Persuasion under Ambiguity

Jonas Hedlund*, T. Florian Kauffeldt†, and Malte Lammert††

First version: July 25, 2019
This version: January 23, 2020

Abstract

This paper introduces a receiver who perceives ambiguity in a binary model of Bayesian persuasion. The sender has a well-defined prior, while the receiver considers an interval of priors and maximizes a convex combination of worst and best expected payoffs (α-maxmin preferences). We characterize the sender's optimal signal and find that the receiver's payoff differences across states given each action (sensitivities), play a fundamental role in the characterization and the comparative statics: If the sender's preferred action is the least (most) sensitive one, then the sender's equilibrium payoff, as well as the sender's preferred degree of receiver ambiguity, is increasing (decreasing) in the receiver's pessimism. We document a tendency for ambiguity-sensitive receivers to be more difficult to persuade.

Keywords: Persuasion, Ambiguity, Signaling, Information transmission

1. Introduction

The recent literature on Bayesian persuasion, pioneered by Kamenica and Gentzkow (2011) (henceforth, K&G), is concerned with situations where a

*Corresponding author
Department of Economics, University of Texas at Dallas, USA.
Email: jonas.hedlund@utdallas.edu
†Department of Economics, Heilbronn University, Germany.
Email: mail@florian-kauffeldt.com
††Department of Economics, University of California, San Diego, USA.
Email: mlammert@ucsd.edu

sender (he) controls the beliefs of a *receiver* (she) by conducting an informative experiment. The bulk of this literature assumes that sender and receiver update beliefs given a well-defined common prior about the payoff-relevant state. The present paper departs from this tradition by studying a persuasion model in which the sender has a well-defined prior on the payoff-relevant state, but the receiver is uncertain about the prior distribution (she perceives *ambiguity*) in the sense of Ellsberg (1961).[1]

Several situations of information control are likely to feature such an asymmetry between the sender's and the receiver's prior beliefs. For example, suppose a doctor (sender) conducts a test to convince a patient (receiver) that she is a suitable candidate for a certain treatment. The prior probability of the patient's suitability may depend on age, symptoms, and medical history, and may be unclear from the point of view of the patient. The doctor, however, is an expert and deals with similar situations on a daily basis, and is more likely to have a clear prior about the patient's suitability. Similarly, a car salesman who controls the conditions under which a customer can test-drive a used car, seems more likely than the customer to have a clear prior about the possibility that the car is a lemon. Can the doctor and the car salesman exploit the receiver's ambiguity to their advantage? Is it easier to persuade a receiver that responds more optimistically to her uncertainty about the prior distribution? Our analysis sheds light on these and related questions.

We examine the effects of receiver ambiguity in a persuasion framework with binary payoff-relevant states and binary receiver actions. The sender's objective is to persuade the receiver to take the "high" action. The sender has a well-defined prior on the payoff-relevant state, whereas the receiver perceives ambiguity and has too little information to form a unique prior. Instead, she considers an interval of priors possible. The receiver observes an outcome of a signal chosen by the sender, updates her prior set pointwise according to Bayes' rule[2], and chooses an action that maximizes her objective function. The receiver has α-maxmin preferences and, therefore, maximizes a convex combination of her worst and best expected payoff with respect

[1] Ellsberg (1961) was the first to point out that when individuals perceive uncertainty about probabilities (*ambiguity*), they frequently do not behave as if they were governed by a well-defined prior. The literature refers to uncertainty about probabilities also as "Knightian Uncertainty," in reference to Knight (1921).

[2] This updating rule is known as "Full Bayesian updating," see Pires (2002).

to the priors in her prior set.[3] The convex combination parameter, α, is the weight that the receiver puts on the worst expected payoff. We may therefore interpret α as a measure of receiver's ambiguity attitude, ranging from full optimism ($\alpha = 0$) to full pessimism ($\alpha = 1$). We take the size of the receiver's prior set as a measure of the magnitude of the receiver's ambiguity and refer to this as the receiver's *perceived ambiguity* (see Ghirardato and Marinacci, 2002).

The main objective of our analysis is to understand the effect of the receiver's ambiguity attitude and perceived ambiguity on the sender's optimal strategy of persuasion and equilibrium expected payoff.

Our first observation is that there are similarities and differences between the fundamental logic of optimal persuasion of a receiver who perceives ambiguity, and the benchmark where the sender and the receiver share a common prior (analyzed by K&G). As in the common prior benchmark, an ambiguity perceiving receiver is persuaded by pushing the prior set up until she becomes indifferent between both actions.[4] The characterization of the receiver's indifference, however, is more complex, and different, when the receiver perceives ambiguity.[5]

Our main results reveal that the effect of the receiver's ambiguity attitude and perceived ambiguity on the sender's optimal strategy of persuasion, is organized by the *sensitivities* of the receiver's actions, i.e., the receiver's payoff differences across states given each action. In particular, the qualitative effect of the receiver's ambiguity attitude on the equilibrium signal, depends only on the sensitivities of the receiver's actions. The sender finds it easier to persuade a more pessimistic receiver if and only if the sender's preferred receiver action is the least sensitive one, i.e., if that action has the smallest payoff difference across states. Intuitively, when the sender's preferred receiver action is the least sensitive one, the receiver perceives less uncertainty

[3] See Hurwicz (1951), Jaffray (1989) and Ghirardato et al. (2004). A prominent special case of α-maxmin preferences are maxmin-preferences, which are axiomatized in Gilboa and Schmeidler (1989).

[4] Another similarity is that the optimal signal under receiver ambiguity collapses the receiver's prior set such that she is certain (i.e., she perceives no ambiguity) about the state, whenever she chooses the low action.

[5] In the common prior case, the optimal signal is characterized by the receiver's cost of type I and type II errors, i.e., the receiver's payoff differences across actions given state. In contrast, in our model, with everything else fixed, the optimal signal is characterized by the cost of type II errors and the receiver's payoff differences across states given action.

given that action. As the receiver becomes more pessimistic, she tends to view the small uncertainty more favorably. This makes it easier to persuade the receiver to take the sender's preferred action.

