# Persuasion with Ambiguous Communication*

Xiaoyu Cheng† Peter Klibanoff‡ Sujoy Mukerji§
Ludovic Renou¶

This Version: November 10, 2026

**Abstract**

We explore whether ambiguous communication can be beneficial to the sender in persuasion games when the receiver is ambiguity averse and dynamically consistent. Splitting of obedient experiments into a convex combination of two Pareto-ranked experiments (those for which both players agree on a payoff ranking) plays a key role. That some obedient experiment can be split in this way is necessary for the sender to benefit. If an optimal Bayesian persuasion experiment can be split in this way, then any not-too-ambiguity-averse sender benefits. The sender never benefits from ambiguous communication when the receiver has only two actions. We identify a class of games in which our conditions fully characterize when the sender benefits and are generically satisfied when there are at least three actions.

*We thank Dilip Abreu, Ian Ball, Piotr Dworczak, Huiyi Guo, Alex Jakobsen, Jian Li, Alessandro Pavan, Eran Shmaya and seminar audiences at Texas A&M, Washington University in St. Louis, University of California, Berkeley, RUD 2024, D-TEA 2024, EC'24, Northwestern University, University of Southern California, Queen Mary University of London, University of Notre Dame, HKUST, HKU Business School, 4th Durham Economic Theory Conference, Arizona State University, NYU, University of Manchester, Princeton University, and CETC 2025 for helpful feedback. A two-page abstract of an earlier version of this paper appeared in EC'24. Part of this work was carried out during Klibanoff's Hallsworth Visiting Professor appointment at the University of Manchester. Mukerji acknowledges financial support from ESRC Grant ES/S015299/1.

†Department of Economics, Florida State University, Tallahassee, FL USA. E-mail: xcheng@fsu.edu.

‡Managerial Economics and Decision Sciences, Kellogg School of Management at Northwestern University, Evanston, IL USA. E-mail: peterk@kellogg.northwestern.edu.

§School of Economics and Finance, Queen Mary University of London, London, UK. E-mail: s.mukerji@qmul.ac.uk.

¶Department of Economics, Arizona State University, Tempe, AZ USA, and CEPR. E-mail: lrenou.econ@gmail.com.

# 1 Introduction

> *"If I seem unduly clear to you, you must have misunderstood what I said."*
>
> Alan Greenspan, Speaking to a Senate Committee in 1987, as quoted in the Guardian Weekly, November 4, 2005.

This paper considers the problem of a sender who wishes to favorably influence, through strategic communication of information, the action taken by a receiver. As in the large literature on Bayesian persuasion following Kamenica and Gentzkow (2011) (see also Rayo and Segal (2010) and surveys by Bergemann and Morris (2019) and Kamenica (2019)), we model the sender as committing to a communication strategy and the receiver as best responding to that strategy. A communication strategy for the sender is usually described as a *statistical experiment*, a function mapping from payoff-relevant states to probability distributions over messages (or signals).[1] Our first key departure from most of the literature is that we enlarge the set of the sender's communication strategies to include ambiguous strategies. These are strategies generating, from the perspective of both players, subjective uncertainty about which statistical experiment will be used to generate the message. Our second key departure is that the receiver (and possibly the sender) treats this uncertainty as ambiguity and is ambiguity averse, and that the receiver is dynamically consistent. The main motivation for dynamic consistency is that we want to study whether and how the sender benefits from ambiguous communication even if the channel of dynamic *in*consistency with the receiver's given ex-ante preferences — the channel at work in nearly all previous literature on mechanism or information design with ambiguity — is shut down.

Would the sender ever benefit from intentionally using an ambiguous communication strategy? We show that even with a consistent receiver the answer can be yes, and provide understanding of the circumstances under which this occurs and the nature of these beneficial strategies.

In which kind of persuasion games can ambiguity benefit the sender? We show that if an optimal Bayesian persuasion experiment can be written as a convex combination of two Pareto-ranked experiments (i.e., two experiments such that both players agree that the first is better than the second assuming action recommendations are followed), this is sufficient for an ambiguity-neutral sender to benefit from ambiguous communication (Theorem 2). These sufficient conditions are also necessary in a class of *belief-threshold*

[1] In Bayesian persuasion, it is without loss of generality to assume that the messages are action recommendations. We show (Proposition 1) that this remains true in our setting.

*games* (Section 6, Theorem 3). For such games, these conditions are easy to check, and are generically satisfied when there are three or more actions. In these games, the sender has state-independent preferences and there is a *default* action that is strictly worse for the sender than any other action. The receiver has a distinct most-preferred action for each state. For all but the default action, an action is expected payoff maximizing for the receiver only when their belief on the corresponding state exceeds a threshold larger than one-half; otherwise, the default action is better. These preferences capture a variety of economic settings studied in the communication literature, for example, seller-buyer and lobbyist-legislator interactions, see e.g., (Lipnowski and Ravid, 2020).

For such games, whenever an ambiguity-neutral sender can benefit from ambiguous communication, a particularly simple ambiguous communication strategy can improve over the optimal Bayesian persuasion strategy: take two Pareto-ranked experiments whose convex combination gives the optimal Bayesian persuasion experiment and make them the support of the ambiguous strategy. The source of the gain is in allowing the Pareto-superior experiment to be used more often than in Bayesian persuasion. The role of the ambiguity is to ensure that the receiver still obeys the recommendations (i.e., the ambiguous experiment is *obedient*). An ambiguity-averse receiver desires to hedge more than an ambiguity-neutral receiver against the possibility that recommendations are coming from the inferior experiment. This allows the superior experiment to be used more often without inducing the receiver to depart from the recommended actions. An additional feature of this ambiguous strategy is that even the receiver's payoff improves on their payoff under the optimal Bayesian persuasion strategy (Cheng et al., 2026, Theorem 4 (ii)).

The sender's benefit from ambiguity beyond such threshold games mentioned above is also linked to Pareto-ranked convex combinations: the existence of a convex splitting of some (unambiguous) obedient experiment into two Pareto-ranked experiments is necessary for ambiguous communication to benefit the sender (Theorem 1). This condition fails in any persuasion game in which the receiver has only two available actions. Thus it is *never* possible for the sender to gain from ambiguous communication in such games (Corollary 1). In the same class of belief-threshold games, these necessary conditions are also sufficient.

We also analyze the sender's optimal persuasion strategy. We find that Pareto-ranked experiments are an essential feature of any optimal persuasion (Theorem 4). Furthermore, we provide a concavification-like characterization (Proposition 3 and Corollary 2) of the optimum.

We close this section with a brief discussion of a few closely related papers. A more extensive discussion can be found in Sections 8.1 and 8.6. Beauchêne et al. (2019) (BLL henceforth) were first to study strategic use of ambiguous communication in persuasion (see also Cheng (2022)). The key difference in assumptions between BLL and our paper is how the receiver best responds given the sender's ambiguous experiment. We assume the receiver chooses an ex-ante optimal message-contingent strategy. Equivalently, we assume the receiver behaves as in a sequentially optimal equilibrium. This corresponds (Hanany et al., 2020, Theorem 1) to the receiver using a simple dynamically consistent belief updating rule and maximizing their interim preferences given each message (see Remark 2 in Section 3.3 for updating details). BLL assume the receiver chooses, given each message, actions maximizing interim preferences formed using a belief updating rule that leads to dynamic inconsistency with their ex-ante preference. Thus, a key contribution of our paper is establishing and analyzing benefits of ambiguous persuasion that do not stem from receiver's behavior that is suboptimal with respect to their given ex-ante preferences. In Section 8.1, we further compare with BLL, including establishing a sharp contrast in when and how the sender can gain from ambiguous communication in a class of belief-threshold games depending on whether the receiver behaves as in our model or as in BLL.[2] The bulk of BLL's analysis imposes the infinitely ambiguity-averse extreme for both the sender and receiver – a polar case of our model, though they show that their approach extends more broadly. Cheng (2025) shows that all benefits from ambiguous communication identified by BLL in the case of such a sender disappear if the receiver is assumed, as in our paper, to maximize their given ex-ante preference. In light of Cheng (2025)'s result, it is essential that we allow at least the sender to be less than infinitely ambiguity averse for benefits from ambiguous communications to possibly exist. Our analysis allows for varying degrees of ambiguity aversion for both the sender and the receiver.

To illustrate our findings and provide intuition, we turn to an introductory example.

# 2 An Introductory Example

There is a sender and a receiver, three actions $a_1, a_2$ and $a_3$, and two payoff-relevant states $\omega_1$ and $\omega_2$, with equal prior probabilities $p = (1/2, 1/2)$. The sender influences the action

[2]For discussion and approaches to dynamic consistency issues in decision-making under ambiguity more broadly see e.g., Hanany and Klibanoff (2007, 2009) and Siniscalchi (2011). Hanany et al. (2020) address dynamic games of incomplete information with ambiguity averse players.

the receiver takes with the release of information. The payoffs are:

| $(u_s, u_r)$ | $a_1$ | $a_2$ | $a_3$ |
|---|---|---|---|
| $\omega_1$ | $1, 1$ | $-1, -1$ | $-4, 2$ |
| $\omega_2$ | $0, 0$ | $2, 2$ | $-4, -4$ |

Table 1: Payoff table (first coordinate is the sender's payoff)

Both the sender and receiver prefer $a_2$ in state $\omega_2$. The receiver prefers $a_3$ in state $\omega_1$, while the sender prefers $a_1$ in that state. This is the conflict of interest in this example. The receiver prefers $a_1$ when their beliefs about $\omega_2$ are intermediate (i.e., in $[1/5, 1/2]$), $a_2$ when their beliefs are higher than $1/2$, and $a_3$ when they are lower than $1/5$. The sender's challenge is to choose a communication strategy to better coordinate the receiver's play of actions $a_1$ and $a_2$ with the state, without inducing action $a_3$.[3]

Under Bayesian persuasion, the sender's optimal payoff is the concavification of their indirect utility at the prior (Kamenica and Gentzkow, 2011). In our example, this value is $5/4$ as illustrated in Figure 1. The figure plots the receiver's expected payoff associated with each of the three actions as dotted lines – each line is labelled with its action; the sender's indirect utility, i.e., the expected payoff the sender obtains when the receiver best responds, as a thick solid curve, and its concavification as a thick dashed curve.

It is immediate to verify that the experiment $\sigma^{BP}$ defined in Figure 1 attains this optimal payoff using the messages "$a_1$" and "$a_2$". The message "$a_2$" reveals that the state is $\omega_2$. Intuitively, since the preferences are perfectly aligned when the state is $\omega_2$, the sender wants the receiver to learn it. At the same time, the sender does not want the receiver to be too pessimistic about $\omega_2$ when the state is $\omega_1$, as the receiver chooses $a_3$ at all beliefs less than $1/5$ on $\omega_2$. The optimal experiment $\sigma^{BP}$ balances these two forces by making the likelihood of message "$a_1$" in state $\omega_2$ just large enough to induce action $a_1$. We note that the experiment $\sigma^{BP}$ is *canonical*, that is, the messages are action recommendations, and, in the example, it is also *obedient*, that is, the receiver finds it optimal to obey the recommendations. In Bayesian persuasion, the restriction to canonical and obedient experiments is without loss, and we prove (Proposition 1) that this remains true in our generalization.

Now, suppose that the sender can design *ambiguous* experiments. These are communication strategies that leave some uncertainty about which statistical experiment will gener-

[3]See the working paper version (Cheng et al., 2026) for an interpretation of this example in terms of bank stress testing.

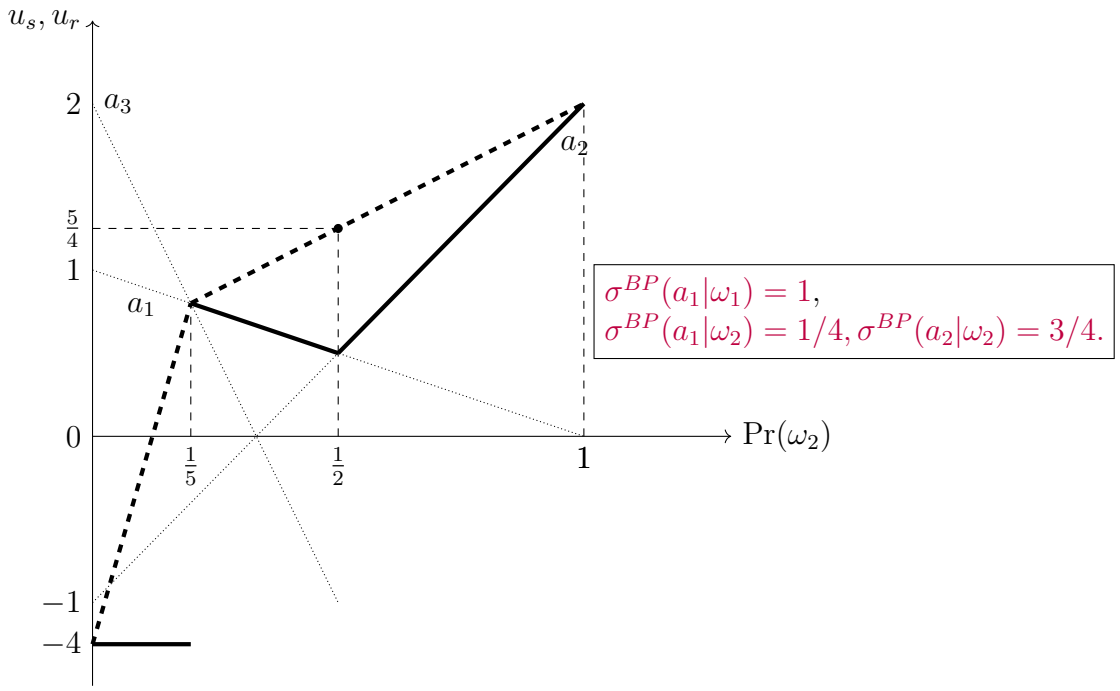



Figure 1: Sender's indirect utility (thick curve) and its concavification (thick dashed curve)

ate the messages, and, in this sense, may not completely pin-down the message likelihoods. We model ambiguous experiments as subjective (finite-support) distributions, $\mu$, over statistical experiments. The key implicit assumption is that the sender and receiver share this common view of the uncertainty over these experiments, would use the weights $\mu$ if they were ambiguity neutral, and that the receiver (and possibly the sender as well) is ambiguity averse and treats this uncertainty about the experiments as ambiguity. Conceptually, for example, such ambiguity could be generated by determining which experiment will be used by methods such as the outcome of artificial devices like draws from an Ellsberg urn, or by a third party who selects and conducts the experiment under public instructions by the sender that leave some room for discretion. By varying the artificial device or the instructions to the third party, the sender can vary $\mu$. Our theory is agnostic about why there is a common view of this uncertainty. Such a view may arise from shared (but limited) historical data or experience with the third party or, particularly in the case of artificially generated ambiguity, from symmetry or other logical considerations. Alternatively, the common-view assumption can simply be regarded as a convenient modeling benchmark.

We now illustrate (see Figure 2) how ambiguous experiments can benefit an ambiguity-

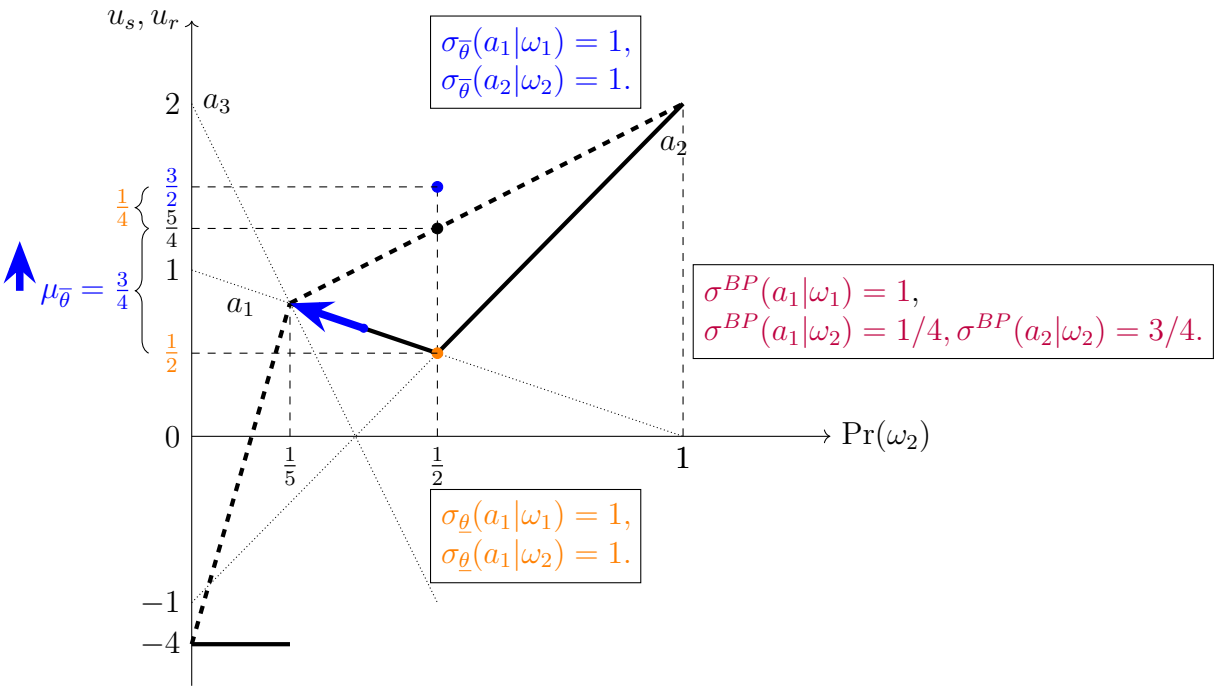



Figure 2: Construction of the ambiguous experiment

neutral (SEU) sender in the example.[4] Consider an ambiguous experiment such that one of two experiments $\sigma_{\underline{\theta}}$ and $\sigma_{\overline{\theta}}$ (see the figure) could be generating the messages, where $\theta$ is used to index experiments in the support of $\mu$. The experiment $\sigma_{\underline{\theta}}$ is uninformative, while the experiment $\sigma_{\overline{\theta}}$ is fully informative. The interpretation of the message $a_2$ is unambiguous: the receiver learns that the state is $\omega_2$. The interpretation of the message $a_1$ is, however, ambiguous: either it means that the state is $\omega_1$ (if $\sigma_{\overline{\theta}}$ generated the message) or it is uninformative (if $\sigma_{\underline{\theta}}$ generated the message). The associated expected payoff profiles for these two experiments assuming action recommendations are obeyed are $(u_s, u_r) = (3/2, 3/2)$ for $\sigma_{\overline{\theta}}$ and $(u_s, u_r) = (1/2, 1/2)$ for $\sigma_{\underline{\theta}}$. Observe that $(3/4)\sigma_{\overline{\theta}} + (1/4)\sigma_{\underline{\theta}} = \sigma^{BP}$ – we call such a configuration a splitting of $\sigma^{BP}$. Thus, if an ambiguity-neutral sender assigns $\mu$-probability more than $3/4$ to the experiment $\sigma_{\overline{\theta}}$ generating the messages, their expected payoff would be strictly higher than their Bayesian persuasion payoff.

