

Persuasion with Ambiguous Communication*

Xiaoyu Cheng[†]

Peter Klibanoff[‡]

Sujoy Mukerji[§]

Ludovic Renou[¶]

This Version: November 7, 2024
Click [here](#) for the newest version.

Abstract

This paper explores whether and to what extent ambiguous communication can be beneficial to the sender in a persuasion problem, when the receiver (and possibly the sender) is ambiguity averse. We provide a concavification-like characterization of the sender's optimal ambiguous communication. The characterization highlights the necessity of using a collection of experiments that form a splitting of an obedient (i.e., incentive compatible) experiment. Some experiments in the collection must be Pareto-ranked in the sense that both players agree on their payoff ranking. The existence of a binary such Pareto-ranked splitting is necessary for ambiguous communication to benefit the sender, and, if an optimal Bayesian persuasion experiment can be split in this way, this is sufficient for an ambiguity-neutral sender as well as the receiver to benefit. Such gains are impossible when the receiver has only two actions. The possibility of gains is substantially robust to (non-extreme) sender ambiguity aversion.

*We thank Ian Ball, Piotr Dworczak, Huiyi Guo, Jian Li, Alessandro Pavan and seminar audiences at Texas A&M, Washington University at 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 for helpful feedback. A two-page abstract of an earlier version of this paper appeared in EC'24.

[†]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.

[¶]School of Economics and Finance, Queen Mary University of London, London, UK, 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). The key departures from most of the literature and the focus of our analysis are that we enlarge the set of the sender’s communication strategies to include ambiguous strategies – strategies for which, from the perspective of both players, the probability that a given statistical experiment will be used to generate the signal is subjectively uncertain, and the receiver (and possibly the sender) is assumed to be ambiguity averse (i.e., averse to this subjective uncertainty about these probabilities). In such an environment, would the sender ever benefit from intentionally using an ambiguous communication strategy? If so, we would like to understand when and why this might occur.

What might it mean in a real-world context for the sender to choose an ambiguous communication strategy? Consider, for example, a pharmaceutical company communicating with a health authority that is responsible for deciding on the approval of a drug. This communication often takes the form of results from clinical trials, and these trials are frequently outsourced to sub-contractors. From the perspectives of both the company and the health authority, this sub-contracting may be viewed as introducing additional uncertainty. In particular, the instructions to the sub-contractors involving what experiments to carry out can be (and in practice are) made contingent to a greater or lesser extent on the knowledge/experience/discretion of the sub-contractors and their specialized expertise. This contingent nature of the instructions to the sub-contractors, though known to all parties, implies that there will be some ambiguity, on the part of both the company and the health authorities, about exactly how to interpret the clinical trial results, driven by the underlying ambiguity that the company and health authority have about

the exact knowledge/experience/discretion of the sub-contractors. Abstracting from standard cost/efficiency motivations for sub-contracting, our theory suggests that there may be purely strategic reasons for the pharmaceutical company to use contingent sub-contracting to inject such ambiguity. Another context in which choice to communicate ambiguously arises is banking regulators' stress testing of financial institutions. In this setting, ambiguity can and often is introduced into communication by banking regulators choosing to use as input to the stress tests "bottom-up" tests conducted by the banks themselves based on their own private in-house models and data.

Like sub-contracting to third parties, the choice to use a particular AI algorithm to generate recommendations or diagnoses can be viewed as a communication strategy that is seen as more or less ambiguous by both players (in addition to possibly varying in overall accuracy). Such algorithms are known to differ in their degrees of interpretability or transparency. These dimensions are acknowledged as relevant in choosing an algorithm.¹ One aspect of interpretability relates to mapping from instances to accuracy properties. For example, a linear regression model is unambiguous in this respect because it is clear how changes in the input characteristics affect the accuracy of the prediction. In contrast, the accuracy of an individual prediction from a deep neural network is more ambiguous in that it is more difficult to determine whether any given instance is one for which the algorithm is likely to predict more or less accurately. This is so even though the neural network may be more accurate on average across inputs than linear regression.

Even though it is often possible to eliminate some or all of the ambiguity if one wishes to, one might benefit from deliberately introducing or maintaining some ambiguity. The examples above suggest sub-contracting to third-parties or delegation to an algorithm as some practical ways to do so, but our analysis and model are agnostic about how such ambiguity might be committed to – we assume an ability to commit to communication strategies, including ambiguous ones, and study when choosing ambiguous strategies is beneficial and why. Some intuition for how it might be beneficial is as follows. When confronted with a host of possible interpretations of the same evidence, ambiguity aversion motivates the receiver to value hedging against variation in expected payoffs across these interpretations. This leads them to best respond as if, compared to an ambiguity-neutral receiver, they overweight the interpretations that give them less favorable expected payoffs. The resulting change in best response may potentially benefit the sender. The more ambiguity averse the receiver is, the more they effectively overweight the less favorable interpreta-

¹See, for example, [Linardatos et al. \(2021\)](#), [Schmitt \(2024\)](#), and [Telus International Website \(2021\)](#).

tions, and the more scope there is for the sender to potentially benefit. While correct, this intuition is quite incomplete – it gives no sense of when this ability to induce different best responses can benefit the sender, nor of the characteristics of the induced ambiguity that will deliver such benefits. Our analysis provides a concavification-like characterization of both the sender’s optimal strategy and when the sender may benefit from the ability to communicate ambiguously. We highlight the necessity of generating ambiguity using a collection of statistical experiments that form a splitting of an experiment whose messages are incentive-compatible action recommendations for the receiver. At least some of the experiments in this collection must be Pareto-ranked in the sense that both the sender and receiver agree on their payoff ranking if their recommendations are followed. The existence of a two-experiment collection forming such a Pareto-ranked splitting is necessary for ambiguous communication to benefit the sender, and, if an optimal Bayesian persuasion experiment can be split in this way, this is sufficient for an ambiguity-neutral sender as well as the receiver to benefit. Surprisingly to us, we prove that this necessary condition is *never* met in problems with binary actions, encompassing many examples in the literature.

The rest of the paper is organized as follows. The next section illustrates the main intuition and some of our results with the help of a simple example and offers a brief discussion of the related literature. (Section 7.2 contains a more extensive discussion.) Section 3 presents the model. Results are in Sections 4 through 6. The main proofs are in the Appendix. Additional material is in the Online Appendix (Cheng et al., 2024).

2 An Introductory Example

We illustrate our main results with the help of a simple example. There is a sender and a receiver, three actions a_1, a_2 and a_3 , and two payoff-relevant states ω_1 and ω_2 , with equal prior probabilities $p = (1/2, 1/2)$.² The sender influences the action the receiver takes with the release of information. The payoffs are:

²The example needs at least three actions since we show (Corollary 3) there is no benefit from ambiguous experiments when the receiver has only two actions.

(u_s, u_r)	a_1	a_2	a_3
ω_1	1, 1	-1, -1	-4, 2
ω_2	0, 0	2, 2	-4, -4

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

The receiver prefers a_3 in state ω_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 ω_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 $1/5$. In the rest of the example, we omit the state when speaking about beliefs – all beliefs are about ω_2 . An interpretation of this example in the context of stress testing and banking regulation is offered in the Online Appendix (Cheng et al., 2024).

We first apply the seminal work of Kamenica and Gentzkow (2011) on Bayesian persuasion to this example. Kamenica and Gentzkow (2011) study a dynamic game between a sender and a receiver, where the sender first designs a statistical experiment $\sigma : \{\omega_1, \omega_2\} \rightarrow \Delta(S)$, the receiver observes the chosen experiment σ and the outcome s , and then chooses an action. In our language, the information design is *unambiguous*, that is, upon observing a signal, the receiver knows the experiment that generated the signal and, therefore, knows how to interpret it. Kamenica and Gentzkow (2011) show that the highest payoff the sender can achieve is the value of the concavification of their indirect utility at the prior p . In our example, the best the sender can do is to induce the beliefs $1/5$ and 1 , resulting in a payoff of $5/4$ – see Figure 1 for a graphical illustration. In Figure 1, we plot the receiver’s expected payoff associated with each of the three actions as dotted lines – each line is labelled with its action. We plot the sender’s indirect utility, i.e., the utility the sender obtains when the receiver chooses an optimal action, as a thick solid curve, and its concavification as a thick dashed curve.

It is immediate to verify that the experiment σ^{BP} implements the splitting of the prior into the beliefs $1/5$ and 1 . The signal “ a_2 ” reveals that the state is ω_2 , while the signal “ a_1 ” leaves some uncertainty. Intuitively, since the preferences are perfectly aligned when the state is ω_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 ω_2 – the receiver chooses a_3 at all beliefs less than $1/5$. The optimal experiment σ^{BP} balances these two forces. We note that the experiment σ^{BP} is canonical, that is, it recommends actions and the receiver finds it optimal to obey the recommendations. This is without loss of generality, and we prove (Proposition




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

1) that this continues to be without loss in our generalization. We therefore restrict attention to canonical experiments in what follows.

Now, suppose that the sender can design *ambiguous* experiments. An experiment is ambiguous when the receiver is unsure how to interpret signals. In other words, the receiver is uncertain about the true signal-generating process, analogously to a statistician uncertain about the true data-generating process.

We model ambiguous experiments as (finitely supported) *distributions* over experiments, that is, an ambiguous experiment is a tuple $(\mu_\theta, \sigma_\theta)_\theta$, with μ_θ the probability of experiment σ_θ . Upon observing a signal from an ambiguous experiment, the receiver does not know which experiment σ_θ generated it and, therefore, is uncertain how to interpret it. We interpret θ as a realization of some ambiguity the sender and receiver perceive in their environment. For example, as discussed in the introduction, outsourcing clinical trials to third parties creates ambiguity for both the health authorities and pharmaceutical companies.

An ambiguity-neutral receiver treats the ambiguous experiment $(\mu_\theta, \sigma_\theta)_\theta$ as equivalent to the unambiguous experiment $\sum_\theta \mu_\theta \sigma_\theta$. In this case, the problem is identical to Bayesian persuasion, and the sender cannot benefit from ambiguity. We assume instead that the re-

ceiver is ambiguity averse and represent their preference with the smooth ambiguity model of [Klibanoff et al. \(2005\)](#). Specifically, let $u_r(\sigma_\theta, \tau^*)$ be the receiver's payoff when the (canonical) experiment is σ_θ and the receiver is obedient.³ The receiver values the ambiguous experiment as $\phi_r^{-1}(\sum_\theta \mu_\theta \phi_r(u_r(\sigma_\theta, \tau^*)))$, where ϕ_r is some strictly increasing, concave and differentiable function. The concavity of ϕ_r captures ambiguity aversion. Greater concavity corresponds to more ambiguity aversion. At one extreme, when the receiver is infinitely ambiguity averse, we have an instance of the maxmin expected utility (MEU) model ([Gilboa and Schmeidler, 1989](#)). At the other, when ϕ_r is affine, we have the expected utility model (implying ambiguity neutrality).

As a preliminary result, we show ([Lemma 1](#)) that the receiver's behavior is equivalent to that of an ambiguity-neutral receiver, but one who uses probabilities $(\nu_\theta)_\theta$ – which we call the effective measure – that generally differ from $(\mu_\theta)_\theta$. This result implies that the ambiguity-averse receiver is obedient if, and only if, the ambiguity-neutral receiver is obedient when facing the unambiguous experiment $\sum_\theta \nu_\theta \sigma_\theta$. We use this result throughout the analysis. At least as importantly, [Lemma 1](#) describes how ambiguity aversion imposes structure on ν_θ and that ν_θ is a function of the entire profile $(\mu_\theta, u_r(\sigma_\theta, \tau^*))_\theta$ as well as the receiver's ambiguity aversion. Thus, the ν_θ are endogenous. Even local changes in the ambiguous experiment, say only changing σ_θ to σ'_θ , impact *all* ν_θ . Furthermore, if $u_r(\sigma_\theta, \tau^*) < u_r(\sigma_{\theta'}, \tau^*)$, then $\nu_\theta/\nu_{\theta'} > \mu_\theta/\mu_{\theta'}$, that is, the effective measure assigns a higher (relative) probability than $(\mu_\theta)_\theta$ to lower payoffs. The more ambiguity averse the receiver, the higher is the (relative) effective probability on the lower payoff. These properties distinguish our model from a model with exogenously fixed heterogeneous priors, e.g., [Alonso and Câmara \(2016\)](#), [Laclau and Renou \(2017\)](#) and [Galperti \(2019\)](#).

We now illustrate how ambiguous experiments can benefit the sender in the example. Consider an ambiguous experiment such that only two experiments $\sigma_{\underline{\theta}}$ and $\sigma_{\bar{\theta}}$ get positive μ -weight. The experiment $\sigma_{\underline{\theta}}$ is uninformative and defined by $\sigma_{\underline{\theta}}(a_1|\omega_1) = \sigma_{\underline{\theta}}(a_1|\omega_2) = 1$, while the experiment $\sigma_{\bar{\theta}}$ is fully informative and defined by $\sigma_{\bar{\theta}}(a_1|\omega_1) = \sigma_{\bar{\theta}}(a_2|\omega_2) = 1$. Observe that the interpretation of the signal/recommendation a_2 is unambiguous: the receiver learns that the state is ω_2 . The interpretation of the signal/recommendation a_1 is, however, ambiguous: either it means that the state is ω_1 (when $\sigma_{\bar{\theta}}$ generated the signal) or it means nothing (when $\sigma_{\underline{\theta}}$ generated the signal). The associated payoff profiles are $(u_s(\sigma_{\underline{\theta}}, \tau^*), u_r(\sigma_{\underline{\theta}}, \tau^*)) = (1/2, 1/2)$ and $(u_s(\sigma_{\bar{\theta}}, \tau^*), u_r(\sigma_{\bar{\theta}}, \tau^*)) = (3/2, 3/2)$. Thus, if $\mu_{\bar{\theta}} > 3/4$, an ambiguity-neutral sender's expected payoff is strictly higher than the

³Obedient in the sense of following the action recommendations. We denote the obedient strategy by τ^* .

Bayesian persuasion payoff of $5/4$.⁴ We now argue that we can simultaneously choose $\mu_{\bar{\theta}} > 3/4$ and guarantee obedience. First, observe that $(1/4)\sigma_{\underline{\theta}} + (3/4)\sigma_{\bar{\theta}} = \sigma^{BP}$ – we call such a configuration a splitting of σ^{BP} . Since the receiver is obedient when facing σ^{BP} , the receiver is obedient when the effective weight $\nu_{\bar{\theta}}$ equals $3/4$. In fact, the receiver continues to be obedient for any effective weight weakly below $3/4$. Second, since $1/2 = u_r(\sigma_{\underline{\theta}}, \tau^*) < u_r(\sigma_{\bar{\theta}}, \tau^*) = 3/2$, $\nu_{\bar{\theta}}$ is strictly lower than $\mu_{\bar{\theta}}$ (unless the receiver is ambiguity neutral) – as mentioned above, this is a consequence of ambiguity aversion. Therefore, since $\nu_{\bar{\theta}} < 3/4$ when $\mu_{\bar{\theta}} = 3/4$, there is room to increase $\mu_{\bar{\theta}}$ above $3/4$ and maintain obedience until the point where $\nu_{\bar{\theta}}$ equals $3/4$.⁵ This construction is illustrated in Figure 2, where the thick arrow moving along the sender’s indirect utility curve indicates the movement of $\nu_{\bar{\theta}}$ towards $3/4$ from below as $\mu_{\bar{\theta}}$ increases above $3/4$ (along the thick arrow next to $\mu_{\bar{\theta}}$). Thus, the ambiguous communication strategy allows the sender to place more weight on the experiment $\sigma_{\bar{\theta}}$ while maintaining obedience, than would be possible with unambiguous communication. This is how ambiguous communication provides benefits.

An important observation is that the experiments $\sigma_{\underline{\theta}}$ and $\sigma_{\bar{\theta}}$ are Pareto-ranked. If they were not, then ambiguity aversion would push the receiver’s effective measure in a direction that would *hurt* rather than help the sender. In fact, we prove that the existence of a two-experiment Pareto-ranked splitting of some (unambiguous) obedient experiment is necessary for ambiguous experiments to benefit the sender over Bayesian persuasion (Theorem 4). Moreover, any ambiguous experiment delivering such benefit must assign positive μ_{θ} -weight to some pair of experiments that are Pareto-ranked and such that the better of the two improves on Bayesian persuasion for the sender, assuming its recommendations were to be followed (Theorem 3). If, as here, σ^{BP} can be split in this way, this is sufficient for an ambiguity-neutral sender (and the receiver too!) to benefit (Theorem 5 and Corollary 5).

We close this section with a brief discussion of a few closely related papers. A more extensive discussion can be found in Section 7.2. [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 signal-contingent strategy (see Remark 1 for an equiv-

⁴The same arguments remain valid as long as the sender is not too ambiguity averse. In particular, the sender continues to benefit even if they are as ambiguity averse as the receiver (and even a bit more so) as long as the sender is not infinitely ambiguity averse.

⁵The effective weight $\nu_{\bar{\theta}}$ is $3/4$ when $\mu_{\bar{\theta}} = \frac{3\phi'_r(1/2)}{3\phi'_r(1/2) + \phi'_r(3/2)}$, with ϕ'_r the derivative of ϕ_r .




Figure 2: Construction of the ambiguous experiment

alent interim implementation). BLL assume the receiver chooses, for each signal, actions maximizing interim preferences formed using a belief updating rule that leads to dynamic inconsistency with their ex-ante preference. Thus, one 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 (see our further discussion in Section 7.2, including the approach to consistency of Pahlke (2023)). 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 (2023) 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 (2023)'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 will allow for varying degrees of ambiguity aversion for both the sender and the receiver.

3 The Persuasion Problem with Ambiguous Communication

We consider a persuasion game between a sender and a receiver, where the sender can choose *ambiguous* experiments if they wish.

3.1 The Model

There is a finite set Ω of payoff-relevant states ω , with common prior probability distribution $p \in \Delta(\Omega)$. 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 ω . A *statistical experiment* is a finite set of messages M and a map σ from Ω to $\Delta(M)$, and we write $\sigma(m|\omega)$ for the probability of m given ω .

We assume that the sender can condition their statistical experiment on the realization of a *source of ambiguity*, which we define as a finite set Θ of *payoff-irrelevant ambiguous events* together with a probability distribution $\mu \in \Delta(\Theta)$. The sender thus chooses a source of ambiguity μ and a collection, $\sigma := (\sigma_\theta)_{\theta \in \Theta}$, specifying a statistical experiment for each event $\theta \in \Theta$.^{6,7} We stress that the set of payoff-irrelevant ambiguous events Θ is not fixed – the sender chooses it. For example, Θ might be a chosen partition of the continuum of values for a parameter about which there is ambiguity. We call a pair (σ, μ) an *ambiguous experiment*. Henceforth, whenever we use the term “experiment” without a modifier, it refers to a standard, unambiguous statistical experiment. Notice that ambiguous experiments are a generalization of experiments in the sense that any experiment can be viewed as an ambiguous experiment with a collection σ that does not vary with θ , i.e., $\sigma_\theta = \sigma$ for all $\theta \in \text{supp}(\mu)$.⁸ Of special interest in some of our later constructions are *binary ambiguous experiments*, those in which σ is a *binary collection* – collections of experiments such that $|\cup_{\theta \in \Theta} \{\sigma_\theta\}| = 2$ (i.e., collections such that exactly two distinct experiments appear in σ).

We analyze the receiver’s behavior from the perspective of their ex-ante preferences,

⁶It is without loss to assume that all statistical experiments in the collection share the same message space.

⁷The sender chooses and commits to μ and σ before θ and ω are realized. Thus, just as in standard Bayesian persuasion where the sender chooses and commits to σ before ω is realized, the sender’s choice of a communication strategy is influenced by their beliefs about how uncertainty may unfold. Since Θ is viewed as ambiguous, any ambiguity aversion on the part of the sender may influence their choice. See Section 7.1 for a discussion of how things would change under the alternative assumption that it is common knowledge that the sender privately learns θ before committing to a strategy.

⁸The notation $\text{supp}(\mu)$ means the support of μ .

that is, we assume that the receiver observes the sender's choice of (σ, μ) and commits to a strategy $\tau : M \rightarrow \Delta(A)$.⁹ In other words, we impose dynamic consistency.¹⁰ The main motivation is that we want to study whether the sender benefits from ambiguous communication even if the channel of dynamic inconsistency — the channel at work in nearly all previous literature on mechanism or information design with ambiguity — is shut down. We refer the interested reader to Section 7.2 for more discussion.

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

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

To isolate the role of intentional ambiguous communication, we work in a stylized environment where ambiguity is not payoff-relevant unless it becomes so by strategic choice of the sender to condition their communication on the realization of ambiguous events. Thus, while the payoff-irrelevant events $\theta \in \Theta$ are viewed as ambiguous, the payoff-relevant events $\omega \in \Omega$ and any randomization over messages induced by a statistical experiment are viewed as unambiguous. It follows that a message m is viewed as ambiguous by the sender and receiver only if the experiments the sender chooses to associate with distinct possible θ 's generate m with different positive likelihoods. These different likelihoods may lead the expected payoff $u_i(\sigma_\theta, \tau)$ to vary with θ and thus itself be viewed as ambiguous.