The comparative statics with respect to the receiver's perceived ambiguity are remarkably rich, given the simplicity of our framework. We evaluate the effect of increased perceived ambiguity by uniformly expanding the receiver's prior set around an arbitrary midpoint. We find that the sender's equilibrium payoff can be decreasing, increasing, or non-monotonic (with a unique maximum), in the receiver's perceived ambiguity. This means that the sender's preferred degree of receiver ambiguity is either zero, maximal, or there is an interior ambiguity bliss point. The distinction between these cases is organized by the receiver's ambiguity attitude α and the sensitivities of her actions. Specifically, suppose the sender's preferred receiver action is the least sensitive one. The sender's preferred degree of receiver ambiguity is then zero if $\alpha \leq 1/2$, continuously increasing in α for intermediate values $\alpha \in (1/2, \hat{\alpha}]$, and maximal for $\alpha > \hat{\alpha}$. That is, the sender's preferred degree of receiver ambiguity is (i) increasing in the receiver's pessimism and (ii) tilted towards a preference for zero receiver ambiguity.

This result can, roughly, be understood in terms of two main effects. First, the smaller payoff uncertainty of the sender's preferred action is more appealing to a more pessimistic receiver. This tendency magnifies as the receiver perceives more ambiguity, i.e., more uncertainty about the relevant probability distribution. This effect is behind the fact that the sender's preferred degree of receiver ambiguity is increasing in the receiver's pessimism (when the sender's preferred receiver action is the least sensitive one). Second, because of a concavity property of Bayes' rule in the prior, it becomes inherently more difficult to persuade the receiver as her prior set expands. The concavity property implies that as the receiver's prior set expands uniformly, the posterior set expands faster at the lower bound than at the upper bound. This makes it more difficult to persuade the receiver to take the high action. This effect is responsible for the tilt of the sender's preferences towards zero receiver ambiguity.

The tilt of the sender's preferences is a robust feature of our model, which arises inevitably as the sender needs to push the receiver's posterior set away from the prior set, to induce the desired action. Consequently, when, in contrast to the case discussed above, the sender's preferred receiver action is the most sensitive one, the sender's preferred degree of receiver ambiguity is decreasing in the receiver's pessimism, but again tilted towards a preference for zero receiver ambiguity.

Finally, an intuitive way of summarizing the comparative statics with respect to the receiver's ambiguity attitude and perceived ambiguity is the following: If the sender's preferred receiver action is the least (most) sensitive one, then the sender's equilibrium payoff as well as the sender's preferred degree of receiver ambiguity are both increasing (decreasing) in the receiver's pessimism, with a tilt in the sender's preferred degree of receiver ambiguity towards zero ambiguity.

The paper is organized as follows. We first review the related literature. Section 2 describes our model in detail. In section 3, we present our results. Finally, Section 4 concludes with a summary. All proofs are in the appendix.

Related literature

Our paper contributes to the literature on Bayesian persuasion, initiated by K&G. Since their pioneering contribution, the Bayesian persuasion model has been extended in many directions, including sender private information (Alonso and Camara, 2018; Hedlund, 2017; Perez-Richet, 2014), receiver private information (Kolotilin, 2018; Kolotilin et al., 2017), multiple senders (Gentzkow and Kamenica 2016 and 2017; Li and Norman, 2018) and multiple receivers (Alonso and Camara, 2016).

A few recent papers (mostly working papers so far) within this literature deal with ambiguity in a K&G persuasion setup. Beauchene et al. (2019) assume that the sender can introduce ambiguity by using an ambiguous communication device, which generates ambiguity for both the sender and the receiver.[6] Hence, in contrast to our model in which ambiguity is exogenously present in the prior distribution, ambiguity in their model is endogenously introduced by the sender through the signal. Sender and receiver are maxmin expected utility maximizers. Beauchene et al. show that it can be beneficial for the sender to use an ambiguous device, even though this generates ambiguity for him.

Hu and Weng (2018) analyze a model in which the receiver is privately informed and the sender perceives ambiguity about the source of the re-

[6]Riedel and Sass (2014) introduce the class of Ellsberg games where players may choose ambiguous strategies (sets of probability distributions over their pure strategies) in addition to classical mixed strategies (single probability distributions over pure strategies). The ambiguous communication device in Beauchene et al. can be considered as such an ambiguous strategy and their setup as an "Ellsberg communication game."

ceiver's private information. As in Beauchene et al., the sender is a maxmin expected utility maximizer. Hu and Weng show that the sender fully discloses the payoff-relevant state whenever he perceives full ambiguity about the receiver's private information.

Kosterina (2018) introduces a model in which the sender perceives ambiguity about the receiver's prior. Again, it is implicitly assumed that the sender has maxmin expected utility preferences. Kosterina characterizes the sender's optimal signal in this setting.

Finally, the paper closest to ours is Laclau and Renou (2016). In Laclau and Renou's model, the sender's payoff depends on a set of receiver beliefs. Laclau and Renou point out that their setup applies both to persuasion of multiple receivers, and to situations where a single receiver perceives ambiguity about the state of the world – as in our model. Whereas our models are related, however, our points of departure are different. Laclau and Renou focus on extending K&G's concavification technique to situations where the sender wishes to simultaneously control multiple beliefs. Their main result result is a characterization of the sender's maximal payoff along these lines. In contrast, our main objective is to conduct comparative statics exercises within the specific context of an ambiguity-sensitive receiver, in order to understand the concrete effects of ambiguity on optimal persuasion.

It is worth emphasizing that the papers reviewed above assume maxmin expected utility preferences. That is, they consider agents who are extremely averse to ambiguity.[7] To the best of our knowledge, our model is the first within this literature to explicitly examine ambiguity-loving behavior. The α-maxmin approach allows us to examine the impact of both different levels of perceived ambiguity and of different ambiguity attitudes on agents' optimal behavior.[8]

[7] Sometimes, it is argued that the size of the prior set can be varied to incorporate different ambiguity attitudes. For instance, a "small" prior set would represent more optimistic responses to ambiguity. In this case, maxmin expected utility preferences would not represent extreme ambiguity aversion. As our analysis shows, however, in a persuasion context important differences between the α-maxmin and maxmin model emerge in the updating stage of beliefs. This provides a justification for our emphasis on the α-maxmin model.

[8] The α-maxmin model explicitly distinguishes between ambiguity and ambiguity attitude by the ambiguity attitude parameter α. Ghirardato et al. (2004) attempted to axiomatize this model. However, as shown by Eichberger et al. (2011), their axiomatization has some flaws if the state space is finite.

Our paper also relates to the small literature on cheap talk communication in the presence of ambiguity. The most related papers in this literature are by Kellner and Le Quement (2017 and 2018). Kellner and Le Quement (2017) consider a simple cheap talk game in which both sender and receiver perceive ambiguity about the state. Kellner and Le Quement (2018) consider a sender facing an ambiguity averse receiver. The sender is able to choose an ambiguous strategy a la Riedel and Sass (2014), see footnote 5.

