^*}(W_A)$.

Optimistic approach: The sender assumes that B selects the best-case distribution for A. In this case, the optimal message for A satisfies

$$t_0 \in \arg\max_{t\in\mathcal{M}} \mathsf{pl}_{g_t^*}(W_A).$$

These approaches capture a spectrum of attitudes toward ambiguity, from pessimistic (robust) to neutral (pignistic) to optimistic.

Equilibria. Let $g \in \mathsf{DS}\big(\mathcal{P}(X) \times \mathcal{P}(\mathcal{P}(X))\big)$ be a Dempster-Shafer mass function representing the (common knowledge) prior joint belief regarding the true state and the sender's winning set. Let σ^A denote the sender's strategy. For any possible message $t \in \mathcal{M}$, define

$$\mathsf{SW}(t, \sigma^A) := (\sigma^A)^{-1}(t) = \{(x, W) \in X \times \mathcal{P}(X) \mid \sigma^A(x, W) = t\}.$$

Let $\mathbf{1}_{\mathsf{SW}(t,\sigma^A)}$ denote the mass function assigning mass 1 to the set $\mathsf{SW}(t, \sigma^A)$ and 0 to every other set. If σ^A is known to B, then, upon receiving message t, player B (the receiver) conditions the prior belief function g on t. Hence, B's updated belief is given by the following mass function:

$$g_t^* = \begin{cases} \big(g \oplus \mathbf{1}_{\mathsf{SW}(t,\sigma^A)}\big)\big|_{\mathcal{P}(X)} & \text{if } \mathsf{pl}_g(\mathsf{SW}(t, \sigma^A)) > 0, \\ g\big|_{\mathcal{P}(X)} & \text{otherwise} \end{cases}$$

where, if m is any mass function, $m\big|_{\mathcal{P}(X)}$ denotes its marginalization to $\mathcal{P}(X)$, and $\oplus$ is the Dempster-Shafer combination operator.

An *equilibrium* of the cheap-talk game with a Dempster-Shafer prior, for the given attitude of the receiver to uncertainty, is a tuple (g, σ^A, τ^B) such that:

1. g is a common-knowledge mass function describing B's describing prior over both the true state and A's winning set;
2. $\tau^B(t) = g_t^* \cdot S^*$, where S^* is the specialization matrix resolving ambiguity;
3. For any $(x, W) \in X \times \mathcal{P}(X)$ and any attitude of A towards uncertainty, $\sigma^A(x, W) \in \mathcal{M}$ must satisfy the optimality condition under that attitude, as discussed above.

Prior on the sender's strategies. Let Strat_A be the set of sender strategies, and $G \in \mathsf{DS}\big(\mathcal{P}(\mathsf{Strat}_A)\big)$ be the receiver's prior over them, known to A. For $t \in \mathcal{M}$ and $\Sigma \subseteq \mathsf{Strat}_A$, let

$$\mathsf{SW}(t, \Sigma) = \bigcup_{\sigma\in\Sigma} \mathsf{SW}(t, \sigma).$$

Each message t induces a mass function G_t^* on $X \times \mathcal{P}(X)$ by transferring mass from Σ to $\mathsf{SW}(t, \Sigma)$:

$$G_t^*(\mathcal{Z}) = \sum\{G(\Sigma) \mid \Sigma \subseteq \mathsf{Strat}_A,\ \mathsf{SW}(t, \Sigma) = \mathcal{Z}\}.$$

The receiver's updated belief is then

$$g_t^* = \begin{cases} (g \oplus G_t^*)\big|_{\mathcal{P}(X)} & \text{if conflict}(g, G_t^*) < 1, \\ g\big|_{\mathcal{P}(X)} & \text{otherwise,} \end{cases}$$

The best strategies of the two players are then determined by this g_t^*, the receiver's specialization matrix S^*, and the sender's attitude to uncertainty, as discussed earlier.

Remark 1. *If we assume that g and G are common knowledge and both players are fully rational, then it is common knowledge that A will only play strategies that are best responses (for a given receiver type) for the given g and G, according to A's attitude toward uncertainty.*

If the receiver knows the sender's attitude $Y \in \{P, R, O\}$ to uncertainty, then $\mathsf{bel}_G(\Sigma^Y_{best}) = 1$, where Σ^Y_{best} is the set of best strategies for A for given Y. However, in many real-life scenarios, players are not fully rational, and this condition may not hold.