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

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

where $\phi_i : \mathbb{R} \rightarrow \mathbb{R}$ 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 subjective expected utility. 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.

⁹For compactness, this notation suppresses the allowed dependence of τ on (σ, μ) .

¹⁰In Remark 1 in Section 3.2.2, we point out a simple receiver's updating rule that guarantees dynamic consistency. Assuming such updating is an equivalent approach to modeling the receiver.

Observe that the only uncertainty treated as ambiguity is that over Θ and this ambiguity matters only to the extent that the expected payoff $u_i(\sigma_\theta, \tau)$ varies with θ . Recall that σ_θ is the θ -th coordinate of the vector σ . Note that we assume that the sender and the receiver share the same μ . We interpret μ as the “best guess” an ambiguity neutral individual would feel comfortable using to evaluate the likelihood of events $\theta \in \Theta$. An ambiguity averse individual, however, values robustness with respect to the uncertainty of this best guess. As we shall see, this translates into the distortion of μ into an *effective* measure – the “as-if” measure the individual uses to make robust decisions. We will also show (Proposition 3) that any benefits from ambiguous communication are robust to the receiver somewhat misperceiving (σ, μ) .

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

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

This MEU case is considered in Section SA.1 of the Online Appendix (Cheng et al., 2024).

Writing $BR(\sigma, \mu)$ for the set of best replies of the receiver (i.e., the maximizers of $U_r(\sigma, \mu, \tau)$ with respect to τ), the sender’s problem is:

$$(\mathcal{P}) = \begin{cases} \max_{(\sigma, \mu, \tau)} U_s(\sigma, \mu, \tau), \\ \text{subject to } \tau \in BR(\sigma, \mu). \end{cases}$$

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:¹¹

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

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

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

¹¹Formally, a collection σ that does not vary with θ , i.e., $\sigma_\theta = \sigma$ for all θ .

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

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

3.2 A Revelation Principle and Incentive Compatibility

3.2.1 A Revelation Principle

Definition 2. *An ambiguous experiment (σ, μ) is **canonical** if $M = A$.*

We write $\tau^* : A \rightarrow \Delta(A)$ 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 (σ, μ) is **obedient** if $\tau^* \in BR(\sigma, \mu)$.*

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 $((\sigma, \mu), \tau)$ such that $\tau \in BR(\sigma, \mu)$, there exists a canonical and obedient ambiguous experiment (σ^*, μ) such that $u_i(\sigma_\theta, \tau) = u_i(\sigma_\theta^*, \tau^*)$ for all $i \in \{s, r\}$ and θ .*

It is well-known that such a revelation principle holds in the persuasion game setting without ambiguity. However, one might have thought 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 τ^* ? 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 choosing a strategy τ to maximize $U_r(\sigma, \mu, \tau)$, their ex-ante payoff from $((\sigma, \mu), \tau)$, which imposes dynamic consistency on the receiver.

This revelation principle result is extremely useful in facilitating our analysis in the remainder of the paper. From here on, we restrict attention to canonical experiments, and represent incentive compatibility via obedience. 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.

3.2.2 Incentive Compatibility and Effective Measure

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

Definition 4. *Given an ambiguous experiment (σ, μ) , the receiver's effective measure $em^{(\sigma, \mu)} \in \Delta(\Theta)$ is given by:*

$$em_{\theta}^{(\sigma, \mu)} := \frac{\mu_{\theta} \phi_r'(u_r(\sigma_{\theta}, \tau^*))}{\sum_{\bar{\theta}} \mu_{\bar{\theta}} \phi_r'(u_r(\sigma_{\bar{\theta}}, \tau^*))}, \text{ for all } \theta \in \Theta. \quad (1)$$

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

The next result states that τ^* is the receiver's best response to the ambiguous experiment (σ, μ) if, and only if, it is a best response to the experiment, σ^* , defined below as the convex combination of the experiments in the collection σ with weights given by the receiver's effective measure.

Lemma 1. *The ambiguous experiment (σ, μ) is obedient if, and only if, the (unambiguous)*

experiment σ^* is obedient, where

$$\sigma^* = \sum_{\theta} em_{\theta}^{(\sigma, \mu)} \sigma_{\theta}.$$

Lemma 1 follows from the first-order conditions of the receiver's maximization problem $\max_{\tau} U_r(\sigma, \mu, \tau)$, evaluated at τ^* . 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,

$$\begin{aligned} 2 &= u_r(\bar{\sigma}, br(\mu\bar{\sigma} + (1 - \mu)\underline{\sigma})) > u_r(\bar{\sigma}, \tau^*) = 3/2 \\ &> u_r(\underline{\sigma}, \tau^*) = 1/2 \\ &> u_r(\underline{\sigma}, br(\mu\bar{\sigma} + (1 - \mu)\underline{\sigma})) = -1, \end{aligned}$$

showing that the obedience strategy τ^* is hedged against the uncertainty about the weight μ more than the strategy $br(\mu\bar{\sigma} + (1 - \mu)\underline{\sigma})$. 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 $(\sigma, em^{(\sigma, \mu)})$ would be obedient.

Remark 1. These properties of the effective measure also give rise to an updating implementation of the receiver's ex-ante optimality – a receiver who updates after observing a message using Bayes' rule with the effective measure in place of μ and adopts this update as their effective posterior will be dynamically consistent.

Finally, for later reference, observe that, fixing σ , we can invert (1) to express μ as a function of the effective measure it generates:

$$\mu_{\theta} = \frac{em_{\theta}^{(\sigma, \mu)} / \phi'_r(u_r(\sigma_{\theta}, \tau^*))}{\sum_{\bar{\theta}} em_{\bar{\theta}}^{(\sigma, \mu)} / \phi'_r(u_r(\sigma_{\bar{\theta}}, \tau^*))}. \quad (2)$$

4 Optimal Persuasion with Ambiguous Communication

In this section, we present a concavification-like characterization of optimal persuasion with ambiguous communication and then use it to derive necessary conditions for optimality. We stress that the concavification-like characterization is not immediate from the

one of [Kamenica and Gentzkow \(2011\)](#). A key complication is the non-separability across recommendations in determining obedience (coming from the appearance of the $u_r(\sigma_\theta, \tau^*)$ terms in the effective measure formula (1)), which is a consequence of ambiguity aversion.

4.1 A Concavification-like Characterization

We need to introduce some additional notation. Let Σ denote the set of all experiments and $\Sigma^* \subseteq \Sigma$ the set of obedient experiments (i.e., $\Sigma^* := \{\sigma \in \Sigma : \tau^* \in br(\sigma)\}$). Notice that both Σ and Σ^* are non-empty convex sets and can be embedded into an $|\Omega| \times (|A| - 1)$ -dimensional Euclidean space since an experiment specifies, for each state $\omega \in \Omega$, a probability distribution over actions in A .

For each scalar $u \in \mathbb{R}$, define the function $\Phi_u : \Sigma \rightarrow \mathbb{R}$ by

$$\Phi_u(\sigma) := \frac{\phi_s(u_s(\sigma, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\sigma, \tau^*))},$$

and consider the following maximization problem:

$$(\Phi^*(u)) := \begin{cases} \max_{(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}} \sum_{\theta \in \Theta} \lambda_\theta \Phi_u(\sigma_\theta), \\ \text{subject to: } \sum_{\theta \in \Theta} \lambda_\theta \sigma_\theta \in \Sigma^*, \sum_{\theta \in \Theta} \lambda_\theta = 1, \lambda_\theta \geq 0, \sigma_\theta \in \Sigma, \forall \theta \in \Theta. \end{cases}$$

Theorem 1 states that the value of the optimal ambiguous persuasion program (\mathcal{P}) is the unique utility level u such that the value of the program $(\Phi^*(u))$ is equal to zero. 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 $|\Omega| \times (|A| - 1) + 1$ experiments.

Theorem 1. *The value of (\mathcal{P}) is u if, and only if, the value of $(\Phi^*(u))$ is 0. Moreover, there exists a solution (σ^*, μ^*) to (\mathcal{P}) such that $|\text{supp}(\mu^*)| \leq |\Omega| \times (|A| - 1) + 1$.*

To understand the relationship between the programs (\mathcal{P}) and $(\Phi^*(u))$, we first note that the definition of U_s , the fact that ϕ_s^{-1} is strictly increasing, and the characterization of obedience in Lemma 1, implies that the value of (\mathcal{P}) is u if, and only if, the value of the program

$$(\hat{\mathcal{P}}) = \begin{cases} \max_{(\mu_\theta, \sigma_\theta)_{\theta \in \Theta}} \sum_{\theta \in \Theta} \mu_\theta \phi_s(u_s(\sigma_\theta, \tau^*)), \\ \text{subject to: } \sum_{\theta \in \Theta} em_\theta^{(\sigma, \mu)} \sigma_\theta \in \Sigma^*, \sigma_\theta \in \Sigma, \forall \theta \in \Theta, \end{cases}$$

is $\phi_s(u)$. Next, we can do a change of variables to maximize over the choice of effective measures and experiments. Formally, if we write λ_θ for $em_\theta^{(\sigma, \mu)}$, we can use (2) to substitute for μ_θ in terms of λ_θ to yield:

$$(\widehat{\mathcal{P}}) = \begin{cases} \max_{(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}} \left(\sum_{\tilde{\theta} \in \Theta} \frac{\lambda_{\tilde{\theta}}}{\phi'_r(u_r(\sigma_{\tilde{\theta}}, \tau^*))} \right)^{-1} \sum_{\theta \in \Theta} \frac{\lambda_\theta}{\phi'_r(u_r(\sigma_\theta, \tau^*))} \phi_s(u_s(\sigma_\theta, \tau^*)), \\ \text{subject to: } \sum_{\theta \in \Theta} \lambda_\theta \sigma_\theta \in \Sigma^*, \sum_{\theta} \lambda_\theta = 1, \lambda_\theta \geq 0, \sigma_\theta \in \Sigma, \forall \theta \in \Theta. \end{cases}$$

Finally, observe that the normalization factor $\left(\sum_{\tilde{\theta} \in \Theta} \frac{\lambda_{\tilde{\theta}}}{\phi'_r(u_r(\sigma_{\tilde{\theta}}, \tau^*))} \right)^{-1}$ makes the objective function highly non-linear in the maximizers $(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}$. This is the motivation for subtracting off $\phi_s(u)$. Indeed, if the value of $(\widehat{\mathcal{P}})$ is $\phi_s(u)$, then

$$\sum_{\theta \in \Theta} \lambda_\theta \frac{\phi_s(u_s(\sigma_\theta, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\sigma_\theta, \tau^*))} = 0.$$

Conversely, if the value of $(\Phi^*(u))$ is zero, then the value of $(\widehat{\mathcal{P}})$ is $\phi_s(u)$. In effect, this reformulation discards the messy (but strictly positive) normalization term without changing the solution.

We can go further. An object appearing in $(\Phi^*(u))$ that proves useful throughout the paper is a *splitting* of an experiment into a convex combination of experiments. The constraint $\sum_{\theta \in \Theta} \lambda_\theta \sigma_\theta \in \Sigma^*$ makes clear that any solution is a splitting of an obedient experiment, where the $(\lambda_\theta)_{\theta \in \Theta}$ are the splitting weights. Importantly, $(\Phi^*(u))$ is linear in these splitting weights.

An implication of this linearity and Theorem 1 is to provide a concavification-like characterization (Aumann and Maschler, 1966, 1995) of the value of optimal persuasion with ambiguous communication. Notice that concavification can be used to compute the value of program $(\Phi^*(u))$: For each $u \in \mathbb{R}$, the program $(\Phi^*(u))$ maximizes over convex combinations of points on the graph of Φ_u , exactly the type of program that concavification characterizes. Specifically, for each $u \in \mathbb{R}$, let $\text{cav}\Phi_u : \Sigma \rightarrow \mathbb{R}$ denote the concavification of Φ_u , that is,

$$\text{cav}\Phi_u(\sigma) = \begin{cases} \max_{(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}} \sum_{\theta \in \Theta} \lambda_\theta \Phi_u(\sigma_\theta), \\ \text{subject to: } \sum_{\theta \in \Theta} \lambda_\theta \sigma_\theta = \sigma, \sum_{\theta} \lambda_\theta = 1, \lambda_\theta \geq 0, \sigma_\theta \in \Sigma, \forall \theta \in \Theta, \end{cases}$$

and the maximum over $\sigma \in \Sigma^*$ of $\text{cav}\Phi_u(\sigma)$ is the value of $(\Phi^*(u))$. Observe that any

such maximum is achieved by a splitting of some obedient experiment, with the splitting weights given by the effective measure. The following immediate corollary of Theorem 1 thus provides a concavification-like characterization of the value of (\mathcal{P}) .

Corollary 1. *The value of (\mathcal{P}) is u if, and only if, $\max_{\sigma \in \Sigma^*} \text{cav} \Phi_u(\sigma) = 0$.*

Algorithmically, we can start with $u_0 = u_s^{BP}$, the payoff the sender obtains at a solution to (\mathcal{P}^{BP}) , which is a lower bound on what the sender can achieve with ambiguous communication. If $\max_{\sigma \in \Sigma^*} \text{cav} \Phi_{u_0}(\sigma) = 0$, then we are done – the sender’s best payoff is u_s^{BP} . If $\max_{\sigma \in \Sigma^*} \text{cav} \Phi_{u_0}(\sigma) > 0$, we can increase u_0 to $u_1 = \max_{a, \omega} u_s(a, \omega)$ and check again. If the solution is zero, we are done. If it is strictly negative, we can then consider the mid-point $u_2 = (1/2)u_0 + (1/2)u_1$. If $\max_{\sigma \in \Sigma^*} \text{cav} \Phi_{u_2}(\sigma) > 0$ (resp., < 0), we can then consider the midpoint $u_3 = (1/2)u_2 + (1/2)u_1$ (resp., $u_3 = (1/2)u_0 + (1/2)u_2$) and repeat the maximization problem, and so on.

We now relate this concavification-like result with the concavification characterization of Bayesian persuasion in [Kamenica and Gentzkow \(2011\)](#) and its extension to allow for exogenously heterogeneous priors in [Alonso and Câmara \(2016\)](#). Both of these characterizations are formulated in terms of splittings of priors, rather than, as in our characterization, splittings on the higher-dimensional space of experiments. 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, i.e., ϕ_s linear. Any ambiguous experiment $(\mu_\theta, \sigma_\theta)_{\theta \in \Theta}$ induces a distribution over the receiver’s *effective* posteriors, that is, the posteriors that the “effective” experiment $\sum_\theta \lambda_\theta \sigma_\theta$ induces, where $\lambda_\theta = em_\theta^{(\sigma, \mu)}$, the effective measure. Thus, the splitting the “effective” experiment $\sum_\theta \lambda_\theta \sigma_\theta$ 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 function that is linear in the distribution over the receiver’s effective posteriors. However, since the relationship between the $em_\theta^{(\sigma, \mu)}$ and $(\mu_\theta, \sigma_\theta)_{\theta \in \Theta}$ is highly non-linear, the desired linearity is impossible. Economically, this non-linearity has its source in the fact that ambiguity aversion causes the effective measure to be proportional to the product (and thus, essentially, the covariance) of the ambiguity neutral probability μ_θ and the marginal utility $\phi'_r(u_r(\sigma_\theta, \tau^*))$ and the latter is non-separable across action recommendations. This explains how the non-separability across action recommendations in determining obedience is what prevents adopting the strategies of [Kamenica and Gentzkow \(2011\)](#) and [Alonso and Câmara \(2016\)](#) to establish our characterization.

4.2 Properties of Optimal Persuasion with Ambiguous Communication

We next use our characterization to derive properties of optimal persuasion. Two experiments are *Pareto-ranked* if the sender and receiver agree on their strict ranking under the assumption of obedience. As we shall see, Pareto-ranking and splittings into Pareto-ranked experiments play a key role in optimal persuasion and, more generally, in the sender benefiting from ambiguous communication – the latter will be focus of the next section.

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

$$u_s(\bar{\sigma}, \tau^*) \geq u_s(\underline{\sigma}, \tau^*) \text{ and } u_r(\bar{\sigma}, \tau^*) \geq u_r(\underline{\sigma}, \tau^*), \quad (3)$$

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

*A **Pareto-ranked splitting** of the experiment σ is a triple $(\bar{\sigma}, \underline{\sigma}, \lambda)$ such that (i) $\lambda\bar{\sigma} + (1 - \lambda)\underline{\sigma} = \sigma$, (ii) $\lambda \in (0, 1)$, and (iii) (3) holds with strict inequalities, i.e., $\bar{\sigma}$ and $\underline{\sigma}$ are Pareto-ranked.*

Our next result shows that these concepts are useful in indicating whether an ambiguous experiment (σ, μ) can be improved (for the sender) by removing or adding some splittings. Part (i) identifies instances where the sender uses too much ambiguity (i.e., splitting in an ineffective manner that should be removed), while parts (ii) and (iii) identify instances in which the sender fails to use some additional and beneficial ambiguity in the form of further Pareto-ranked splitting. Part (i) of the result says that if two experiments in the support of μ bracket the sender's payoff from the ambiguous experiment, the sender can strictly improve whenever they are not weakly Pareto-ranked. The proof shows improvement can be achieved by merging the two experiments. Parts (ii) and (iii) give conditions under which the introduction of additional ambiguity through further Pareto-ranked splittings that bracket the sender's payoff from the ambiguous experiment help the sender. When such Pareto-ranked splittings exist, these conditions are always satisfied for an ambiguity-neutral sender.

Theorem 2. *Assume that ϕ_r is strictly concave.¹² Let (σ, μ) be an obedient ambiguous experiment and $U_s(\sigma, \mu, \tau^*)$ the corresponding sender's payoff. If either*

¹²As the proof in the Appendix makes clear, the only role of this assumption is to simplify the statement of

(i) for some $\theta, \theta' \in \text{supp}(\mu)$ such that $u_s(\sigma_\theta, \tau^*) \geq U_s(\boldsymbol{\sigma}, \mu, \tau^*) \geq u_s(\sigma_{\theta'}, \tau^*)$, σ_θ and $\sigma_{\theta'}$ are not weakly Pareto-ranked,

or,

(ii) for some $\theta \in \text{supp}(\mu)$, there exists a Pareto-ranked splitting of σ_θ , $(\bar{\sigma}, \underline{\sigma}, \lambda)$, such that $u_s(\bar{\sigma}, \tau^*) \geq u_s(\sigma_\theta, \tau^*) > U_s(\boldsymbol{\sigma}, \mu, \tau^*) \geq u_s(\underline{\sigma}, \tau^*)$, and

$$\frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_s(u_s(\underline{\sigma}, \tau^*))} > \frac{\phi'_r(u_r(\bar{\sigma}, \tau^*))}{\phi'_r(u_r(\sigma_\theta, \tau^*))}, \quad (M_+)$$

or

(iii) for some $\theta \in \text{supp}(\mu)$, there exists a Pareto-ranked splitting of σ_θ , $(\bar{\sigma}, \underline{\sigma}, \lambda)$, such that $u_s(\bar{\sigma}, \tau^*) \geq U_s(\boldsymbol{\sigma}, \mu, \tau^*) \geq u_s(\sigma_\theta, \tau^*) \geq u_s(\underline{\sigma}, \tau^*)$, and

$$\frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_s(u_s(\underline{\sigma}, \tau^*))} > \frac{\phi'_r(u_r(\sigma_\theta, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))}, \quad (M_-)$$

then there exists an obedient ambiguous experiment $(\hat{\boldsymbol{\sigma}}, \hat{\mu})$ that is strictly better than $(\boldsymbol{\sigma}, \mu)$ for the sender.

Theorem 2 describes properties that indicate when an ambiguous experiment $(\boldsymbol{\sigma}, \mu)$ is not exploiting ambiguous communication optimally. To gain intuition for part (i), first observe that if such σ_θ and $\sigma_{\theta'}$ are not weakly Pareto-ranked, then the receiver must get a strictly higher expected payoff from $\sigma_{\theta'}$ than from σ_θ , while the reverse is true for the sender. Ambiguity aversion then implies that the receiver's effective measure places more weight on σ_θ relative to $\sigma_{\theta'}$ than the ambiguity neutral weights do, i.e., $\lambda_\theta/\lambda_{\theta'} > \mu_\theta/\mu_{\theta'}$. If σ_θ and $\sigma_{\theta'}$ are the only two experiments in the support of μ , the sender can merge them into the (unambiguous) experiment $\lambda_\theta\sigma_\theta + \lambda_{\theta'}\sigma_{\theta'}$. By construction, the receiver would continue to be obedient, and the sender would strictly benefit from this merging – a profitable deviation. When σ_θ and $\sigma_{\theta'}$ are not the only two experiments in the support of μ , however, this is not the complete story as this merging may also impact the weighting of the merged experiment relative to the other experiments. Part of the additional insight of the proof is that when $u_s(\sigma_\theta, \tau^*) \geq U_s(\boldsymbol{\sigma}, \mu, \tau^*) \geq u_s(\sigma_{\theta'}, \tau^*)$ holds, this impact is not unfavorable to the sender.

the theorem. Without it, one needs to add conditions checking if $\phi'_r(u_r(\sigma_\theta, \tau^*)) \neq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$ to each part of the theorem.