2. Setup

There are two players, the *sender* and the *receiver*, and two payoff-relevant states, ω_L and ω_H. The sender has a prior $\mu_0 \in (0,1)$ that the state is ω_H. The receiver has too little information to form a unique prior, and instead has a prior set, each element of which is a belief that the state is ω_H. We assume that the receiver's prior set is a closed interval contained in $(0,1)$, with arbitrary midpoint m and size 2ε. That is, the receiver's prior set is an interval $[m-\varepsilon, m+\varepsilon] \subset (0,1)$, where $m \in (0,1)$ and $\varepsilon \in [0,\bar{\varepsilon})$, $\bar{\varepsilon} := \min\{m, 1-m\}$.[9] We interpret the size of the prior set 2ε as the receiver's *perceived ambiguity.*[10] Notice that we allow for the case $m = \mu_0$, where the receiver perceives ambiguity around the sender's prior.

A *signal* $\pi = (\pi(\cdot|\omega_L), \pi(\cdot|\omega_H))$ consists of a pair of conditional probability distributions over a binary set of outcomes $S = \{s_0, s_1\}$. For notational simplicity, we abbreviate $\pi_{ij} := \pi(s_i|\omega_j)$ and (without loss) require $\pi_{1L} \leq \pi_{1H}$ (so s_1 is evidence in favor of ω_H). The sender chooses a signal to modify the receiver's prior set. The receiver observes the sender's choice of signal and an outcome $s \in S$. The receiver then updates her prior set to a set of posterior beliefs, by applying Bayes' rule to each prior in her prior set. Given a signal

[9]We parameterize the receiver's prior set in this way because it is convenient for our comparative statics analysis. This is discussed in more detail in Section 3.3. The assumption that the limits of the prior set are different from 0 and 1 is essentially without loss. We make this assumption because there are sometimes discontinuities at $\varepsilon = m$, which makes the analysis more tedious (but with nearly identical results). Furthermore, we assume a closed and convex set of priors. As the analysis shows, only the extreme points of the set of priors matter. Hence, the results would be the same for any closed set of priors. However, the underlying axiomatizations of the multiple prior model (Gilboa and Schmeidler, 1989; Ghirardato et al., 2004) imply that the set of priors is also convex.

[10]For convenience of exposition, in the comparative statics analysis below, we sometimes refer to the parameter ε as the receiver's perceived ambiguity.

π and an outcome s_i, the receiver updates each prior $\mu \in [m-\varepsilon, m+\varepsilon]$ to a posterior

$$\beta_i(\pi,\mu) := \frac{\pi_{iH}\mu}{\pi_{iH}\mu + \pi_{iL}(1-\mu)}.$$

Because Bayes' rule is increasing in the prior, the receiver's posterior set given π and s_i equals $[\beta_i(\pi, m-\varepsilon), \beta_i(\pi, m+\varepsilon)]$. Notice also that our normalization $\pi_{1L} \leq \pi_{1H}$ implies that $\beta_0(\pi,\mu) \leq \mu \leq \beta_1(\pi,\mu)$ for any prior $\mu \in (0,1)$.

After updating her beliefs, the receiver takes an action $a \in A := \{0,1\}$. The receiver's preferences over actions and states are represented by a payoff function $u(a,\omega)$. We assume that the receiver prefers action $a = 0$ when the state is known to be ω_L and action $a = 1$ when the state is known to be ω_H. That is, $u(0,\omega_L) > u(1,\omega_L)$ and $u(0,\omega_H) < u(1,\omega_H)$. The only substantial assumption with respect to the receiver's payoff function is that the receiver does not have an action that is dominant across both states of the world (in which case the problem becomes trivial).

We assume that the receiver has α-maxmin preferences. This means that her objective function is a convex combination of the lowest and the highest expected payoff across the beliefs in her posterior set. Formally, the receiver's preferences over actions given a (arbitrary) posterior set $[\mu_1,\mu_2] \subset [0,1]$ are represented by the function

$$U_{[\mu_1,\mu_2]}(a) := \alpha \min_{\mu\in[\mu_1,\mu_2]} \{\mu u(a,\omega_H) + (1-\mu)u(a,\omega_L)\} \\ + (1-\alpha) \max_{\mu\in[\mu_1,\mu_2]} \{\mu u(a,\omega_H) + (1-\mu)u(a,\omega_L)\},$$

where $\alpha \in [0,1]$ is the weight that the receiver puts on the worst-case scenario (i.e., the lowest expected payoff across the beliefs in the posterior set). Therefore, α can be interpreted as the degree of pessimism of the receiver. We refer to α as the receiver's *ambiguity attitude*, ranging from full optimism ($\alpha = 0$), to full pessimism ($\alpha = 1$). We say that the receiver is *pessimistic* (*optimistic*) if $\alpha \geq (\leq) 1/2$ and *ambiguity neutral* if $\alpha = 1/2$.[11]

The sender's payoff depends only on the receiver's action, and is given by a strictly increasing payoff function $v(a)$. This means that the goal of the

[11] An $\alpha = 1/2$ receiver is ambiguity neutral in our model in the sense that her behavior under her prior set does not respond to our "midpoint preserving" expansions of the prior set, $[m-\varepsilon, m+\varepsilon]$. In fact, the utility $U_{[m-\varepsilon,m+\varepsilon]}(a)$ of such a receiver is constant in ε for each action a.

sender is to persuade the receiver to take action $a = 1$. For simplicity (and without loss), we normalize the sender's payoff function such that $v(0) = 0$ and $v(1) = 1$.

In an equilibrium of the game, the receiver's choice of action maximizes her α-maxmin expected utility $U_{[\beta_i(\pi,m-\varepsilon),\beta_i(\pi,m+\varepsilon)]}(a)$ for every possible signal π and outcome s_i, and the sender's choice of signal π maximizes sender's expected payoff (given receiver's maximizing behavior). For simplicity, we assume that whenever the receiver is indifferent between both actions, she chooses $a = 1$.[12] We refer to an equilibrium signal as an *optimal signal*.