We now argue that the sender can simultaneously assign $\mu$-probability above $3/4$ to $\sigma_{\overline{\theta}}$ *and* guarantee that an ambiguity-averse receiver is willing to obey the action recom-

[4]Similar arguments remain valid as long as the sender is not too ambiguity averse. See the discussion of sufficient conditions for the sender to benefit from ambiguous communication in Section 5, including a demonstration that the sender may benefit even if they are more ambiguity averse than the receiver.

mendations. If the receiver were ambiguity-neutral, they would evaluate the ambiguous experiment using $\mu$, just like the sender. Ambiguity aversion, however, leads the receiver to behave as if they shift some of the weight on $\sigma_{\overline{\theta}}$ to the lower expected payoff experiment $\sigma_{\underline{\theta}}$. We refer to this "as if" weighting as the receiver's *effective measure* on the experiments. This relative pessimism of the effective measure reflects the value an ambiguity-averse receiver places on robustness across probabilities over the experiments. As a result, the receiver chooses a strategy whose expected payoff is less sensitive across experiments than the ambiguity-neutral best response. The more ambiguity averse the receiver, the stronger this relative pessimism.

Since the receiver is obedient when facing $\sigma^{BP}$, the receiver is obedient when their effective measure places weight $3/4$ on $\sigma_{\overline{\theta}}$. In the figure, the thick arrow moving along the sender's indirect utility curve indicates the movement of the receiver's effective weight on $\sigma_{\overline{\theta}}$ towards $3/4$ from below as the probability $\mu_{\overline{\theta}}$ on $\sigma_{\overline{\theta}}$ increases above $3/4$ (along the thick arrow next to $\mu_{\overline{\theta}}$). Thus, the ambiguous communication strategy allows the sender to place more weight on the better experiment $\sigma_{\overline{\theta}}$, while maintaining obedience, than would be possible with unambiguous communication. This is how ambiguous communication provides benefits.

An important observation is that $\sigma_{\overline{\theta}}$ and $\sigma_{\underline{\theta}}$ are Pareto-ranked – both players prefer $\sigma_{\overline{\theta}}$ assuming action recommendations are followed. If they were not Pareto-ranked, then ambiguity aversion would push the receiver's effective measure in a direction that would *hurt* rather than help the sender – for example, if the receiver thought instead that $\sigma_{\underline{\theta}}$ were the better experiment this would cause their effective weight on $\sigma_{\overline{\theta}}$ to *exceed* $\mu_{\overline{\theta}}$, leading any $\mu_{\overline{\theta}} \geq 3/4$ to cause the receiver to defect to the action $a_3$ when the message is $a_1$. The ambiguity generated using Pareto-ranked experiments serves to beneficially misalign the (endogenous) effective beliefs of the sender and receiver.[5]

The remainder of the paper is organized as follows. The next section presents the model and two key preliminary results – a revelation principle and an incentive-compatibility lemma. Main results are in Sections 4 through 7. Section 8 contains further discussion. Proofs are in the Appendix. Additional material and proofs of auxiliary results from the Appendix may be found in the working paper version (Cheng et al., 2026).

[5] A player's effective measure is endogenous in the sense that it depends on their strategy's expected payoff profile across the possible experiments in addition to their ambiguity aversion. This endogeneity of the players' misalignment distinguishes our model from models with exogenously fixed heterogeneous priors, e.g., Alonso and Câmara (2016), Laclau and Renou (2017) and Galperti (2019).

# 3 Modeling Persuasion with Ambiguous Communication

## 3.1 The Model

We consider a persuasion game between a sender and a receiver, where the sender can choose *ambiguous* experiments if they wish. There is a finite set $\Omega$ of payoff-relevant states $\omega$, with common full-support prior $p \in \Delta(\Omega)$. There is a finite set $A$ of actions the receiver can choose from. If the receiver chooses $a \in A$, the payoff to the sender (resp., receiver) is $u_s(a, \omega) \in \mathbb{R}$ (resp., $u_r(a, \omega) \in \mathbb{R}$), when the state is $\omega$. A *statistical experiment* is a map $\sigma$ from $\Omega$ to $\Delta(M)$ where $M$ is a finite set of messages, and we write $\sigma(m|\omega)$ for the probability of $m$ given $\omega$. Let $\Sigma$ denote the set of all statistical experiments.

We model an *ambiguous experiment* as a finite-support distribution over statistical experiments, and denote it by $\mu \in \Delta_f(\Sigma)$, where the uncertainty described by $\mu$ about which statistical experiment is generating the messages is viewed as ambiguity by the players. We index the experiments in the support of $\mu$ (denoted $\text{supp}(\mu)$), and use $\theta$ as a generic index. Thus $\sigma_\theta$ is an experiment in $\text{supp}(\mu)$ and $\mu_\theta$ is the probability $\mu$ assigns to $\sigma_\theta$. The sender may choose any ambiguous experiment.[6,7] Henceforth, whenever we use the term "experiment" without a modifier, it refers to a standard, unambiguous statistical experiment. Our model enlarges the sender's strategy space relative to the standard Bayesian persuasion model in the sense that any experiment $\sigma$ can be viewed as an ambiguous experiment with a degenerate $\mu$ concentrated on $\sigma$, i.e., $\mu = \delta_\sigma$. Of special interest in some of our later constructions are *binary ambiguous experiments*, those in which $\text{supp}(\mu)$ is a pair of experiments. For $\kappa \in (0, 1)$ and $\overline{\sigma}, \underline{\sigma} \in \Sigma$, we denote the binary ambiguous experiment $\kappa\delta_{\overline{\sigma}} + (1 - \kappa)\delta_{\underline{\sigma}}$ by $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$.

**Remark 1.** What is a possible foundation for the availability of these ambiguous experiments $\mu$ that are treated as commonly perceived ambiguity? One possibility is that the desired $\mu$ results from the sender conditioning the statistical experiment on the realization of a suitably chosen finite partition of a *source of ambiguity*. The idea is that the sender commits to allowing the realization of external, payoff-irrelevant ambiguous events to determine which experiment will generate the messages. Formally, a source of ambiguity is a

[6] It is without loss to assume that all statistical experiments in $\text{supp}(\mu)$ share the same message space.

[7] The sender chooses and commits to the ambiguous experiment *before* the statistical experiment and $\omega$ are realized. This parallels standard Bayesian persuasion where the sender chooses and commits to $\sigma$ before $\omega$ is realized. See Section 8.2 for a discussion of how things would change under the alternative assumption that it is common knowledge that the sender privately learns the realization of $\mu$ before committing to it.

continuum, $\mathcal{A}$, of payoff-irrelevant ambiguous states, $\alpha$, along with a continuous $\rho \in \Delta(\mathcal{A})$ that represents the common view of uncertainty over $\mathcal{A}$ that an ambiguity-neutral (SEU) player would use to compute their expected utility. Let $B$ be any finite partition of $\mathcal{A}$ and define $\tilde{\mu} \in \Delta(B)$ by $\tilde{\mu}(b) \equiv \int_{\alpha \in b} \rho(\alpha) d\alpha$ for each $b \in B$. By richness of $\mathcal{A}$ and continuity of $\rho$, *any* finitely-supported distribution $\mu$ over experiments can be induced by the choice of some finite partition $B$ together with some mapping from $B$ to experiments.

More informally, and more generally, our model of ambiguous experiments is intended to capture any setting in which the sender's instructions give rise to commonly-viewed ambiguity about which experiment will generate the messages (i.e., about the likelihood-implication of the messages) and in which, by varying their choice of instructions, the sender has the ability to influence that common view. There might be some tension between the sender's ability to finely tune the uncertainty about which experiment will be used while still viewing that uncertainty as ambiguity. However, we think that both features arise, for example, in settings where the sender instructs a third party to generate the message, while leaving them with some discretion to adapt the experiment based on their own private information. By varying the extent and direction of this discretion, the sender controls both the degree and nature of the resulting ambiguity. □

We analyze the receiver's behavior from the perspective of their ex-ante preferences, that is, we assume that the receiver observes the sender's choice of $\mu$ and then chooses a strategy $\tau : M \to \Delta(A)$ to maximize the receiver's ex-ante preference.[8] We further assume that $\tau$ is indeed carried out after the message is observed. This follows from either assuming that the receiver can commit to $\tau$, or that the receiver is dynamically consistent.[9] Recall from the Introduction that this allows us to shut down the channel of dynamic *in*consistency that is at work in nearly all previous literature on mechanism or information design with ambiguity. We refer the interested reader to Section 8.1 for more discussion on this point.

We write $u_i(\sigma, \tau)$ for the expected payoff of player $i \in \{s, r\}$ when the experiment generating the message $m$ is $\sigma$ and the receiver's strategy is $\tau$, that is,

$$u_i(\sigma, \tau) = \sum_{\omega, m, a} p(\omega)\sigma(m|\omega)\tau(a|m)u_i(a, \omega). \tag{1}$$

To isolate and clarify the role of intentional ambiguous communication, we work in a stylized environment where ambiguity is not payoff-relevant unless it becomes so through

[8] For compactness, this notation suppresses the allowed dependence of $\tau$ on $\mu$.

[9] For a simple receiver's updating rule that guarantees dynamic consistency, see Remark 2 in Section 3.3.

the sender's strategic choice of a non-degenerate $\mu$. In particular, the payoff-relevant events $\omega$ and any randomization over messages induced by an experiment are viewed as unambiguous.[10] It follows that a message $m$ is viewed as ambiguous by the sender and receiver only if experiments in the support of $\mu$ generate $m$ with different positive likelihoods. These different likelihoods may lead the expected payoff $u_i(\sigma_\theta, \tau)$ to vary with $\theta$.

How does such ambiguity enter the sender's and receiver's preferences? We assume, as in the smooth ambiguity model (Klibanoff et al., 2005), that player $i$ evaluates the strategy profile $(\mu, \tau)$ as

$$U_i(\mu, \tau) = \phi_i^{-1}\left(\sum_\theta \mu_\theta \phi_i(u_i(\sigma_\theta, \tau))\right), \tag{2}$$

where $\phi_i : \mathbb{R} \to \mathbb{R}$ is a weakly concave and strictly increasing and differentiable function. An affine $\phi_i(\cdot)$ corresponds to ambiguity neutrality, in which case the preferences reduce to SEU with belief $\mu$. Greater concavity of $\phi_i(\cdot)$ corresponds to greater ambiguity aversion. We do not consider ambiguity loving behavior, as this would build-in a direct preference benefit from ambiguous communication, while with ambiguity aversion, ambiguous communication can only be valuable if it has a strategic benefit.

When the expected payoff $u_i(\sigma_\theta, \tau)$ varies with $\theta$, an ambiguity-averse player $i$ responds to $\mu$ *as if* they are ambiguity neutral and using an *effective* measure over $\Sigma$ that is more pessimistic than $\mu$. While (2) delivers a particularly tractable formula for this effective measure (Section 3.3), the key qualitative properties of comparative pessimism and endogeneity would apply to the effective measure under many models of ambiguity averse preferences.

The following maxmin expected utility (MEU) objective can be seen as an appropriate limit of (2) as ambiguity aversion goes to infinity (Klibanoff et al., 2005, Proposition 3):

$$U_i^{MEU}(\mu, \tau) = \min_{\sigma_\theta \in \text{supp}(\mu)} u_i(\sigma_\theta, \tau).$$

Such an MEU receiver together with an ambiguity-neutral sender is considered in Appendix B of the working paper version (Cheng et al., 2026).[11]

Writing $BR(\mu)$ for the set of best replies of the receiver (i.e., the maximizers of $U_r(\mu, \tau)$

[10] Section 8.4 discusses extensions to environments with pre-existing ambiguity about $\omega$.

[11] Key differences from our main analysis are that (1) the sender never needs more than binary ambiguous experiments to approach their optimal payoff, and (2) the sender can approach their optimal payoff without providing the receiver any benefits from communication, and this may lead the receiver's payoff to drop discontinuously when passing to the MEU limit.

with respect to $\tau$), the sender's problem is:

$$\begin{array}{l} \max_{(\mu,\tau)} U_s(\mu,\tau), \\ \text{subject to } \tau \in BR(\mu). \end{array} \tag{$\mathcal{P}$}$$

Observe that the sender's Bayesian persuasion problem (Kamenica and Gentzkow, 2011) corresponds to the special case of our model where the sender is restricted to choosing an experiment:

$$\begin{array}{l} \max_{(\sigma,\tau)} u_s(\sigma,\tau), \\ \text{subject to } \tau \in br(\sigma), \end{array} \tag{$\mathcal{P}^{BP}$}$$

where $br(\sigma)$ denotes the set of best replies to $\sigma$, i.e., the maximizers of $u_r(\sigma,\tau)$ with respect to $\tau$. Let $u_s^{BP}$ denote the value of ($\mathcal{P}^{BP}$), i.e., the sender's payoff at a solution to ($\mathcal{P}^{BP}$).

**Definition 1.** ***Ambiguous communication benefits the sender*** *if the value of* ($\mathcal{P}$) *is strictly higher than* $u_s^{BP}$.

Our analysis will focus on when and how ambiguous communication may benefit the sender compared to the standard, unambiguous case of Bayesian persuasion. We also present results on optimal persuasion with ambiguous communication (the solution to ($\mathcal{P}$)) and its properties.

We next present two preliminary results – a revelation principle and a characterization of incentive compatibility for ambiguous experiments – that play a key role in our analysis.

## 3.2 A Revelation Principle

**Definition 2.** *An ambiguous experiment* $\mu$ *is* ***canonical*** *if* $M = A$.

We write $\tau^* : A \to \Delta(A)$ for the receiver's obedient strategy, that is, $\tau^*(a|a) = 1$ for all $a$. We will refer to any canonical ambiguous experiment that induces such obedience as itself obedient.

**Definition 3.** *A canonical ambiguous experiment* $\mu$ *is* ***obedient*** *if* $\tau^* \in BR(\mu)$.

We start with a preliminary observation: a revelation principle holds – for payoff purposes, it is without loss of generality to restrict attention to canonical and obedient ambiguous experiments.

**Proposition 1.** *For any $(\mu, \tau)$ such that $\tau \in BR(\mu)$, there exists a canonical and obedient ambiguous experiment $\mu^*$ such that $U_i(\mu, \tau) = U_i(\mu^*, \tau^*)$ for all $i \in \{s, r\}$. Furthermore, for each $\sigma_\theta \in$ supp$(\mu)$, $\mu(\{\sigma \in \Sigma : u_i(\sigma, \tau) = u_i(\sigma_\theta, \tau)\}) = \mu^*(\{\sigma \in \Sigma : u_i(\sigma, \tau^*) = u_i(\sigma_\theta, \tau)\})$ for all $i \in \{s, r\}$.*

It is well-known that such a revelation principle holds in the persuasion setting without ambiguity. However, one might think of at least two reasons why the same might not be true in our environment. First, an ambiguity averse receiver might strictly prefer a mixed strategy to any pure strategy for hedging reasons in the face of ambiguity. How can the receiver's desire to mix be reconciled with the revelation principle, which states that it is without loss of generality to have the receiver play the pure strategy $\tau^*$? The answer is that any mixing the receiver might desire to do can always be emulated through the use of experiments that mix over action recommendations. It is the standard Bayesian persuasion assumption of sender's commitment that guarantees that this emulation is always possible. Second, dynamic inconsistency, generated by ambiguity aversion together with assumptions on updating, is the main channel leading to the failure of such a revelation principle in existing literature. As previously mentioned, we shut down this channel by modeling the receiver as maximizing $U_r(\mu, \tau)$, their ex-ante payoff from $(\mu, \tau)$, or, equivalently, using dynamically consistent updating (see Remark 2 in Section 3.3).