The intuition for part (ii) is similar. Suppose the sender constructs $\hat{\sigma}$ from σ by adding ambiguity, replacing σ_θ with $\bar{\sigma}$ and $\underline{\sigma}$ and choosing $\hat{\mu}$ such that the total weight on $\bar{\sigma}$ and $\underline{\sigma}$ is μ_θ , i.e., $\hat{\mu}_{\bar{\sigma}} + \hat{\mu}_{\underline{\sigma}} = \mu_\theta$. (All other weights remain the same.) If σ_θ is the only experiment in the support of μ , choosing

$$\hat{\mu}_{\bar{\sigma}} = \frac{\lambda \phi'_r(u_r(\underline{\sigma}, \tau^*))}{\lambda \phi'_r(u_r(\underline{\sigma}, \tau^*)) + (1 - \lambda) \phi'_r(u_r(\underline{\sigma}, \tau^*))}$$

guarantees that the effective measure places weight λ on $\bar{\sigma}$ and, therefore, that the receiver remains obedient (since $(\bar{\sigma}, \underline{\sigma}, \lambda)$ is a Pareto-ranked splitting of σ_θ). Since $\hat{\mu}_{\bar{\sigma}} > \lambda$, an ambiguity neutral sender would strictly benefit from adding ambiguity. Note that if the sender is ambiguity neutral, the condition (M_+) is automatically satisfied, since the left-hand side is one, while the right-hand side is strictly less than one (because the experiments are Pareto-ranked). However, if the sender is ambiguity averse, introducing some additional ambiguity comes at a cost. The condition (M_+) guarantees that that the gain outweighs the cost. Lastly, if there is more than one experiment in the support of μ , a similar argument continues to work. Since part (iii) is the mirror image of part (ii), the same intuition applies to it, with (M_-) playing the role of (M_+) .

In Theorem 2, the conditions refer to pairs of experiments for which the sender's payoffs bracket $U_s(\sigma, \mu, \tau^*)$. Intuition for why similar conclusions may not apply when the pairs involved in the Pareto-ranking or the Pareto-ranked splitting lie on the same side of $U_s(\sigma, \mu, \tau^*)$ is related to how the receiver's ambiguity aversion, as reflected in properties of ϕ_r , connects μ with the effective measure $em^{(\sigma, \mu)}$ via (1). In particular, when there are more than two experiments in σ , splitting or merging experiments on the same side of $U_s(\sigma, \mu, \tau^*)$ may shift their combined weights in the effective measure relative to the other experiments in a manner unfavorable to the sender. In the Appendix, we see that concavity (resp. convexity) of $1/\phi'_r$ is sufficient to extend the conclusions to pairs of experiments on a particular side of $U_s(\sigma, \mu, \tau^*)$, and assuming linearity of $1/\phi'_r$ leads to the following simpler necessary conditions for optimal persuasion:

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

$$\phi_r(x) = c \ln(ax + b) + d \tag{4}$$

for some $a, b, c, d \in \mathbb{R}$ where $a, c > 0$ and $ax + b > 0$ for all $x \in [\min_{a, \omega} u_r(a, \omega), \max_{a, \omega} u_r(a, \omega)]$. Then, all experiments are weakly Pareto-ranked, that is, for all $\theta, \theta' \in \text{supp}(\mu)$, σ_θ and $\sigma_{\theta'}$

are weakly Pareto-ranked.

If, in addition, the sender is ambiguity neutral, no Pareto-ranked splitting of σ_θ exists for any $\theta \in \text{supp}(\mu)$.

Note that (4) may be interpreted as constant relative ambiguity aversion (see [Klibanoff et al. \(2005\)](#)). The result that all experiments used must be weakly Pareto-ranked is reminiscent of a key Pareto-ranking result ([Rayo and Segal, 2010](#), p. 959, Lemma 2) in an entirely different persuasion setting (one in which ambiguity plays no role). Mathematically, the common source of both results is the maximization of the product of an increasing function of the sender's expected payoff and an increasing function of the receiver's expected payoff. Indeed, under (4), $1/\phi'_r$ is linear and the objective function Φ_u in our characterizations has the product form

$$\Phi_u(\sigma) = (a/c) \times u_r(\sigma, \tau^*) \times (\phi_s(u_s(\sigma, \tau^*)) - \phi_s(u)).$$

In [Rayo and Segal \(2010\)](#), their sender receives profit only when the receiver “accepts” (i.e., takes the higher of two actions) and this occurs with probability equal to the conditional expected gross payoff (normalized to $[0,1]$) of the receiver when accepting. Thus their sender maximizes an expected payoff that is, signal-by-signal, equal to the product of their conditional expected profit and the receiver's conditional expected gross payoff. As far as we know, there is no obvious analogue in the setting of [Rayo and Segal \(2010\)](#) of our problem with general concave ϕ_r and the corresponding partial Pareto-ranking result in Part (i) of our Theorem 2.

We now solve our introductory example for a ϕ_r satisfying (4):

Example 1 (Introductory Example Continued). *Suppose $\phi_r(x) = \ln(x+5)$ and $\phi_s(x) = x$. Then a sender's optimal persuasion strategy is the $((\bar{\sigma}, \underline{\sigma}), \mu)$ described in Figure 2 with $\mu_{\bar{\sigma}} = 39/50$. The payoffs from this optimal persuasion are as follows:*

$$U_s((\bar{\sigma}, \underline{\sigma}), \mu, \tau^*) = 39/50 \times 3/2 + 11/50 \times 1/2 = 1.28,$$

$$U_r((\bar{\sigma}, \underline{\sigma}), \mu, \tau^*) = e^{(39/50 \ln(13/2) + 11/50 \ln(11/2))} - 5 \approx 1.265.$$

Thus both the sender and receiver do better than the payoff of $5/4$ they would each obtain under Bayesian persuasion.

So far, the analysis was devoted to the characterization of optimal communication strategies when ambiguous experiments are allowed. However, it does not directly tell

us whether the sender would benefit from introducing ambiguity into their communication. We now turn to this issue, which we view as a primary focus of the paper.

5 When Does Ambiguous Communication Benefit The Sender?

5.1 A Concavification Characterization

A characterization of when ambiguous communication benefits the sender can be derived from our characterization of optimal persuasion (Theorem 1 and Corollary 1). More specifically, Lemma A.2 in the proof of Theorem 1 shows that ambiguous communication gives the sender a strictly higher payoff than u if, and only if, the value of the program $(\Phi^*(u))$ is strictly positive. By letting $u = u_s^{BP}$, we obtain the following.

Corollary 2. *Ambiguous communication benefits the sender if, and only if, the value of $(\Phi^*(u_s^{BP}))$ is strictly positive, or, equivalently, $\max_{\sigma \in \Sigma^*} \text{cav} \Phi_{u_s^{BP}}(\sigma) > 0$.*

We next derive some necessary conditions for the sender to benefit from ambiguous communication that emphasize the role of Pareto-ranked experiments.

5.2 Necessary Conditions for Ambiguity to Benefit The Sender

We show that Pareto-ranked experiments continue to be key in determining whether ambiguous communication is better for the sender than unambiguous communication. The following theorem shows having Pareto-ranked experiments in the collection (in particular, better and worse ones having sender's expected payoffs bracketing u_s^{BP}) is necessary for an ambiguous experiment to benefit the sender.

Theorem 3. *If an obedient ambiguous experiment (σ, μ) benefits the sender, then there exist $\theta, \theta' \in \text{supp}(\mu)$ such that σ_θ and $\sigma_{\theta'}$ are Pareto-ranked, with $u_s(\sigma_\theta, \tau^*) > u_s^{BP} \geq u_s(\sigma_{\theta'}, \tau^*)$.*

Comparing with part (i) of Theorem 2, we see that while optimal persuasion requires *weak* Pareto-ranking of experiments that bracket the sender's payoff from that ambiguous experiment, Theorem 3 says that any improvement over Bayesian persuasion requires some Pareto-ranked experiments (and thus strictly ranked) bracketing u_s^{BP} for the sender.

We next present two equivalent sets of necessary conditions for ambiguity to benefit the sender, and show that these conditions imply that ambiguous communication can never

benefit the sender when the receiver has only two available actions – a common assumption in many examples and applications in the literature. Whereas Theorem 3 described a necessary property of any sender’s strategy that improves on Bayesian persuasion, these next conditions relate the possibility of ambiguity benefiting the sender in a given persuasion game to the existence of Pareto-ranked experiments with certain properties.

Theorem 4. *Ambiguous communication benefits the sender only if ϕ_r is not affine, and*

- (a) *there exists a Pareto-ranked splitting, $(\bar{\sigma}, \underline{\sigma}, \lambda)$, of an obedient experiment $\hat{\sigma}$ such that $u_s(\bar{\sigma}, \tau^*) > u_s^{BP}$; or, equivalently,*
- (b) *there exist Pareto-ranked experiments, σ and σ^* such that: (i) $\text{supp } \sigma(\cdot|\omega) = \text{supp } \sigma^*(\cdot|\omega)$ for all ω , (ii) $u_s(\sigma, \tau^*) > u_s^{BP}$, and (iii) $\tau^* \in \text{br}(\sigma^*) \setminus \text{br}(\sigma)$.*

Example 2 (Introductory Example Continued). *Recall that for the collection $\sigma = (\bar{\sigma}, \underline{\sigma})$ constructed in Figure 2 of the introductory example, $(\bar{\sigma}, \underline{\sigma}, \frac{3}{4})$ is a Pareto-ranked splitting of σ^{BP} . Thus, for this example, the conditions in part (a) of the theorem are satisfied for $\hat{\sigma} = \sigma^{BP}$.*

Remark 2 (Construction of Pareto-ranked experiments in (b)). The argument that the conditions in part (b) of the theorem are necessary is constructive, and the effective measure plays a key role. Suppose that there exists a solution (σ^*, μ^*) to the sender’s program (\mathcal{P}) that benefits the sender. Construct σ and σ^* in (b) by letting $\sigma = \sum_{\theta} \mu_{\theta}^* \sigma_{\theta}^*$ and $\sigma^* = \sum_{\theta} \text{em}_{\theta}^{(\sigma^*, \mu^*)} \sigma_{\theta}^*$.

Remark 3 (Not necessary for a Pareto-ranked splitting of σ^{BP} to exist). The reader might wonder if a stronger version of necessary condition (a) that requires the Pareto-ranked splitting to be of σ^{BP} is also necessary. This is false. There are examples in which the sender benefits from ambiguous communication even though no Pareto-ranked splitting of σ^{BP} exists (as is true, for instance, whenever all σ^{BP} are efficient). In such cases, it is splittings of some other obedient experiment that generate the gains over Bayesian persuasion for the sender.

The conditions in Theorem 4 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 3. *If the receiver has only two actions, the sender cannot benefit from ambiguous communication.*

The intuition is as follows. From Theorem 4, we have that part of a necessary condition for ambiguity helping the sender is the existence of an experiment σ that strictly improves the receiver's expected payoff compared to some other experiment σ^* , with the added property that obedience of σ is not optimal, i.e., $\tau^* \notin br(\sigma)$. Intuitively, such an improvement is possible only when σ 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 extra 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.

5.3 Robust Benefits

So far, we have assumed that if the sender designs the ambiguous experiment (σ, μ) , the receiver perceives it correctly. More realistically, the receiver might have a somewhat different perception of the experiment than the one the sender intends to convey. After all, conveying the exact specifications of an experiment is a complex task, let alone of an ambiguous one. Yet, we show that if the sender benefits from ambiguous communication, they continue to benefit even if the receiver somewhat misperceives the intended experiment.

Proposition 3. *Suppose ambiguous communication benefits the sender and that the set of obedient experiments has a non-empty interior. Then, there exists a non-empty open set of obedient ambiguous experiments that benefit the sender.*

A sketch of the argument is that the sender benefits from ambiguous communication if, and only if, $\Phi^*(u_s^{BP}) > 0$, and the problem is sufficiently continuous to guarantee the existence of an open set of obedient ambiguous experiments that benefit the sender. In fact, this continues to be true even under small perturbations of ϕ_r , so that the existence of benefits does not rest on exact knowledge of the receiver's ϕ_r .

Corollary 4. *Suppose ambiguous communication benefits the sender and that the set of obedient experiments has a non-empty interior. Then, there exists an ambiguous experiment that benefits the sender, and continues to do so under small perturbations of ϕ_r .*

6 Benefits from Binary Ambiguous Communication

This section restricts attention to binary ambiguous experiments. This restriction is not without loss of generality because there are examples in which the sender benefitting from ambiguous communication requires ambiguous experiments with σ containing more than two distinct experiments (see Section SA.3 in the Online Appendix (Cheng et al., 2024)). Nonetheless, this restriction allows us to derive sufficient conditions for the sender to benefit from ambiguous communication and how these conditions vary with the extent of the sender's and/or receiver's ambiguity aversion. It also allows us to see that binary ambiguous communication may also improve the receiver's payoff.

If a binary ambiguous experiment benefits the sender compared to Bayesian persuasion, it follows from Theorem 3 that the experiments must be Pareto-ranked. We therefore focus on Pareto-ranked binary ambiguous experiments in what follows.

The next theorem, Theorem 5, provides necessary and sufficient conditions for a binary ambiguous experiment based on a Pareto-ranked splitting of any obedient experiment σ^* to (a) strictly improve the receiver's payoff compared to σ^* , and (b) strictly improve the sender's payoff compared to σ^* . We later apply the theorem to the case in which $u_s(\sigma^*, \tau^*) = u_s^{BP}$, thereby obtaining sufficient conditions for the sender to benefit from ambiguous communication (see Corollary 5). Proposition 4 provides conditions on the primitives sufficient for existence of a Pareto-ranked splitting of a given experiment.

The theorem uses the following notion of *probability premium*.

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

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

The probability premium $\rho^{\phi, u}((\bar{\sigma}, \underline{\sigma}), \lambda)$ is exactly the ϕ -payoff difference between σ^* and the ambiguous experiment $((\bar{\sigma}, \underline{\sigma}), \lambda)$, 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) (see Eeckhoudt and Laeven (2015) for a graphical representation of Pratt's concept).

Thus, if we let

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

be the probability of $\bar{\sigma}$, then

$$\begin{aligned} U((\bar{\sigma}, \underline{\sigma}), \mu, \tau^*) &= \phi^{-1}\left((\lambda + \rho^{\phi, u}((\bar{\sigma}, \underline{\sigma}), \lambda))\phi(u(\bar{\sigma}, \tau^*)) + (1 - \lambda - \rho^{\phi, u}((\bar{\sigma}, \underline{\sigma}), \lambda))\phi(u(\underline{\sigma}, \tau^*))\right) \\ &= \phi^{-1}\left(\phi(u(\sigma^*, \tau^*))\right) = u(\sigma^*, \tau^*), \end{aligned}$$

meaning that the premium $\rho^{\phi, u}((\bar{\sigma}, \underline{\sigma}), \lambda)$ is exactly the increase in $\mu_{\bar{\sigma}}$ above λ needed to make the player indifferent between the ambiguous experiment $((\bar{\sigma}, \underline{\sigma}), \mu)$ and σ^* . Thus, assuming obedience, this premium is the smallest increase in the μ -probability of the higher payoff experiment required to compensate for exposure to the ambiguous experiment:

Lemma 2. *Let $\bar{\sigma}$ and $\underline{\sigma}$ be experiments such that $u_i(\bar{\sigma}, \tau^*) > u_i(\underline{\sigma}, \tau^*)$. For all $\mu_{\bar{\sigma}}, \lambda \in [0, 1]$, $U_i((\bar{\sigma}, \underline{\sigma}), \mu, \tau^*) > u_i(\lambda\bar{\sigma} + (1 - \lambda)\underline{\sigma}, \tau^*)$ if, and only if, player i 's $((\bar{\sigma}, \underline{\sigma}), \lambda)$ -probability premium is strictly less than $\mu_{\bar{\sigma}} - \lambda$.*

As a consequence, we have the following result:

Theorem 5. *Let σ^* be an obedient experiment. Suppose that $(\bar{\sigma}, \underline{\sigma}, \lambda)$ is a Pareto-ranked splitting of σ^* satisfying $u_s(\bar{\sigma}, \tau^*) > u_s(\underline{\sigma}, \tau^*)$. The binary ambiguous experiment (σ, μ) , with $\sigma = (\bar{\sigma}, \underline{\sigma})$ and*

$$\mu_{\bar{\sigma}} = \frac{\lambda\phi'_r(u_r(\underline{\sigma}, \tau^*))}{\lambda\phi'_r(u_r(\underline{\sigma}, \tau^*)) + (1 - \lambda)\phi'_r(u_r(\bar{\sigma}, \tau^*))}, \quad (5)$$

satisfies the following properties:

- (i) (σ, μ) is obedient,
- (ii) $U_r(\sigma, \mu, \tau^*) > u_r(\sigma^*, \tau^*)$ if, and only if, $\mu_{\bar{\sigma}} > \lambda$,
- (iii) $U_s(\sigma, \mu, \tau^*) > u_s(\sigma^*, \tau^*)$ if, and only if, the sender's $((\bar{\sigma}, \underline{\sigma}), \lambda)$ -probability premium is strictly less than $\mu_{\bar{\sigma}} - \lambda$.

Furthermore, the sender's $((\bar{\sigma}, \underline{\sigma}), \lambda)$ -probability premium is increasing in the sender's ambiguity aversion, and $\mu_{\bar{\sigma}}$ is increasing in the receiver's ambiguity aversion.

That a μ satisfying (5) ensures that the obedience of σ^* extends to the binary ambiguous experiment (σ, μ) as in (i) is a straightforward consequence of Lemma 1. The necessary

and sufficient conditions in (ii) for the receiver to be better off when the sender communicates ambiguously using (σ, μ) rather than unambiguously using σ^* require some elaboration. First, the condition $\mu_{\bar{\sigma}} > \lambda$ is equivalent to $\phi'_r(u_r(\underline{\sigma}, \tau^*)) > \phi'_r(u_r(\bar{\sigma}, \tau^*))$, i.e., the receiver is, within this range of payoffs, not everywhere ambiguity neutral. In particular, this condition is always satisfied if ϕ_r is strictly concave. Second, the condition $\mu_{\bar{\sigma}} > \lambda$ can be shown to be equivalent to the receiver's $((\bar{\sigma}, \underline{\sigma}), \lambda)$ -probability premium being strictly less than $\mu_{\bar{\sigma}} - \lambda$, which, by Lemma 2, characterizes when the receiver is better off under (σ, μ) than under σ^* . The necessary and sufficient conditions in (iii) for (σ, μ) to be better for the sender than σ^* follow directly from Lemma 2. For an ambiguity neutral sender, the probability premium is zero, and thus the condition in (iii) reduces to $\mu_{\bar{\sigma}} > \lambda$, as in (ii). Thus, for an ambiguity neutral sender facing a strictly ambiguity averse receiver, the ambiguity introduced in (σ, μ) improves on σ^* for both sender and receiver.

The source of the economic gain from ambiguous communication, for both sender and receiver, is the greater use, as measured by $\mu_{\bar{\sigma}} - \lambda$, of the Pareto-better experiment $\bar{\sigma}$. 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. The fact that $\mu_{\bar{\sigma}}$ is constructed to respect obedience taking into account the receiver's ambiguity aversion but not the sender's, is what explains why the condition for this net gain to be positive can be simplified for the receiver, but not the sender.

The comparative static about $\mu_{\bar{\sigma}}$ in the final section of the theorem, when combined with (ii) and (iii), shows that the payoff difference between the ambiguous experiment (σ, μ) and the unambiguous σ^* satisfies single-crossing with respect to the receiver's ambiguity aversion for both the sender and receiver. Similarly, the comparative static in the sender's probability premium, together with (iii), shows that the negative of this payoff difference for the sender satisfies single-crossing with respect to the sender's ambiguity aversion.

Starting from any given obedient experiment, Theorem 5 provides necessary and sufficient conditions for binary ambiguous communication to strictly improve the sender's payoff, and thus sufficient conditions for *some* ambiguous communication to do so. Thus, if we apply Theorem 5 to the case where σ^* is an optimal Bayesian persuasion, we obtain sufficient conditions for ambiguity to benefit the sender:

Corollary 5. *Let σ^{BP} be an obedient experiment such that $u_s(\sigma^{BP}, \tau^*) = u_s^{BP}$. If there exists a Pareto-ranked splitting of σ^{BP} , $((\bar{\sigma}, \underline{\sigma}), \lambda)$, for which $\rho^{\phi_s, u_s}((\bar{\sigma}, \underline{\sigma}), \lambda) < \mu_{\bar{\sigma}} - \lambda$, with $\mu_{\bar{\sigma}}$ given by equation (5), then ambiguous communication benefits the sender.*

Therefore, whenever a Pareto-ranked splitting of a σ^{BP} exists, an ambiguity neutral sender

benefits from ambiguous communication as long as the receiver is not completely ambiguity neutral over the payoff range of the splitting.¹³

Theorem 5 and Corollary 5 require the existence of a Pareto-ranked splitting. This existence is not guaranteed. For instance, if the obedient experiment σ^* induces an efficient payoff profile, no Pareto-ranked splitting of it exists. The next result provides sufficient conditions on the primitives for the existence of a Pareto-ranked splitting of σ^* .