3. Analysis

3.1. Preliminaries

It is convenient to conduct the analysis in terms of the following payoff differences:

$$\begin{aligned} \Delta_1 &:= u(1, \omega_H) - u(1, \omega_L) \qquad & \Delta_L &:= u(0, \omega_L) - u(1, \omega_L) \\ \Delta_0 &:= u(0, \omega_L) - u(0, \omega_H) \qquad & \Delta_H &:= u(1, \omega_H) - u(0, \omega_H). \end{aligned}$$

We refer to $|\Delta_a|$, i.e., the absolute value of the receiver's payoff difference across state given action a, as the *sensitivity* of action a. We say that a is *more sensitive* than a' if $|\Delta_a| \geq |\Delta_{a'}|$.

The term Δ_L (Δ_H) can be interpreted as the cost of a type I (type II) error. Since (by assumption) $\Delta_L > 0$ and $\Delta_H > 0$ and (by definition) $\Delta_L + \Delta_H = \Delta_0 + \Delta_1$, it follows that $\Delta_0 + \Delta_1 > 0$. As a result, action a is more sensitive than action a' if and only if $\Delta_a \geq \Delta_{a'}$.

The signs of Δ_0 and Δ_1 are important in the analysis, because they determine the beliefs at which the receiver evaluates each of her actions. Since Δ_0 and Δ_1 cannot be negative at the same time, there are three possible cases:

(A) $\Delta_0, \Delta_1 \geq 0$: the receiver payoff is larger for action $a = 0$ in state ω_L and for action $a = 1$ in state ω_H.

[12] This assumption is without loss as long as the receiver is restricted to pure strategies, and then merely implies that we do not need to specify the receiver's behavior as a part of our equilibrium characterizations. We discuss the implications of allowing the receiver to choose probabilistically in the concluding remarks.

(B) $\Delta_0 \geq 0 \geq \Delta_1$: the receiver's payoff is larger in state ω_L than in state ω_H regardless of the action (ω_L is the receiver's "favorite" state).

(C) $\Delta_0 \leq 0 \leq \Delta_1$: the receiver's payoff is larger in state ω_H than in state ω_L regardless of the action (ω_H is the receiver's "favorite" state).

In cases (B) and (C), where the receiver has a favorite state, we say that the receiver's actions are *comonotonic*.[13] In case (A), we say that the receiver's actions are *non-comonotonic*. In case (B) (case (C)) action $a = 0$ ($a = 1$) is the more sensitive one, while in case (A) either action could be the more sensitive one.

3.2. The optimal signal

The goal of the sender is to persuade the receiver to take action $a = 1$. This means that the sender must construct a signal such that the receiver is willing to take action $a = 1$ for at least one of the posterior sets induced by that signal. In order to characterize such posterior sets, it is useful to define the function

$$G(\mu_1, \mu_2) = U_{[\mu_1,\mu_2]}(1) - U_{[\mu_1,\mu_2]}(0).$$

$G(\mu_1, \mu_2)$ gives the receiver's payoff difference between action $a = 1$ and $a = 0$ if her posterior set is $[\mu_1, \mu_2]$. That is, the receiver is willing to take action $a = 1$ given posterior set $[\mu_1, \mu_2]$ if and only if $G(\mu_1, \mu_2) \geq 0$.

An explicit formula for $G(\mu_1, \mu_2)$ can be obtained straightforwardly. First notice that it is possible to write

$$\begin{aligned} U_{[\mu_1,\mu_2]}(1) &= u(1, \omega_L) + \alpha \min_{\mu \in [\mu_1,\mu_2]} \mu\Delta_1 + (1-\alpha) \max_{\mu \in [\mu_1,\mu_2]} \mu\Delta_1 \\ U_{[\mu_1,\mu_2]}(0) &= u(0, \omega_L) - \alpha \max_{\mu \in [\mu_1,\mu_2]} \mu\Delta_0 - (1-\alpha) \min_{\mu \in [\mu_1,\mu_2]} \mu\Delta_0. \end{aligned}$$

Plugging these expressions into G and solving the min and max terms case by case yields

[13] "Comonotonic" stands for "common monotonicity": Two actions a and a' are said to be comonotonic if there are no states ω, ω' such that $u(a, \omega) < u(a, \omega')$ and $u(a', \omega) > u(a', \omega')$.

$$G(\mu_1, \mu_2) = \begin{cases} \mu'\Delta_0 + \mu''\Delta_1 - \Delta_L & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \geq 0 \quad \text{(A)} \\ \mu'(\Delta_0 + \Delta_1) - \Delta_L & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \leq 0 \quad \text{(B)} \\ \mu''(\Delta_0 + \Delta_1) - \Delta_L & \text{if } \Delta_0 \leq 0 \text{ and } \Delta_1 \geq 0 \quad \text{(C)} \end{cases} \tag{1}$$

where $\mu' = (1-\alpha)\mu_1 + \alpha\mu_2$ and $\mu'' = \alpha\mu_1 + (1-\alpha)\mu_2$.

The expression for G reveals a qualitative difference between the conditions under which the receiver chooses $a = 1$ when her actions are comonotonic (B and C) and when her actions are non-comonotonic (A). The reason for this difference is that when the receiver's actions are comonotonic, the worst case belief and the best case belief are constant across actions. The receiver, therefore, evaluates both actions using the same α-mixture of the upper and lower bound of her posterior set, and chooses $a = 1$ whenever this α-mixture is above a threshold. When the receiver's actions are non-comonotonic, the worst case and the best case beliefs switch across actions, and the receiver, therefore, uses different α-mixtures to evaluate the different actions. Since there are two α-mixtures of posterior beliefs, the receiver's choice of $a = 1$ is no longer characterized in terms of a single threshold for a single α-mixture of beliefs.

Since $\Delta_0 + \Delta_1 > 0$, it follows that G is strictly increasing in both its arguments. That is, the receiver becomes more inclined to take action $a = 1$ when the lower bound or the upper bound of her posterior set moves up. This is intuitive, given that the receiver prefers $a = 1$ when the state is known to be ω_H, and also aligns with the common prior counterpart of our model, where the receiver becomes more inclined to $a = 1$ as the common prior moves up.

If $G(m-\varepsilon, m+\varepsilon) \geq 0$, then the receiver chooses action $a = 1$ at her prior set, which means that there is no need for the sender to persuade the receiver. Formally, the sender maximizes his expected payoff by using an *uninformative signal*, i.e., a signal such that receiver's posterior set equals the prior set with probability one.[14] The following result, therefore, follows immediately, and is left without proof.

Lemma 1. *If $G(m-\varepsilon, m+\varepsilon) \geq 0$, then any uninformative signal is optimal.*

[14] It follows from a straightforward calculation that a signal is uninformative if and only if $\pi_{1L} = \pi_{1H}$.