From here on, we restrict attention to canonical experiments, and represent incentive compatibility via obedience. When writing payoffs, $U_i$ and $u_i$, if we omit the second argument, it should be understood that $\tau = \tau^*$.

Given the prominent role obedient experiments play, understanding when obedience holds is important. We next present a characterization of such incentive compatibility for ambiguous experiments.

## 3.3 Incentive Compatibility and Effective Measure

We present a result we repeatedly use linking the obedience of an *ambiguous* experiment to the obedience of an *unambiguous* experiment derived from that ambiguous experiment. To state the result, we need the following definition:

**Definition 4.** *Given an ambiguous experiment $\mu$, the receiver's* ***effective measure*** $em^\mu \in$

$\Delta_f(\Sigma)$ *is given by:*

$$em^{\mu}_{\theta} := \frac{\mu_\theta \phi_r^{'}(u_r(\sigma_\theta))}{\sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi_r^{'}(u_r(\sigma_{\tilde{\theta}}))}, \text{ for all } \sigma_\theta \in \text{supp}(\mu). \tag{3}$$

The effective measure $em^{\mu}$ is a probability measure with the same support as $\mu$. It is equal to $\mu$ when the receiver is ambiguity neutral (i.e., $\phi_r$ is affine), and is more pessimistic than $\mu$ for an ambiguity averse receiver (i.e., $\phi_r$ concave). Pessimism here means shifting weight toward $\sigma_\theta$ yielding lower expected receiver's payoffs, i.e., if $u_r(\sigma_\theta) < u_r(\sigma_{\theta'})$, then $em^{\mu}_{\theta}/em^{\mu}_{\theta'} > \mu_\theta/\mu_{\theta'}$. Notice also that the effective measure of a given $\sigma_\theta$ depends on the specification of the ambiguous experiment for all $\sigma_\theta \in \text{supp}(\mu)$.

The next result states that $\tau^*$ is the receiver's best response to the ambiguous experiment $\mu$ if, and only if, it is a best response to the experiment, $\sigma^*$, defined below as the convex combination of experiments with weights given by the receiver's effective measure. In such a case, we say that these experiments form a *splitting* of $\sigma^*$ with splitting weights given by the effective measure.

**Lemma 1.** *The ambiguous experiment $\mu$ is obedient if, and only if, the (unambiguous) experiment $\sigma^*$ is obedient, where*

$$\sigma^* = \sum_{\theta} em^{\mu}_{\theta} \sigma_\theta.$$

Lemma 1 follows from the first-order conditions of the receiver's maximization problem $\max_\tau U_r(\mu, \tau)$, evaluated at $\tau^*$. Some intuition is that obedience will differ from the best response to the experiment $\sum_\theta \mu_\theta \sigma_\theta$ in that it will be better hedged against uncertainty about the weights on the experiments. In our introductory example, for instance,

$$\begin{aligned} 2 = u_r(\sigma_{\overline{\theta}}, br(\mu_{\overline{\theta}}\sigma_{\overline{\theta}} + (1-\mu_{\overline{\theta}})\sigma_{\underline{\theta}})) &> u_r(\sigma_{\overline{\theta}}) = 3/2 \\ &> u_r(\sigma_{\underline{\theta}}) = 1/2 > u_r(\sigma_{\underline{\theta}}, br(\mu_{\overline{\theta}}\sigma_{\overline{\theta}} + (1-\mu_{\overline{\theta}})\sigma_{\underline{\theta}})) = -1, \end{aligned}$$

showing that the obedience strategy $\tau^*$ is hedged against the uncertainty about $\mu$ more than the strategy $br(\mu_{\overline{\theta}}\sigma_{\overline{\theta}} + (1-\mu_{\overline{\theta}})\sigma_{\underline{\theta}})$. Thus the relative pessimism of the effective measure reflects the fact that an ambiguity-averse receiver values such hedging.

Lemma 1 gives rise to the following interpretation of the receiver's effective measure: It is an "ambiguity-neutral measure supporting obedience" in the sense that *if* the receiver

were ambiguity neutral, the ambiguous experiment $em^\mu$ would be obedient.[12]

**Remark 2** (Dynamically consistent updating)**.** These properties of the effective measure also give rise to an updating implementation of the receiver's ex-ante optimality. A receiver who forms their effective posterior after observing a message by updating using Bayes' rule but with the effective measure in place of $\mu$ in the Bayes' formula will be dynamically consistent. Formally, such a receiver updates $em^\mu_\theta p(\omega)\sigma_\theta(a|\omega)$, their effective prior over $\Sigma \times \Omega \times A$, to $\frac{em^\mu_\theta p(\omega)\sigma_\theta(a|\omega)}{\sum_{\tilde{\theta},\tilde{\omega}} em^\mu_{\tilde{\theta}} p(\tilde{\omega})\sigma_{\tilde{\theta}}(a|\tilde{\omega})}$, their effective posterior over $\Sigma \times \Omega$ after observing message $a$. □

Finally, for later reference, observe that we can invert (3) to express $\mu$ as a function of the effective measure it generates:

$$\mu_\theta = \frac{em^\mu_\theta / \phi_r'(u_r(\sigma_\theta))}{\sum_{\tilde{\theta}} em^\mu_{\tilde{\theta}} / \phi_r'(u_r(\sigma_{\tilde{\theta}}))}. \tag{4}$$

# 4 A Necessary Condition for Ambiguous Communication to Benefit The Sender

This section establishes a necessary condition for ambiguous communication benefits the sender, and shows that this condition implies no benefit for the sender when the receiver has only two available actions – a common assumption in many examples and applications in the literature. The necessary condition uses the concept of Pareto-ranked experiments and splittings into such experiments. Two experiments are *Pareto-ranked* if the sender and receiver agree on their strict ranking under the assumption of obedience.

**Definition 5.** *Two experiments $\overline{\sigma}$ and $\underline{\sigma}$ are **weakly Pareto-ranked** if either the two inequalities*

$$u_s(\overline{\sigma}) \geq u_s(\underline{\sigma}) \text{ and } u_r(\overline{\sigma}) \geq u_r(\underline{\sigma}), \tag{5}$$

*hold or both reversed inequalities hold. They are **Pareto-ranked** if the same holds true with strict inequalities.*

[12] For the benefits of working with distributions over statistical experiments rather than distributions over posteriors, see Section 8.5.

*A **Pareto-ranked splitting** of the experiment $\sigma$ is a triple $(\overline{\sigma}, \underline{\sigma}, \lambda)$ such that (i) $\lambda\overline{\sigma} + (1-\lambda)\underline{\sigma} = \sigma$, (ii) $\lambda \in (0,1)$, and (iii) (5) holds with strict inequalities, i.e., $\overline{\sigma}$ and $\underline{\sigma}$ are Pareto-ranked.*

The following theorem shows that the existence of Pareto-ranked experiments (in particular, ones having sender's expected payoffs bracketing $u_s^{BP}$) splitting an obedient experiment is necessary for ambiguous communication to benefit the sender:

**Theorem 1.** *Ambiguous communication benefits the sender only if $\phi_r$ is not affine and there exists a Pareto-ranked splitting, $(\overline{\sigma}, \underline{\sigma}, \lambda)$, of an obedient experiment $\hat{\sigma}$ such that $u_s(\overline{\sigma}) > u_s^{BP} > u_s(\underline{\sigma})$.*

**Example 1** (Verifying that the introductory example satisfies this necessary condition)**.** *In Figure 2 of the introductory example, $(\sigma_{\overline{\theta}}, \sigma_{\underline{\theta}}, 3/4)$ is a Pareto-ranked splitting of $\sigma^{BP}$ and $u_s(\sigma_{\overline{\theta}}) > u_s^{BP} > u_s(\sigma_{\underline{\theta}})$.*

**Remark 3** (Not necessary for $\hat{\sigma}$ to be an optimal Bayesian persuasion experiment)**.** The reader might wonder if a stronger version of the theorem that requires $\hat{\sigma}$ to be an optimal Bayesian persuasion experiment holds. This is false. There are examples in which the sender benefits from ambiguous communication even though no Pareto-ranked splitting of any optimal Bayesian persuasion experiment exists (as is true, for instance, whenever all such experiments are efficient). □

The conditions in Theorem 1 are deceptively powerful: From these conditions alone, strict benefit from ambiguity can be ruled out for a simple yet important class of problems – those in which the receiver has a binary action space.

**Corollary 1.** *If the receiver has only two actions, the sender cannot benefit from ambiguous communication.*

An important step in the proof of Theorem 1 is showing that its existence condition is equivalent to the existence of Pareto-ranked experiments, $\sigma$ and $\sigma^*$ such that: (i) $\operatorname{supp}\sigma(\cdot|\omega) = \operatorname{supp}\sigma^*(\cdot|\omega)$ for all $\omega$, (ii) $u_s(\sigma, \tau^*) > u_s^{BP}$, and (iii) $\tau^* \in br(\sigma^*) \setminus br(\sigma)$. Given $\overline{\sigma}$, $\underline{\sigma}$, and $\hat{\sigma}$ from Theorem 1, letting $\sigma^* = \hat{\sigma}$ and $\sigma$ be a well-chosen mixture of $\overline{\sigma}$ and $\hat{\sigma}$ satisfies (i) - (iii). Conversely, given $\sigma$ and $\sigma^*$ satisfying (i) - (iii) there exists a Pareto-ranked splitting of the obedient $\sigma^*$ with $\overline{\sigma} = \sigma$ that satisfies $u_s(\overline{\sigma}) > u_s^{BP} > u_s(\underline{\sigma})$. Necessity is then established by the following argument: Suppose that there exists a solution $\tilde{\mu}$ to the sender's program ($\mathcal{P}$) that benefits the sender. Construct $\sigma$ and $\sigma^*$ satisfying (i) - (iii) by letting $\sigma = \sum_\theta \tilde{\mu}_\theta \tilde{\sigma}_\theta$ and $\sigma^* = \sum_\theta em_\theta^{\tilde{\mu}} \tilde{\sigma}_\theta$.

The intuition for Corollary 1 is as follows. From the above, ambiguous communication benefiting the sender implies the existence of a non-obedient experiment $\sigma$ that is Pareto-ranked higher than some obedient experiment $\sigma^*$. Intuitively, such an improvement for the receiver is possible only when $\sigma$ is more informative for the receiver and the benefit of this extra information outweighs the cost of not best responding. When there are only two actions, taking advantage of information requires best responding. To see this, note that not best responding implies either taking the same action always (and thus ignoring any information) or always doing the opposite of what is optimal for the receiver (which hurts more when there is more information). In contrast, when there are three or more actions, it becomes possible to have some beneficial responsiveness to information without going all the way to best responding. As we saw in the introductory example, this indeed can leave scope for possible improvements.

# 5 A Sufficient Condition for Ambiguous Communication to Benefit the Sender

This section shows that a strengthening of the necessary conditions described in Theorem 1 is sufficient for ambiguous communication to benefit the sender. Given strict ambiguity aversion of the receiver, existence of a Pareto-ranked splitting of an optimal Bayesian persuasion experiment, $\sigma^{BP}$, is sufficient for an ambiguity-neutral sender to benefit. This sufficiency extends to ambiguity-averse senders when paired with a condition giving a tight upper bound on a measure of the sender's ambiguity aversion. The measure of ambiguity aversion is the following notion of *probability premium*.

**Definition 6.** *Given $\phi$, $u$, and experiments $\overline{\sigma}$ and $\underline{\sigma}$ such that $u(\overline{\sigma}) > u(\underline{\sigma})$, the $((\overline{\sigma}, \underline{\sigma}), \lambda)$-**probability premium** required to compensate for replacing the unambiguous experiment $\sigma^* := \lambda\overline{\sigma} + (1-\lambda)\underline{\sigma}$ by the binary ambiguous experiment $\mu^{((\overline{\sigma},\underline{\sigma}),\lambda)}$, assuming obedience, is:*

$$\rho^{\phi,u}((\overline{\sigma}, \underline{\sigma}), \lambda) := \frac{\phi(u(\sigma^*)) - \lambda\phi(u(\overline{\sigma})) - (1-\lambda)\phi(u(\underline{\sigma}))}{\phi(u(\overline{\sigma})) - \phi(u(\underline{\sigma}))}.$$

The probability premium $\rho^{\phi,u}((\overline{\sigma}, \underline{\sigma}), \lambda)$ is exactly the $\phi$-payoff difference between $\sigma^*$ and the ambiguous experiment $\mu^{((\overline{\sigma},\underline{\sigma}),\lambda)}$, normalized to lie in $[0, 1]$. This premium is non-negative under ambiguity aversion, and is zero under ambiguity neutrality. Similar notions

of probability premium in the context of risk go back to at least Pratt (1964). Letting

$$\kappa = \lambda + \rho^{\phi,u}((\overline{\sigma},\underline{\sigma}),\lambda) \in [0,1],$$

$$U(\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}) = \phi^{-1}\Big(\kappa\phi(u(\overline{\sigma})) + (1-\kappa)\phi(u(\underline{\sigma}))\Big) = \phi^{-1}\Big(\phi(u(\sigma^*))\Big) = u(\sigma^*),$$

meaning that the premium $\rho^{\phi,u}((\overline{\sigma},\underline{\sigma}),\lambda)$ is exactly the increase in $\kappa$ above $\lambda$ needed to make the player indifferent between the ambiguous experiment $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$ and $\sigma^*$. Thus, assuming obedience, this premium is the smallest increase in the probability on $\overline{\sigma}$ required to compensate for exposure to the ambiguous experiment.

**Theorem 2.** *Let $\sigma^{BP}$ be an obedient experiment such that $u_s(\sigma^{BP}) = u_s^{BP}$. If there exists a Pareto-ranked splitting of $\sigma^{BP}$, $(\overline{\sigma},\underline{\sigma},\lambda)$, for which $\rho^{\phi_s,u_s}((\overline{\sigma},\underline{\sigma}),\lambda) < \kappa - \lambda$ with*

$$\kappa = \frac{\lambda\phi_r'(u_r(\underline{\sigma}))}{\lambda\phi_r'(u_r(\underline{\sigma})) + (1-\lambda)\phi_r'(u_r(\overline{\sigma}))}, \tag{6}$$

*then ambiguous communication benefits the sender.*

*Furthermore, $\rho^{\phi_s,u_s}((\overline{\sigma},\underline{\sigma}),\lambda)$ is increasing in the sender's ambiguity aversion, and $\kappa$ given by (6) is increasing in the receiver's ambiguity aversion.*

Therefore, whenever a Pareto-ranked splitting of an optimal Bayesian persuasion experiment exists, an ambiguity-neutral sender benefits from ambiguous communication if $0 < \kappa - \lambda$. The latter is equivalent to $\phi_r'(u_r(\underline{\sigma})) > \phi_r'(u_r(\overline{\sigma}))$, i.e., the receiver is, within this range of payoffs, not everywhere ambiguity neutral, which is always satisfied if $\phi_r$ is strictly concave.

The key to Theorem 2 is to consider the binary ambiguous experiment $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$. First, observe that (6) is equivalent to $em_{\overline{\theta}}^{\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}} = \lambda$, ensuring that the obedience of $\sigma^{BP}$ extends to the ambiguous experiment $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$ by Lemma 1. Second, the inequality $\rho^{\phi_s,u_s}((\overline{\sigma},\underline{\sigma}),\lambda) < \kappa - \lambda$ characterizes exactly when the sender strictly prefers using the ambiguous experiment $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$ rather than $\sigma^{BP}$. The source of the economic gain from this ambiguous communication is the greater use, as measured by $\kappa - \lambda$, of the Pareto-better experiment $\overline{\sigma}$ while maintaining obedience. This gain has to be netted-off against the probability premium, which encapsulates the cost due to the player's own ambiguity aversion of the exposure to ambiguity from the ambiguous experiment.