Proposition 4. *Given any experiment σ^* , fix, for each $\omega \in \Omega$, $a_\omega \in \text{supp}(\sigma^*(\cdot|\omega))$ and consider the following set of vectors,*

$$\left\{ \begin{bmatrix} p(\omega)(u_s(a, \omega) - u_s(a_\omega, \omega)) \\ p(\omega)(u_r(a, \omega) - u_r(a_\omega, \omega)) \end{bmatrix} : a \in \text{supp}(\sigma^*(\cdot|\omega)), \omega \in \Omega \right\}. \quad (6)$$

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

Remark 4 (Relation with efficiency of σ^*). The spanning condition in Proposition 4 is stronger than the statement that σ^* , assuming obedience, is not efficient. The reason for this is that existence of a Pareto-ranked splitting needs not only a more efficient experiment, but one that for all $\omega \in \Omega$, never generates an action recommendation that could not have come from σ^* in that ω . Arieli et al. (2024) argue that Bayesian persuasion solutions are typically inefficient¹⁴ and provide a necessary condition for their efficiency. This condition, $\sum_{\omega \in \Omega} |\text{supp}(\sigma^*(\cdot|\omega))| \leq |\Omega| + 1$, is necessarily violated when the spanning condition holds. Observe that whenever σ^* randomizes in at least two states in the support of p , the spanning condition holds for a generic specification of the payoffs, $u_i(a, \omega)$.

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

While the optimal persuasion does depend on ϕ_r and ϕ_s , i.e., the ambiguity attitudes, we next show that the possibility of strict improvement for the sender from using binary ambiguous experiments is robust in several respects. First, the same ambiguous experiment remains beneficial to any less ambiguity averse sender. Second, it is robust to the sender underestimating the extent of ambiguity aversion of the receiver. In other words, if an

¹³Recall from Remark 3 that existence of a Pareto-ranked splitting of σ^{BP} is not necessary to benefit.

¹⁴Though Ichihashi (2019) proves it is always efficient when the receiver has only two actions.

improvement is possible when facing a given receiver, it is also possible when facing a more ambiguity-averse receiver. While we show that the same collection σ can be used to generate the improvement for all more ambiguity-averse receivers, in general, the μ guaranteeing improvement may need to change. Part (iii) of the result shows that adding the requirement that $\underline{\sigma}$ is obedient allows a stronger robustness: the same μ that generates an improvement for the sender when facing a receiver with ϕ_r also generates an improvement when facing any more ambiguity averse receiver (more concave ϕ_r).

Theorem 6. *Suppose there exists a binary collection $\sigma = (\bar{\sigma}, \underline{\sigma})$ and a non-degenerate μ such that (σ, μ) is obedient and benefits the sender (compared to σ^{BP}). Then:*

- (i) (σ, μ) also benefits all weakly less ambiguity averse senders, and
- (ii) for any weakly more ambiguity averse receiver, there exists some $\tilde{\mu}$ such that $(\sigma, \tilde{\mu})$ benefits all weakly less ambiguity averse senders, and
- (iii) if $\tau^* \in br(\underline{\sigma})$, then $\tilde{\mu}$ in (ii) can be set equal to μ .

7 Further Discussion

7.1 What if the Sender Learns θ in Advance?

We have assumed that the sender commits to an ambiguous experiment not knowing which θ obtains. Consider the alternative assumption that the sender privately observes θ before committing to a strategy (and this is common knowledge). As in cheap-talk games, this would imply that the sender must be indifferent between all experiments σ_θ in any equilibrium of this modified game.¹⁵ It follows that the unambiguous experiment $\sum_\theta em_\theta^{(\sigma, \mu)} \sigma_\theta$ would also give the same payoff to the sender. (Since the sender commits to experiments, we can still restrict attention to obedient canonical experiments.) This alternative assumption thus implies that the sender cannot benefit from ambiguous communication. If the sender wants to benefit, they must commit to not learn θ prior to committing to an experiment. Delegation to third-parties might be one way to achieve this in practice.

¹⁵More precisely, this must be true for all experiments σ_θ for $\theta \in \text{supp}(\mu)$.

7.2 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 information a receiver might learn.¹⁶ [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 α -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.

We conclude by returning to the discussion of BLL begun in the introduction. Broadly speaking, the gains we identify work through key properties, such as Pareto-ranking of (at least some) of the experiments in the collection σ chosen by the sender. Such properties contrast sharply with the “synonym” constructions emphasized in BLL that lead to

¹⁶[Dworczak and Pavan \(2022\)](#)’s model is not restricted to single-receiver persuasion settings.

collections in which each experiment yields the same expected payoff to the sender. BLL and our approach also lead to different outcomes. For example, our Corollary 3 shows that ambiguous communication never benefits the sender when the receiver has only two actions. In contrast, BLL find gains from ambiguous communication in such cases, including their main example. Conversely, there are examples in which there is no benefit for the sender according to BLL’s approach (even when extended to include sender preferences less extremely ambiguity averse than U_s^{MEU}), in which the sender benefits from ambiguous communication in our approach.

As previously mentioned, the benefits from ambiguous communication in BLL involve in an essential way the receiver behaving suboptimally with respect to their ex-ante preferences as specified by BLL. Pahlke (2023) uses constructions based on rectangularity (Epstein and Schneider, 2003) to construct alternative ex-ante MEU preferences (different from BLL and from U_r^{MEU}) that are consistent with the receiver’s interim behavior in BLL. When there are gains in BLL from ambiguous communication, some of the measures in Pahlke’s construction must reflect correlation between Ω and which experiment from the ambiguous collection generates the messages. This is the manifestation of the dynamic inconsistency in BLL within the dynamically consistent reformulation of Pahlke (2023).¹⁷

References

- ALONSO, R., AND O. CÂMARA (2016): “Bayesian persuasion with heterogeneous priors,” *Journal of Economic Theory*, 165, 672–706.
- ARIELI, I., Y. BABICHENKO, A. JAIN, AND R. SMORODINSKY (2024): “Are Bayesian persuasion outcomes efficient.”
- AUMANN, R. J., AND M. MASCHLER (1966): “Game theoretic aspects of gradual disarmament,” *Report of the US Arms Control and Disarmament Agency*, 80, 1–55.
- (1995): *Repeated games with incomplete information*: MIT press.
- BEAUCHÊNE, D., J. LI, AND M. LI (2019): “Ambiguous persuasion,” *Journal of Economic Theory*, 179, 312–365.

¹⁷For 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). See Hanany et al. (2020) on dynamic games of incomplete information with ambiguity averse players.

- BERGEMANN, D., AND S. MORRIS (2019): “Information design: A unified perspective,” *Journal of Economic Literature*, 57, 44–95.
- BOSE, S., AND L. RENO (2014): “Mechanism design with ambiguous communication devices,” *Econometrica*, 82, 1853–1872.
- CHENG, X. (2022): “Relative maximum likelihood updating of ambiguous beliefs,” *Journal of Mathematical Economics*, 99, 102587.
- (2023): “Ambiguous persuasion: An ex-ante formulation,” *arXiv preprint arXiv:2010.05376*.
- CHENG, X., P. KLIBANOFF, S. MUKERJI, AND L. RENO (2024): “Supplemental Appendix to Persuasion with Ambiguous Communication.”
- DI TILLIO, A., N. KOS, AND M. MESSNER (2017): “The design of ambiguous mechanisms,” *The Review of Economic Studies*, 84, 237–276.
- DÜTTING, P., M. FELDMAN, D. PERETZ, AND L. SAMUELSON (2024): “Ambiguous contracts,” *arXiv preprint arXiv:2302.07621*.
- DWORCZAK, P., AND A. PAVAN (2022): “Preparing for the worst but hoping for the best: Robust (Bayesian) persuasion,” *Econometrica*, 90, 2017–2051.
- ECKHOUDT, L. R., AND R. J. LAEVEN (2015): “The probability premium: A graphical representation,” *Economics Letters*, 136, 39–41.
- EPSTEIN, L. G., AND M. SCHNEIDER (2003): “Recursive multiple-priors,” *Journal of Economic Theory*, 113, 1–31.
- GALPERTI, S. (2019): “Persuasion: The Art of Changing Worldviews,” *American Economic Review*, 109, 996–1031, [10.1257/aer.20161441](https://doi.org/10.1257/aer.20161441).
- GHIRARDATO, P., F. MACCHERONI, AND M. MARINACCI (2004): “Differentiating ambiguity and ambiguity attitude,” *Journal of Economic Theory*, 118, 133–173.
- GILBOA, I., AND D. SCHMEIDLER (1989): “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18, 141–153.
- 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., AND P. KLIBANOFF (2007): “Updating preferences with multiple priors,” *Theoretical Economics*, 2, 261–298.
- (2009): “Updating ambiguity averse preferences,” *The BE Journal of Theoretical Economics*, 9.
- HANANY, E., P. KLIBANOFF, AND S. MUKERJI (2020): “Incomplete information games with ambiguity averse players,” *American Economic Journal: Microeconomics*, 12, 135–87.
- HEDLUND, J., T. F. KAUFFELDT, AND M. LAMMERT (2020): “Persuasion under ambiguity,” *Theory and Decision*, 1–28.
- ICHIHASHI, S. (2019): “Limiting Sender’s information in Bayesian persuasion,” *Games and Economic Behavior*, 117, 276–288.
- KAMENICA, E. (2019): “Bayesian persuasion and information design,” *Annual Review of Economics*, 11, 249–272.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian persuasion,” *American Economic Review*, 101, 2590–2615.
- KELLNER, C., AND M. T. LE QUEMENT (2017): “Modes of ambiguous communication,” *Games and Economic Behavior*, 104, 271–292.
- (2018): “Endogenous ambiguity in cheap talk,” *Journal of Economic Theory*, 173, 1–17.
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): “A smooth model of decision making under ambiguity,” *Econometrica*, 73, 1849–1892.
- KOSTERINA, S. (2022): “Persuasion with unknown beliefs,” *Theoretical Economics*, 17, 1075–1107.
- LACLAU, M., AND L. RENO (2017): “Public persuasion,” *mimeo*.
- LINARDATOS, P., V. PAPASTEFANOPOULOS, AND S. KOTSIANTIS (2021): “Explainable AI: A Review of Machine Learning Interpretability Methods,” *Entropy*, 23.
- NIKZAD, A. (2021): “Persuading a pessimist: Simplicity and robustness,” *Games and Economic Behavior*, 129, 144–157, <https://doi.org/10.1016/j.geb.2021.05.007>.

PAHLKE, M. (2023): “Dynamic Consistency and Ambiguous Communication,” *mimeo*.

PRATT, J. W. (1964): “Risk Aversion in the Small and in the Large,” *Econometrica*, 32, 122–136, <http://www.jstor.org/stable/1913738>.

RAYO, L., AND I. SEGAL (2010): “Optimal information disclosure,” *Journal of Political Economy*, 118, 949–987.

SCHMITT, M. (2024): “Interpretable Machine Learning,” <https://towardsdatascience.com/interpretable-machine-learning-45b467dbelaf>, accessed May 7 2024.

SINISCALCHI, M. (2011): “Dynamic choice under ambiguity,” *Theoretical Economics*, 6, 379–421.

TANG, R., AND M. ZHANG (2021): “Maxmin implementation,” *Journal of Economic Theory*, 194, 105250.

TELUS INTERNATIONAL WEBSITE (2021): <https://www.telusinternational.com/insights/ai-data/article/how-to-select-the-right-machine-learning-algorithm>, accessed May 7 2024.

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

A Proofs of Main Results in Order of Appearance

Proof of Proposition 1. Fix (σ, μ) and $\tau \in BR(\sigma, \mu)$. Construct a canonical ambiguous experiment (σ^*, μ) as follows: For each θ , define $\sigma_\theta^*(a|\omega) = \sum_m \tau(a|m)\sigma_\theta(m|\omega)$, for all (a, ω) , and let $\sigma^* = (\sigma_\theta^*)_\theta$. For $i \in \{s, r\}$, $u_i(\sigma_\theta, \tau) = \sum_{\omega, a} p(\omega)u_i(a, \omega) \sum_m \tau(a|m)\sigma_\theta(m|\omega) = u_i(\sigma_\theta^*, \tau^*)$. Therefore $U_i(\sigma^*, \mu, \tau^*) = U_i(\sigma, \mu, \tau)$, for $i \in \{s, r\}$.

Suppose (σ^*, μ) is not obedient. Then $\exists \delta : A \rightarrow \Delta(A)$ such that $U_r(\sigma^*, \mu, \tau^*) < U_r(\sigma^*, \mu, \delta)$. Then, τ' defined by $\tau'(a|m) = \sum_{a'} \delta(a|a')\tau(a'|m)$ for all (a, m) , yields

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

Thus, $U_r(\boldsymbol{\sigma}, \mu, \tau') = U_r(\boldsymbol{\sigma}^*, \mu, \delta) > U_r(\boldsymbol{\sigma}^*, \mu, \tau^*) = U_r(\boldsymbol{\sigma}, \mu, \tau)$, so $\tau \notin BR(\boldsymbol{\sigma}, \mu)$. \square

Proof of Lemma 1. IF. Let $(\boldsymbol{\sigma}, \mu)$ be an ambiguous experiment. We argue that if $\tau^* \in br(\boldsymbol{\sigma}^*)$, where $\boldsymbol{\sigma}^* = \sum_{\theta} em_{\theta}^{(\boldsymbol{\sigma}, \mu)} \boldsymbol{\sigma}_{\theta}$, then $\tau^* \in BR(\boldsymbol{\sigma}, \mu)$. Since $\tau^* \in br(\boldsymbol{\sigma}^*)$, $\forall b, a \in A$,

$$\sum_{\omega} p(\omega) \boldsymbol{\sigma}^*(a|\omega) u_r(a, \omega) \geq \sum_{\omega} p(\omega) \boldsymbol{\sigma}^*(a|\omega) u_r(b, \omega). \quad (\text{A.1})$$

For any strategy τ , $\exists \delta \in \mathbb{R}^{|A| \times |A|}$ such that $\tau = \tau^* + \delta$, which satisfies the following:

$$\forall a \neq b \in A, \delta(b|a) \geq 0, \delta(a|a) \leq 0, \text{ and } \sum_{\tilde{a} \in A} \delta(\tilde{a}|a) = 0 \quad (\text{A.2})$$

The concavity of ϕ_r implies that $\phi_r(U_r(\boldsymbol{\sigma}, \mu, \tau))$ is concave in τ . Hence, for all δ ,

$$\phi_r(U_r(\boldsymbol{\sigma}, \mu, \tau^* + \delta)) \leq \phi_r(U_r(\boldsymbol{\sigma}, \mu, \tau^*)) + \sum_{b, a \in A} \left. \frac{\partial \phi_r(U_r(\boldsymbol{\sigma}, \mu, \tau))}{\partial \tau(b|a)} \right|_{\tau=\tau^*} \delta(b|a).$$

Sufficient for τ^* to be a solution to the receiver's program is that for all δ satisfying (A.2),

$$\sum_{b, a \in A} \left. \frac{\partial \phi_r(U_r(\boldsymbol{\sigma}, \mu, \tau))}{\partial \tau(b|a)} \right|_{\tau=\tau^*} \delta(b|a) \leq 0. \quad (\text{A.3})$$

To show (A.3) holds, note that $\left. \frac{\partial \phi_r(U_r(\boldsymbol{\sigma}, \mu, \tau))}{\partial \tau(b|a)} \right|_{\tau=\tau^*} = \sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\boldsymbol{\sigma}_{\tilde{\theta}}, \tau^*)) \sum_{\omega} p(\omega) \boldsymbol{\sigma}^*(a|\omega) u_r(b, \omega)$.

Then by (A.2), $-\delta(a|a) = \sum_{b \neq a} \delta(b|a)$, and we have:

$$\begin{aligned} & \frac{\sum_{b, a \in A} \left. \frac{\partial \phi_r(U_r(\boldsymbol{\sigma}, \mu, \tau))}{\partial \tau(b|a)} \right|_{\tau=\tau^*} \delta(b|a)}{\sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\boldsymbol{\sigma}_{\tilde{\theta}}, \tau^*))} = \sum_{b, a \in A} \sum_{\omega} p(\omega) \boldsymbol{\sigma}^*(a|\omega) u_r(b, \omega) \delta(b|a) \\ & = \sum_{a \in A} \sum_{b \neq a} \delta(b|a) \left(\sum_{\omega} p(\omega) \boldsymbol{\sigma}^*(a|\omega) u_r(b, \omega) - \sum_{\omega} p(\omega) \boldsymbol{\sigma}^*(a|\omega) u_r(a, \omega) \right) \leq 0, \end{aligned}$$

where the last inequality follows from $\delta(b|a) \geq 0$ for all $b \neq a$ and (A.1). This implies (A.3) as $(\sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\boldsymbol{\sigma}_{\tilde{\theta}}, \tau^*))) > 0$. Therefore, we have shown that $\tau^* \in BR(\boldsymbol{\sigma}, \mu)$.

ONLY IF. The proof is nearly identical to the if direction and left to the reader. \square

Proof of Theorem 1. First, we show that there is a unique u that solves $\Phi^*(u) = 0$.

Lemma A.1. $\Phi^*(u)$ satisfies single-crossing, i.e., for any $u > u'$, $\Phi^*(u) \geq 0 \Rightarrow \Phi^*(u') > 0$. Thus, there exists a unique $u \in \mathbb{R}$ such that $\Phi^*(u) = 0$.

Proof of Lemma A.1. $\Phi^*(u) \geq 0$ and ϕ_s, ϕ_r strictly increasing implies $\exists(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}$ such that $\sum_{\theta \in \Theta} \lambda_\theta \frac{\phi_s(u_s(\sigma_\theta, \tau^*))}{\phi_r'(u_r(\sigma_\theta, \tau^*))} \geq \sum_{\theta \in \Theta} \lambda_\theta \frac{\phi_s(u)}{\phi_r'(u_r(\sigma_\theta, \tau^*))} > \sum_{\theta \in \Theta} \lambda_\theta \frac{\phi_s(u')}{\phi_r'(u_r(\sigma_\theta, \tau^*))}$, which implies $\Phi^*(u') > 0$ for $u' < u$. Therefore, there is at most one solution to $\Phi^*(u) = 0$. Since ϕ_s is continuous, Φ_u is continuous in (σ, u) and, thus, by Berge's Maximum Theorem, $\Phi^*(u)$ is continuous. Since $\Phi^*(u) < 0$ for $u > \max\{u_s(a, \omega)\}$ and $\Phi^*(u) > 0$ for $u < \min\{u_s(a, \omega)\}$, $\exists u$ such that $\Phi^*(u) = 0$. \square

By Proposition 1, we can rewrite (\mathcal{P}) as

$$(\mathcal{P}) = \begin{cases} \max_{(\sigma, \mu)} U_s(\sigma, \mu, \tau^*), \\ \text{subject to } \tau^* \in BR(\sigma, \mu). \end{cases}$$

We can then show the conclusion using the following lemma.

Lemma A.2. For each $u \in \mathbb{R}$, $(\mathcal{P}) > u$ if, and only if, $\Phi^*(u) > 0$.

Proof of Lemma A.2. IF. Suppose there exists a solution $(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}$ such that $\Phi^*(u) > 0$. Let $\sigma = (\sigma_\theta)_{\theta \in \Theta}$ and μ be defined by $\mu_\theta := \frac{\lambda_\theta / \phi_r'(u_r(\sigma_\theta, \tau^*))}{\sum_j \lambda_j / \phi_r'(u_r(\sigma_{\theta^j}, \tau^*))}$ for each $\theta \in \Theta$. By construction, (σ, μ) satisfies $em_\theta^{(\sigma, \mu)} = \lambda_\theta$ and $\sum_{\theta \in \Theta} \lambda_\theta \sigma_\theta \in \Sigma^*$. Lemma 1 implies $\tau^* \in BR(\sigma, \mu)$. Thus, $(\mathcal{P}) > u$ since $U_s(\sigma, \mu, \tau^*) > u$ as

$$\phi_s(U_s(\sigma, \mu, \tau^*)) - \phi_s(u) = \frac{1}{\sum_{j \in I} \lambda_j / \phi_r'(u_r(\sigma_{\theta^j}, \tau^*))} \sum_{\theta \in \Theta} \lambda_\theta \frac{\phi_s(u_s(\sigma_\theta, \tau^*)) - \phi_s(u)}{\phi_r'(u_r(\sigma_\theta, \tau^*))} > 0,$$

ONLY IF. $(\mathcal{P}) > u$ implies there exists an obedient (σ, μ) such that $U_s(\sigma, \mu, \tau^*) > u$. Let $\lambda_\theta = em_\theta^{(\sigma, \mu)}$. Lemma 1 implies that $\sum_\theta \lambda_\theta \sigma_\theta \in \Sigma^*$. Thus, $\Phi^*(u) > 0$ as $\sum_\theta \lambda_\theta \Phi_u(\sigma_\theta) = (\phi_s(U_s(\sigma, \mu, \tau^*)) - \phi_s(u)) \sum_\theta \lambda_\theta / \phi_r'(u_r(\sigma_\theta, \tau^*)) > 0$ since $U_s(\sigma, \mu, \tau^*) > u$. \square

We now complete the proof by showing that $(\mathcal{P}) = u$ if, and only if, $\Phi^*(u) = 0$. Suppose $(\mathcal{P}) = u$. Then for all $u' < u$, $(\mathcal{P}) > u'$ and thus $\Phi^*(u') > 0$ by Lemma A.2. By Lemma A.1, there exists a unique \hat{u} such that $\Phi^*(\hat{u}) = 0$. If $\hat{u} > u$, then $\Phi^*(u) > 0$ by Lemma A.1, and thus $(\mathcal{P}) > u$ by Lemma A.2, a contradiction. Thus, $\Phi^*(u) = 0$.