If instead $G(m-\varepsilon, m+\varepsilon) < 0$, the receiver chooses $a = 1$ only if the sender provides sufficiently compelling evidence in favor of that action. In this case, because G is strictly increasing in both arguments, and because $\beta_0(\pi, \mu) \leq \mu \leq \beta_1(\pi, \mu)$ for any $\mu \in (0, 1)$, we have that $G(\beta_0(\pi, m-\varepsilon), \beta_0(\pi, m+\varepsilon)) < 0$. In other words, sender can only persuade receiver to take action $a = 1$ for one of the outcomes of the signal. Let $\Pi = \{(\pi_{1L}, \pi_{1H}) \in [0, 1]^2 : \pi_{1L} \leq \pi_{1H}\}$. It follows that a signal π is optimal if and only if it solves

$$\begin{aligned} &\max_{(\pi_{1L}, \pi_{1H}) \in \Pi} \pi_{1L}(1 - \mu_0) + \pi_{1H}\mu_0 \\ &\text{s.t. } G(\beta_1(\pi, m - \varepsilon), \beta_1(\pi, m + \varepsilon)) \geq 0. \end{aligned}$$

That is, an optimal signal maximizes the probability of producing outcome s_1, subject to outcome s_1 persuading the receiver to choose action $a = 1$. Since the sender's objective function is strictly increasing in π_{1L} and π_{1H}, and since the constraint is violated if $\pi_{1L} = \pi_{1H} = 1$, the constraint is binding at the optimum. Since, in addition, G is strictly increasing in both arguments and $\beta_1(\pi, m - \varepsilon)$ and $\beta_1(\pi, m + \varepsilon)$ are strictly increasing in π_{1H}, it follows that $\pi_{1H} = 1$ at the optimum. The optimal value of π_{1L}, then, is characterized by the equality of the constraint in the sender's optimization problem, given $\pi_{1H} = 1$. This means that we have the following characterization of the optimal signal whenever the receiver would choose $a = 0$ at her prior set.

Proposition 1. *Suppose $G(m - \varepsilon, m + \varepsilon) < 0$. A signal π is then optimal if and only if $\pi_{1H} = 1$ and for cases (A), (B), and (C), we have, respectively, that equations (*), (**), and (***) below hold.*

(A) $\Delta_0 \geq 0$, $\Delta_1 \geq 0$ *and*

$$\gamma\Delta_0 + \eta\Delta_1 - \Delta_L = 0 \tag{*}$$

(B) $\Delta_0 \geq 0$, $\Delta_1 \leq 0$ *and*

$$\gamma(\Delta_0 + \Delta_1) - \Delta_L = 0 \tag{**}$$

(C) $\Delta_0 \leq 0$, $\Delta_1 \geq 0$ *and*

$$\eta(\Delta_0 + \Delta_1) - \Delta_L = 0 \tag{***}$$

where

$$\gamma = (1-\alpha)\frac{m-\varepsilon}{m-\varepsilon+(1-m+\varepsilon)\pi_{1L}} + \alpha\frac{m+\varepsilon}{m+\varepsilon+(1-m-\varepsilon)\pi_{1L}} \quad \textit{and}$$

$$\eta = \alpha\frac{m-\varepsilon}{m-\varepsilon+(1-m+\varepsilon)\pi_{1L}} + (1-\alpha)\frac{m+\varepsilon}{m+\varepsilon+(1-m-\varepsilon)\pi_{1L}}$$

Further, there is a unique optimal signal.

Notice that the formulae in (*)-(***) correspond to $G(\beta_1(\pi, m-\varepsilon), \beta_1(\pi, m+\varepsilon)) = 0$ evaluated at $\pi_{1H} = 1$, for each of the cases (A)-(C).

A couple of remarks are in place. First, at the optimal signal, the receiver is certain about the state when she chooses $a = 0$. That is, outcome s_0 collapses the receiver's prior set to $\{0\}$. On the other hand, outcome s_1 pushes the receiver's prior set up to the point where she is indifferent between both actions. This fundamental logic is analogous to the common prior counterpart of our model (analyzed by K&G), and follows because the receiver becomes more willing to choose action $a = 1$ as the upper or lower bound of the posterior set goes up. Persuading the receiver to choose action $a = 1$, then, amounts to pushing the posterior set up. Since the probability of producing the outcome that pushes the posterior set up becomes smaller as the posterior set is pushed further up, the optimal signal leaves the receiver indifferent between both actions when she chooses $a = 1$.

Second, formulae (*)-(***) present a suitable opportunity to compare the main predictions of the common prior benchmark with those of our model. When the sender and the receiver share common prior m, the receiver's indifference between both actions is characterized by the equation $\beta_1(\pi, m) = \Delta_L/(\Delta_L + \Delta_H)$. This can be seen by setting $\varepsilon = 0$ in (*)-(***). Therefore, the sender's optimal signal under a common prior simply sets the receiver's posterior equal to the cutoff value $\Delta_L/(\Delta_L + \Delta_H)$, which is characterized by the ratio of the cost of type II and type I errors, Δ_H/Δ_L. The explicit formula for the optimal signal under a common prior can be calculated as

$$\pi_{1L} = \frac{m}{1-m}\frac{\Delta_H}{\Delta_L}.$$

Comparative statics with respect to the common prior follow from straightforward differentiation. As (*)-(***) make clear, the characterization of the optimal signal under ambiguity, as well as the comparative statics, depend on α-mixtures of receiver beliefs and are more complex. In fact, the explicit formula for π_{2L} is too cumbersome to be provided in the main text, and is given in Appendix A for the interested reader.

Third, the characterization of the optimal signal under ambiguity is different depending on whether the receiver's actions are comonotonic or non-comonotonic. In the case of comonotonic actions (B and C), the receiver's indifference is again characterized by the ratio $\Delta_L/(\Delta_0+\Delta_1) = \Delta_L/(\Delta_L+\Delta_H)$, and, therefore, by the ratio of the cost of type II and type I errors Δ_H/Δ_L. This ratio now determines the cutoff for the α-mixture of the receiver's posteriors such that the receiver is indifferent between both actions. It follows that in the case of comonotonic actions, the main difference between our model and the common prior case, is that here the equilibrium value of π_{1L} is determined such that an α-mixture of posteriors, rather than a single posterior, equals $\Delta_L/(\Delta_L + \Delta_H)$.