Notice that the gain here is not a story of a less ambiguity averse sender gaining

by insuring a more ambiguity averse receiver. In particular, it is possible to satisfy the sufficient conditions in Theorem 2 even when the sender is more ambiguity averse than the receiver, as measured by their respective probability premiums. To see this, suppose that $\phi_r$ is strictly concave so that $\kappa > \lambda$ and let $(\overline{\sigma}, \underline{\sigma}, \lambda)$ be a Pareto-ranked splitting of $\sigma^{BP}$. Algebra and concavity can be used to show that $\rho^{\phi_r,u_r}((\overline{\sigma}, \underline{\sigma}), \lambda) < \kappa - \lambda$, and thus for any sender such that $\rho^{\phi_s,u_s}((\overline{\sigma}, \underline{\sigma}), \lambda) \in (\rho^{\phi_r,u_r}((\overline{\sigma}, \underline{\sigma}), \lambda), \kappa - \lambda)$ the sufficient conditions in Theorem 2 are satisfied while the sender has the higher probability premium. For instance, in the introductory example, setting $\phi_r(x) = (x+4)^{1/2}$ and $\phi_s(x) = (x+4)^{1/4}$, and using the Pareto-ranked splitting $(\sigma_{\overline{\theta}}, \sigma_{\underline{\theta}}, 3/4)$ of $\sigma^{BP}$ as in Figure 2, yields $\rho^{\phi_r,u_r}((\sigma_{\overline{\theta}}, \sigma_{\underline{\theta}}), 3/4) < 0.01 < \rho^{\phi_s,u_s}((\sigma_{\overline{\theta}}, \sigma_{\underline{\theta}}), 3/4) < \kappa - 3/4 \approx 0.018$, so that the sufficient conditions in Theorem 2 are satisfied while the sender is more ambiguity averse than the receiver.

The comparative statics in the final part of the theorem deliver robustness in several respects, including to the sender underestimating the extent of ambiguity aversion of the receiver. Formally, let $(\overline{\sigma}, \underline{\sigma}, \lambda)$ be a Pareto-ranked splitting of $\sigma^{BP}$ satisfying the conditions in Theorem 2. Then ambiguous communication using $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$ also benefits all less ambiguity averse senders. If the receiver is more ambiguity averse than stated, then if $\underline{\sigma}$ is obedient, $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$ continues to benefit the sender (while otherwise only $\kappa$ may need adjustment to ensure a benefit).

We next study a class of persuasion games where Theorem 2 is particularly powerful.

# 6 Belief-Threshold Games

In this section, we characterize when ambiguous communication benefits an ambiguity-neutral sender in a class of persuasion games called *belief-threshold games* (Aybas and Turkel, 2026). This characterization takes the form of a simple and direct condition on payoffs that is easy to satisfy when the receiver has more than two actions (and, of course, is never satisfied when the receiver has only two actions).

In a belief-threshold game, there are $n+1$ states and actions indexed by $i = 0, \ldots, n$. For $i \geq 1$, action $a_i$ maximizes the receiver's expected payoff if and only if the probability of state $\omega_i$ exceeds a threshold $T_i > 1/2$. When none of these thresholds are met, action $a_0$, called the default action, maximizes the receiver's expected payoff. The sender's payoff is state-independent and ranks $a_0$ below all the other actions. Formally:

**Definition 7.** *Fix $n \geq 1$, and let $\Omega = \{\omega_0, \omega_1, \ldots, \omega_n\}$ and $A = \{a_0, a_1, \ldots, a_n\}$. A persuasion game is a* ***belief-threshold game*** *if the following conditions hold:*

*(i) The sender's payoff is state-independent and satisfies*

$$u_s(a_i, \omega) = v_i > v_0 = u_s(a_0, \omega) \quad \textit{for all } i = 1, \ldots, n \textit{ and } \omega \in \Omega. \tag{7}$$

*(ii) There exist unique thresholds $T_1, \ldots, T_n \in (1/2, 1)$ such that, for every $q \in \Delta(\Omega)$,*

$$a_i \in \arg\max_{a \in A} \sum_{\omega \in \Omega} q(\omega) u_r(a, \omega) \iff q(\omega_i) \geq T_i, \textit{ for all } i = 1, \ldots, n,$$

*and*

$$a_0 \in \arg\max_{a \in A} \sum_{\omega \in \Omega} q(\omega) u_r(a, \omega) \iff q(\omega_i) \leq T_i \textit{ for all } i = 1, \ldots, n.$$

Without loss of generality, we normalize $v_0$ to 0 throughout.

Such preferences are well-suited for a variety of economic settings studied in the communication literature. Some of the seminal motivating examples in the persuasion literature, including the prosecutor-judge example in Kamenica and Gentzkow (2011), are belief-threshold games. The combination of sender state-independence (called *transparent motives* following Lipnowski and Ravid (2020)) and receiver threshold preferences models settings where the receiver seeks to match the action to an underlying state, while the sender benefits from inducing any action other than the default. These include seller–buyer and lobbyist–legislator interactions. For example, Chakraborty and Harbaugh (2010) consider a seller persuading a buyer to purchase one of several products, with no purchase as the default action, while Lipnowski and Ravid (2020) study advice-seeking environments, for instance, one in which an industry lobbyist is trying to persuade a politician to adopt one of several policy reforms rather than maintain the status quo.

In the seller-buyer application, the state $\omega$ represents characteristics of the products, with $\omega_i$ the characteristics making product $i$ the best option for the buyer, while $\omega_0$ are those indicating that the buyer should walk away. The seller makes more profit from any sale than no sale. In lobbyist-legislator, the state $\omega$ represents information about the effects of each policy on the legislator's constituents, with $\omega_i$ indicating that policy $i$ is the best option for them, while $\omega_0$ indicates that the status quo is the best option. The lobbyist is

paid a (possibly policy-dependent) bonus if they convince the legislator to support any of the new policies, and thus has state-independent payoffs that rank the status quo last.

The following lemma establishes that a simple parametric representation of the receiver's payoffs characterizes belief-threshold games.

**Lemma 2.** *A persuasion game is a belief-threshold game if and only if the sender's payoffs satisfy* (7) *and the receiver's payoffs satisfy*

$$u_r(a_i, \omega_j) - u_r(a_0, \omega_j) = k_i(\mathbf{1}_{i=j} - T_i), \quad \text{for all } j \in \{0, 1, \ldots, n\}, \tag{8}$$

*for some* $k_i > 0$, $T_i \in (1/2, 1)$, $i \in \{1, \ldots, n\}$.

Our next result establishes a simple-to-check and generically satisfied necessary and sufficient condition on payoffs for ambiguous communication to benefit an ambiguity-neutral sender in belief-threshold games for which the prior satisfies $\sum_{i=1}^{n} p(\omega_i)/T_i < 1$. The proof additionally establishes that in these games, the sufficient conditions from Theorem 2 are also necessary, and the necessary conditions from Theorem 1 are also sufficient.

**Theorem 3.** *Fix any belief-threshold game with* $\sum_{i=1}^{n} p(\omega_i)/T_i < 1$ *and suppose* $\phi_r$ *is strictly concave. Then ambiguous communication benefits an ambiguity-neutral sender if and only if there exist* $i, j \in \{1, \ldots, n\}$ *such that*

$$\frac{v_i}{k_i T_i} \neq \frac{v_j}{k_j T_j}. \tag{9}$$

Condition (9) requires that the ratio of the sender's payoff from $a_i$ compared to $a_0$ and the receiver's payoff from the same in any state where $a_i$ is not optimal is not constant across all non-default actions. In the special case where the thresholds and $k$'s are the same for all actions, (9) reduces to the condition that the sender has a strict preference between two non-default actions. Similarly, if the sender is indifferent between all non-default actions, (9) reduces to the condition that the receiver's $k_i T_i$ varies across non-default actions. In general, (9) is generically satisfied when $n \geq 2$ in the sense that it is satisfied for all but a measure-zero set of payoff parameters for this class of games.[13]

What is the role of the condition $\sum_{i=1}^{n} p(\omega_i)/T_i < 1$? First, observe that it ensures that the prior does not reach the threshold for any $a_i, i = 1, \ldots, n$, and thus, absent persuasion

[13] For a similar class of games, Arieli et al. (2025) show that $\sigma^{BP}$ is generically inefficient when $n \geq 2$. This inefficiency is a prerequisite for the existence of a Pareto-ranked splitting of $\sigma^{BP}$.

the receiver would take the worst action for the sender, $a_0$, creating room for the sender to possibly gain from persuasion. Second, and the reason why it cannot be replaced by the weaker condition that the receiver would take $a_0$ absent persuasion, the proof shows that it ensures that $\sigma^{BP}$ recommends all actions with positive probability in state $\omega_0$ (and recommends $a_i$ with probability one in state $\omega_i$ for $i > 0$). This property is what allows the condition in Lemma A.5 ensuring the existence of a Pareto-ranked splitting to be checked using the payoff-difference vectors from state $\omega_0$ and hence allows the necessary and sufficient condition to be expressed as the simple ratio condition (9). See (Cheng et al., 2026, Example 2) for an example showing that relaxing $\sum_{i=1}^{n} p(\omega_i)/T_i < 1$ leads the characterization in Theorem 3 to fail.

How can Theorem 3 be extended to the case of an ambiguity-averse sender? First, the "only if" direction of the theorem goes through unchanged, as the argument makes no use of the sender's ambiguity neutrality. Second, the part of the "if" direction establishing the existence of a Pareto-ranked splitting of $\sigma^{BP}$ also goes through unchanged. The only adjustments are needed in ensuring that the splitting satisfies the inequality $\rho^{\phi_s,u_s}((\overline{\sigma},\underline{\sigma}),\lambda) < \kappa - \lambda$ from the sufficient conditions in Theorem 2. One approach would be to argue that, since $\kappa$ increases with the concavity of $\phi_r$ and the other terms are unaffected by $\phi_r$, fixing the splitting and $\phi_s$, there always exists a critical coefficient of absolute ambiguity aversion $\hat{\alpha} > 0$ such that the inequality is satisfied when facing any receiver with $\phi_r$ more concave than $-e^{-\hat{\alpha}x}$. See (Cheng et al., 2026, Proposition 2) for a result establishing this fact.

# 7 Optimal Persuasion with Ambiguous Communication

Just as Pareto-ranked splittings were key to understanding when and how ambiguous communication can benefit the sender, our next result shows that Pareto-ranked experiments are fundamental in describing properties of *optimal* ambiguous communication.

**Theorem 4.** *Suppose $\mu$ is a solution to ($\mathcal{P}$), and that $\phi_r$ is strictly concave.*[14] *Then*

*(i) for all $\sigma_\theta, \sigma_{\theta'} \in supp(\mu)$ such that $u_s(\sigma_\theta) \geq U_s(\mu) \geq u_s(\sigma_{\theta'})$ with at least one inequality strict, $\sigma_\theta$ and $\sigma_{\theta'}$ are weakly Pareto-ranked, and*

*(ii) if $\phi_s$ is affine, then for all $\sigma_\theta \in supp(\mu)$, there does not exist a Pareto-ranked splitting of $\sigma_\theta$, $(\overline{\sigma},\underline{\sigma},\lambda)$, such that $u_s(\overline{\sigma}) \geq U_s(\mu) \geq u_s(\underline{\sigma})$.*

[14] The only role of strict concavity of $\phi_r$ is to simplify the statement of the theorem. Without it, one needs to add conditions checking if $\phi_r'(u_r(\sigma_\theta)) \neq \phi_r'(u_r(\sigma_{\theta'}))$ to each part of the theorem.

Part (i) of the result says that if two experiments in the support of $\mu$ bracket the sender's payoff from an optimal ambiguous experiment, they must be weakly Pareto-ranked. Part (ii) shows there must not exist any opportunities to introduce additional ambiguity through Pareto-ranked splittings that bracket the sender's payoff from the ambiguous experiment. Unlike Part (i), Part (ii) requires the sender to be ambiguity neutral. However, Remark A.2 in the Appendix explains why a similar result holds whenever the sender is not too ambiguity averse compared to the receiver.

To gain intuition for part (i), first observe that if such $\sigma_\theta$ and $\sigma_{\theta'}$ are not weakly Pareto-ranked, then the receiver must get a strictly higher expected payoff from $\sigma_{\theta'}$ than from $\sigma_\theta$, while the reverse is true for the sender. Ambiguity aversion then implies that the receiver's effective measure places more weight on $\sigma_\theta$ relative to $\sigma_{\theta'}$ than the ambiguity neutral weights do, i.e., $em^\mu_\theta / em^\mu_{\theta'} > \mu_\theta / \mu_{\theta'}$. If $\sigma_\theta$ and $\sigma_{\theta'}$ are the only two experiments in the support of $\mu$, the sender can merge them into the (unambiguous) experiment $em^\mu_\theta \sigma_\theta + em^\mu_{\theta'} \sigma_{\theta'}$. By construction, the receiver would continue to be obedient, and the sender would strictly benefit from this merging – a profitable deviation. When $\sigma_\theta$ and $\sigma_{\theta'}$ are not the only two experiments in the support of $\mu$, however, this is not the complete story. Splitting or merging experiments may shift their combined weights in the effective measure relative to the other experiments. Part of the additional insight of the proof is that when $u_s(\sigma_\theta) \geq U_s(\mu) \geq u_s(\sigma_{\theta'})$ holds, the merging nevertheless benefits the sender. The intuition for part (ii) is similar.

Lemmas A.9 and A.10 in the Appendix show that concavity (resp. convexity) of $1/\phi_r'$ is sufficient to extend the conclusions to pairs of experiments on a particular side of $U_s(\mu)$, and assuming linearity of $1/\phi_r'$ leads to the following simpler necessary conditions for optimal persuasion:

**Proposition 2.** *Suppose $\mu$ is a solution to ($\mathcal{P}$), and*

$$\phi_r(x) = \ln(x + b) \tag{10}$$

*for some $b \in \mathbb{R}$ where $x + b > 0$ for $x \in [\min_{a,\omega} u_r(a,\omega), \max_{a,\omega} u_r(a,\omega)]$. Then, all experiments are weakly Pareto-ranked, that is, for all $\sigma_\theta, \sigma_{\theta'} \in supp(\mu)$, $\sigma_\theta$ and $\sigma_{\theta'}$ are weakly Pareto-ranked. If, in addition, the sender is ambiguity neutral, no Pareto-ranked splitting of any $\sigma_\theta \in supp(\mu)$ exists.*

The result that all experiments used must be weakly Pareto-ranked is reminiscent of a key Pareto-ranking result (Rayo and Segal, 2010, Lemma 2) in an entirely different per-

suasion setting (one in which ambiguity plays no role). The more general results of our Theorem 4 have no obvious analogue in the setting of Rayo and Segal (2010).

The proofs of Theorem 4 and Proposition 2 make use of a concavification-like characterization (Aumann and Maschler, 1966) of the value of optimal persuasion with ambiguous communication. Since concavification plays a prominent role in the existing persuasion literature, we explicitly describe this characterization here.

Let $\Sigma^* \subseteq \Sigma$ be the set of obedient experiments. Both $\Sigma$ and $\Sigma^*$ are non-empty convex sets and can be embedded into an $|\Omega|(|A|-1)$-dimensional Euclidean space since an experiment specifies, for each $\omega \in \Omega$, a probability distribution over $A$. For each scalar $u \in \mathbb{R}$, define

$$\Phi_u(\sigma) := \frac{\phi_s(u_s(\sigma)) - \phi_s(u)}{\phi_r'(u_r(\sigma))},$$

and consider the following maximization problem:

$$\begin{aligned} &\max_{\lambda \in \Delta_f(\Sigma)} \textstyle\sum_\theta \lambda_\theta \Phi_u(\sigma_\theta), \\ &\text{subject to } \textstyle\sum_\theta \lambda_\theta \sigma_\theta \in \Sigma^*. \end{aligned} \tag{$\Phi_u^*$}$$

Let $V(\Phi_u^*)$ denote the value of the program ($\Phi_u^*$).

Proposition 3 states that the value of the ambiguous persuasion program ($\mathcal{P}$) is the unique utility level $u$ such that $V(\Phi_u^*) = 0$. An optimal ambiguous persuasion strategy can be directly constructed from a solution to ($\Phi_u^*$), and there always exists such an optimal strategy that makes use of no more than $|\Omega|(|A|-1)+1$ experiments.

**Proposition 3.** *The value of ($\mathcal{P}$) is $u$ if and only if $V(\Phi_u^*) = 0$. Moreover, there exists a solution $\mu^*$ to ($\mathcal{P}$) such that $|supp\ (\mu^*)| \leq |\Omega|(|A|-1)+1$.*

Concavification can be used to compute $V(\Phi_u^*)$: For each $u \in \mathbb{R}$, the program ($\Phi_u^*$) maximizes over convex combinations of points on the graph of $\Phi_u$, exactly the type of program that concavification characterizes. Specifically, for each $u \in \mathbb{R}$, let $\text{cav}\Phi_u : \Sigma \to \mathbb{R}$ denote the concavification of $\Phi_u$, that is,

$$\text{cav}\Phi_u(\sigma) = \begin{cases} \max_{\lambda \in \Delta_f(\Sigma)} \sum_\theta \lambda_\theta \Phi_u(\sigma_\theta), \\ \text{subject to } \sum_\theta \lambda_\theta \sigma_\theta = \sigma, \end{cases}$$

and the maximum over $\sigma \in \Sigma^*$ of $\text{cav}\Phi_u(\sigma)$ is $V(\Phi_u^*)$. Any such maximum is achieved by a splitting of some obedient experiment, with the splitting weights given by the effective mea-

sure. The following immediate corollary of Proposition 3 thus provides a concavification-like characterization of the value of ($\mathcal{P}$):[15]

**Corollary 2.** *The value of* ($\mathcal{P}$) *is* $u$ *if, and only if,* $\max_{\sigma \in \Sigma^*} \text{cav}\Phi_u(\sigma) = 0$.