For the other direction, suppose $\Phi^*(u) = 0$. If $(\mathcal{P}) > u$, then Lemma A.2 implies $\Phi^*(u) > 0$, a contradiction. If $(\mathcal{P}) < u$, then there exists $u' < u$ such that $(\mathcal{P}) = u'$. Then by the previous direction, $\Phi^*(u') = 0$, contradicting Lemma A.1. Thus, $(\mathcal{P}) = u$.

Recall that Σ is a convex subset of $\mathbb{R}^{|\Omega| \times (|A|-1)}$. As a result, the graph of Φ_u is a subset of $\mathbb{R}^{(|\Omega| \times (|A|-1) + 1)}$. Suppose that $(\mathcal{P}) = u$. Then $\Phi^*(u) = 0$. Let $(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}$ be such that $\sum_{\theta \in \Theta} \lambda_\theta \Phi_u(\sigma_\theta) = 0$ and $\sum_{\theta \in \Theta} \lambda_\theta \sigma_\theta = \sigma^* \in \Sigma^*$. Thus, $(\sigma^*, 0)$ is on the boundary of, and thus an element of a supporting hyperplane of, the convex hull of the graph of Φ_u . The intersection of this hyperplane and the set forms a face and thus has dimension at most $|\Omega| \times (|A| - 1)$. Extreme points of the face are also extreme points of the convex hull of the graph of Φ_u . Any such extreme point has the form $(\sigma, \Phi_u(\sigma))$ for some $\sigma \in \Sigma$. By Caratheodory's theorem applied to the face, $(\sigma^*, 0)$ can be written as a convex combination of at most $(|\Omega| \times (|A| - 1) + 1)$ such extreme points. Denote the coefficients in the convex combination and the experiments corresponding to these extreme points by $(\hat{\lambda}_{\hat{\theta}}, \hat{\sigma}_{\hat{\theta}})_{\hat{\theta} \in \hat{\Theta}}$ with $|\hat{\Theta}| \leq (|\Omega| \times (|A| - 1) + 1)$. Thus, there exists a solution to (\mathcal{P}) with $\sigma^* = (\hat{\sigma}_{\hat{\theta}})_{\hat{\theta} \in \hat{\Theta}}$ and μ^* with $\text{supp}(\mu^*) \subseteq \hat{\Theta}$ so that $|\text{supp}(\mu^*)| \leq (|\Omega| \times (|A| - 1) + 1)$. \square

To prove Theorem 2 and Proposition 2, we first state two auxiliary results whose proofs are in the Online Appendix (Cheng et al., 2024).

Definition A.1. For any $u \in \mathbb{R}$, let $\Sigma_+(u) = \{\sigma \in \Sigma : u_s(\sigma, \tau^*) > u\}$ and $\Sigma_-(u) = \{\sigma \in \Sigma : u_s(\sigma, \tau^*) \leq u\}$.

Theorem A.1. Let $(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}$ be a solution to $(\Phi^*(u))$. Then:

- (i) For all θ and θ' such that $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_+(u) \times \Sigma_-(u)$, if $\phi'_r(u_r(\sigma_\theta, \tau^*)) \neq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$, then they are Pareto-ranked.
- (ii) If $1/\phi'_r$ is concave, then $\forall \theta, \theta'$ such that $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_+(u) \times \Sigma_+(u)$, if $u_s(\sigma_\theta, \tau^*) \neq u_s(\sigma_{\theta'}, \tau^*)$ and $\phi'_r(u_r(\sigma_\theta, \tau^*)) \neq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$, they are Pareto-ranked.
- (iii) If $1/\phi'_r$ is convex, then $\forall \theta, \theta'$ such that $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_-(u) \times \Sigma_-(u)$, if $u_s(\sigma_\theta, \tau^*) \neq u_s(\sigma_{\theta'}, \tau^*)$ and $\phi'_r(u_r(\sigma_\theta, \tau^*)) \neq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$, they are Pareto-ranked.

Theorem A.2. Let $(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}$ be a solution to $(\Phi^*(u))$. Then:

- (i) $\forall \sigma_\theta, \nexists$ a Pareto-ranked splitting, $(\bar{\sigma}, \underline{\sigma}, \lambda)$ with $(\bar{\sigma}, \underline{\sigma}) \in \Sigma_+(u) \times \Sigma_-(u)$, $\phi'_r(u_r(\bar{\sigma}, \tau^*)) < \phi'_r(u_r(\underline{\sigma}, \tau^*))$, $\frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_s(u_s(\underline{\sigma}, \tau^*))} > \frac{\phi'_r(u_r(\bar{\sigma}, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))}$, if $\sigma_\theta \in \Sigma_+(u)$, and $\frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_s(u_s(\underline{\sigma}, \tau^*))} > \frac{\phi'_r(u_r(\sigma_\theta, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))}$, if $\sigma_\theta \in \Sigma_-(u)$.

(ii) If $1/\phi'_r$ is concave, then $\forall \sigma_\theta$, \exists a Pareto-ranked splitting, $(\bar{\sigma}, \underline{\sigma}, \lambda)$ with $(\bar{\sigma}, \underline{\sigma}) \in \Sigma_-(u) \times \Sigma_-(u)$, $\phi'_r(u_r(\bar{\sigma}, \tau^*)) < \phi'_r(u_r(\underline{\sigma}, \tau^*))$, and $\frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_s(u_s(\underline{\sigma}, \tau^*))} > \frac{\phi'_r(u_r(\bar{\sigma}, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))}$.

(iii) If $1/\phi'_r$ is convex, then $\forall \sigma_\theta$, \exists a Pareto-ranked splitting, $(\bar{\sigma}, \underline{\sigma}, \lambda)$ with $(\bar{\sigma}, \underline{\sigma}) \in \Sigma_+(u) \times \Sigma_+(u)$, $\phi'_r(u_r(\bar{\sigma}, \tau^*)) < \phi'_r(u_r(\underline{\sigma}, \tau^*))$, and $\frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_s(u_s(\underline{\sigma}, \tau^*))} > \frac{\phi'_r(u_r(\bar{\sigma}, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))}$.

Versions of Theorems A.1 and A.2 with $\Sigma_+(u)$ and $\Sigma_-(u)$ defined by swapping the strict and weak inequalities in Definition A.1 also hold.

Proof of Theorem 2. Suppose (σ, μ) is obedient and ϕ_r is strictly concave. Let $\hat{u} := U_s(\sigma, \mu, \tau^*)$ and $\lambda_\theta = em_\theta^{(\sigma, \mu)}$ for all θ . Then $\sum_i \lambda_\theta \frac{\phi_s(u_s(\sigma_\theta, \tau^*)) - \phi_s(\hat{u})}{\phi'_r(u_r(\sigma_\theta, \tau^*))} = 0$. If there exists $(\hat{\lambda}_\theta, \hat{\sigma}_\theta)$ such that $\sum_\theta \hat{\lambda}_\theta \frac{\phi_s(u_s(\hat{\sigma}_\theta, \tau^*)) - \phi_s(\hat{u})}{\phi'_r(u_r(\hat{\sigma}_\theta, \tau^*))} > 0$, then the ambiguous experiment $(\hat{\sigma}, \hat{\mu})$ with $em_\theta^{(\hat{\sigma}, \hat{\mu})} := \hat{\lambda}_\theta$ will strictly improve upon (σ, μ) . The existence of such $(\hat{\lambda}_\theta, \hat{\sigma}_\theta)$ under the conditions in (i) and (ii) of Theorem 2 follows from (i) of Theorems A.1 and A.2, respectively (and their alternative versions). \square

Proof of Proposition 2. When both ϕ_s and $1/\phi'_r$ are linear, notice all the conditions in Theorem A.1 and Theorem A.2 (and the variations of them using definitions of $\Sigma_+(u)$ and $\Sigma_-(u)$ that swap the strict and weak inequalities) are satisfied. \square

Proof of Theorem 3. Suppose that $U_s(\sigma, \mu, \tau^*) > u_s^{BP}$. Observe that $\Sigma_+(u_s^{BP})$ and $\Sigma_-(u_s^{BP})$ have non-empty intersections with the support of μ since, if $\Sigma_+(u_s^{BP})$ did not, then (σ, μ) could not benefit the sender, while if $\Sigma_-(u_s^{BP})$ did not, then $\tau^* \notin BR(\sigma, \mu)$, contradicting obedience. Thus, there exists $\theta, \theta' \in \text{supp}(\mu)$ such that $\sigma_\theta \in \Sigma_+(u_s^{BP})$ and $\sigma_{\theta'} \in \Sigma_-(u_s^{BP})$. Define $\sigma^* := \sum_{\hat{\theta}} em_{\hat{\theta}}^{(\sigma, \mu)} \sigma_{\hat{\theta}}$. Then,

$$\sigma^* = \sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_+(u_s^{BP})} \frac{\mu_{\hat{\theta}} \phi'_r(u_r(\sigma_{\hat{\theta}}, \tau^*))}{\sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\sigma_{\tilde{\theta}}, \tau^*))} \sigma_{\hat{\theta}} + \sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_-(u_s^{BP})} \frac{\mu_{\hat{\theta}} \phi'_r(u_r(\sigma_{\hat{\theta}}, \tau^*))}{\sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\sigma_{\tilde{\theta}}, \tau^*))} \sigma_{\hat{\theta}}.$$

By Lemma 1, we have σ^* is obedient. By definition of u_s^{BP} , this implies

$$\phi_s(u_s(\sigma^*, \tau^*)) \leq \phi_s(u_s^{BP}). \quad (\text{A.4})$$

Suppose that for all $\theta, \theta' \in \text{supp}(\mu)$ such that $\sigma_\theta \in \Sigma_+(u_s^{BP})$ and $\sigma_{\theta'} \in \Sigma_-(u_s^{BP})$, σ_θ and $\sigma_{\theta'}$ are not Pareto-ranked. This is equivalent to

$$u_r(\sigma_\theta, \tau^*) \leq u_r(\sigma_{\theta'}, \tau^*). \quad (\text{A.5})$$

The remainder of the proof shows that this contradicts (A.4). From (A.5) and concavity of ϕ_r , $\phi'_r(u_r(\sigma_\theta, \tau^*)) \geq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$. Observe that,

$$\begin{aligned}
& (\phi_s(u_s(\sigma^*, \tau^*)) - \phi_s(u_s^{BP})) \sum_{\hat{\theta}} \mu_{\hat{\theta}} \phi'_r(u_r(\sigma_{\hat{\theta}}, \tau^*)) \\
& \geq \left(\sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_+(u_s^{BP})} \frac{\mu_{\hat{\theta}} \phi'_r(u_r(\sigma_{\hat{\theta}}, \tau^*))}{\sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\sigma_{\tilde{\theta}}, \tau^*))} (\phi_s(u_s(\sigma_{\hat{\theta}})) - \phi_s(u_s^{BP})) \right. \\
& \quad \left. + \sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_-(u_s^{BP})} \frac{\mu_{\hat{\theta}} \phi'_r(u_r(\sigma_{\hat{\theta}}, \tau^*))}{\sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\sigma_{\tilde{\theta}}, \tau^*))} (\phi_s(u_s(\sigma_{\hat{\theta}})) - \phi_s(u_s^{BP})) \right) \sum_{\tilde{\theta}} \mu_{\tilde{\theta}} \phi'_r(u_r(\sigma_{\tilde{\theta}}, \tau^*)) \\
& = \sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_+(u_s^{BP})} \mu_{\hat{\theta}} \phi'_r(u_r(\sigma_{\hat{\theta}}, \tau^*)) (\phi_s(u_s(\sigma_{\hat{\theta}}, \tau^*)) - \phi_s(u_s^{BP})) \\
& \quad + \sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_-(u_s^{BP})} \mu_{\hat{\theta}} \phi'_r(u_r(\sigma_{\hat{\theta}}, \tau^*)) (\phi_s(u_s(\sigma_{\hat{\theta}}, \tau^*)) - \phi_s(u_s^{BP})) \\
& \geq \phi'_r \left(\max_{\hat{\theta} \in \Sigma_+(u_s^{BP})} u_r(\sigma_{\hat{\theta}}, \tau^*) \right) \sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_+(u_s^{BP})} \mu_{\hat{\theta}} (\phi_s(u_s(\sigma_{\hat{\theta}}, \tau^*)) - \phi_s(u_s^{BP})) \\
& \quad + \phi'_r \left(\min_{\hat{\theta} \in \Sigma_-(u_s^{BP})} u_r(\sigma_{\hat{\theta}}, \tau^*) \right) \sum_{\hat{\theta}: \sigma_{\hat{\theta}} \in \Sigma_-(u_s^{BP})} \mu_{\hat{\theta}} (\phi_s(u_s(\sigma_{\hat{\theta}}, \tau^*)) - \phi_s(u_s^{BP})) \\
& \geq \phi'_r \left(\min_{\hat{\theta} \in \Sigma_-(u_s^{BP})} u_r(\sigma_{\hat{\theta}}, \tau^*) \right) \left[\left(\sum_{\hat{\theta}} \mu_{\hat{\theta}} \phi_s(u_s(\sigma_{\hat{\theta}}, \tau^*)) \right) - \phi_s(u_s^{BP}) \right] > 0,
\end{aligned}$$

implying $\phi_s(u_s(\sigma^*, \tau^*)) > \phi_s(u_s^{BP})$, contradicting (A.4). The first inequality follows from substituting for σ^* and concavity of ϕ_s , the second from the definitions of $\Sigma_+(u_s^{BP})$ and $\Sigma_-(u_s^{BP})$ and concavity of ϕ_r , the third since (A.5) implies $\max_{\hat{\theta} \in \Sigma_+(u_s^{BP})} u_r(\sigma_{\hat{\theta}}, \tau^*) \leq \min_{\hat{\theta} \in \Sigma_-(u_s^{BP})} u_r(\sigma_{\hat{\theta}}, \tau^*)$, and the final one since (σ, μ) benefits the sender. \square

Proof of Theorem 4. Suppose that there exists a solution $(\sigma^*, \mu^*, \tau^*)$ to the maximization problem (P) that benefits the sender. Let $\sigma := \sum_{\theta} \mu_{\theta}^* \sigma_{\theta}^*$ and $\sigma^* := \sum_{\theta} em_{\theta}^{(\sigma^*, \mu^*)} \sigma_{\theta}^*$. These experiments satisfy the conditions in part (b) of the theorem: Since $em^{(\sigma^*, \mu^*)}$ and μ^* have the same support on Θ , $\text{supp } \sigma(\cdot | \omega) = \text{supp } \sigma^*(\cdot | \omega)$ for all ω . From Lemma 1, since $\tau^* \in BR(\sigma^*, \mu^*)$, $\tau^* \in br(\sigma^*)$. Since the sender benefits from ambiguous communication, we have that $u_s^{BP} < \phi_s^{-1}(\sum_{\theta} \mu_{\theta}^* \phi_s(u_s(\sigma_{\theta}^*, \tau^*))) \leq \sum_{\theta} \mu_{\theta}^* u_s(\sigma_{\theta}^*, \tau^*) = u_s(\sigma, \tau^*)$. This further implies that $\tau^* \notin br(\sigma)$, and thus $em^{(\sigma^*, \mu^*)} \neq \mu^*$ implying that ϕ_r is not affine. Since $\tau^* \in br(\sigma^*)$, we have that $u_s^{BP} \geq u_s(\sigma^*, \tau^*)$ and, thus, $u_s(\sigma, \tau^*) > u_s^{BP} \geq u_s(\sigma^*, \tau^*)$.

Since $em^{(\sigma^*, \mu^*)} \neq \mu^*$, there must exist a pair (θ, θ') in the support of μ^* such that

$u_r(\sigma_\theta, \tau^*) > u_r(\sigma_{\theta'}, \tau^*)$ and $\phi'_r(u_r(\sigma_\theta, \tau^*)) < \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$. We next use the following lemma and concavity of ϕ_r to show $u_r(\sigma, \tau^*) > u_r(\sigma^*, \tau^*)$.

Lemma A.3. *Fix any two monotonic sequences $x_1 \geq x_2 \geq \dots \geq x_n$, $0 < y_1 \leq y_2 \leq \dots \leq y_n$, and a probability $\mu \in \Delta(\{1, 2, \dots, n\})$. Assume that there exist indices $i^* < j^*$ such that $\mu_{i^*} > 0$, $\mu_{j^*} > 0$, $x_{i^*} > x_{j^*}$ and $y_{i^*} < y_{j^*}$. Then $\sum_{i=1}^n x_i \frac{\mu_i y_i}{\sum_{j=1}^n \mu_j y_j} < \sum_{i=1}^n x_i \mu_i$.*

(See Online Appendix (Cheng et al., 2024) for a proof of Lemma A.3.)

To prove $u_r(\sigma, \tau^*) > u_r(\sigma^*, \tau^*)$, we apply the lemma to the decreasing rearrangement of the sequence $(u_r(\sigma_\theta^*, \tau^*))_\theta$ (the x_i 's) and the increasing rearrangement of $(\phi'_r(u_r(\sigma_\theta^*, \tau^*)))_\theta$ (the y_i 's). Since ϕ_r is strictly increasing and concave, we have that $\phi'_r(u_r(\sigma_\theta^*, \tau^*)) > 0$, and $u_r(\sigma_\theta^*, \tau^*) \geq u_r(\sigma_{\theta'}^*, \tau^*)$ implies that $\phi'_r(u_r(\sigma_\theta^*, \tau^*)) \leq \phi'_r(u_r(\sigma_{\theta'}^*, \tau^*))$. There exists $i^* < j^*$ such that $\mu_{i^*} > 0$, $\mu_{j^*} > 0$, $x_{i^*} > x_{j^*}$ and $y_{i^*} < y_{j^*}$ since there exists a pair (θ, θ') in the support of μ such that $u_r(\sigma_\theta, \tau^*) > u_r(\sigma_{\theta'}, \tau^*)$ and $\phi'_r(u_r(\sigma_\theta, \tau^*)) < \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$. Applying the lemma and the definition of σ and σ^* yields $u_r(\sigma, \tau^*) > u_r(\sigma^*, \tau^*)$. This completes the proof that the conditions in (b) are necessary for ambiguity to benefit the sender. To establish that (b) implies (a), we rely on the following sufficient condition for the existence of a Pareto-ranked splitting of an experiment σ :

Lemma A.4. *Let σ be an experiment. If there exists $\hat{\sigma}$ such that $u_s(\hat{\sigma}, \tau^*) > u_s(\sigma, \tau^*)$, $u_r(\hat{\sigma}, \tau^*) > u_r(\sigma, \tau^*)$ and for all $\omega \in \Omega$, $\text{supp}(\hat{\sigma}(\cdot|\omega)) \subseteq \text{supp}(\sigma(\cdot|\omega))$, then there exists a Pareto-ranked splitting of σ , $(\bar{\sigma}, \underline{\sigma}, \lambda)$ with $\bar{\sigma} = \hat{\sigma}$.*

(See the Online Appendix (Cheng et al., 2024) for a proof of Lemma A.4.)