We should also remark that except for the special cases of $\alpha = 0$ and $\alpha = 1$, there is not an analogy in the sense that behavior under ambiguity with comonotonic actions can be replicated by an expected utility model with an appropriately defined receiver prior. In particular, for generic values of α, m, and ε, one cannot define a receiver prior $\mu(\alpha, m, \varepsilon)$ such that the optimal signal is the same function of Δ_L and Δ_H under both ambiguity and expected utility. This follows because the functional form of $\beta_0(\pi, m)$ is different from the functional forms of the α-mixtures γ and η. In fact, solving for the optimal value of π_{1L} in the case of a common prior requires solving a linear equation, while a quadratic equation must be solved in the case of comonotonic actions. We refer the reader to Appendix A for the explicit solution to this equation.

In the case of non-comonotonic actions (A), the result is qualitatively different. The fact that the receiver evaluates action $a = 0$ and $a = 1$ using the opposite α-mixtures of beliefs, translates to a more complex characterization of the optimal signal. There are now three parameters, Δ_0, Δ_1 and Δ_L, in the expression determining the optimal value of π_{1L}. In particular, it is necessary to know the ratio of the sensitivities Δ_0/Δ_1 and the ratio of the cost of type I errors to the sensitivity of action $a = 1$ (or $a = 0$), Δ_L/Δ_1, in order to pin down the value of π_{1L}. This implies that one can specify a set of receiver payoff functions that would all produce the same prediction in the common prior case, but across which the predictions would vary in our model. Notice also that in the common prior case, the costs of type I and type II errors are the crucial parameters. In contrast, in the case of non-comonotonic actions and an ambiguity-sensitive receiver, the sensitivities of the receiver's actions play an important role.[15]

The optimal value of π_{1L} is an inverse measure of the informativeness of the sender's optimal signal. As the optimal value of π_{1L} increases, the optimal signal becomes less informative (in the sense of Blackwell, 1951). This makes the receiver worse off in the sense that her choice of action $a = 1$ will coincide with state ω_H less frequently (regardless of the "objective" prior probability of state ω_H). At the same time, as the optimal value of π_{1L} increases, the sender's equilibrium payoff increases, because action $a = 1$ is induced with a higher probability. We now proceed to examine the effect of the parameters in the model on the optimal value of π_{1L}.

3.3. Comparative statics

We now proceed to our comparative statics analysis, which examines the effect of the receiver's ambiguity attitude α and the receiver's perceived ambiguity, measured by ε, on the optimal value of π_{1L}. Interestingly, while the receiver's payoff function is separately linear in both α and ε, the comparative statics with respect to these parameters are qualitatively different. This follows because when the receiver is exposed to information, ε goes into Bayes rule, which creates a non-linear response to changes in ε. This can be gleaned from formulae (*)-(***) in Proposition 1 and is reflected in our comparative statics results below.

We first introduce some notation. Fix $\pi_{1H} = 1$ and let $\beta(\pi_{1L}, \mu) := \beta_1(\pi, \mu)$. That is, $\beta(\pi_{1L}, \mu)$ denotes the posterior belief given outcome s_1 as a function of π_{1L} and the prior μ, with $\pi_{1H} = 1$ fixed. Second, define the function $\tilde{\pi}_{1L} : [0, \bar{\varepsilon}] \times [0, 1] \to [0, 1]$, such that $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is the optimal value of π_{1L} given (ε, α) and with $\pi_{1H} = 1$ fixed. If $G(m - \varepsilon, m + \varepsilon) < 0$, then $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is defined by $G(\beta(\tilde{\pi}_{1L}(\varepsilon, \alpha), m - \varepsilon), \beta(\tilde{\pi}_{1L}(\varepsilon, \alpha), m + \varepsilon)) = 0$, and otherwise, if the optimal signal is uninformative, $\tilde{\pi}_{1L}(\varepsilon, \alpha) = 1$. Our object of study in this section is $\tilde{\pi}_{1L}(\varepsilon, \alpha)$.

Since the explicit formulae for $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ obtained by solving equations (*)-(***) are rather complex, our analysis relies to a large extent on differentiating the equation $G(\beta(\pi_{1L}, m - \varepsilon), \beta(\pi_{1L}, m + \varepsilon)) = 0$ implicitly. The following straightforward observation is useful in this regard.

[15] Notice, however, that in the special cases of an ambiguity neutral receiver ($\alpha = 1/2$) and equal sensitivities ($\Delta_0 = \Delta_1$), the characterization under non-comonotonic actions aligns with the common prior case.

Remark 1. The function $G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))$ is twice continuously differentiable in Δ_0, Δ_1, α, ε and π_{1L}.

3.3.1. Ambiguity attitude

We first examine the comparative statics effects of the receiver's ambiguity attitude α on the equilibrium signal. Because $G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))$ is linear in α, the comparative statics with respect to α are straightforward. In particular, differentiating $G(\mu_1, \mu_2)$ with respect to α (at an arbitrary posterior set) yields

$$\frac{\partial G(\mu_1,\mu_2)}{\partial \alpha} = \begin{cases} (\mu_2-\mu_1)(\Delta_0-\Delta_1) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \geq 0 \quad \text{(A)} \\ (\mu_2-\mu_1)(\Delta_0+\Delta_1) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \leq 0 \quad \text{(B)} \\ -(\mu_2-\mu_1)(\Delta_0+\Delta_1) & \text{if } \Delta_0 \leq 0 \text{ and } \Delta_1 \geq 0 \quad \text{(C)} \end{cases}.$$

Since case (B) implies $\Delta_0 - \Delta_1 > 0$ and case (C) implies $\Delta_0 - \Delta_1 < 0$, the sign of $\partial G(\mu_1, \mu_2)/\partial\alpha$ coincides with the sign of $\Delta_0 - \Delta_1$. Therefore, as α goes up, the sender's constraint becomes more relaxed if $\Delta_0 - \Delta_1 > 0$ and less relaxed if $\Delta_0 - \Delta_1 < 0$. Since, in addition, the sender's objective function is increasing in π_{1L}, it follows that $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\alpha$ shares sign with $\Delta_0 - \Delta_1$ and we have the following result.