We can apply Corollary 2 and Proposition 2 to the Introductory Example:

**Example 2** (Optimal Persuasion in the Introductory Example)**.** *Suppose* $\phi_r(x) = \ln(x + 5)$ *and* $\phi_s(x) = x$*. These results allow us to show the sender's optimal strategy in the Introductory Example is the binary ambiguous experiment described in Figure 2 with* $\mu_{\overline{\theta}} = 39/50$*. Both the sender and receiver do better than they would under Bayesian persuasion.*

# 8 Further Discussion

## 8.1 Comparison with BLL and persuasion of dynamically inconsistent ambiguity-averse receivers

To facilitate comparison with BLL's approach, while maintaining our modeling setting as much as possible, consider the following dynamically inconsistent model of the receiver's interim preference formation: The receiver applies Bayesian updating using $\mu$ rather than the effective measure (recall the dynamically consistent updating described in Remark 2 in Section 3.3), and chooses their message-contingent strategy to maximize these interim preferences. Formally, fixing an ambiguous experiment $\mu$, and any message $m \in M$ that it could generate, let

$$q_\theta(\omega|m) = \frac{p(\omega)\sigma_\theta(m|\omega)}{\sum_{\omega'} p(\omega')\sigma_\theta(m|\omega')}$$

denote this receiver's posterior belief over $\Omega$ conditional on message $m$ under experiment $\sigma_\theta$. Furthermore, let

$$\eta(\sigma_\theta|m) = \frac{\mu(\sigma_\theta)\sum_\omega p(\omega)\sigma_\theta(m|\omega)}{\sum_{\theta'}\sum_\omega \mu(\sigma_{\theta'})p(\omega)\sigma_{\theta'}(m|\omega)}$$

[15] See Section 8.5 for a discussion comparing this result with concavification in Bayesian persuasion.

denote this receiver's posterior belief over $\Sigma$ conditional on message $m$. Then such a receiver chooses their strategy $\tau$ such that, for each $m$, $\tau(m) \in \Delta(A)$ maximizes

$$\sum_{\theta} \eta(\sigma_\theta|m)\phi_r\left(\sum_{\omega,a} q_\theta(\omega|m)\tau(a|m)u_r(a,\omega)\right).$$

We refer to this as an *inconsistent* receiver.[16] Similarly, we refer to a receiver who behaves as in our model as a *consistent* receiver.

One key difference between inconsistent and consistent receivers relates to opting out of persuasion games with ambiguous communication. It is easy to construct examples where the sender's optimal ambiguous experiment makes an inconsistent receiver worse off than they would be under no communication (see, e.g., Section 2 of BLL). Such a receiver has an incentive to opt out of the persuasion game ex ante. In contrast, a consistent receiver will never be worse off under ambiguous communication than under no communication.

Turning to when and how the sender benefits from ambiguous communication when facing a consistent versus an inconsistent receiver, Belief-threshold games provide a useful setting for this comparison. Consider belief-threshold games with $\sum_{i=1}^{n} p(\omega_i)/T_i < 1$. Theorem 3 told us that an ambiguity-neutral sender benefits from ambiguous communication when facing a consistent receiver if and only if $v_i/k_iT_i \neq v_j/k_jT_j$ for some $i,j \in \{1,\ldots,n\}$, i.e., (9) in Theorem 3 holds. Suppose that the receiver's payoff to some action is *safe* in the sense that it is independent of the state. Note that the safety of any specified action is compatible with any choice of the $v_i$'s, $k_i$'s, and $T_i$'s, and thus does not affect whether (9) holds. The next result shows that the sender benefits from ambiguous communication when facing a sufficiently ambiguity-averse inconsistent receiver if and only if this safe action is not $a_0$.[17]

**Theorem 5.** *Fix any belief-threshold game with $\sum_{i=1}^{n} p(\omega_i)/T_i < 1$ and $u_r(a_{\bar{i}},\omega)$ constant across $\omega \in \Omega$ for some $\bar{i}$. If $\bar{i} > 0$, then the sender benefits from ambiguous communication with any inconsistent receiver having sufficiently concave $\phi_r$, and if $\bar{i} = 0$, then the sender does not benefit from ambiguous communication with any inconsistent receiver.*

A key construction in helping understand this result is the following ambiguous experiment: given an optimal Bayesian persuasion experiment $\sigma^{BP}$, let ambiguous experiment $\mu$

[16] Being precise, the inconsistent receiver's interim preferences are inconsistent with their given ex-ante preferences. As Pahlke (2023) points out, the interim preferences in BLL are consistent with *some* ex-ante preferences, though these may be very different from the given ones.

[17] Belief-threshold games cannot have multiple safe actions for the receiver.

be such that $\mu(\sigma_1) = \mu(\sigma_2) = 1/2$, $\sigma_1$ and $\sigma_2$ send messages in $\{m_0, m_j\} \cup A$, and

$$\sigma_1(m_0|\omega_0) = \sigma_2(m_j|\omega_0) = \sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_j|\omega_0),$$
$$\sigma_1(m_j|\omega_j) = \sigma_2(m_0|\omega_j) = \sigma^{BP}(a_j|\omega_j),$$

where $j = \bar{i}$ if $\bar{i} > 0$ and $j = 1$ otherwise, all other entries for $m_0$, $m_j$, $a_0$, and $a_j$ are zero, and for all other actions, $\sigma_1 = \sigma_2 = \sigma^{BP}$. Starting from $\sigma^{BP}$, this ambiguous experiment $\mu$ injects ambiguity through the messages $m_0$ and $m_j$. Observe that the Bayesian posterior beliefs over $\Omega$ conditional on $m_0$ under $\sigma_1$ are the same as those conditional on $m_j$ under $\sigma_2$, and vice versa. These two messages are an example of what BLL call strong synonyms, and such synonyms play an important role in their analysis. Intuitively, the ambiguity generated by these synonyms drives the inconsistent receiver toward the safe action $a_{\bar{i}}$ in states $\omega_0$ and $\omega_j$. When $\bar{i} > 0$, this results in $a_{\bar{i}}$ being played more often and $a_0$ less often than under $\sigma^{BP}$, which helps the sender, and is how that case of the theorem is proved. When $\bar{i} = 0$, however, this results in $a_0$ being played more often and $a_j$ less often than under $\sigma^{BP}$, which hurts the sender. So in this case, ambiguity generated by strong synonyms doesn't benefit the sender. The proof shows that this remains true for any other ambiguous experiment.

The combination of Theorems 3 and 5 provide a simple recipe for constructing examples that contrast the sender's ability to benefit from ambiguous communication with a consistent versus an inconsistent receiver. First, consider a binary-action belief-threshold game and make action $a_1$ safe. Corollary 1 rules out any sender benefit with a consistent receiver, while Theorem 5 shows that the sender can benefit with an inconsistent receiver having sufficiently concave $\phi_r$. Conversely, consider a three-action belief-threshold game and make $a_0$ safe. This places no restrictions on the $k_i, T_i$, or $v_i$, so consider $k_1 T_1 = k_2 T_2$ and $v_1 \neq v_2$. Thus, condition (9) in Theorem 3 is satisfied, implying an ambiguity-neutral sender can benefit with a consistent receiver whenever $\phi_r$ is strictly concave, while the safety of $a_0$ rules out any benefit with an inconsistent receiver by Theorem 5.

More generally, these results show that the determinants of whether the sender can benefit from ambiguous communication are entirely different when facing a consistent versus an inconsistent receiver. Furthermore, when benefit is possible, it can be obtained through very different forms of ambiguous communication: Pareto-ranked splittings of $\sigma^{BP}$ when facing a consistent receiver, and augmenting $\sigma^{BP}$ with strong synonyms when facing an inconsistent receiver.

Why doesn't a Pareto-ranked splitting of $\sigma^{BP}$ benefit the sender when facing an in-

consistent receiver in these games? An inconsistent receiver can be thought of as using a different effective measure for each message. Recall from Section 6 that $\sigma^{BP}$ for such games recommends all actions with positive probability in state $\omega_0$ and recommends $a_i$ with probability one in state $\omega_i$ for $i > 0$. Thus, in any Pareto-ranked splitting of $\sigma^{BP}$, the recommendation $a_i$ may be produced only in states $\omega_i$ and $\omega_0$. Fix $i > 0$. From Lemma 2 and that $a_{\bar{i}}$ is a safe action, $u_r(a_i, \omega_i) - u_r(a_i, \omega_0) \geq 0$. Therefore, upon receiving the recommendation $a_i$, the inconsistent receiver's effective measure will place at least as much weight as an ambiguity-neutral receiver on the $\theta$ such that $\sigma_\theta(a_i|\omega_0)$ is higher (because this will give a lower interim payoff to the receiver). For the better experiment in the splitting, it must be that enough $a_i$'s for $i > 0$ get more weight in state $\omega_0$ than they do in that state in the worse experiment (otherwise it would not be better for the sender). But it will be exactly for these $a_i$'s that the inconsistent receiver's effective measure fails to move incentives in the direction that would allow the sender to place more weight on the better experiment while maintaining obedience.

Correspondingly, why doesn't the strong synonym construction benefit the sender when facing a consistent receiver? Observe that this modification of $\sigma^{BP}$ can be described in two steps. First, form an experiment $\sigma^{BP}_m$ by merging the two messages $a_0$ and $a_j$ into a single message $m$. Second, inject ambiguity by splitting $m$ into two messages $m_0$ and $m_j$, where under $\theta_1$, $m_0$ substitutes for $m$ in state $\omega_0$ and $m_j$ substitutes for $m$ in state $\omega_j$, while under $\theta_2$, the pattern of substitution is reversed. By the symmetry of these substitutions, the receiver's optimal response is the same to $m_0$ as to $m_j$. Thus, a consistent receiver has the same expected payoff from $\sigma_{\theta_1}$ as from $\sigma_{\theta_2}$, and therefore behaves as if facing $\sigma^{BP}_m$, and in particular plays $a_0$ in response to $m$, which is worse for the sender than $\sigma^{BP}$. In this sense, a consistent receiver "sees through" the ambiguity generated by the strong synonyms.

**Remark 4.** There are games outside the belief-threshold class in which the sender might benefit from using a Pareto-ranked splitting of $\sigma^{BP}$ when facing an inconsistent receiver. The introductory example is such a game. The splitting used there has the property that for every message in the support of both experiments in the splitting, the receiver's expected payoff is higher under $\sigma_{\overline{\theta}}$ than under $\sigma_{\underline{\theta}}$, and thus all of the receiver's interim effective measures move in a direction favorable for the sender. Even so, belief-threshold games make clear that the existence of Pareto-ranked splittings with this strong property is not a necessary condition for the sender to benefit from ambiguous communication when facing an inconsistent receiver. □

## 8.2 What if the sender knows the resolution of ambiguity in advance?

We have assumed that the sender commits to an ambiguous experiment $\mu$ without knowing in advance how this ambiguity will resolve. In Section 8.2 of the working paper version (Cheng et al., 2026), we consider the alternative assumption that the sender privately observes the resolution of the ambiguity, i.e., which experiment will generate the message, before committing to $\mu$ (and this is common knowledge). We show that, in any perfect Bayesian equilibria of the modified game, the sender does *not* benefit from ambiguous communication. One lesson is that to benefit from ambiguous communication, the sender needs to use a source of ambiguity the resolution of which they do not already know at the time of committing to the ambiguous experiment.

## 8.3 Robustness to Small Misperceptions

We have assumed that if the sender designs the ambiguous experiment $\mu$, the receiver perceives it correctly, and that the sender knows $\phi_r$ exactly. In Section 8.3 of the working paper version (Cheng et al., 2026), we show that if the sender can benefit from ambiguous communication, this property is robust to small misperceptions of the experiment or of $\phi_r$.

## 8.4 Robustness to Pre-Existing Ambiguity about Payoff-Relevant States

To isolate the role of ambiguous communication, we have assumed that there is no pre-existing ambiguity about the payoff-relevant states. A full analysis of persuasion with ambiguous communication given arbitrary pre-existing ambiguity over states is beyond the scope of this paper and is an interesting topic for future research. Nonetheless, many results and concepts emphasized in this paper remain relevant. In Section 8.4 of the working paper version, we establish (Cheng et al., 2026, Theorem 9) the continued role of Pareto-ranked splittings of obedient experiments in generating benefits from ambiguous communication even under pre-existing ambiguity. Additionally, the conditions under which the sender benefits from ambiguous communication are robust to the introduction of a small amount of pre-existing ambiguity (Cheng et al., 2026, Theorem 10).

## 8.5 Why experiments rather than distributions over posteriors?

Our analysis works directly with statistical experiments rather than the corresponding Bayes-plausible distributions over posteriors frequently used in the literature. Experiments prove

more useful in our setting: convex combinations of experiments, which we exploit repeatedly (e.g., in defining effective measure and Pareto-ranked splittings), need not correspond to a convex combination of the corresponding distributions over posteriors, because the latter ignores any overlap in the messages sent by the experiments in the support.

Similarly, our concavification result in Section 7 is another instance where experiments prove more amenable than distributions. The splittings in Corollary 2 are on the higher-dimensional space of experiments, rather than on the space of priors as in Kamenica and Gentzkow (2011), and its extension to allow for exogenously heterogeneous priors in Alonso and Câmara (2016). We now explain why this is the case.

Suppose we try to write a program in which the sender maximizes with respect to splittings of the prior. Consider the simplest case of an ambiguity-neutral sender. Any ambiguous experiment $\mu$ induces a distribution over the receiver's *effective* posteriors, that is, the posteriors that the "effective" experiment, $\sum_\theta em_\theta^\mu \sigma_\theta$, induces. Thus, the splitting $\sum_\theta em_\theta^\mu \sigma_\theta$ induces may differ from the splitting the experiment $\sum_\theta \mu_\theta \sigma_\theta$ induces. The latter is the one the ambiguity-neutral sender uses to evaluate their payoff. To be amenable to a concavification approach on this space, the sender's objective function would therefore need to be, as in Alonso and Câmara (2016), an increasing transformation of a linear function of the distribution over the receiver's effective posteriors. However, the highly non-linear relationship between $em_\theta^\mu$ and $\mu$ makes this impossible. Economically, this non-linearity has its source in the fact that ambiguity aversion causes the effective measure to be proportional to the product of the ambiguity neutral probability $\mu_\theta$ and the marginal utility $\phi_r'(u_r(\sigma_\theta))$ and the latter is non-separable across action recommendations. Thus, it is this non-separability across action recommendations in determining obedience under ambiguity aversion that prevents adopting the strategies of Kamenica and Gentzkow (2011) and Alonso and Câmara (2016).

## 8.6 Related Literature

In addition to the papers cited in the introduction, the following are also at the intersection of Bayesian persuasion (BP) and ambiguity. Kosterina (2022) studies BP when an MEU sender is ambiguous about the receiver's prior, while in Dworczak and Pavan (2022) an MEU sender (who also has a preference for selecting among MEU-optimal strategies those that perform best under a baseline conjecture) is ambiguous about the exogenous informa-

tion a receiver might learn.[18] Nikzad (2021) studies BP when the receiver is MEU and has ambiguity about the prior over states. Hedlund et al. (2020) studies BP in problems with two states and two actions, when the receiver has $\alpha$-MEU preferences (Ghirardato et al., 2004) and considers an interval of priors and the sender has state-independent preferences over the action taken by the agent and is ambiguity neutral. In all of these papers, the sender is limited to standard, unambiguous experiments, and thus any ambiguity is exogenous. This stands in contrast to the endogeneity of ambiguity in our setting, where it becomes payoff-relevant only through the intentional communication choices of the sender.

Kellner and Le Quement (2018) study cheap talk communication assuming that the receiver has MEU preferences and the sender can choose to communicate ambiguously. The key difference between cheap talk and persuasion is the sender's inability to commit to a communication strategy. Their receiver uses the same dynamically inconsistent update rule as in BLL. They find that both sender and receiver may benefit from the sender choosing to communicate ambiguously. Kellner and Le Quement (2017) studies cheap talk communication with purely exogenous ambiguity.

Papers studying mechanism design with ambiguity include Bose and Renou (2014), Wolitzky (2016), Di Tillio et al. (2017), Guo (2019) and Tang and Zhang (2021), among others. All but Wolitzky (2016) consider ambiguity that arises intentionally through design of the mechanism. Dütting et al. (2024) allow a principal to offer ambiguous contracts to an MEU agent and show how the principal may benefit and that optimal contracts have a simple form. All gains from ambiguous contracting disappear in their model if the agent can hedge against ambiguity by randomizing.