Observe that the conditions in (b) of Theorem 4 imply the conditions in Lemma A.4. Therefore a Pareto-ranked splitting of σ^* , $(\bar{\sigma}, \underline{\sigma}, \lambda)$ exists with $\bar{\sigma} = \sigma$. Then by condition (ii) in (b) of Theorem 4, $u_s(\bar{\sigma}, \tau^*) > u_s^{BP}$. This shows that (b) implies (a). It remains to show (a) implies (b). Suppose there exists an obedient experiment $\hat{\sigma}$ satisfying the conditions in (a). Since $u_s(\bar{\sigma}, \tau^*) > u_s^{BP}$, there exists a $\underline{\gamma} \in (\lambda, 1)$ such that for all $\gamma \in (\underline{\gamma}, 1)$, $\gamma u_s(\bar{\sigma}, \tau^*) + (1 - \gamma)u_s(\underline{\sigma}, \tau^*) > u_s^{BP}$, and thus $\tau^* \notin \text{br}(\gamma\bar{\sigma} + (1 - \gamma)\underline{\sigma})$. Fix any such γ and let $\sigma = \gamma\bar{\sigma} + (1 - \gamma)\underline{\sigma}$. Let $\sigma^* = \hat{\sigma}$. It follows that $\tau^* \in \text{br}(\sigma^*)$ since $\hat{\sigma}$ is obedient, and $u_s^{BP} \geq u_s(\sigma^*, \tau^*)$, hence $\gamma > \lambda$. Since both σ and σ^* are strict mixtures of $\bar{\sigma}$ and $\underline{\sigma}$, they satisfy the common support condition (i) of (b). By the definition of Pareto-ranked splitting and $\gamma > \lambda$, $u_r(\sigma, \tau^*) > u_r(\sigma^*, \tau^*)$. This establishes that (a) implies (b), completing the proof. \square

Proof of Corollary 3. Suppose that there exist experiments, σ and σ^* , satisfying the conditions in (b) of Theorem 4. Since condition (iii) of (b) implies $\tau^* \notin br(\sigma)$, either

$$\sum_{\omega} u_r(a_1, \omega) \sigma(a_1|\omega) p(\omega) < \sum_{\omega} u_r(a_2, \omega) \sigma(a_1|\omega) p(\omega), \text{ or} \quad (\text{A.6})$$

$$\sum_{\omega} u_r(a_2, \omega) \sigma(a_2|\omega) p(\omega) < \sum_{\omega} u_r(a_1, \omega) \sigma(a_2|\omega) p(\omega). \quad (\text{A.7})$$

Assume (A.6). From (iii) of (b), $\tau^* \in br(\sigma^*)$, implying $\sum_{\omega} u_r(a_1, \omega) \sigma^*(a_1|\omega) p(\omega) \geq \sum_{\omega} u_r(a_2, \omega) \sigma^*(a_1|\omega) p(\omega)$. It follows that: $\sum_{\omega} u_r(a_1, \omega) [\sigma(a_1|\omega) - \sigma^*(a_1|\omega)] p(\omega) < \sum_{\omega} u_r(a_2, \omega) [\sigma(a_1|\omega) - \sigma^*(a_1|\omega)] p(\omega) = \sum_{\omega} u_r(a_2, \omega) [\sigma^*(a_2|\omega) - \sigma(a_2|\omega)] p(\omega)$. Therefore, $\sum_{\omega} u_r(a_2, \omega) [\sigma^*(a_2|\omega) - \sigma(a_2|\omega)] p(\omega) + \sum_{\omega} u_r(a_1, \omega) [\sigma^*(a_1|\omega) - \sigma(a_1|\omega)] p(\omega) > 0$, contradicting (ii) of (b). The argument from (A.7) is analogous. \square

Proof of Lemma 2.

$$\begin{aligned} & \frac{\phi_i(u_i(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - (\lambda\phi_i(u_s(\bar{\sigma}, \tau^*)) + (1-\lambda)\phi_i(u_s(\underline{\sigma}, \tau^*)))}{\phi_i(u_i(\bar{\sigma}, \tau^*)) - \phi_i(u_i(\underline{\sigma}, \tau^*))} < \mu - \lambda \\ \Leftrightarrow & \phi_i(u_i(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) < \mu\phi_i(u_i(\bar{\sigma}, \tau^*)) + (1-\mu)\phi_i(u_i(\underline{\sigma}, \tau^*)) \end{aligned}$$

Thus, equivalently, $u_i(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*) < U_i(\sigma, \mu, \tau^*)$. \square

Proof of Theorem 5. Fix an obedient σ^* and let $(\bar{\sigma}, \underline{\sigma}, \lambda)$ be a Pareto-ranked splitting of σ^* satisfying $u_i(\bar{\sigma}, \tau^*) > u_i(\underline{\sigma}, \tau^*)$ for $i \in \{s, r\}$.

Observe that for $i \in \{s, r\}$, $U_i(\sigma, \mu, \tau^*) = \phi_i^{-1}(\mu\phi_i(u_i(\bar{\sigma}, \tau^*)) + (1-\mu)\phi_i(u_i(\underline{\sigma}, \tau^*)))$, and $u_i(\sigma^*, \tau^*) = \lambda u_i(\bar{\sigma}, \tau^*) + (1-\lambda)u_i(\underline{\sigma}, \tau^*)$. (5) implies that $em_{\theta_1}^{(\sigma, \mu)} = \lambda$, and $em_{\theta_2}^{(\sigma, \mu)} = 1 - \lambda$. Lemma 1 then implies that (σ, μ) is obedient since σ^* is. This proves part (i).

By Lemma 2, $U_r(\sigma, \mu, \tau^*) > u_r(\sigma^*, \tau^*)$ if and only if the receiver's $(\{\bar{\sigma}, \underline{\sigma}\}, \lambda)$ -probability premium is strictly less than $\mu - \lambda$. The latter is equivalent to:

$$\begin{aligned} & \frac{\phi_r(u_r(\sigma^*, \tau^*)) - \lambda\phi_r(u_r(\bar{\sigma}, \tau^*)) - (1-\lambda)\phi_r(u_r(\underline{\sigma}, \tau^*))}{\phi_r(u_r(\bar{\sigma}, \tau^*)) - \phi_r(u_r(\underline{\sigma}, \tau^*))} + \lambda < \frac{\lambda\phi_r'(u_r(\underline{\sigma}, \tau^*))}{\lambda\phi_r'(u_r(\underline{\sigma}, \tau^*)) + (1-\lambda)\phi_r'(u_r(\bar{\sigma}, \tau^*))} \\ \Leftrightarrow & \frac{\phi_r(u_r(\sigma^*, \tau^*)) - \phi_r(u_r(\underline{\sigma}, \tau^*))}{\phi_r(u_r(\sigma^*, \tau^*)) - \phi_r(u_r(\underline{\sigma}, \tau^*)) + \phi_r(u_r(\bar{\sigma}, \tau^*)) - \phi_r(u_r(\sigma^*, \tau^*))} \\ & < \frac{\phi_r'(u_r(\underline{\sigma}, \tau^*))(u_r(\sigma^*, \tau^*) - u_r(\underline{\sigma}, \tau^*))}{\phi_r'(u_r(\underline{\sigma}, \tau^*))(u_r(\sigma^*, \tau^*) - u_r(\underline{\sigma}, \tau^*)) + \phi_r'(u_r(\bar{\sigma}, \tau^*))(u_r(\bar{\sigma}, \tau^*) - u_r(\sigma^*, \tau^*))} \end{aligned}$$

$$\Leftrightarrow \frac{1}{1 + \frac{\phi_r(u_r(\bar{\sigma}, \tau^*)) - \phi_r(u_r(\sigma^*, \tau^*))}{\phi_r(u_r(\sigma^*, \tau^*)) - \phi_r(u_r(\underline{\sigma}, \tau^*))}} < \frac{1}{1 + \frac{\phi'_r(u_r(\bar{\sigma}, \tau^*))(u_r(\bar{\sigma}, \tau^*) - u_r(\sigma^*, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))(u_r(\sigma^*, \tau^*) - u_r(\underline{\sigma}, \tau^*))}}$$

$$\Leftrightarrow \phi'_r(u_r(\bar{\sigma}, \tau^*)) < \phi'_r(u_r(\underline{\sigma}, \tau^*)) \Leftrightarrow \lambda < \frac{\lambda \phi'_r(u_r(\underline{\sigma}, \tau^*))}{\lambda \phi'_r(u_r(\underline{\sigma}, \tau^*)) + (1 - \lambda) \phi'_r(u_r(\bar{\sigma}, \tau^*))} = \mu.$$

where the first equivalence uses the fact that $\lambda = \frac{u_r(\sigma^*, \tau^*) - u_r(\underline{\sigma}, \tau^*)}{u_r(\bar{\sigma}, \tau^*) - u_r(\underline{\sigma}, \tau^*)}$, the second is algebra, the third follows from concavity and differentiability of ϕ_r , and the fourth from the strict positivity of ϕ'_r .¹⁸ This proves part (ii). Part (iii) follows directly from Lemma 2.

Within the smooth ambiguity model, an increase in ambiguity aversion corresponds to ϕ becoming more concave (Klibanoff et al., 2005). Given differentiability, $\tilde{\phi}$ more concave than ϕ means that $\tilde{\phi} := \varphi \circ \phi$ for some strictly increasing, concave, and differentiable φ . To see that the sender's probability premium increases in the sender's ambiguity aversion, observe that $\rho^{\tilde{\phi}, u_s}((\bar{\sigma}, \underline{\sigma}), \lambda) + \lambda$ is equal to

$$\frac{\varphi(\phi(u(\sigma^*, \tau^*))) - \varphi(\phi(u(\underline{\sigma}, \tau^*)))}{\varphi(\phi(u(\bar{\sigma}, \tau^*))) - \varphi(\phi(u(\sigma^*, \tau^*))) + \varphi(\phi(u(\sigma^*, \tau^*))) - \varphi(\phi(u(\underline{\sigma}, \tau^*)))}$$

$$\geq \frac{\varphi'(\phi(u(\sigma^*, \tau^*)))(\phi(u(\sigma^*, \tau^*)) - \phi(u(\underline{\sigma}, \tau^*)))}{\varphi'(\phi(u(\sigma^*, \tau^*)))(\phi(u(\bar{\sigma}, \tau^*)) - \phi(u(\sigma^*, \tau^*))) + \varphi'(\phi(u(\sigma^*, \tau^*)))(\phi(u(\sigma^*, \tau^*)) - \phi(u(\underline{\sigma}, \tau^*)))}$$

$$= \frac{\phi(u(\sigma^*, \tau^*)) - \phi(u(\underline{\sigma}, \tau^*))}{\phi(u(\bar{\sigma}, \tau^*)) - \phi(u(\underline{\sigma}, \tau^*))} = \rho^{\phi, u_s}(\{\bar{\sigma}, \underline{\sigma}\}, \lambda) + \lambda,$$

where the inequality follows from the concavity of φ . The inequality is strict if and only if $\varphi'(\phi(u(\bar{\sigma}, \tau^*))) < \varphi'(\phi(u(\underline{\sigma}, \tau^*)))$, as the strict inequality on φ' implies that either $\varphi'(\phi(u(\bar{\sigma}, \tau^*))) < \varphi'(\phi(u(\sigma^*, \tau^*)))$ or $\varphi'(\phi(u(\sigma^*, \tau^*))) < \varphi'(\phi(u(\underline{\sigma}, \tau^*)))$ or both.

It remains to show that the r.h.s. of (5) is increasing in the receiver's ambiguity aversion. This follows since it is increasing in $\frac{\phi'_r(u_r(\underline{\sigma}, \tau^*))}{\phi'_r(u_r(\bar{\sigma}, \tau^*))}$, and

$$\frac{\tilde{\phi}'(u_r(\underline{\sigma}, \tau^*))}{\tilde{\phi}'(u_r(\bar{\sigma}, \tau^*))} = \frac{\varphi'(\phi(u(\underline{\sigma}, \tau^*)))\phi'(u_r(\underline{\sigma}, \tau^*))}{\varphi'(\phi(u(\bar{\sigma}, \tau^*)))\phi'(u_r(\bar{\sigma}, \tau^*))} \geq \frac{\phi'(u_r(\underline{\sigma}, \tau^*))}{\phi'(u_r(\bar{\sigma}, \tau^*))},$$

by concavity of φ , and strict if and only if $\varphi'(\phi(u(\bar{\sigma}, \tau^*))) < \varphi'(\phi(u(\underline{\sigma}, \tau^*)))$. \square

Proof of Proposition 4. Let $\Sigma_{\sigma^*} \subseteq \Sigma$ denote the set of experiments that, for all $\omega \in \Omega$, have the same support as σ^* . For each $\omega \in \Omega$, fix any $a_\omega \in \text{supp}(\sigma^*(\cdot|\omega))$. For any $\sigma \in \Sigma_{\sigma^*}$, by

¹⁸If ϕ_r were concave, but not differentiable at $u_r(\sigma^*, \tau^*)$, the third equivalence would fail in one direction since we could then have a linear piece from $u_r(\underline{\sigma}, \tau^*)$ to $u_r(\sigma^*, \tau^*)$ with slope $\phi'_r(u_r(\underline{\sigma}, \tau^*))$ and another one from $u_r(\sigma^*, \tau^*)$ to $u_r(\bar{\sigma}, \tau^*)$ with slope $\phi'_r(u_r(\bar{\sigma}, \tau^*))$.

substituting $\sigma(a_\omega|\omega) = 1 - \sum_{a \neq a_\omega} \sigma(a|\omega)$, we have, for $i \in \{s, r\}$,

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

where $\Omega_{\sigma^*} \subseteq \Omega$ denotes the set of ω such that $| \text{supp}(\sigma^*(\cdot|\omega)) | > 1$.

Given any $\sigma \in \Sigma_{\sigma^*}$, we use $\tilde{\sigma} \in \mathbb{R}^{\sum_{\omega \in \Omega_{\sigma^*}} |\text{supp}(\sigma^*(\cdot|\omega)) - 1|} =: \mathbb{R}^{\sigma^*}$ to denote the vector of those components of $\sigma(a|\omega)$ with $\omega \in \Omega_{\sigma^*}$ and $a \in \text{supp}(\sigma^*(\cdot|\omega)) \setminus \{a_\omega\}$. 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 (6) are exactly the 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 linear mapping from \mathbb{R}^{σ^*} to \mathbb{R}^2 defined from the right hand side above is surjective. Since $\{\tilde{\sigma} : \sigma \in \Sigma_{\sigma^*}\} \ni \tilde{\sigma}^*$ is open in \mathbb{R}^{σ^*} , 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.4, there exists a Pareto-ranked splitting of σ^* with $\bar{\sigma} = \hat{\sigma}$. \square

Online Supplemental Appendix to “Persuasion with Ambiguous Communication”

Xiaoyu Cheng Peter Klibanoff Sujoy Mukerji Ludovic Renou

SA.1 A Polar Case: A Maxmin Receiver and Ambiguity Neutral Sender

This section analyzes the case of an ambiguity-neutral sender and an infinitely ambiguity averse receiver, represented by the maxmin preferences U_r^{MEU} . As the comparative statics statements in Theorem 5 suggest, this is the most favorable case for the sender to benefit from ambiguous communication. In fact, as we shall see, the sender can extract nearly all the surplus.

More precisely, we show that in this case (a) binary ambiguous experiments are sufficient to exhaust all gains from persuasion, and (b) the sender can attain a payoff arbitrarily close to their best feasible payoff subject to the receiver getting at least the payoff they would obtain if no information were disclosed. A conclusion we draw is that assuming an infinitely ambiguity averse receiver is very powerful and, in our view, unrealistically so, further motivating the analysis in the rest of the paper which allows for more moderate levels of aversion.

Before turning to the analysis, we remark that the opposite cases, of either an infinitely ambiguity averse sender with payoffs U_s^{MEU} or an ambiguity neutral receiver, preclude any benefit from ambiguous persuasion. The latter case follows from Theorem 4, while Cheng (2023) shows that in the former case the sender never benefits from ambiguous communication.

The following lemma relates obedience for an ambiguous experiment to obedience for an experiment. It is thus the analogue of Lemma 1 for a receiver with preferences U_r^{MEU} :

Lemma SA.1.1. *(σ, μ) is obedient if, and only if, the experiment σ^* is obedient, where*

$$\sigma^* := \frac{\sum_{\substack{\theta \in \arg \min_{\theta \in \text{supp}(\mu)} u_r(\sigma_\theta, \tau^*)}} \mu_\theta \sigma_\theta}{\sum_{\substack{\theta \in \arg \min_{\theta \in \text{supp}(\mu)} u_r(\sigma_\theta, \tau^*)}} \mu_\theta}.$$

Proof of Lemma SA.1.1. Fix (σ, μ) . Obedience requires that $\min_{\theta \in \Theta} u_r(\sigma_\theta, \tau^*) \geq \min_{\theta \in \Theta} u_r(\sigma_\theta, \tau)$. From the minmax theorem, this is equivalent to the existence of $\mu^* \in \Delta(\Theta)$ such that $\sum_{\theta \in \Theta} \mu_\theta u_r(\sigma_\theta, \tau^*) \geq \sum_{\theta \in \Theta} \mu_\theta^* u_r(\sigma_\theta, \tau^*) \geq \sum_{\theta \in \Theta} \mu_\theta^* u_r(\sigma_\theta, \tau)$, for all (μ, τ) . The result follows since μ^* is a minimizer of $\sum_{\theta \in \Theta} \mu_\theta u_r(\sigma_\theta, \tau^*)$. \square

Observe that when the argmin in Lemma SA.1.1 is a singleton, σ^* equals the receiver's payoff-minimizing experiment from σ . More generally, it is a convex combination of the possibly multiple minimizing experiments in σ with relative weights inherited from μ . Thus the analogue of the effective measure here may have a smaller support than μ (something that never happens for a smooth ambiguity receiver). Lemma SA.1.1 says that *only* those payoff-minimizing experiments affect obedience of (σ, μ) . Thus, the sender is free to include in σ and arbitrarily weight *any other experiments* as long as they don't disrupt the receiver's minimum.

Since the receiver can always ignore any recommendations made, they can guarantee themselves the payoff

$$\underline{u}_r^* := \max_{a \in A} \sum_{\omega} p(\omega) u_r(a, \omega),$$

which is the payoff they would obtain if no information were disclosed. The consequence of the great flexibility available to the sender given Lemma SA.1.1 is the next theorem, which states that the sender's optimal payoff approaches their highest feasible payoff subject to the receiver getting at least \underline{u}_r^* . The corresponding communication strategy uses a binary ambiguous experiment with the μ -weight on the better experiment approaching 1, and the worse experiment an obedient one holding the receiver to \underline{u}_r^* .

Theorem SA.1.1. *Suppose there exists $\hat{\sigma}$ such that $u_r(\hat{\sigma}, \tau^*) > \underline{u}_r^*$. The value of the following program is the supremum of the payoff that an ambiguity neutral sender can obtain when the receiver has maxmin preferences U_r^{MEU} :*

$$\begin{aligned} & \max_{\sigma} u_s(\sigma, \tau^*), \\ & \text{s.t. } u_r(\sigma, \tau^*) \geq \underline{u}_r^*. \end{aligned}$$

Proof of Theorem SA.1.1. Let $\bar{\sigma}$ attain the value of the program in the theorem for the sender and $\underline{\sigma}$ be an obedient experiment with $u_r(\underline{\sigma}, \tau^*) = \underline{u}_r^*$. Assume $u_r(\bar{\sigma}, \tau^*) > u_r(\underline{\sigma}, \tau^*)$. If the sender chooses $((\bar{\sigma}, \underline{\sigma}), (\mu, 1 - \mu))$, the receiver is obedient since the worst payoff is $u_r(\underline{\sigma}, \tau^*)$. As μ approaches 1, the sender's payoff approaches $u_s(\bar{\sigma}, \tau^*)$. The

sender cannot do better than this, since the receiver's payoff is at least \underline{u}_r^* . If $u_r(\bar{\sigma}, \tau^*) = u_r(\underline{\sigma}, \tau^*)$, mix $\bar{\sigma}$ with $\hat{\sigma}$. For any $\varepsilon > 0$, $u_r((1-\varepsilon)\bar{\sigma} + \varepsilon\hat{\sigma}, \tau^*) > u_r(\underline{\sigma}, \tau^*)$. As ε approaches 0 and μ approaches 1, the payoff for the sender approaches $u_s(\bar{\sigma}, \tau^*)$. \square

There is a sense in which Theorem SA.1.1 could be argued to overstate what the sender can achieve. For MEU, the “effective” experiment σ^* could have a smaller support than (σ, μ) . Lemma SA.1.1 treats action recommendations that could occur under (σ, μ) but not under σ^* as zero probability events. However, observing such action recommendations would reveal to the receiver that $\theta \notin \arg \min_{\theta \in \text{supp}(\mu)} u_r(\sigma_\theta, \tau^*)$. In this case, the receiver may no longer be indifferent between obeying or not. Therefore, Theorem SA.1.1 could be seen as forcing the receiver to be obedient in such situations.

This issue can be addressed by strengthening obedience to further require that σ^* always has the same support as (σ, μ) (which was always true for smooth ambiguity receivers). This strengthening does not substantially change the conclusions of Theorem SA.1.1, as it only replaces the program in Theorem SA.1.1 by:

$$\begin{aligned} & \sup_{\sigma} u_s(\sigma, \tau^*), \\ & \text{s.t. } u_r(\sigma, \tau^*) > \underline{u}_r^* \text{ and } \text{supp}(\sigma) \subseteq A_0, \end{aligned}$$

where A_0 is the set of all actions which can be best responses for the receiver to some probability distribution over the states in the support of the prior p . The corresponding communication strategies would be binary with μ -weight on the better one approaching 1 as before, but with the worse experiment now adjusted to have full support on A_0 by mixing it with an arbitrarily small amount of an obedient experiment with full support on A_0 that yields the receiver more than \underline{u}_r^* (such an experiment exists under the assumption of Theorem SA.1.1).