Proposition 2. *The optimal signal $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is continuous in α, and whenever $\tilde{\pi}_{1L}(\varepsilon, \alpha) < 1$, it is is also differentiable in α with*

$$\frac{\partial \tilde{\pi}_{1L}(\varepsilon,\alpha)}{\partial \alpha} \begin{cases} < 0 & \textit{if } \Delta_0 < \Delta_1 \\ = 0 & \textit{if } \Delta_0 = \Delta_1 \\ > 0 & \textit{if } \Delta_0 > \Delta_1 \end{cases}.$$

There is a natural intuition for Proposition 2. Recall that $\Delta_a \geq \Delta_{a'} \Leftrightarrow |\Delta_a| \geq |\Delta_{a'}|$. Proposition 2 states that if the sender's preferred action $a = 1$ is more sensitive than $a = 0$, then the optimal signal becomes more informative as the receiver becomes more pessimistic. This means that the sender's equilibrium payoff also declines as the receiver becomes more pessimistic. The logic is that when $a = 1$ is the more sensitive action, the receiver perceives more uncertainty given $a = 1$ than given $a = 0$. This uncertainty is viewed less favorably by the receiver as she becomes more pessimistic, and it then becomes harder to persuade the receiver to take action $a = 1$. The opposite logic holds when $a = 0$ is the more sensitive action. In this case, the receiver perceives more uncertainty given $a = 0$, and it then becomes easier to persuade her to take action $a = 1$ as she becomes more pessimistic and views uncertainty less favorably.[16]

3.3.2. Perceived ambiguity

We now examine the effects of the receiver's perceived ambiguity (i.e., the size of the receiver's prior set $[m - \varepsilon, m + \varepsilon]$), on the sender's optimal signal. That is, we examine the dependence of $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ on ε. Our exercise is constructed such that it evaluates the effect on the optimal signal as the prior set expands uniformly at the upper and lower bound, keeping the location m of the prior set fixed. This means that the results that we obtain are valid regardless of the location of the prior set.[17]

If we set $m = \mu_0$, then the comparative statics with respect to ε provides a comparison between the common prior counterpart of our model and a uniform expansion of the receiver's prior set around the common prior. This gives an idea of whether an expert would be interested in "convincing" an ambiguity perceiving client about his prior, before exposing the client to more information (e.g., whether it would be in the interest of a doctor to share background information with an ambiguity perceiving patient, before subjecting the patient to a test).

The comparative statics with respect to ε are characterized in a lemma and two propositions. Our first result reveals that the dependence of $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ on ε has an intuitive overall structure.

Lemma 2. *Suppose $I = \{\varepsilon \in [0, \bar{\varepsilon}) : \tilde{\pi}_{1L}(\varepsilon, \alpha) < 1\} \neq \varnothing$. Then $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is continuous in ε, the set I is a right-open interval, and $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is differentiable in ε on I with either*

1. *$\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon \leq 0$ for all $\varepsilon \in I$ and $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon < 0$ for $\varepsilon \in I$ such that $\varepsilon > 0$,*

[16] Our emphasis on sensitivity as the driving force is not artificial. If one differentiates $G(\mu_1, \mu_2)$ with respect to α without imposing our assumption $\Delta_L, \Delta_H > 0$, it explicitly follows that the derivative shares sign with $|\Delta_0| - |\Delta_1|$.

[17] On the other hand, our exercise only gives a partial notion of the effect of expanding the upper and lower bound of the prior set at different rates. An alternative, more exhaustive, exercise, would trace level curves in terms of the upper and lower bound, such that the optimal value of π_{1L} is kept constant. The non-linearity of Bayes' rule and the relatively large number of parameters in our model, however, make this exercise exceedingly complex, to the point that it becomes difficult to reach conclusions beyond the ones we obtain with the approach taken here.

2. $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon > 0$ *for all* $\varepsilon \in I$, *or*

3. $I = [0,\bar{\varepsilon})$ *and there is some* $\varepsilon^* > 0$ *such that*

$$\frac{\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)}{\partial\varepsilon} \begin{cases} > 0 & \text{if } \varepsilon < \varepsilon^* \\ = 0 & \text{if } \varepsilon = \varepsilon^* \\ < 0 & \text{if } \varepsilon > \varepsilon^* \end{cases} .$$

Lemma 2 implies that the equilibrium signal as a function of the receiver's perceived ambiguity, and, therefore, the sender's preferences over the receiver's perceived ambiguity, can be understood in terms of three cases. In case (i), because of the continuity of $\tilde{\pi}_{1L}(\varepsilon,\alpha)$ and the condition on the partial derivative, the equilibrium value of π_{1L} (and the sender's equilibrium payoff) is monotonically decreasing in the receiver's perceived ambiguity ε. In this case, $\tilde{\pi}_{1L}(\varepsilon,\alpha) = 1$ on some (possibly empty) interval $[0,\hat{\varepsilon}]$ and $\tilde{\pi}_{1L}(\varepsilon,\alpha)$ is strictly decreasing on $I = (\hat{\varepsilon},\bar{\varepsilon})$. If $\tilde{\pi}_{1L}(0,\alpha) < 1$, then the sender's payoff is strictly decreasing in the receiver's perceived ambiguity over $I = [0,\bar{\varepsilon})$. In case (ii), $\tilde{\pi}_{1L}(\varepsilon,\alpha)$ (and the sender's equilibrium payoff) is instead strictly increasing on some interval $I = [0,\hat{\varepsilon})$ and $\tilde{\pi}_{1L}(\varepsilon,\alpha) = 1$ on a (possibly empty) interval $[\hat{\varepsilon},\bar{\varepsilon})$. In case (iii), $\tilde{\pi}_{1L}(\varepsilon,\alpha)$ (and the sender's equilibrium payoff) is non-monotonic in the receiver's perceived ambiguity. In this case, $\tilde{\pi}_{1L}(\varepsilon,\alpha)$ is first strictly increasing and then strictly decreasing, with a unique ambiguity bliss point $\varepsilon^* \in (0,\bar{\varepsilon})$. Cases (i)-(iii) are illustrated in Figure 1.

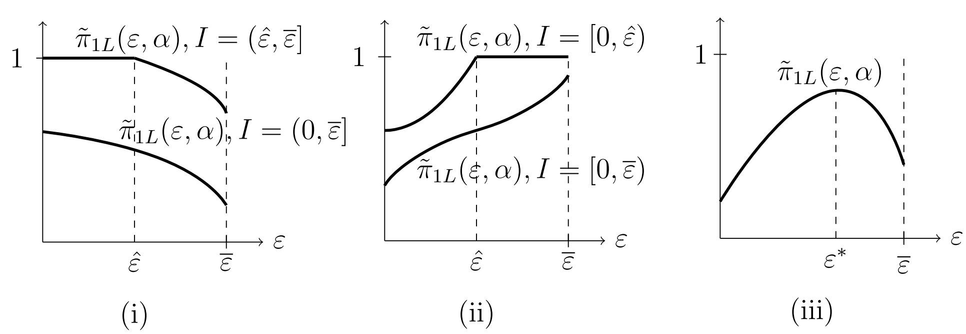

Figure 1: Examples of $\tilde{\pi}_{1L}(\varepsilon,\alpha)$ as a function of ε for cases (i)-(iii) in Lemma 2. Subfigures (i) and (ii) contain two examples of the curve.