# Appendices

## A Proofs of Main Results

### A.1 Proofs for Section 3

*Proof of Proposition 1.* Fix $\mu$ and $\tau \in BR(\mu)$. For each $\sigma_\theta \in \text{supp}(\mu)$, define $\sigma^*_\theta$ by $\sigma^*_\theta(a|\omega) = \sum_m \tau(a|m)\sigma_\theta(m|\omega)$ for all $a \in A$ and $\omega \in \Omega$. For each $\theta$ and $i \in \{s, r\}$, $u_i(\sigma_\theta, \tau) = \sum_{\omega,a} p(\omega)u_i(a, \omega) \sum_m \tau(a|m)\sigma_\theta(m|\omega) = u_i(\sigma^*_\theta, \tau^*)$. Define the canonical

[18] Dworczak and Pavan (2022)'s model is not restricted to single-receiver persuasion settings.

ambiguous experiment $\mu^* = \sum_\theta \mu_\theta \delta_{\sigma_\theta^*}$. Then $U_i(\mu^*, \tau^*) = U_i(\mu, \tau)$ for each $i \in \{s, r\}$.

Suppose $\mu^*$ is not obedient. Then there exists $\tilde{\tau} : A \to \Delta(A)$ such that $U_r(\mu^*, \tau^*) < U_r(\mu^*, \tilde{\tau})$. Then, define $\tau'(a|m) = \sum_{a'} \tilde{\tau}(a|a')\tau(a'|m)$, and note that for all $\theta$,

$$\begin{aligned} u_r(\sigma_\theta, \tau') &= \sum_{\omega,a} p(\omega) u_r(a, \omega) \sum_{a'} \tilde{\tau}(a|a') \sum_m \tau(a'|m)\sigma_\theta(m|\omega) \\ &= \sum_{\omega,a,a'} p(\omega) u_r(a, \omega)\tilde{\tau}(a|a')\sigma_\theta^*(a'|\omega) = u_r(\sigma_\theta^*, \tilde{\tau}). \end{aligned}$$

Thus, $U_r(\mu, \tau') = U_r(\mu^*, \tilde{\tau}) > U_r(\mu^*, \tau^*) = U_r(\mu, \tau)$, implying $\tau \notin BR(\mu)$. □

*Proof of Lemma 1.* Let $\mu$ be an ambiguous experiment. We show that $\tau^* \in br(\sigma^*)$, where $\sigma^* = \sum_\theta em_\theta^\mu \sigma_\theta$, if and only if $\tau^* \in BR(\mu)$. $\tau^* \in br(\sigma^*)$ is equivalent to

$$u_r(\sigma^*, \tau) - u_r(\sigma^*, \tau^*) \leq 0, \quad \forall \tau. \tag{A.1}$$

The concavity and differentiability of $\phi_r$ implies that $\phi_r(U_r(\mu, \tau))$ is concave and differentiable in $\tau$. Thus, $\tau^* \in BR(\mu)$ is equivalent to

$$\langle \nabla_\tau \phi_r(U_r(\mu, \tau^*)), \tau - \tau^* \rangle \leq 0, \quad \forall \tau. \tag{A.2}$$

Using the formula (3) for the effective measure $em_\theta^\mu$, we have

$$\langle \nabla_\tau \phi_r(U_r(\mu, \tau^*)), \tau - \tau^* \rangle = \left( \sum_\theta \mu_\theta \phi_r'(u_r(\sigma_\theta)) \right) (u_r(\sigma^*, \tau) - u_r(\sigma^*, \tau^*)).$$

By the positivity of $\sum_\theta \mu_\theta \phi_r'(u_r(\sigma_\theta))$, (A.2) if and only if (A.1). □

## A.2 Proofs for Section 4

*Proof of Theorem 1.* We prove the theorem by first showing the following lemma and then establishing that the conditions in Theorem 1 are equivalent to the conditions in the lemma.

**Lemma A.1.** *Ambiguous communication benefits the sender only if $\phi_r$ is not affine and there exist Pareto-ranked experiments, $\sigma$ and $\sigma^*$ such that: (i) $\operatorname{supp} \sigma(\cdot|\omega) = \operatorname{supp} \sigma^*(\cdot|\omega)$ for all $\omega$, (ii) $u_s(\sigma) > u_s^{BP}$, and (iii) $\tau^* \in br(\sigma^*) \setminus br(\sigma)$.*

*Proof of Lemma A.1.* Let $\mu^*$ solve ($\mathcal{P}$) and benefit the sender. Define $\sigma := \sum_\theta \mu_\theta^* \sigma_\theta$ and

$\sigma^* := \sum_\theta em_\theta^{\mu^*} \sigma_\theta$. Since $\mu^*$ and $em^{\mu^*}$ have the same support, $\sigma$ and $\sigma^*$ satisfy (i). Lemma 1 gives $\tau^* \in br(\sigma^*)$, and concavity of $\phi_s$ yields $u_s(\sigma^*) \leq u_s^{BP} < U_s(\mu^*) \leq \sum_\theta \mu_\theta^* u_s(\sigma_\theta) = u_s(\sigma)$. Thus $\tau^* \notin br(\sigma)$, proving (ii) and (iii). Also, $\sigma \neq \sigma^*$ implies $em^{\mu^*} \neq \mu^*$ and hence $\phi_r$ is not affine. To establish $u_r(\sigma) > u_r(\sigma^*)$, we use the following lemma.

**Lemma A.2.** *Fix two monotonic sequences $x_1 \geq \cdots \geq x_n$, $0 < y_1 \leq \cdots \leq y_n$, and $\mu \in \Delta(\{1, 2, \ldots, n\})$. If there exist indices $i^* < j^*$ such that $\mu_{i^*} > 0$, $\mu_{j^*} > 0$, $x_{i^*} > x_{j^*}$ and $y_{i^*} < y_{j^*}$. Then $(\sum_{i=1}^n x_i \mu_i y_i)/(\sum_{j=1}^n \mu_j y_j) < \sum_{i=1}^n x_i \mu_i$.*[19]

Order $u_r(\sigma_\theta)$ decreasingly. Since $\phi_r$ is strictly increasing and concave, the corresponding derivatives $\phi_r'(u_r(\sigma_\theta))$ are positive and nondecreasing. They are nonconstant on $\text{supp}(\mu^*)$ because $em^{\mu^*} \neq \mu^*$. Hence some $\sigma_\theta, \sigma_{\theta'} \in \text{supp}(\mu^*)$ satisfy $u_r(\sigma_\theta) > u_r(\sigma_{\theta'})$ and $\phi_r'(u_r(\sigma_\theta)) < \phi_r'(u_r(\sigma_{\theta'}))$. Lemma A.2 therefore gives $u_r(\sigma) > u_r(\sigma^*)$. □

We show conditions in Lemma A.1 imply those in Theorem 1 with the following lemma:

**Lemma A.3.** *Fix $\sigma$. If there exists $\hat{\sigma}$ such that $u_s(\hat{\sigma}) > u_s(\sigma)$, $u_r(\hat{\sigma}) > u_r(\sigma)$ and for all $\omega \in \Omega$, supp$(\hat{\sigma}(\cdot|\omega)) \subseteq$ supp$(\sigma(\cdot|\omega))$, then there exists a Pareto-ranked splitting of $\sigma$, $(\overline{\sigma}, \underline{\sigma}, \lambda)$ with $\overline{\sigma} = \hat{\sigma}$.*[20]

Applying Lemma A.3 to the experiments in Lemma A.1 gives a Pareto-ranked splitting $(\overline{\sigma}, \underline{\sigma}, \lambda)$ of the obedient $\sigma^*$ with $\overline{\sigma} = \sigma$. Conditions (ii) and (iii) then give $u_s(\overline{\sigma}) > u_s^{BP} \geq u_s(\sigma^*) > u_s(\underline{\sigma})$, as required.

Finally, we show the conditions in Theorem 1 imply those in Lemma A.1. Suppose there exists an obedient experiment $\hat{\sigma}$ satisfying the conditions in Theorem 1. Because $u_s(\overline{\sigma}) > u_s^{BP}$, there exists $\underline{\gamma} \in (\lambda, 1)$ such that for all $\gamma \in (\underline{\gamma}, 1)$, $\gamma u_s(\overline{\sigma}) + (1-\gamma) u_s(\underline{\sigma}) > u_s^{BP}$, and thus $\tau^* \notin br(\gamma\overline{\sigma} + (1-\gamma)\underline{\sigma})$. Fix any such $\gamma$ and let $\sigma = \gamma\overline{\sigma} + (1-\gamma)\underline{\sigma}$. Let $\sigma^* = \hat{\sigma}$. It follows that $\tau^* \in br(\sigma^*)$ since $\hat{\sigma}$ is obedient, and thus $u_s^{BP} \geq u_s(\sigma^*)$. Since both $\sigma$ and $\sigma^*$ are strict mixtures of $\overline{\sigma}$ and $\underline{\sigma}$, they satisfy the common support condition (i). By the definition of Pareto-ranked splitting and $\gamma > \lambda$, $u_r(\sigma) > u_r(\sigma^*)$. □

*Proof of Corollary 1.* Suppose $\sigma$ and $\sigma^*$ satisfy the conditions in Lemma A.1. Since $\tau^* \notin br(\sigma)$, assume without loss of generality that obedience fails after recommendation $a_1$. Together with $\tau^* \in br(\sigma^*)$, this gives $\sum_\omega p(\omega)\sigma(a_1|\omega)[u_r(a_1, \omega) - u_r(a_2, \omega)] < 0$ and $\sum_\omega p(\omega)\sigma^*(a_1|\omega)[u_r(a_1, \omega) - u_r(a_2, \omega)] \geq 0$. Using $\sigma(a_2|\omega) - \sigma^*(a_2|\omega) = \sigma^*(a_1|\omega) -$

[19] This lemma seems like a known result, but we could not locate a reference and so include a proof in the Appendix E of the working paper version (Cheng et al., 2026) for completeness.

[20] The proof of Lemma A.3 is in Appendix E of the working paper version (Cheng et al., 2026).

$\sigma(a_1|\omega)$, we obtain $u_r(\sigma)-u_r(\sigma^*) = \sum_\omega p(\omega)[u_r(a_1,\omega)-u_r(a_2,\omega)][\sigma(a_1|\omega)-\sigma^*(a_1|\omega)] < 0$, contradicting $u_r(\sigma) > u_r(\sigma^*)$. □

## A.3 Proofs for Section 5

The following lemma is used in proving Theorem 2:

**Lemma A.4.** *Let $\overline{\sigma}$ and $\underline{\sigma}$ be experiments such that $u_i(\overline{\sigma}) > u_i(\underline{\sigma})$. For all $\kappa, \lambda \in [0,1]$, $U_i(\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}) > u_i(\lambda\overline{\sigma} + (1-\lambda)\underline{\sigma})$ if and only if player $i$'s $((\overline{\sigma},\underline{\sigma}),\lambda)$-probability premium is strictly less than $\kappa - \lambda$.*

*Proof of Lemma A.4.*

$$\frac{\phi_i(u_i(\lambda\overline{\sigma} + (1-\lambda)\underline{\sigma})) - (\lambda\phi_i(u_i(\overline{\sigma})) + (1-\lambda)\phi_i(u_i(\underline{\sigma})))}{\phi_i(u_i(\overline{\sigma})) - \phi_i(u_i(\underline{\sigma}))} < \kappa - \lambda$$
$$\iff \phi_i(u_i(\lambda\overline{\sigma} + (1-\lambda)\underline{\sigma})) < \kappa\phi_i(u_i(\overline{\sigma})) + (1-\kappa)\phi_i(u_i(\underline{\sigma}))$$

Thus, equivalently, $u_i(\lambda\overline{\sigma} + (1-\lambda)\underline{\sigma}) < U_i(\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)})$. □

*Proof of Theorem 2.* Let $\sigma^{BP}$ be obedient and satisfy $u_s(\sigma^{BP}) = u_s^{BP}$. Suppose that $(\overline{\sigma},\underline{\sigma},\lambda)$ is a Pareto-ranked splitting of $\sigma^{BP}$, $u_i(\overline{\sigma}) > u_i(\underline{\sigma})$ for $i \in \{s,r\}$. For $\kappa$ given by (6), consider the binary ambiguous experiment $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$. Its effective weight on $\overline{\sigma}$ is exactly $\lambda$, Lemma 1 then implies $\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}$ is obedient since $\sigma^{BP}$ is. Then $\rho^{\phi_s,u_s}((\overline{\sigma},\underline{\sigma}),\lambda) < \kappa - \lambda$ implies $U_s(\mu^{((\overline{\sigma},\underline{\sigma}),\kappa)}) > u_s^{BP}$ by Lemma A.4, thus the sender benefits.

An increase in ambiguity aversion corresponds to replacing $\phi$ by $\tilde{\phi} = \varphi \circ \phi$, where $\varphi$ is strictly increasing, concave, and differentiable (Klibanoff et al., 2005). For the sender, concavity of $\varphi$ gives

$$\frac{\tilde{\phi}_s(u_s(\sigma^{BP})) - \tilde{\phi}_s(u_s(\underline{\sigma}))}{\tilde{\phi}_s(u_s(\overline{\sigma})) - \tilde{\phi}_s(u_s(\underline{\sigma}))} \geq \frac{\phi_s(u_s(\sigma^{BP})) - \phi_s(u_s(\underline{\sigma}))}{\phi_s(u_s(\overline{\sigma})) - \phi_s(u_s(\underline{\sigma}))}.$$

The two sides equal the probability premiums under $\tilde{\phi}_s$ and $\phi_s$, respectively, plus $\lambda$.

Finally, $\kappa$ in (6) is strictly increasing in the ratio $\phi_r'(u_r(\underline{\sigma}))/\phi_r'(u_r(\overline{\sigma}))$, and

$$\frac{\tilde{\phi}_r'(u_r(\underline{\sigma}))}{\tilde{\phi}_r'(u_r(\overline{\sigma}))} = \frac{\varphi'(\phi_r(u_r(\underline{\sigma})))}{\varphi'(\phi_r(u_r(\overline{\sigma})))}\frac{\phi_r'(u_r(\underline{\sigma}))}{\phi_r'(u_r(\overline{\sigma}))} \geq \frac{\phi_r'(u_r(\underline{\sigma}))}{\phi_r'(u_r(\overline{\sigma}))}.$$

Thus $\kappa$ increases with the receiver's ambiguity aversion. □

## A.4 Proofs for Section 6

*Proof of Lemma 2.* **Only If.** Fix any $i > 0$. By part (ii) of Definition 7, the receiver is indifferent between $a_i$ and $a_0$ at any belief $q$ with $q(\omega_i) = T_i$. This indifference must hold in particular for $q = T_i\delta_{\omega_i} + (1 - T_i)\delta_{\omega_j}$ for every $j \neq i$. This is equivalent to

$$T_i(u_r(a_i, \omega_i) - u_r(a_0, \omega_i)) + (1 - T_i)(u_r(a_i, \omega_j) - u_r(a_0, \omega_j)) = 0.$$

Since $a_i$ is strictly optimal at $\omega_i$, (8) holds with $k_i$ defined as

$$k_i := \frac{u_r(a_i, \omega_i) - u_r(a_0, \omega_i)}{1 - T_i} = -\frac{u_r(a_i, \omega_j) - u_r(a_0, \omega_j)}{T_i} > 0.$$

**If.** Fix $k_i > 0$, $T_i \in (1/2, 1)$, $i > 0$ and suppose the receiver's payoffs satisfy (8). Then observing that maximizing $\sum_{\omega\in\Omega} q(\omega)u_r(a, \omega)$ over $a \in A$ is equivalent to maximizing $\sum_{\omega\in\Omega} q(\omega)(u_r(a, \omega) - u_r(a_0, \omega))$, and substituting (8) yields Definition 7. □

*Proof of Theorem 3.* **If.** For any obedient experiment $\sigma$,

$$q(\omega_i|a_i) = \frac{p(\omega_i)\sigma(a_i|\omega_i)}{\sum_\omega p(\omega)\sigma(a_i|\omega)} \geq T_i.$$

Since $\sigma(a_i|\omega_i) \leq 1$, it follows that $\sum_\omega p(\omega)\sigma(a_i|\omega) \leq p(\omega_i)/T_i$. Therefore, every obedient experiment $\sigma$ satisfies $u_s(\sigma, \tau^*) = \sum_{i=1}^n v_i \sum_\omega p(\omega)\sigma(a_i|\omega) \leq \sum_{i=1}^n v_i p(\omega_i)/T_i$. When $\sum_{i=1}^n p(\omega_i)/T_i < 1$, this upper bound can be achieved by $\sigma^{BP}$ defined as follows:

$$\sigma^{BP}(a_i|\omega_i) = 1, \quad \sigma^{BP}(a_i|\omega_0) = \frac{p(\omega_i)(1 - T_i)}{p(\omega_0)T_i}, \quad \text{for all } i \geq 1,$$
$$\sigma^{BP}(a_0|\omega_0) = 1 - \sum_{i=1}^n \frac{p(\omega_i)(1 - T_i)}{p(\omega_0)T_i},$$

and all other probabilities are zero. To see this is feasible, notice that

$$\sum_{i=1}^n \frac{p(\omega_i)(1 - T_i)}{p(\omega_0)T_i} < 1 \iff \sum_{i=1}^n \frac{p(\omega_i)}{T_i} - \sum_{i=1}^n p(\omega_i) < p(\omega_0) \iff \sum_{i=1}^n \frac{p(\omega_i)}{T_i} < 1.$$

Moreover, $\sigma^{BP}$ is obedient because $q(\omega_i|a_i) = T_i$ for all $i \geq 1$, and satisfies

$$u_s(\sigma^{BP}) = \sum_{i=1}^{n} v_i \sum_{\omega} p(\omega)\sigma^{BP}(a_i|\omega) = \sum_{i=1}^{n} v_i \frac{p(\omega_i)}{T_i}.$$

We next use the following lemma to show that a Pareto-ranked splitting of $\sigma^{BP}$ exists.