Formally, a stronger notion of obedience that does not allow positive μ weight on experiments that recommend actions outside the support of the effective measure weighted experiment of a maxmin receiver is the following:

Lemma SA.1.2. *(σ, μ) is (strongly) obedient if, and only if, the experiment σ^* is obedient, where*

$$\sigma^* := \frac{\sum_{\substack{\theta \in \arg \min \\ \theta \in \text{supp}(\mu)} u_r(\sigma_\theta, \tau^*)} \mu_\theta \sigma_\theta}{\sum_{\substack{\theta \in \arg \min \\ \theta \in \text{supp}(\mu)} u_r(\sigma_\theta, \tau^*)} \mu_\theta},$$

and $\text{supp}(\sigma_\theta) \subseteq \text{supp}(\sigma^*)$ for all $\theta \in \text{supp}(\mu)$.

Define \underline{u}_r^* as before

$$\underline{u}_r^* := \max_{a \in A} \sum_{\omega} p(\omega) u_r(a, \omega).$$

Further define A_0 as the set of all actions which can be best responses for the receiver:

$$A_0 := \{a \in A : \exists q \in \Delta(\Omega) \text{ s.t. } a \in \arg \max_{a' \in A} \sum_{\omega} q(\omega) u_r(a', \omega)\}.$$

The following is the version of Theorem SA.1.1 using the stronger obedience notion:

Theorem SA.1.2. *Suppose there exists $\hat{\sigma}$ such that $u_r(\hat{\sigma}, \tau^*) > \underline{u}_r^*$. The value of the following program is the supremum of the payoff that an ambiguity neutral sender can obtain when the receiver has maxmin preferences U_r^{MEU} and the version of obedience in Lemma SA.1.2 is used:*

$$\begin{aligned} & \sup_{\sigma} u_s(\sigma, \tau^*), \\ & \text{s.t. } u_r(\sigma, \tau^*) > \underline{u}_r^* \text{ and, } \text{supp}(\sigma) \subseteq A_0, \end{aligned}$$

Proof of Theorem SA.1.2. Let $\bar{\sigma}$ attain the value for the sender of the following program

$$\begin{aligned} & \max_{\sigma} u_s(\sigma, \tau^*), \\ & \text{s.t. } u_r(\sigma, \tau^*) \geq \underline{u}_r^* \text{ and, } \text{supp}(\sigma) \subseteq A_0. \end{aligned}$$

Let $\underline{\sigma}$ be an obedient uninformative experiment, so that $u_r(\underline{\sigma}, \tau^*) = \underline{u}_r^*$. Observe that $\text{supp}(\underline{\sigma}) \subseteq A_0$. Assume that $u_r(\bar{\sigma}, \tau^*) > u_r(\underline{\sigma}, \tau^*)$. There exists an obedient experiment $\tilde{\sigma}$ such that $\text{supp}(\tilde{\sigma}) = A_0$ and $u_r(\tilde{\sigma}, \tau^*) > \underline{u}_r^*$.¹⁹ Define a sequence of experiments $\underline{\sigma}_n = (1 - \epsilon_n)\underline{\sigma} + \epsilon_n\tilde{\sigma}$ where $\epsilon_n > 0$ with $\epsilon_n \rightarrow 0$ as n goes to infinity. If the sender offers the ambiguous experiment $((\bar{\sigma}, \underline{\sigma}_n), (\mu, 1 - \mu))$ for small enough ϵ_n , the receiver is strongly obedient since the worst payoff is $u_r(\underline{\sigma}_n, \tau^*)$, obedience is preserved under con-

¹⁹For any $a \in A_0$, fix some $q_a \in \Delta(\text{supp}(p))$ under which a is optimal for the receiver. There exists a $\beta_a \in (0, 1)$ and a $q \in \Delta(\text{supp}(p))$ such that $p = \beta_a q_a + (1 - \beta_a)q'_a$. Let $a'_a \in A_0$ denote an action that is optimal for the receiver under q'_a . Applying this argument to all $a \in A_0$ to construct a set $\cup_{a \in A_0} \{q_a, q'_a\}$ whose convex hull contains p in its interior. Since each probability distribution in the set can be thought of as a Bayesian posterior, this interior convex combination is a Bayes plausible distribution over the posteriors and thus, by Kamenica and Gentzkow (2011), corresponds to an obedient $\tilde{\sigma}$ with $\text{supp}(\tilde{\sigma}) = A_0$. Finally, since $u_r(\hat{\sigma}, \tau^*) > \underline{u}_r^*$, there exists an $a \in A_0$ such that $\sum_{\omega} q_a(\omega) u_r(a, \omega) > \underline{u}_r^*$. Thus, $u_r(\tilde{\sigma}, \tau^*) > \underline{u}_r^*$.

vex combinations of experiments, and $\text{supp}(\bar{\sigma}) \subseteq A_0 = \text{supp}(\underline{\sigma}_n)$. As we can choose μ arbitrarily close to 1 and ϵ_n arbitrarily close to 0, we approach the value of the program in Theorem SA.1.2. Furthermore, it is not possible for the sender to do better than this (i.e., have a higher supremum), since the receiver’s payoff from any obedient experiment (and thus from any obedient ambiguous experiment) is at least \underline{u}_r^* and the strong version of obedience requires that all experiments in the support of μ recommend actions in A_0 .

If $u_r(\bar{\sigma}, \tau^*) = u_r(\underline{\sigma}, \tau^*)$, we need to slightly modify the construction to guarantee obedience. The idea is to mix $\bar{\sigma}$ with a bit of $\tilde{\sigma}$ to guarantee a unique worst payoff, i.e., $u_r((1-2\epsilon_n)\bar{\sigma} + 2\epsilon_n\tilde{\sigma}, \tau^*) > u_r(\underline{\sigma}_n, \tau^*)$ for all $\epsilon_n > 0$. As ϵ_n approaches 0 and μ approaches 1, the payoff for the sender approaches the value of the program in the theorem. \square

SA.2 An Interpretation of the Introductory Example

We present an interpretation of our introductory example from Section 2. This helps motivate the example and illustrate the strategic choices in our model in a stylized real-world setting.

Think of the sender as a banking regulatory authority (“the regulator”) who must design, conduct and communicate the results of stress testing of the banking sector (“the bank”). The receiver can be thought of as a typical investor (“the investor”) choosing among alternative investments whose payoff depends on the realization of the state ω . The regulator’s choice of communication strategy can be seen as coming from a combination of the announced specifications of the stress tests and the protocol for communicating the results.²⁰ Think of the ω as investment-relevant information about the health of the banking sector, with ω_1 and ω_2 associated with “bad” and “good” health, respectively. Actions a_1 and a_2 are socially-productive investments (i.e., productive from the viewpoint of the economy as a whole, a viewpoint that we assume the regulator adopts). Action a_3 is a socially-detrimental purely speculative investment. The regulator’s challenge is to design and communicate stress tests so as to better coordinate investment behavior with the health of the banking sector, without diverting investments to the speculative activity.

How can we think about an ambiguous communication strategy in this context? One realistic channel through which ambiguity could be introduced into communication arises

²⁰See [Bergemann and Morris \(2016\)](#) for another example of Bayesian persuasion in the context of stress testing.

from the fact²¹ that some banking regulators use as input to the stress tests “bottom-up” tests conducted by the banks based on their own in-house models and data. For example, the EU-wide stress tests conducted by the European Banking Authority have to-date been almost entirely based on such bottom-up inputs. By choosing in the announced specifications of the stress tests how much leeway to leave to the banks regarding the exact models/tests the banks will use, the regulator can manipulate the degree to which test results are subject to ambiguous interpretation. If they grant no discretion on this dimension, then, at least in principle, the interpretation of test results coming from a model with known stochastic properties is clear and unambiguous. If, in contrast, substantial leeway is granted – for example, by prescribing the model/test run by the bank to be contingent on a parameter value to be calculated based on data private to the bank – the proper interpretation of any given announced result may then vary substantially with the (unobserved at the time the specifications are announced and committed to) private information of the bank. This “model variation” as a function of the bank’s private information would thus generate uncertainty for parties outside the bank about the precise interpretation of any given bank test results.

How can we think about choosing μ in the banking context? Suppose the regulator requires the bank to carry out different tests depending on the bank’s exposure to a particular asset class and the calculation of the exposure parameter requires input of data that is privately known to the bank. Let the parameter space for this private information be the interval $[0, 1]$.²² Suppose the belief common to the regulator and investor about this parameter is a continuous probability measure on this interval. Then, by prescribing different partitions of $[0, 1]$ to serve as the contingencies under which the various models/tests will be run by the bank, the regulator may vary the μ associated with the stress test. For instance, the regulator may require the bank to undertake a more sensitive and detailed simulation exercise if the exposure parameter is above a particular threshold and a less sensitive and comprehensive test otherwise.²³

²¹See e.g., Table 1 in [Dent et al. \(2016\)](#).

²²This parameter by itself does not signal a particular state ω . It is only through complex interactions with many other exposures, the wider domestic and world economy, and imperfectly observed activities of the banking sector that these exposures may influence banking health. Only testing can reveal the particular strength or vulnerability.

²³For instance, in recent EU stress tests (see [European Banking Authority \(2023\)](#), Section 2.4.4.): “Banks with *significant* foreign currency exposure are required to take into account the altered creditworthiness of their respective obligors, given the FX development under the baseline and adverse scenarios. In particular, banks are only required to evaluate this impact if the exposures of certain asset classes in foreign currencies are above certain thresholds.”

SA.3 The insufficiency of binary ambiguous experiments

Proposition SA.3.1. *It is not always sufficient to consider only binary ambiguous experiments in searching for either a strict benefit from ambiguity or optimal ambiguous persuasion.*

We provide a detailed sketch of the proof. The full proof is available in [Cheng et al. \(2024\)](#).

Sketch of Proof of Proposition SA.3.1. The proof is by construction. We first show an example in which the only optimal ambiguous experiments are more than binary. A modification of this example is then used to provide an example in which the sender may strictly benefit from ambiguous communication even when no binary ambiguous experiment benefits the sender.

Example in which all optimal ambiguous experiments are more than binary.

Suppose $\phi_s(x) = x$ and $\phi_r(x) = \ln(x + 5)$. Let $\Omega = \{\omega_1, \omega_2\}$, with equal prior probabilities $p = (1/2, 1/2)$. There are five actions $\{a_1, a_2, b_1, b_2, b_3\}$ and the payoff matrix is

(u_s, u_r)	ω_1	ω_2
a_1	3, 3	0, 0
a_2	-1, -1	3, 3
b_1	0, 4	-1, -2
b_2	0, 2	1, 2
b_3	-2, -4	1, 4

The optimal Bayesian persuasion experiment is

$$\begin{aligned} \sigma_a(a_1|\omega_1) &= 4/5, & \sigma_a(a_2|\omega_1) &= 1/5; \\ \sigma_a(a_1|\omega_2) &= 2/5, & \sigma_a(a_2|\omega_2) &= 3/5. \end{aligned}$$

Notice that

$$u_s(\sigma_a, \tau^*) = u_r(\sigma_a, \tau^*) = 2.$$

Let $\sigma_{11}, \sigma_{12}, \sigma_{21}$ and σ_{22} denote the extreme experiments where σ_{ij} recommends a_i and a_j deterministically in states ω_1 and ω_2 , respectively. Notice that these extreme experiments

are all Pareto-ranked:

$$\begin{aligned} u_s(\sigma_{11}, \tau^*) &= u_r(\sigma_{11}, \tau^*) = 3/2; \\ u_s(\sigma_{12}, \tau^*) &= u_r(\sigma_{12}, \tau^*) = 3; \\ u_s(\sigma_{21}, \tau^*) &= u_r(\sigma_{21}, \tau^*) = -1/2; \\ u_s(\sigma_{22}, \tau^*) &= u_r(\sigma_{22}, \tau^*) = 1. \end{aligned}$$

Consider the following splitting of σ_a ,

$$\sigma_a = \frac{1}{5}\sigma_{11} + \frac{3}{5}\sigma_{12} + \frac{1}{5}\sigma_{21}.$$

It can be verified that for $\hat{\sigma} = (\sigma_{11}, \sigma_{12}, \sigma_{21})$ and $\hat{\mu}$ such that $\sum_{\theta} em_{\theta}^{(\hat{\sigma}, \hat{\mu})} \sigma_{\theta} = \sigma_a$, $(\hat{\sigma}, \hat{\mu})$ is an obedient ambiguous experiment yielding the sender a payoff of $159/70 = 2.27143$, strictly higher than Bayesian persuasion of 2. Therefore, any optimal ambiguous experiment must involve ambiguity and thus be at least binary.

As $\phi_s(x) = x$ and $\phi_r(x) = \ln(x + 5)$, by Proposition 2, in any optimal ambiguous experiment, there cannot exist any further Pareto-ranked splitting of any experiment in the collection.

Observe that σ_a is the only incentive-compatible experiment that never recommends any of the b actions. Furthermore, σ_a cannot be split into a convex combination of two extreme experiments. Thus, any binary splitting of σ_a must involve at least one non-extreme experiment. However, since all these extreme experiments are Pareto-ranked, there must exist a Pareto-ranked splitting of any such non-extreme experiment (into extreme experiments). Therefore, any binary ambiguous experiment constructed from splittings of σ_a cannot be optimal.

The proof goes on to show that an optimal ambiguous experiment in this example also cannot be a binary ambiguous experiment that is constructed from a splitting of any other incentive-compatible experiment (in particular, any recommending a b action with a positive probability).

Example in which ambiguous communication benefits the sender, but does not do so when restricted to binary ambiguous experiments

Suppose $\phi_s(x) = x$ and $\phi_r(x) = \ln(x + 5)$. Let $\Omega = \{\omega_1, \omega_2\}$ and the prior p be uniform. There are seven actions $\{a_1, a_2, b_1, b_2^+, b_2^-, b_3, c\}$. Let the payoff matrix be, for some $x > 2$,

(u_s, u_r)	ω_1	ω_2
a_1	3, 3	0, 0
a_2	-1, -1	3, 3
b_1	0, 4	0, -2
b_2^-	0, 5/2	0, 1
b_2^+	0, 5/4	0, 9/4
b_3	0, -4	0, 4
c	$x, 7/4$	$x, 7/4$

The only differences from the previous example are the addition of c and the replacement of b_2 by b_2^- and b_2^+ . Let σ_c denote the experiment that recommends action c deterministically in both states. Because $x > 2$, the optimal Bayesian persuasion experiment is σ_c , yielding the sender a payoff of x . The proof then shows the existence of $x > 2$ such that the sender's payoff from $(\hat{\sigma}, \hat{\mu})$ is strictly higher than x but the sender's payoff from any binary ambiguous experiment is lower than x . The replacement of b_2 by b_2^- and b_2^+ serves to make σ_c obedient only at the prior p , which helps simplify the calculations in the proof. \square

SA.4 Proofs of auxiliary results from “Persuasion with Ambiguous Communication”

SA.4.1 Proof of Theorem A.1

Proof of Theorem A.1. Fix any $u \in \mathbb{R}$. Let $(\sigma_\theta, \lambda_\theta)_{\theta \in \Theta}$ be feasible for the maximization problem $\Phi^*(u)$. Suppose that there exists a pair $(\sigma_\theta, \sigma_{\theta'})$ with $\lambda_\theta > 0$ and $\lambda_{\theta'} > 0$ and such that there exists a $\lambda \in (0, 1)$ for which, $\Phi_u(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}) > \lambda\Phi_u(\sigma_\theta) + (1-\lambda)\Phi_u(\sigma_{\theta'})$.

Then, $(\sigma_\theta, \lambda_\theta)_{\theta \in \Theta}$ cannot be a solution to the maximization problem $\Phi^*(u)$. This can be seen from the following construction of a strict improvement satisfying the constraints in that problem: If $\frac{\lambda_\theta}{\lambda} \leq \frac{\lambda_{\theta'}}{1-\lambda}$, then replacing σ_θ by the merged experiment $\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}$ and replacing λ_θ by $\hat{\lambda}_\theta = \frac{\lambda_\theta}{\lambda}$ and $\lambda_{\theta'}$ by $\hat{\lambda}_{\theta'} = \lambda_{\theta'} - (1-\lambda)\frac{\lambda_\theta}{\lambda}$ yields such an improvement. If instead $\frac{\lambda_\theta}{\lambda} > \frac{\lambda_{\theta'}}{1-\lambda}$, then replacing $\sigma_{\theta'}$ by the merged experiment $\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}$ and replacing $\lambda_{\theta'}$ by $\hat{\lambda}_{\theta'} = \frac{\lambda_{\theta'}}{1-\lambda}$ and λ_θ by $\hat{\lambda}_\theta = \lambda_\theta - \lambda\frac{\lambda_{\theta'}}{1-\lambda}$ is such an improvement.

To prove (i), towards a contradiction, suppose in the solution there exists $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_+(u) \times \Sigma_-(u)$ with $\phi'_r(u_r(\sigma_\theta, \tau^*)) \neq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$ and they are not Pareto-ranked, i.e., $u_s(\sigma_\theta, \tau^*) > u \geq u_s(\sigma_{\theta'}, \tau^*)$, and $u_r(\sigma_\theta, \tau^*) < u_r(\sigma_{\theta'}, \tau^*)$. Then there exists $\lambda \in (0, 1)$

such that $\phi'_r(u_r(\sigma_\theta, \tau^*)) > \phi'_r(u_r(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*))$ (by differentiability of ϕ_r), and $\Phi_u(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}) > \lambda\Phi_u(\sigma_\theta) + (1-\lambda)\Phi_u(\sigma_{\theta'})$. To see the last point, notice that

$$\begin{aligned}\Phi_u(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}) &= \frac{\phi_s(u_s(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*))} \\ &\geq \frac{\lambda[\phi_s(u_s(\sigma_\theta, \tau^*)) - \phi_s(u)] + (1-\lambda)[\phi_s(u_s(\sigma_{\theta'}, \tau^*)) - \phi_s(u)]}{\phi'_r(u_r(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*))} \\ &= \frac{\lambda\phi'_r(u_r(\sigma_\theta, \tau^*))}{\phi'_r(u_r(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*))} \Phi_u(\sigma_\theta) + \frac{(1-\lambda)\phi'_r(u_r(\sigma_{\theta'}, \tau^*))}{\phi'_r(u_r(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*))} \Phi_u(\sigma_{\theta'}) \\ &> \lambda\Phi_u(\sigma_\theta) + (1-\lambda)\Phi_u(\sigma_{\theta'}),\end{aligned}$$

where the first inequality follows from concavity of ϕ_s , and the second inequality follows from concavity of ϕ_r and $\Phi_u(\sigma_\theta) > 0$ and $\Phi_u(\sigma_{\theta'}) \leq 0$, and $\phi'_r(u_r(\sigma_\theta, \tau^*)) > \phi'_r(u_r(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*))$.

To prove (ii), towards a contradiction, suppose in the solution there exists $(\sigma_\theta, \sigma_{\theta'}) \in \Sigma_+(u) \times \Sigma_+(u)$ with $u_s(\sigma_\theta, \tau^*) \neq u_s(\sigma_{\theta'}, \tau^*)$ and $\phi'_r(u_r(\sigma_\theta, \tau^*)) \neq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$, but they are not Pareto-ranked, i.e., $u_s(\sigma_\theta, \tau^*) > u_s(\sigma_{\theta'}, \tau^*) > u$, and $u_r(\sigma_\theta, \tau^*) < u_r(\sigma_{\theta'}, \tau^*)$. Then for any $\lambda \in (0, 1)$, we can show $\Phi_u(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}) > \lambda\Phi_u(\sigma_\theta) + (1-\lambda)\Phi_u(\sigma_{\theta'})$. To see this, observe that

$$\begin{aligned}\Phi_u(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}) &= \frac{\phi_s(u_s(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*))} \\ &\geq \lambda \frac{\phi_s(u_s(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\sigma_\theta, \tau^*))} + (1-\lambda) \frac{\phi_s(u_s(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\sigma_{\theta'}, \tau^*))} \\ &\geq \lambda \frac{\lambda\phi_s(u_s(\sigma_\theta, \tau^*)) + (1-\lambda)\phi_s(u_s(\sigma_{\theta'}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\sigma_\theta, \tau^*))} \\ &\quad + (1-\lambda) \frac{\lambda\phi_s(u_s(\sigma_\theta, \tau^*)) + (1-\lambda)\phi_s(u_s(\sigma_{\theta'}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\sigma_{\theta'}, \tau^*))} \tag{SA.4.1} \\ &= \lambda\Phi_u(\sigma_\theta) + \lambda(1-\lambda) \frac{\phi_s(u_s(\sigma_{\theta'}, \tau^*)) - \phi_s(u_s(\sigma_\theta, \tau^*))}{\phi'_r(u_r(\sigma_\theta, \tau^*))} \\ &\quad + (1-\lambda)\Phi_u(\sigma_{\theta'}) + (1-\lambda)\lambda \frac{\phi_s(u_s(\sigma_\theta, \tau^*)) - \phi_s(u_s(\sigma_{\theta'}, \tau^*))}{\phi'_r(u_r(\sigma_{\theta'}, \tau^*))} \\ &> \lambda\Phi_u(\sigma_\theta) + (1-\lambda)\Phi_u(\sigma_{\theta'}),\end{aligned}$$

where the first inequality follows from concavity of $1/\phi'_r$ and positivity of $\phi_s(u_s(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*)) - \phi_s(u)$, the second inequality follows from concavity of ϕ_s , and the last strict inequality follows from $u_r(\sigma_\theta, \tau^*) < u_r(\sigma_{\theta'}, \tau^*)$ and $\phi'_r(u_r(\sigma_\theta, \tau^*)) \neq \phi'_r(u_r(\sigma_{\theta'}, \tau^*))$.