3 Games with Asymmetric Priors

In the previous section, we assumed that for any message t, the mass function g^*_t is common knowledge. In many settings, however, the sender A may not know the exact mass function describing the receiver B's beliefs about the true state and A's winning set. We now consider scenarios in which A has only partial or approximate information about g^*_t.

When the receiver has additional private information. So far, ambiguity has been represented *within* a mass function, via uncertainty about the true state encoded in g^*_t. We now introduce a second layer of uncertainty concerning *the mass function itself.*

Specifically, the sender may not know the exact g^*_t describing the receiver's beliefs. Instead, for each message t, a coarser prior h^*_t is common knowledge. This captures baseline information before the receiver incorporates additional (possibly private) evidence, represented by a specialization matrix T:

$$g^*_t = h^*_t \cdot T.$$

The sender does not observe T, only that such a transformation determines g^*_t.

We focus on a robust attitude toward this higher-order uncertainty: the sender chooses t to maximize the minimal payoff over all specializations consistent with h^*_t. Other attitudes can be treated analogously. For any set X we use $\mathsf{Matr}(X)$ to denote the set of all sepcialization matrices on X.

Lemma 1. *For sender A, given h^* the best choice of message $t_0 \in \mathcal{M}$ and resulting expected payoff are indicated in the following table.*

approach	best strategy	payoff
robust	$t_0 \in \arg\max_{t\in\mathcal{M}} \min_{S\in\mathsf{Matr}(X)} \mathsf{bel}_{h^*_t\cdot S}(W_A)$ $= \arg\max_{t\in\mathcal{M}} \mathsf{bel}_{h^*_t}(W_A)$	$\max_{t\in\mathcal{M}} \mathsf{bel}_{h^*_t}(W_A)$
pignistic	$t_0 \in \arg\max_{t\in\mathcal{M}} \min_{S\in\mathsf{Matr}(X)} bet_P(h^*_t \cdot S)(W_A)$ $= \arg\max_{t\in\mathcal{M}} \mathsf{bel}_{h^*_t}(W_A)$	$\max_{t\in\mathcal{M}} \mathsf{bel}_{h^*_t}(W_A)$
optimistic	$t_0 \in \arg\max_{t\in\mathcal{M}} \min_{S\in\mathsf{Matr}(X)} \mathsf{pl}_{h^*_t\cdot S}(W_A)$ $= \arg\max_{t\in\mathcal{M}} \mathsf{bel}_{h^*_t}(W_A)$	$\max_{t\in\mathcal{M}} \mathsf{bel}_{h^*_t}(W_A)$

Thus, the sender can achieve the guaranteed payoff $\max_{t\in\mathcal{M}} \mathsf{bel}_{h^*_t\cdot S}(W_A)$ by adapting the best strategy under the robust approach.

When the sender's approximation is unfaithful. So far, we assumed that the sender A has correct knowledge of the receiver B's prior. In many settings, however, A's beliefs about B's prior may be inaccurate, and B may in turn hold beliefs about A's beliefs. Such mismatches generate higher-order uncertainty and richer strategic behavior. An example of such scenarios is that of *double bluffing*. In a simple two-state setting $X = \{p_1, p_2\}$ with messages $\mathcal{M} = \{p_1, p_2\}$, if A knows that B believes A to report the state she wants B to believe, then A may optimally send the opposite message, knowing that B will discount it. This strategic reversal leads B to update beliefs in a way that can favor A, even when A's message is literally false. Thus, double bluffing emerges naturally from misaligned beliefs about priors and strategies: optimal communication hinges not only on the message itself, but on how it is expected to be interpreted.

4 Conclusions and Future Directions

We introduced a generalization of cheap-talk games in which priors are modeled by Dempster–Shafer mass functions rather than probabilities. We analyze both common- and asymmetric-prior settings and characterize optimal strategies under different attitudes toward ambiguity. The framework captures not only beliefs about states, but also uncertainty about those beliefs. Several interesting examples of strategic communication can be formalized using this framework highlighting many important features of strategic communication:

- **Strategic revelation.** Even truthful communication can be persuasive: the sender can selectively reveal information to steer the receiver's beliefs.
- **Discounting and credibility.** Optimal strategies depend on how much the receiver trusts messages. Greater trust allows more aggressive signaling, while lower trust favors conservative, prior-aligned messages.
- **Strategic lying to be caught.** A sender may use detectable lies to manipulate higher-order beliefs, shaping how subsequent messages are interpreted.