**Lemma A.5.** *Given any experiment $\sigma^*$, fix, for each $\omega \in \Omega$, $a_\omega \in supp(\sigma^*(\cdot|\omega))$ and consider the following set of vectors,*

$$\left\{ \begin{bmatrix} u_s(a,\omega) - u_s(a_\omega,\omega) \\ u_r(a,\omega) - u_r(a_\omega,\omega) \end{bmatrix} : a \in supp(\sigma^*(\cdot|\omega)), \omega \in \Omega \right\}. \tag{A.3}$$

*If this set spans $\mathbb{R}^2$, then there exists a Pareto-ranked splitting of $\sigma^*$.*

Let $a_{\omega_0} = a_0$, the payoff-difference vectors in state $\omega_0$ are positive multiples of

$$\begin{bmatrix} v_i \\ -k_i T_i \end{bmatrix} \text{ for } i \geq 1.$$

Two such vectors are linearly independent and hence span $\mathbb{R}^2$ if and only if there exist $i, j$ such that $v_i/k_iT_i \neq v_j/k_jT_j$. Therefore, a Pareto-ranked splitting of $\sigma^{BP}$ exists. Note that such existence requires $n \geq 2$, so there is no conflict with Corollary 1. Since $\phi_r$ is strictly concave, the sufficient conditions in Theorem 2 are satisfied. Therefore, ambiguous communication benefits the sender.

**Only if.** Suppose $v_i/k_iT_i = c$ for all $i > 0$. Consider any non-obedient experiment $\hat{\sigma}$ with $u_s(\hat{\sigma}) > u_s^{BP}$. Then using $v_i/k_iT_i = c$, we have

$$\sum_{i=1}^{n} \sum_{\omega} p(\omega)\hat{\sigma}(a_i|\omega)v_i > \sum_{i=1}^{n} \frac{p(\omega_i)}{T_i} v_i \implies \sum_{i=1}^{n} k_i T_i \sum_{\omega} p(\omega)\hat{\sigma}(a_i|\omega) > \sum_{i=1}^{n} p(\omega_i)k_i.$$

We next show $u_r(\hat{\sigma})$ is strictly worse than always taking action $a_0$:

$$\begin{aligned} u_r(\hat{\sigma}) - \sum_{\omega} p(\omega)u_r(a_0,\omega) &= \sum_{i=1}^{n} \sum_{\omega} p(\omega)\hat{\sigma}(a_i|\omega)(u_r(a_i,\omega) - u_r(a_0,\omega)) \\ &\leq \sum_{i=1}^{n} p(\omega_i)k_i - \sum_{i=1}^{n} k_i T_i \sum_{\omega} p(\omega)\hat{\sigma}(a_i|\omega) < 0. \end{aligned}$$

Since the receiver can always choose action $a_0$, any obedient $\sigma$ must satisfy $u_r(\sigma) \geq \sum_\omega p(\omega)u_r(a_0,\omega)$. Thus no experiment giving the sender more than $u_s^{BP}$ can Pareto-dominate an obedient experiment, violating the necessary condition in Theorem 1. □

*Proof of Lemma A.5.* Let $\Sigma_{\sigma^*} \subseteq \Sigma$ denote the set of experiments having the same support as $\sigma^*$ for each $\omega \in \Omega$. For each $\omega \in \Omega$, fix any $a_\omega \in supp(\sigma^*(\cdot|\omega))$. For any $\sigma \in \Sigma_{\sigma^*}$, by substituting $\sigma(a_\omega|\omega) = 1 - \sum_{a \neq a_\omega} \sigma(a|\omega)$, we have, for $i \in \{s, r\}$,

$$u_i(\sigma,\tau^*) = \sum_{\omega,a} p(\omega)\sigma(a|\omega)u_i(a,\omega) = \sum_{\omega\in\Omega_{\sigma^*},a} p(\omega)\sigma(a|\omega)(u_i(a,\omega) - u_i(a_\omega,\omega)) + \sum_\omega p(\omega)u_i(a_\omega,\omega),$$

where $\Omega_{\sigma^*} \subseteq \Omega$ denotes the set of $\omega$ such that $|\, supp(\sigma^*(\cdot|\omega))| > 1$. Given any $\sigma \in \Sigma_{\sigma^*}$, we use $\tilde{\sigma} \in \mathbb{R}^{\sum_{\omega\in\Omega_{\sigma^*}}(|\, supp(\sigma^*(\cdot|\omega))|-1)} =: \mathbb{R}^{\sigma^*}$ to denote the vector of those components of $\sigma(a|\omega)$ with $\omega \in \Omega_{\sigma^*}$ and $a \in supp(\sigma^*(\cdot|\omega)) \setminus \{a_\omega\}$. Thus, we can write

$$\begin{bmatrix} u_s(\sigma,\tau^*) \\ u_r(\sigma,\tau^*) \end{bmatrix} = \begin{bmatrix} p(\omega)(u_s(a,\omega) - u_s(a_\omega,\omega)) & \cdots \\ p(\omega)(u_r(a,\omega) - u_r(a_\omega,\omega)) & \cdots \end{bmatrix} \tilde{\sigma} + \begin{bmatrix} \sum_\omega p(\omega)u_s(a_\omega,\omega) \\ \sum_\omega p(\omega)u_r(a_\omega,\omega) \end{bmatrix}$$

Notice that any non-zero vectors in (A.3) are positive multiples of non-zero columns of the first matrix on the right-hand side above. When the former set spans $\mathbb{R}^2$, the latter matrix has full rank and thus the affine mapping from $\mathbb{R}^{\sigma^*}$ to $\mathbb{R}^2$ defined above is surjective. Since $\{\tilde{\sigma} : \sigma \in \Sigma_{\sigma^*}\} \ni \tilde{\sigma^*}$ is open in $\mathbb{R}^{\sigma^*}$, by the open mapping theorem, $\{(u_s(\sigma,\tau^*), u_r(\sigma,\tau^*)) : \sigma \in \Sigma_{\sigma^*}\}$ is open. Thus, there exists $\hat{\sigma} \in \Sigma_{\sigma^*}$ such that $u_i(\hat{\sigma},\tau^*) > u_i(\sigma^*,\tau^*)$ for $i \in \{s, r\}$. By Lemma A.3, there exists a Pareto-ranked splitting of $\sigma^*$ with $\overline{\sigma} = \hat{\sigma}$. □

**Remark A.1** (Non-necessity)**.** That the spanning condition in Lemma A.5 is not necessary for the existence of a Pareto-ranked splitting can be seen from our introductory example. For $\sigma^* = \sigma^{BP}$, (A.3) does not span $\mathbb{R}^2$, but there are, as depicted in Figure 2, Pareto-ranked splittings of $\sigma^{BP}$. □

## A.5 Proofs for Section 7

We prove Proposition 3 first and then use it to prove Theorem 4.

*Proof of Proposition 3.* There is a unique $u$ such that $V(\Phi_u^*) = 0$:

**Lemma A.6.** *$V(\Phi_u^*)$ satisfies single-crossing, i.e., for any $u > u'$, if $V(\Phi_u^*) \geq 0$, then $V(\Phi_{u'}^*) > 0$. Thus, there exists a unique $u \in \mathbb{R}$ such that $V(\Phi_u^*) = 0$.*

*Proof of Lemma A.6.* $V(\Phi_u^*) \geq 0$ and $\phi_s, \phi_r$ strictly increasing imply $V(\Phi_{u'}^*) > 0$ for $u' < u$. Thus, there is a unique $u$ such that $V(\Phi_u^*) = 0$. By Caratheodory's theorem, the maximum in ($\Phi_u^*$) can be achieved within a compact set, continuity of $\Phi_u$ and Berge's Maximum Theorem imply $V(\Phi_u^*)$ is continuous. Finally, $V(\Phi_u^*) < 0$ for $u > \max u_s(a, \omega)$ and $V(\Phi_u^*) > 0$ for $u < \min u_s(a, \omega)$ imply the existence of $u$ such that $V(\Phi_u^*) = 0$. □

By Proposition 1, ($\mathcal{P}$) is equivalent to $\max_\mu U_s(\mu)$ subject to $\tau^* \in BR(\mu)$.

**Lemma A.7.** *For each $u \in \mathbb{R}$, the value of ($\mathcal{P}$) is greater than $u$ if and only if $V(\Phi_u^*) > 0$.*

*Proof of Lemma A.7.* By Lemma 1 and (4), $\lambda = em^\mu$ gives a one-to-one correspondence between obedient $\mu$ and feasible $\lambda$ for ($\Phi_u^*$). For each corresponding pair,

$$\phi_s(U_s(\mu)) - \phi_s(u) = \frac{\sum_\theta \lambda_\theta \Phi_u(\sigma_\theta)}{\sum_\theta \lambda_\theta / \phi_r'(u_r(\sigma_\theta))}.$$

Thus $U_s(\mu) > u$ if and only if $\sum_\theta \lambda_\theta \Phi_u(\sigma_\theta) > 0$. Existence of such feasible choices therefore gives the value of ($\mathcal{P}$) above $u$ if and only if $V(\Phi_u^*) > 0$. □

Let $u$ be the unique zero of $V(\Phi_u^*)$. By Lemmas A.6 and A.7, for every $u' \in \mathbb{R}$, the value of ($\mathcal{P}$) is greater than $u' \iff V(\Phi_{u'}^*) > 0 \iff u > u'$. Hence the value of ($\mathcal{P}$) is $u$ if and only if $V(\Phi_u^*) = 0$.

Let $u$ be the value of ($\mathcal{P}$) and choose a maximizer $\lambda$ of ($\Phi_u^*$) with $\sum_\theta \lambda_\theta \sigma_\theta = \sigma^* \in \Sigma^*$. Since $V(\Phi_u^*) = 0$, $(\sigma^*, 0)$ is a boundary point of the convex hull of the graph of $\Phi_u$ in $\mathbb{R}^{|\Omega|(|A|-1)+1}$. Every graph point receiving positive weight under $\lambda$ lies on a supporting hyperplane through $(\sigma^*, 0)$. Caratheodory's theorem applied within this hyperplane therefore gives a representation using at most $|\Omega|(|A|-1)+1$ graph points. Applying (4) yields an optimal $\mu^*$ with $|\text{supp}(\mu^*)| \leq |\Omega|(|A|-1)+1$. □

*Proof of Theorem 4.* Suppose $\mu$ is optimal and $\phi_r$ is strictly concave. Let $\hat{u} := U_s(\mu)$ and $\lambda := em^\mu$. By the preceding characterization, $\sum_\theta \lambda_\theta \Phi_{\hat{u}}(\sigma_\theta) = 0 = V(\Phi_{\hat{u}}^*)$, so $\lambda$ is a solution. Since $\lambda$ and $\mu$ have the same support, strict concavity of $\phi_r$ and the corresponding parts of the following lemma (and its swapped version) imply parts (i) and (ii).

**Definition A.1.** *Let $\Sigma_+(u) = \{\sigma \in \Sigma : u_s(\sigma) > u\}$ and $\Sigma_-(u) = \{\sigma \in \Sigma : u_s(\sigma) \leq u\}$.*

**Lemma A.8.** *Let $\lambda \in \Delta_f(\Sigma)$ be a solution to ($\Phi_u^*$). Then the following hold:*

*(i) For all $\sigma_\theta, \sigma_{\theta'} \in$ supp$(\lambda)$ such that $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_+(u) \times \Sigma_-(u)$, if $\phi_r'(u_r(\sigma_\theta)) \neq \phi_r'(u_r(\sigma_{\theta'}))$, then they are Pareto-ranked,*

*(ii) If $\phi_s$ is affine, then for all $\sigma_\theta \in supp(\lambda)$, there does not exist a Pareto-ranked splitting of $\sigma_\theta$, $(\overline{\sigma}, \underline{\sigma}, \alpha)$ with $(\overline{\sigma}, \underline{\sigma}) \in \Sigma_+(u) \times \Sigma_-(u)$ and $\phi_r'(u_r(\overline{\sigma})) < \phi_r'(u_r(\sigma_\theta)) < \phi_r'(u_r(\underline{\sigma}))$.*

A swapped version of Lemma A.8 with $\Sigma_+(u)$ and $\Sigma_-(u)$ defined by swapping the strict and weak inequalities in Definition A.1 also holds, and the proof is identical. □

*Proof of Lemma A.8.* Fix $u \in \mathbb{R}$. Let $\lambda \in \Delta_f(\Sigma)$ be feasible for ($\Phi_u^*$). To prove (i), suppose there exist $\sigma_\theta, \sigma_{\theta'} \in \text{supp}(\lambda)$ and $\alpha \in (0,1)$ such that $\Phi_u(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}) > \alpha\Phi_u(\sigma_\theta) + (1-\alpha)\Phi_u(\sigma_{\theta'})$. Then, $\lambda$ cannot be a solution to ($\Phi_u^*$). To see this, if $\lambda_\theta/\alpha \leq \lambda_{\theta'}/(1-\alpha)$, then replacing $\sigma_\theta$ by $\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}$, letting $\hat{\lambda}_\theta = \lambda_\theta/\alpha$, and $\hat{\lambda}_{\theta'} = \lambda_{\theta'} - (1-\alpha)\lambda_\theta/\alpha$ yields a strict improvement. The other case is symmetric.

Towards a contradiction, suppose in the solution there exists $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_+(u) \times \Sigma_-(u)$ with $\phi_r'(u_r(\sigma_\theta)) \neq \phi_r'(u_r(\sigma_{\theta'}))$ and they are not weakly Pareto-ranked, i.e., $u_s(\sigma_\theta) > u \geq u_s(\sigma_{\theta'})$, and $u_r(\sigma_\theta) < u_r(\sigma_{\theta'})$. Then there exists $\alpha \in (0,1)$ such that $\phi_r'(u_r(\sigma_\theta)) > \phi_r'(u_r(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}))$ (by differentiability of $\phi_r$), and $\Phi_u(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}) > \alpha\Phi_u(\sigma_\theta) + (1-\alpha)\Phi_u(\sigma_{\theta'})$. To see the last point, notice that

$$\begin{aligned}
\Phi_u(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}) &= \frac{\phi_s(u_s(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'})) - \phi_s(u)}{\phi_r'(u_r(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}))} \\
&\geq \frac{\alpha[\phi_s(u_s(\sigma_\theta)) - \phi_s(u)] + (1-\alpha)[\phi_s(u_s(\sigma_{\theta'})) - \phi_s(u)]}{\phi_r'(u_r(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}))} \\
&> \alpha\Phi_u(\sigma_\theta) + (1-\alpha)\Phi_u(\sigma_{\theta'}),
\end{aligned}$$

where the first inequality follows from concavity of $\phi_s$, the second follows from concavity of $\phi_r$, $\Phi_u(\sigma_\theta) > 0$, $\Phi_u(\sigma_{\theta'}) \leq 0$, and $\phi_r'(u_r(\sigma_\theta)) > \phi_r'(u_r(\alpha\sigma_\theta + (1-\alpha)\sigma_{\theta'}))$.

To prove (ii), if there exist $\sigma_\theta \in \text{supp}(\lambda)$ and two experiments $\sigma$ and $\sigma'$ such that $\sigma_\theta = \alpha\sigma + (1-\alpha)\sigma'$ for some $\alpha \in (0,1)$ and $\Phi_u(\alpha\sigma + (1-\alpha)\sigma') < \alpha\Phi_u(\sigma) + (1-\alpha)\Phi_u(\sigma')$, then $\lambda$ cannot be a solution to ($\Phi_u^*$). This follows by noting that splitting $\sigma_\theta$ into $\sigma$ with probability $\alpha\lambda_\theta$ and $\sigma'$ with probability $(1-\alpha)\lambda_\theta$ induce a strict improvement.