The proof of (iii) is the same as the proof of (ii), except that the first inequality in the chain (SA.4.1) now follows from convexity of $1/\phi'_r$ and $\phi_s(u_s(\lambda\sigma_\theta + (1-\lambda)\sigma_{\theta'}, \tau^*)) - \phi_s(u) < 0$. \square

SA.4.2 Proof of Theorem A.2

Proof of Theorem A.2. Fix any $u \in \mathbb{R}$. Let $(\sigma_\theta, \lambda_\theta)_{\theta \in \Theta}$ be feasible for $(\Phi^*(u))$. Suppose that there exist σ_θ satisfying $\lambda_\theta > 0$ and two experiments σ and σ' such that $\sigma_\theta = \lambda\sigma + (1-\lambda)\sigma'$ for some $\lambda \in (0, 1)$ and $\Phi_u(\lambda\sigma + (1-\lambda)\sigma') < \lambda\Phi_u(\sigma) + (1-\lambda)\Phi_u(\sigma')$, then $(\sigma_\theta, \lambda_\theta)_{\theta \in \Theta}$ cannot be a solution to $(\Phi^*(u))$. This follows by noting that splitting σ_θ into σ with probability $\lambda\lambda_\theta$ and σ' with probability $(1-\lambda)\lambda_\theta$ induces a strict improvement.

To prove (i), towards a contradiction, suppose there exists such a Pareto-ranked splitting $(\bar{\sigma}, \underline{\sigma}, \lambda)$ with $\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma} \in \Sigma_+(u)$, then we have

$$\begin{aligned}
& \lambda\Phi_u(\bar{\sigma}) + (1-\lambda)\Phi_u(\underline{\sigma}) - \Phi_u(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}) \\
&= \lambda \left(\frac{\phi_s(u_s(\bar{\sigma}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\bar{\sigma}, \tau^*))} - \frac{\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))} \right) \\
&\quad + (1-\lambda) \left(\frac{\phi_s(u_s(\underline{\sigma}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\underline{\sigma}, \tau^*))} - \frac{\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))} \right) \\
&\geq \frac{\lambda}{\phi'_r(u_r(\bar{\sigma}, \tau^*))} (\phi_s(u_s(\bar{\sigma}, \tau^*)) - \phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))) \\
&\quad + \frac{1-\lambda}{\phi'_r(u_r(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))} (\phi_s(u_s(\underline{\sigma}, \tau^*)) - \phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))) \\
&\geq \frac{\lambda}{\phi'_r(u_r(\bar{\sigma}, \tau^*))} \phi'_s(u_s(\bar{\sigma}, \tau^*)) (u_s(\bar{\sigma}, \tau^*) - \lambda u_s(\bar{\sigma}, \tau^*) - (1-\lambda)u_s(\underline{\sigma}, \tau^*)) \\
&\quad + \frac{1-\lambda}{\phi'_r(u_r(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))} \phi'_s(u_s(\underline{\sigma}, \tau^*)) (u_s(\underline{\sigma}, \tau^*) - \lambda u_s(\bar{\sigma}, \tau^*) - (1-\lambda)u_s(\underline{\sigma}, \tau^*)) \\
&= \lambda(1-\lambda)(u_s(\bar{\sigma}, \tau^*) - u_s(\underline{\sigma}, \tau^*)) \left(\frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_r(u_r(\bar{\sigma}, \tau^*))} - \frac{\phi'_s(u_s(\underline{\sigma}, \tau^*))}{\phi'_r(u_r(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))} \right) > 0,
\end{aligned}$$

where the first inequality follows from $u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}) > u \geq u_s(\underline{\sigma}, \tau^*)$, $\phi'_r(u_r(\bar{\sigma}, \tau^*)) \leq \phi'_r(u_r(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) \leq \phi'_r(u_r(\underline{\sigma}, \tau^*))$, the second inequality follows from concavity of ϕ_s , and the last inequality follows from the supposition. The other case can be shown similarly.

To show (ii), if $1/\phi'_r(\cdot)$ is concave, towards a contradiction, we have

$$\begin{aligned}
\Phi_u(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}) &= \frac{\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*))} \\
&\leq \lambda \frac{\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\bar{\sigma}, \tau^*))} + (1-\lambda) \frac{\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - \phi_s(u)}{\phi'_r(u_r(\underline{\sigma}, \tau^*))} \\
&= \lambda \Phi_u(\bar{\sigma}) + \lambda \frac{\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - \phi_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_r(u_r(\bar{\sigma}, \tau^*))} \\
&\quad + (1-\lambda) \Phi_u(\underline{\sigma}) + (1-\lambda) \frac{\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) - \phi_s(u_s(\underline{\sigma}, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))} \\
&\leq \lambda \Phi_u(\bar{\sigma}) + \lambda(1-\lambda) \frac{\phi'_s(u_s(\bar{\sigma}, \tau^*))}{\phi'_r(u_r(\bar{\sigma}, \tau^*))} (u_s(\underline{\sigma}, \tau^*) - u_s(\bar{\sigma}, \tau^*)) \\
&\quad + (1-\lambda) \Phi_u(\underline{\sigma}) + \lambda(1-\lambda) \frac{\phi'_s(u_s(\underline{\sigma}, \tau^*))}{\phi'_r(u_r(\underline{\sigma}, \tau^*))} (u_s(\bar{\sigma}, \tau^*) - u_s(\underline{\sigma}, \tau^*)) \\
&< \lambda \Phi_u(\bar{\sigma}) + (1-\lambda) \Phi_u(\underline{\sigma}),
\end{aligned}$$

where the first inequality follows from concavity of $1/\phi'_r$ and $\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) \leq \phi_s(u)$, the second inequality from concavity of ϕ_s and the third inequality from the supposition.

The proof of (iii) is the same as the proof of (ii), except that the first inequality now follows from convexity of $1/\phi'_r$ and $\phi_s(u_s(\lambda\bar{\sigma} + (1-\lambda)\underline{\sigma}, \tau^*)) > \phi_s(u)$. \square

SA.4.3 Proof of Lemma A.3

Proof of Lemma A.3. Define, for all integers $k \in [1, n]$,

$$\mathcal{S}_k = \sum_{i=1}^k x_i \mu_i \left[\sum_{j \neq i; j=1}^k \mu_j (y_i - y_j) \right].$$

Notice that when $k = n$, we have

$$\begin{aligned}
\mathcal{S}_n &= \sum_{i=1}^n x_i \mu_i \left[\sum_{j=1; j \neq i}^n \mu_j (y_i - y_j) \right] \\
&= \sum_{i=1}^n x_i \mu_i \left[(1 - \mu_i) y_i - \sum_{j=1; j \neq i}^n \mu_j y_j \right]
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^n x_i \mu_i \left[y_i - \sum_{j=1}^n \mu_j y_j \right] \\
&= \left(\sum_{j=1}^n \mu_j y_j \right) \left(\sum_{i=1}^n x_i \frac{\mu_i y_i}{\sum_{j=1}^n \mu_j y_j} - \sum_{i=1}^n x_i \mu_i \right).
\end{aligned}$$

Since $\left(\sum_{j=1}^n \mu_j y_j \right) > 0$, it suffices to show $\mathcal{S}_n < 0$. We prove this by induction. Observe that $\mathcal{S}_1 = 0$. For $k \geq 1$,

$$\begin{aligned}
\mathcal{S}_{k+1} &= \sum_{i=1}^{k+1} x_i \mu_i \left[\sum_{j=1; j \neq i}^{k+1} \mu_j (y_i - y_j) \right] \\
&= \sum_{i=1}^k x_i \mu_i \left[\sum_{j \neq i; j=1}^k \mu_j (y_i - y_j) \right] + \sum_{i=1}^k x_i \mu_i [\mu_{k+1} (y_i - y_{k+1})] + \\
&\quad x_{k+1} \mu_{k+1} \sum_{j=1}^k [\mu_j (y_{k+1} - y_j)] \\
&= \sum_{i=1}^k x_i \mu_i \left[\sum_{j \neq i; j=1}^k \mu_j (y_i - y_j) \right] + \sum_{i=1}^k \mu_i \mu_{k+1} (x_i - x_{k+1}) (y_i - y_{k+1}) \\
&= \mathcal{S}_k + \underbrace{\sum_{i=1}^k \mu_i \mu_{k+1} (x_i - x_{k+1}) (y_i - y_{k+1})}_{\leq 0}.
\end{aligned}$$

For $k = j^* - 1$,

$$\sum_{i=1}^{j^*-1} \mu_i \mu_{j^*} (x_i - x_{j^*}) (y_i - y_{j^*}) \leq \mu_{i^*} \mu_{j^*} (x_{i^*} - x_{j^*}) (y_{i^*} - y_{j^*}) < 0.$$

Therefore, $0 = \mathcal{S}_1 > \mathcal{S}_{j^*} \geq \mathcal{S}_n$. □

SA.4.4 Proof of Lemma A.4

Proof of Lemma A.4. Define

$$\sigma^\lambda = \frac{1}{1-\lambda} \sigma - \frac{\lambda}{1-\lambda} \hat{\sigma} \tag{SA.4.2}$$

where $\lambda \in (0, 1)$. Observe that if $\underline{\sigma}^\lambda$ is a well-defined experiment, then $\lambda\hat{\sigma} + (1-\lambda)\underline{\sigma}^\lambda = \sigma$, and $u_s(\underline{\sigma}^\lambda, \tau^*) < u_s(\sigma, \tau^*)$, $u_r(\underline{\sigma}^\lambda, \tau^*) < u_r(\sigma, \tau^*)$, so that $(\hat{\sigma}, \underline{\sigma}^\lambda, \lambda)$ is a Pareto-ranked splitting of σ .

It remains to show that there exists $\lambda \in (0, 1)$ such that $\underline{\sigma}^\lambda$ is indeed an experiment. In other words, for each ω , $\underline{\sigma}^\lambda(\cdot|\omega)$ must be a probability distribution over actions.

If $|\text{supp}(\sigma(\cdot|\omega))| = 1$, then $\text{supp}(\hat{\sigma}(\cdot|\omega)) \subseteq \text{supp}(\sigma(\cdot|\omega))$ implies $\text{supp}(\hat{\sigma}(\cdot|\omega)) = \text{supp}(\sigma(\cdot|\omega))$. It follows that $\underline{\sigma}^\lambda(\cdot|\omega) = \sigma(\cdot|\omega)$ for all $\lambda \in (0, 1)$, and is thus a distribution over actions.

If $|\text{supp}(\sigma(\cdot|\omega))| > 1$, embed $\sigma(\cdot|\omega)$ into the Euclidean space $\mathbb{R}^{|\text{supp}(\sigma(\cdot|\omega))|}$ and notice that $\sigma(\cdot|\omega)$ is in the relative interior of the probability simplex $\Delta(\text{supp}(\sigma(\cdot|\omega)))$. Thus there exists $\epsilon_\omega > 0$ such that for all $x \in \mathbb{R}^{|\text{supp}(\sigma(\cdot|\omega))|}$ with $\sum_i x_i = 1$, if $\|x - \sigma(\cdot|\omega)\| < \epsilon_\omega$, then $x \in \Delta(\text{supp}(\sigma(\cdot|\omega)))$. Since $\text{supp}(\hat{\sigma}(\cdot|\omega)) \subseteq \text{supp}(\sigma(\cdot|\omega))$, one has $\hat{\sigma}(\cdot|\omega) \in \Delta(\text{supp}(\sigma(\cdot|\omega)))$ as well. Then, for all $\lambda \in (0, 1)$, $\underline{\sigma}^\lambda(\cdot|\omega) \in \mathbb{R}^{|\text{supp}(\sigma(\cdot|\omega))|}$ and $\sum_a \underline{\sigma}^\lambda(a|\omega) = 1$. Moreover, there exists $\lambda_\omega > 0$ such that for all $\lambda \in (0, \lambda_\omega)$, $\|\underline{\sigma}^\lambda(\cdot|\omega) - \sigma(\cdot|\omega)\| < \epsilon_\omega$, and thus $\underline{\sigma}^\lambda(\cdot|\omega) \in \Delta(\text{supp}(\sigma(\cdot|\omega)))$, making it a distribution over actions.

Because Ω is finite, $\underline{\lambda}(\hat{\sigma}, \sigma) \equiv \min_{\omega: |\text{supp}(\sigma(\cdot|\omega))| > 1} \lambda_\omega > 0$. Therefore, for all $\lambda \in (0, \underline{\lambda}(\hat{\sigma}, \sigma))$, $\underline{\sigma}^\lambda$ is a well-defined experiment. \square

SA.4.5 Proofs of Proposition 3 and Corollary 4

Proof of Proposition 3. Throughout, we view a splitting $(\lambda_\theta, \sigma_\theta)_{\theta \in \Theta}$ of some $\sigma \in \Sigma^*$ as a finitely supported distribution in $\Delta(\Sigma)$. With a slight abuse of notation, we write λ for the distribution, $\Delta_{\text{simple}}(\Sigma)$ for the set of finitely supported distributions on Σ , and \mathbb{E}_λ for the expectation operator with respect to λ . Since ambiguous communication benefits the sender, $\{\lambda \in \Delta_{\text{simple}}(\Sigma) : \mathbb{E}_\lambda[\sigma] \in \Sigma^*, \mathbb{E}_\lambda[\Phi_{u_s^{BP}}(\sigma)] \in (0, \infty)\} \neq \emptyset$ by Corollary 2. Since Σ^* is a convex set and has a non-empty interior, any point in Σ^* can be approached by points in the interior of Σ^* . As the expectations in the above set are continuous in λ , this implies that $\{\lambda \in \Delta_{\text{simple}}(\Sigma) : \mathbb{E}_\lambda[\sigma] \in \text{int } \Sigma^*, \mathbb{E}_\lambda[\Phi_{u_s^{BP}}(\sigma)] \in (0, \infty)\} \neq \emptyset$. Furthermore, since $\text{int } \Sigma^* \times (0, \infty)$ is open in the natural product topology, this set is open. \square

Proof of Corollary 4. By Proposition 3, there exists a non-empty open set of ambiguous experiments that benefit the sender. Fix one. Since $\mathbb{E}_\lambda[\Phi_{u_s^{BP}}(\sigma)]$ is continuous in ϕ'_r , this experiment continues to benefit the sender under small perturbations of ϕ'_r . Finally, since ϕ_r is concave, small perturbations of ϕ_r imply small perturbations of ϕ'_r (Rockafellar, 1970, Theorem 25.7, p. 248). \square

SA.4.6 Proof of Theorem 6

Proof of Theorem 6. That (σ, μ) benefits the sender means that

$$\phi_s^{-1}(\mu\phi_s(u_s(\bar{\sigma}, \tau^*)) + (1 - \mu)\phi_s(u_s(\underline{\sigma}, \tau^*))) > u_s^{BP}. \quad (\text{SA.4.3})$$

Any less ambiguity averse sender will have $\tilde{\phi}_s$ weakly less concave than ϕ_s , weakly increasing the left-hand side of (SA.4.3), while leaving the right-hand side unchanged. This proves (i).

Let $\sigma^* = em_{\theta_1}^{(\sigma, \mu)}\bar{\sigma} + (1 - em_{\theta_1}^{(\sigma, \mu)})\underline{\sigma}$. By Lemma 1, $\tau^* \in br(\sigma^*)$. In light of (SA.4.3) and Theorem 3, $\bar{\sigma}$ and $\underline{\sigma}$ must be Pareto-ranked. Without loss of generality, assume that $\bar{\sigma}$ is the better one. A weakly more ambiguity averse receiver will have a $\tilde{\phi}_r = \varphi \circ \phi_r$ for some increasing, differentiable and concave φ , resulting in an effective measure $e\tilde{m}^{(\sigma, \mu)}$ such that $e\tilde{m}_{\theta_1}^{(\sigma, \mu)} \leq em_{\theta_1}^{(\sigma, \mu)}$.

If, as in (iii), $\tau^* \in br(\underline{\sigma})$, then since $\tau^* \in br(\sigma^*) = br(em_{\theta_1}^{(\sigma, \mu)}\bar{\sigma} + (1 - em_{\theta_1}^{(\sigma, \mu)})\underline{\sigma})$, $\tau^* \in br(e\tilde{m}_{\theta_1}^{(\sigma, \mu)}\bar{\sigma} + (1 - e\tilde{m}_{\theta_1}^{(\sigma, \mu)})\underline{\sigma})$. By Lemma 1, this implies $\tau^* \in BR(\sigma, \mu)$ for a receiver with any such $\tilde{\phi}_r$, proving that σ together with μ continues to benefit the sender (and, by the same argument as for (i), any less ambiguity averse senders as well). This proves (iii).

Finally, if $\tau^* \notin br(\underline{\sigma})$, define $\tilde{\mu}$ by

$$\tilde{\mu}(\theta_1) = \frac{em_{\theta_1}^{(\sigma, \mu)}\tilde{\phi}'_r(u_r(\underline{\sigma}, \tau^*))}{em_{\theta_1}^{(\sigma, \mu)}\tilde{\phi}'_r(u_r(\underline{\sigma}, \tau^*)) + (1 - em_{\theta_1}^{(\sigma, \mu)})\tilde{\phi}'_r(u_r(\bar{\sigma}, \tau^*))}, \text{ and } \tilde{\mu}(\theta_2) = 1 - \tilde{\mu}(\theta_1),$$

so that $e\tilde{m}^{(\sigma, \tilde{\mu})} = em^{(\sigma, \mu)}$. Lemma 1 and $\tau^* \in br(\sigma^*)$ then implies $\tau^* \in BR(\sigma, \tilde{\mu})$ for a receiver with $\tilde{\phi}_r$. Since $(u_s(\bar{\sigma}, \tau^*) > u_s(\underline{\sigma}, \tau^*))$, to show that $(\sigma, \tilde{\mu})$ benefits the sender (and all less ambiguity averse senders) it suffices to show that $\tilde{\mu}_{\theta_1} \geq \mu$. Indeed,

$$\begin{aligned} \tilde{\mu}_{\theta_1} &= \frac{em_{\theta_1}^{(\sigma, \mu)}\phi'_r(u_r(\underline{\sigma}, \tau^*))}{em_{\theta_1}^{(\sigma, \mu)}\phi'_r(u_r(\underline{\sigma}, \tau^*)) + (1 - em_{\theta_1}^{(\sigma, \mu)})\frac{\varphi'(\phi_r(u_r(\bar{\sigma}, \tau^*)))}{\varphi'(\phi_r(u_r(\underline{\sigma}, \tau^*)))}\phi'_r(u_r(\bar{\sigma}, \tau^*))} \\ &\geq \frac{em_{\theta_1}^{(\sigma, \mu)}\phi'_r(u_r(\underline{\sigma}, \tau^*))}{em_{\theta_1}^{(\sigma, \mu)}\phi'_r(u_r(\underline{\sigma}, \tau^*)) + (1 - em_{\theta_1}^{(\sigma, \mu)})\phi'_r(u_r(\bar{\sigma}, \tau^*))} = \mu, \end{aligned}$$

where the inequality follows from $\varphi'(\phi_r(u_r(\underline{\sigma}, \tau^*))) \geq \varphi'(\phi_r(u_r(\bar{\sigma}, \tau^*)))$. \square

References

- Bergemann, Dirk and Stephen Morris (2016) “Information design, Bayesian persuasion and Bayes correlated equilibrium,” *American Economic Review: Papers & Proceedings*, 106 (5), 586–591.
- Cheng, Xiaoyu (2023) “Ambiguous persuasion: An ex-ante formulation,” *arXiv preprint arXiv:2010.05376*.
- Cheng, Xiaoyu, Peter Klibanoff, Sujoy Mukerji, and Ludovic Renou (2024) “Non-binary ambiguous experiments may improve upon Bayesian persuasion when binary ambiguous experiments do not,” <https://www.dropbox.com/scl/fi/baq9nposf7hhow63b27es/Non-binary-improvement.pdf?rlkey=zzqlr6l6pu5epsx0a89rjrf7f&dl=0>.
- Dent, Kieran, Ben Westwood, and Miguel Segoviano Basurto (2016) “Stress testing of banks: an introduction,” *Bank of England Quarterly Bulletin*, Q3.
- European Banking Authority (2023) “Methodological Note,” 20.
- Kamenica, Emir and Matthew Gentzkow (2011) “Bayesian persuasion,” *American Economic Review*, 101 (6), 2590–2615.
- Rockafellar, Ralph Tyrell (1970) *Convex Analysis*, Princeton: Princeton University Press.