This paper is part of a line of research aimed at modelling how agents extract reliable information from narratives and communicative interactions in which speakers may strategically distort the truth [6]. Directions for future research within this line include the computational complexity of optimal strategies and the existence of equilibria under Dempster–Shafer priors, which merit further study, especially in the light of recent algorithmic work on cheap talk and Bayesian persuasion [3, 12]. Second, extending the framework to repeated or multi-stage settings would connect it to models of long or dynamic cheap talk and sequential persuasion [2, 23, 20]. Third, exploring alternative belief aggregation rules beyond Dempster's rule may yield different strategic and equilibrium outcomes [24].

These questions are also relevant for multi-agent artificial intelligence, where agents interact under uncertainty, trust, and potential deception [9, 15, 14]. Frameworks for belief representation and aggregation, such as the Dempster–Shafer approach, can provide normative tools for analyzing communication in such environments [27].

References

[1] Ricardo Alonso and Odilon Câmara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.

[2] Robert J Aumann and Sergiu Hart. Long cheap talk. *Econometrica*, 71(6):1619–1660, 2003.

[3] Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi. Algorithmic cheap talk. *arXiv preprint arXiv:2311.09011*, 2023.

[4] Dorian Beauchêne, Jian Li, and Ming Li. Ambiguous persuasion. *Journal of Economic Theory*, 179:312–365, 2019.

[5] Subir Bose and Ludovic Renou. Mechanism design with ambiguous communication devices. *Econometrica*, 82(5):1853–1872, 2014.

[6] Willem Conradie, Krishna Manoorkar, Alessandra Palmigiano, Apostolos Tzimoulis, and Nachoem Wijnberg. The tiglath pileser principle and models of persuasion. Manuscript in preparation.

[7] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, NJ, 2 edition, 2006. Section 2.7: Log-Sum Inequality and Its Applications.

[8] Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, 1982.

[9] Pedro MP Curvo. The traitors: Deception and trust in multi-agent language model simulations. *arXiv preprint arXiv:2505.12923*, 2025.

[10] Geoffroy De Clippel and Xu Zhang. Non-bayesian persuasion. *Journal of Political Economy*, 130(10):2594–2642, 2022.

[11] Arthur P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics*, 38(2):325–339, 1967.

[12] Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 412–425, 2016.

[13] Kfir Eliaz and Ran Spiegler. A model of competing narratives. *American Economic Review*, 110(12):3786–3816, 2020.

[14] Meta Fundamental AI Research Diplomacy Team (FAIR)†, Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, 2022.

[15] Thilo Hagendorff. Deception abilities emerged in large language models. *Proceedings of the National Academy of Sciences*, 121(24):e2317967121, 2024.

[16] Jonas Hedlund, T Florian Kauffeldt, and Malte Lammert. Persuasion under ambiguity. *Theory and Decision*, 90(3):455–482, 2021.

[17] Ju Hu and Xi Weng. Robust persuasion of a privately informed receiver. *Economic Theory*, 72(3):909–953, 2021.

[18] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[19] Christian Kellner and Mark T Le Quement. Endogenous ambiguity in cheap talk. *Journal of Economic Theory*, 173:1–17, 2018.

[20] Fei Li and Peter Norman. Sequential persuasion. *Theoretical Economics*, 16(2):639–675, 2021.

[21] Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Persuasion via weak institutions. *Journal of Political Economy*, 130(10):2705–2730, 2022.

[22] John Morgan and Phillip C. Stocken. An analysis of stock recommendations. *RAND Journal of Economics*, 34(1):183–203, 2003.

[23] Serkan Sarıtaş, Serdar Yüksel, and Sinan Gezici. Dynamic signaling games with quadratic criteria under nash and stackelberg equilibria. *Automatica*, 115:108883, 2020.

[24] KARI SENTZ and SCOTT FERSON. Combination of evidence in dempster-shafer theory. Technical report, Sandia National Labs., Albuquerque, NM (US); Sandia National Labs., Livermore, CA (US), 04 2002.

[25] Glenn Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.

[26] Philippe Smets and Robert Kennes. The transferable belief model. *Artificial intelligence*, 66(2):191–234, 1994.

[27] Budhitama Subagdja, Ah-Hwee Tan, and Yilin Kang. A coordination framework for multi-agent persuasion and adviser systems. *Expert Systems with Applications*, 116:31–51, 2019.