Towards a contradiction, suppose there exists such a Pareto-ranked splitting $(\overline{\sigma}, \underline{\sigma}, \alpha)$ with $\alpha\overline{\sigma} + (1-\alpha)\underline{\sigma} \in \Sigma_+(u)$ (the other case is symmetric), then we have

$$\begin{aligned}
&\alpha\Phi_u(\overline{\sigma}) + (1-\alpha)\Phi_u(\underline{\sigma}) - \Phi_u(\sigma_\theta) \\
\geq &\frac{\alpha}{\phi_r'(u_r(\overline{\sigma}))}\left[\phi_s(u_s(\overline{\sigma})) - \phi_s(u_s(\sigma_\theta))\right] + \frac{1-\alpha}{\phi_r'(u_r(\sigma_\theta))}\left[\phi_s(u_s(\underline{\sigma})) - \phi_s(u_s(\sigma_\theta))\right]
\end{aligned}$$

$$\geq \alpha(1-\alpha)\big(u_s(\overline{\sigma}) - u_s(\underline{\sigma})\big)\left[\frac{\phi_s'(u_s(\overline{\sigma}))}{\phi_r'(u_r(\overline{\sigma}))} - \frac{\phi_s'(u_s(\underline{\sigma}))}{\phi_r'(u_r(\sigma_\theta))}\right] > 0,$$

where the first inequality follows from $u_s(\sigma_\theta) > u \geq u_s(\underline{\sigma})$ and concavity of $\phi_r$, the second from concavity of $\phi_s$ and $\sigma_\theta = \alpha\overline{\sigma} + (1-\alpha)\underline{\sigma}$, and the last from linearity of $\phi_s$ and $\phi_r'(u_r(\overline{\sigma})) < \phi_r'(u_r(\sigma_\theta))$. □

**Remark A.2.** From the proof of Lemma A.8 (ii), linearity of $\phi_s$ can be relaxed to $\frac{\phi_s'(u_s(\overline{\sigma},\tau^*))}{\phi_s'(u_s(\underline{\sigma},\tau^*))} > \frac{\phi_r'(u_r(\overline{\sigma},\tau^*))}{\phi_r'(u_r(\sigma_\theta,\tau^*))}$, if $\sigma_\theta \in \Sigma_+(u)$, and $\frac{\phi_s'(u_s(\overline{\sigma},\tau^*))}{\phi_s'(u_s(\underline{\sigma},\tau^*))} > \frac{\phi_r'(u_r(\sigma_\theta,\tau^*))}{\phi_r'(u_r(\underline{\sigma},\tau^*))}$, if $\sigma_\theta \in \Sigma_-(u)$. □

*Proof of Proposition 2.* Proposition 2 is implied by the following two auxiliary lemmas that extend Lemma A.8. The proofs of these lemmas are similar to that of Lemma A.8 and may be found in Appendix E of the working paper version (Cheng et al., 2026).

**Lemma A.9.** *Let $\lambda \in \Delta_f(\Sigma)$ be a solution to ($\Phi_u^*$). Then:*

*(i) If $1/\phi_r'$ is concave, then for all $\sigma_\theta, \sigma_{\theta'} \in$ supp$(\lambda)$ such that $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_+(u)\times\Sigma_+(u)$, if $u_s(\sigma_\theta) \neq u_s(\sigma_{\theta'})$ and $\phi_r'(u_r(\sigma_\theta)) \neq \phi_r'(u_r(\sigma_{\theta'}))$, they are Pareto-ranked.*

*(ii) If $1/\phi_r'$ is convex, then for all $\sigma_\theta, \sigma_{\theta'} \in$ supp$(\lambda)$ such that $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_-(u) \times \Sigma_-(u)$, if $u_s(\sigma_\theta) \neq u_s(\sigma_{\theta'})$ and $\phi_r'(u_r(\sigma_\theta)) \neq \phi_r'(u_r(\sigma_{\theta'}))$, they are Pareto-ranked.*

**Lemma A.10.** *Let $\lambda \in \Delta_f(\Sigma)$ be a solution to ($\Phi_u^*$). Then:*

*(i) If $1/\phi_r'$ is concave, then for all $\sigma_\theta \in$ supp$(\lambda)$, there does not exist a Pareto-ranked splitting of $\sigma_\theta$, $(\overline{\sigma}, \underline{\sigma}, \alpha)$ with $(\overline{\sigma}, \underline{\sigma}) \in \Sigma_-(u) \times \Sigma_-(u)$, $\phi_r'(u_r(\overline{\sigma})) < \phi_r'(u_r(\underline{\sigma}))$, and $\phi_s'(u_s(\overline{\sigma}))/\phi_s'(u_s(\underline{\sigma})) > \phi_r'(u_r(\overline{\sigma}))/\phi_r'(u_r(\underline{\sigma}))$.*

*(ii) If $1/\phi_r'$ is convex, then for all $\sigma_\theta \in$ supp$(\lambda)$, there does not exist a Pareto-ranked splitting of $\sigma_\theta$, $(\overline{\sigma}, \underline{\sigma}, \alpha)$ with $(\overline{\sigma}, \underline{\sigma}) \in \Sigma_+(u) \times \Sigma_+(u)$, $\phi_r'(u_r(\overline{\sigma})) < \phi_r'(u_r(\underline{\sigma}))$, and $\phi_s'(u_s(\overline{\sigma}))/\phi_s'(u_s(\underline{\sigma})) > \phi_r'(u_r(\overline{\sigma}))/\phi_r'(u_r(\underline{\sigma}))$.*

With their swapped versions, when $1/\phi_r'$ is affine, Lemma A.8 (i) and A.9 imply the first conclusion. If $\phi_s$ is also affine, Lemmas A.8 (ii) and A.10 imply the second conclusion. □

## A.6 Proofs for Section 8

*Proof of Theorem 5.* Fix belief-threshold game with $\sum_{i=1}^n p(\omega_i)/T_i < 1$. Let $\sigma^{BP}$ be defined as in the proof of Theorem 3. Let $a_{\bar{i}}$ be the safe action and normalize $u_r(a_{\bar{i}}, \omega) = d$.

**If $\bar{i} > 0$.** Without loss of generality, let $\bar{i} = 1$. We show the sender benefits by using the ambiguous experiment $\mu$ such that $\mu(\sigma_{\theta_1}) = \mu(\sigma_{\theta_2}) = 1/2$, where

$$\begin{aligned}
&\sigma_{\theta_1}(a_i|\omega_j) = \sigma_{\theta_2}(a_i|\omega_j) = \sigma^{BP}(a_i|\omega_j), \quad \text{for all } i = 2, \dots, n, \text{ and } j = 0, \dots, n,\\
&\sigma_{\theta_1}(m_0|\omega_0) = \sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0), \quad \sigma_{\theta_1}(m_1|\omega_1) = 1,\\
&\sigma_{\theta_2}(m_1|\omega_0) = \sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0), \quad \sigma_{\theta_2}(m_0|\omega_1) = 1,
\end{aligned}$$

and all other probabilities are zero.

The receiver's posterior over $\omega$ conditional on $a_i$ for $i > 1$ is the same as $\sigma^{BP}$. Thus, the receiver's optimal action remains the corresponding $a_i$. By symmetry, the receiver's optimal action is the same after observing $m_0$ or $m_1$. Since an interim best response can put positive probability on at most one $a_i$ with $i > 0$ (see the proof of Lemma A.11) and $a_0$ dominates other $a_i$ with $i > 1$ in states $\omega_0$ and $\omega_1$, the receiver chooses $\alpha a_1 + (1-\alpha)a_0$ for some $\alpha \in [0, 1]$. Thus, the sender benefits if $\alpha$ is strictly greater than

$$\alpha^{BP} := \frac{p(\omega_1) + p(\omega_0)\sigma^{BP}(a_1|\omega_0)}{p(\omega_1) + p(\omega_0)(\sigma^{BP}(a_1|\omega_0) + \sigma^{BP}(a_0|\omega_0))}.$$

The receiver's posterior beliefs over $\omega$ conditional on message $m_0$ and $\sigma_\theta$ are given by $q_{\theta_1}(\omega_0|m_0) = 1$ and $q_{\theta_2}(\omega_1|m_0) = 1$, with posterior beliefs over $\Sigma$ given by

$$\eta(\sigma_{\theta_1}|m_0) = \frac{p(\omega_0)(\sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0))}{p(\omega_0)(\sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0)) + p(\omega_1)} \text{ and } \eta(\sigma_{\theta_2}|m_0) = 1 - \eta(\sigma_{\theta_1}|m_0).$$

The receiver's expected payoff from taking $\alpha a_1 + (1-\alpha)a_0$ is given by

$$\begin{aligned}
V(\alpha) :=& \frac{p(\omega_0)(\sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0))}{p(\omega_0)(\sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0)) + p(\omega_1)}\phi_r(\alpha d + (1-\alpha)(d + k_1 T_1))\\
&+ \frac{p(\omega_1)}{p(\omega_0)(\sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0)) + p(\omega_1)}\phi_r(\alpha d + (1-\alpha)(d - k_1(1-T_1))).
\end{aligned}$$

Since $V(\alpha)$ is concave in $\alpha$, it holds that the optimal choice of $\alpha$ is strictly greater than $\alpha^{BP}$ if the first-order condition at $\alpha^{BP}$ is strictly positive, i.e., $V'(\alpha^{BP}) > 0$, which says

$$\begin{aligned}
\frac{\phi_r'(\alpha^{BP} d + (1-\alpha^{BP})(d - k_1(1-T_1)))}{\phi_r'(\alpha^{BP} d + (1-\alpha^{BP})(d + k_1 T_1))} &> \frac{p(\omega_0)(\sigma^{BP}(a_0|\omega_0) + \sigma^{BP}(a_1|\omega_0))}{p(\omega_1)}\frac{T_1}{1-T_1}\\
&= 1 + \frac{T_1}{p(\omega_1)(1-T_1)}\left(1 - \sum_{i=1}^{n}\frac{p(\omega_i)}{T_i}\right) > 1.
\end{aligned}$$

This is satisfied when $\phi_r$ is sufficiently concave in the relevant range.

**If $\bar{i} = 0$.** We first establish the following lemma.

**Lemma A.11.** *Suppose $u_r(a_0, \omega)$ is constant across $\omega \in \Omega$. If an inconsistent receiver's interim best response after message $m$ puts positive probability on action $a_i$ for some $i > 0$, then the average posterior probability of $\omega_i$ conditional on $m$ must be at least $T_i$, i.e., $\overline{q}(\omega_i|m) := \sum_\theta \eta(\sigma_\theta|m) q_\theta(\omega_i|m) \geq T_i$.*

*Proof of Lemma A.11.* Any interim best response $\tau(m)$ is optimal under an effective posterior belief $\tilde{q} \in \Delta(\Omega)$. Assigning positive probability to $a_i$, $i > 0$, therefore requires $\tilde{q}(\omega_i) \geq T_i > 1/2$, which can hold for at most one $i$.

Suppose $\tau(m) = \alpha a_i + (1-\alpha)a_0$ is an interim best response for some $i > 0$ and $\alpha \in (0, 1]$, and let $u_r(a_0, \omega) = d$. Then

$$\phi_r(d) \leq \sum_\theta \eta(\sigma_\theta|m) \phi_r\big(d + \alpha k_i (q_\theta(\omega_i|m) - T_i)\big) \leq \phi_r\big(d + \alpha k_i (\overline{q}(\omega_i|m) - T_i)\big),$$

where the first inequality follows from optimality relative to $a_0$, and the second from Jensen's inequality and concavity of $\phi_r$. Since $\phi_r$ is strictly increasing and $\alpha k_i > 0$, it follows that $\overline{q}(\omega_i|m) \geq T_i$. □

Fix any ambiguous experiment $\mu$ and the receiver's interim best response $\tau$. For each action $a_i$ with $i > 0$, let $M_i = \{m \in M : \tau(a_i|m) > 0\}$. By Lemma A.11, for each $m \in M_i$, it holds that $\overline{q}(\omega_i|m) \geq T_i$. Let $\pi(m) = \sum_\theta \mu_\theta \sum_\omega p(\omega)\sigma_\theta(m|\omega)$ denote the average total probability of message $m$. Then, by Bayes plausibility,

$$p(\omega_i) = \sum_{m \in M} \pi(m) \overline{q}(\omega_i|m) \geq \sum_{m \in M_i} \pi(m) \overline{q}(\omega_i|m) \geq T_i \sum_{m \in M_i} \pi(m).$$

Thus, $\sum_{m \in M_i} \pi(m) \leq p(\omega_i)/T_i$. Notice the sender's payoff is bounded above by

$$\phi_s^{-1}\left(\sum_\theta \mu_\theta \phi_s\left(\sum_{i=1}^n v_i \sum_{m \in M_i} \sum_\omega p(\omega)\sigma_\theta(m|\omega)\right)\right) \leq \sum_{i=1}^n v_i \sum_{m \in M_i} \pi(m) \leq \sum_{i=1}^n v_i \frac{p(\omega_i)}{T_i} = u_s^{BP}.$$

Therefore, no ambiguous experiment can give the sender a strictly higher payoff than $u_s^{BP}$, and the sender does not benefit from ambiguous communication. □

# References


Alonso, R., Câmara, O. 2016. Bayesian persuasion with heterogeneous priors. Journal of Economic Theory. 165, 672–706.

Arieli, I., Babichenko, Y., Jain, A., Smorodinsky, R. 2025. Efficiency in games with incomplete information. arXiv preprint arXiv:2510.12508.

Aumann, R. J., Maschler, M. 1966. Game theoretic aspects of gradual disarmament. Report of the US Arms Control and Disarmament Agency. 80, 1–55.

Aybas, Y. C., Turkel, E. 2026. Persuasion with coarse communication. Economic Theory.

Beauchêne, D., Li, J., Li, M. 2019. Ambiguous persuasion. Journal of Economic Theory. 179, 312–365.

Bergemann, D., Morris, S. 2019. Information design: A unified perspective. Journal of Economic Literature. 57 (1), 44–95.

Bose, S., Renou, L. 2014. Mechanism design with ambiguous communication devices. Econometrica. 82 (5), 1853–1872.

Chakraborty, A., Harbaugh, R. 2010. Persuasion by cheap talk. American Economic Review. 100 (5), 2361–2382.

Cheng, X. 2022. Relative maximum likelihood updating of ambiguous beliefs. Journal of Mathematical Economics. 99, 102587.

Cheng, X. 2025. Ambiguous persuasion: An ex-ante formulation. Games and Economic Behavior. 154, 149–158.

Cheng, X., Klibanoff, P., Mukerji, S., Renou, L. 2026. Persuasion with ambiguous communication. arXiv preprint arXiv:2410.05504.

Di Tillio, A., Kos, N., Messner, M. 2017. The design of ambiguous mechanisms. The Review of Economic Studies. 84 (1), 237–276.

Dütting, P., Feldman, M., Peretz, D., Samuelson, L. 2024. Ambiguous contracts. Econometrica. 92 (6), 1967–1992.

Dworczak, P., Pavan, A. 2022. Preparing for the worst but hoping for the best: Robust (bayesian) persuasion. Econometrica. 90 (5), 2017–2051.

Galperti, S. 2019. Persuasion: The art of changing worldviews. American Economic Review. 109 (3), 996–1031.

Ghirardato, P., Maccheroni, F., Marinacci, M. 2004. Differentiating ambiguity and ambiguity attitude. Journal of Economic Theory. 118 (2), 133–173.

Guo, H. 2019. Mechanism design with ambiguous transfers: An analysis in finite dimensional naive type spaces. Journal of Economic Theory. 183, 76–105.

Hanany, E., Klibanoff, P. 2007. Updating preferences with multiple priors. Theoretical Economics. 2 (3), 261–298.

Hanany, E., Klibanoff, P. 2009. Updating ambiguity averse preferences. The BE Journal of Theoretical Economics. 9 (1).

Hanany, E., Klibanoff, P., Mukerji, S. 2020. Incomplete information games with ambiguity averse players. American Economic Journal: Microeconomics. 12 (2), 135–87.

Hedlund, J., Kauffeldt, T. F., Lammert, M. 2020. Persuasion under ambiguity. Theory and Decision. 1–28.

Kamenica, E. 2019. Bayesian persuasion and information design. Annual Review of Economics. 11, 249–272.

Kamenica, E., Gentzkow, M. 2011. Bayesian persuasion. American Economic Review. 101 (6), 2590–2615.

Kellner, C., Le Quement, M. T. 2017. Modes of ambiguous communication. Games and Economic Behavior. 104, 271–292.

Kellner, C., Le Quement, M. T. 2018. Endogenous ambiguity in cheap talk. Journal of Economic Theory. 173, 1–17.

Klibanoff, P., Marinacci, M., Mukerji, S. 2005. A smooth model of decision making under ambiguity. Econometrica. 73 (6), 1849–1892.

Kosterina, S. 2022. Persuasion with unknown beliefs. Theoretical Economics. 17 (3), 1075–1107.

Laclau, M., Renou, L. 2017. Public persuasion. mimeo. Paris School of Economics.

Lipnowski, E., Ravid, D. 2020. Cheap talk with transparent motives. Econometrica. 88 (4), 1631–1660.

Nikzad, A. 2021. Persuading a pessimist: Simplicity and robustness. Games and Economic Behavior. 129, 144–157.

Pahlke, M. 2023. Dynamic consistency and ambiguous communication. mimeo. Corvinus University of Budapest.

Pratt, J. W. 1964. Risk aversion in the small and in the large. Econometrica. 32 (1/2), 122–136.

Rayo, L., Segal, I. 2010. Optimal information disclosure. Journal of Political Economy. 118 (5), 949–987.

Siniscalchi, M. 2011. Dynamic choice under ambiguity. Theoretical Economics. 6 (3), 379–421.

Tang, R., Zhang, M. 2021. Maxmin implementation. Journal of Economic Theory. 194, 105250.

Wolitzky, A. 2016. Mechanism design with maxmin agents: Theory and an application to bilateral trade. Theoretical Economics. 11 (3), 971–1004.