The gist of the proof of Lemma 2 is to establish a quasiconcavity property of $\tilde{\pi}_{1L}(\varepsilon,\alpha)$: Whenever $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon = 0$, we have that $\partial^2\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon^2 < 0$.

This property, in turn, results because $G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))$ is decreasing in π_{1L} and concave in ε. The concavity of G is a consequence of Bayes' rule being concave in the prior whenever the likelihood ratio π_{1L}/π_{1H} is smaller than one. This condition holds here because $\pi_{1H} = 1$, and, more generally, because the sender needs to push the receiver's prior set upwards to persuade her to take his desired action.[18] The concavity of G in ε is a fundamental aspect of our model, which tends to produce a sender preference for degenerate receiver prior sets, in a way discussed in more detail below.

We now discuss the conditions under which each of cases (i)-(iii) in Lemma 2 occurs. Our next result completely characterizes the conditions such that $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is decreasing in ε, i.e., such that the sender prefers the receiver to have a unique prior m rather than any prior set $[m-\varepsilon, m+\varepsilon]$. Because we allow $m = \mu_0$, the result also characterizes the situations in which the sender would like to "clarify" things for the receiver and have her adopt his prior, if this was possible.

Proposition 3. *Suppose* $I = \{\varepsilon \in [0, \bar{\varepsilon}) : \tilde{\pi}_{1L}(\varepsilon, \alpha) < 1\} \neq \varnothing$. *Then* $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon \leq 0$ *for all* $\varepsilon \in I$ *if and only if either (a)* $\Delta_0 = \Delta_1$*, (b)* $\alpha = 1/2$*, (c)* $\Delta_0 > \Delta_1$ *and* $\alpha < 1/2$*, or (d)* $\Delta_0 < \Delta_1$ *and* $\alpha > 1/2$*.*

Conditions (c) and (d) of Proposition 3 align with Proposition 2 and have a related intuition. These conditions roughly state that if the receiver responds negatively to ambiguity, from the sender's point of view, then the sender prefers the receiver to perceive less ambiguity. Specifically, part (d) states that if the receiver is pessimistic ($\alpha > 1/2$) and the sender's preferred action is the more sensitive one ($\Delta_0 < \Delta_1$), then $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ (and the sender's equilibrium payoff) is decreasing in the receiver's perceived ambiguity ε. Intuitively, a pessimistic receiver dislikes the greater payoff uncertainty given $a = 1$. This preference becomes stronger as her prior set expands and she perceives more ambiguity. This makes it more difficult to persuade the receiver to take action $a = 1$. A similar intuition applies to part (c).

Conditions (a) and (b) of Proposition 3 state that if both receiver actions are equally sensitive ($\Delta_0 = \Delta_1$) or if the receiver is ambiguity neutral ($\alpha = 1/2$), then the sender prefers the receiver to perceive less ambiguity. While these cases are razor edge, and should not be overemphasized, they manifest

[18] The concavity of Bayes' rule in the prior in this case is well known and can be verified by a simple calculation.

a tilt in our model toward sender preferences for a degenerate receiver prior set. Although both cases are perfectly neutral, the sender (often strictly) prefers zero receiver perceived ambiguity.

To understand these results, it is helpful to notice that condition (a) and (b) both collapse the sender's constraint to[19]

$$G(\beta(\pi, m-\varepsilon), \beta(\pi, m+\varepsilon)) =$$
$$\frac{1}{2}\left[\frac{m-\varepsilon}{m-\varepsilon+(1-m+\varepsilon)\pi_{1L}}+\frac{m+\varepsilon}{m+\varepsilon+(1-m-\varepsilon)\pi_{1L}}\right](\Delta_0+\Delta_1)-\Delta_L$$

At $\varepsilon = 0$, the effect on both posteriors (in brackets) of raising ε is the same, but of opposite sign, and the effect on the constraint is therefore null. As ε goes further up, however, the concavity in the prior of posteriors that are updated upwards, implies that the negative effect on the smaller of the two posteriors becomes larger, while the positive effect on the larger posterior tapers off. As a result, the net effect is negative, and the sender prefers ε to remain at zero.[20] Interestingly, the concavity of G in ε that drives this result is an unavoidable consequence of the sender's need to push the receiver's posterior set away from the prior set, to induce his desired action. That is, the sender's inclination for zero receiver perceived ambiguity is a fundamental feature of our binary environment, and is not the product of any particular assumption that we make within this environment.[21]

Our final comparative statics result concerns the conditions under which

[19] In passing, notice that in these cases, the optimal signal is again characterized by the ratio of the costs of type I and type II errors, as in the common prior counterpart to our model.

[20] The apparent paradox in the sender strictly preferring an ambiguity neutral receiver to perceive no ambiguity, is due to the updating of beliefs. An ambiguity neutral receiver does not respond to our "midpoint preserving" expansions of the prior set when acting under her prior set. However, as explained here, once the change in the prior set goes through Bayes rule, it becomes "midpoint reducing" and then affects the ambiguity neutral receiver to the sender's detriment.

[21] Naturally, the type of conclusion we obtain depends on the fact that our exercise consists of expanding the upper and lower bounds of the receiver's prior set uniformly. If we allowed the upper and lower bounds to expand at different rates, the conclusion would be of a similar flavor, however. In this case, given $\alpha = 1/2$ or $\Delta_0 = \Delta_1$, the upper bound would need to change quicker than the lower bound for the sender to be indifferent as the receiver's prior set expands. This also reflects, in a related sense, a sender inclination for smaller receiver prior sets.

cases (ii) and (iii) in Lemma 2 occur, i.e., where $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is either increasing or non-monotonic in ε. Proposition 3 implies that these cases occur when $\Delta_0 > \Delta_1$ and $\alpha > 1/2$ or $\Delta_0 < \Delta_1$ and $\alpha < 1/2$. The distinction between cases (ii) and (iii), however, is more complex. We find it useful to characterize the comparative statics in terms of the receiver's ambiguity attitude. Let $\varepsilon(\alpha)$ denote the interior ambiguity bliss point in case (iii), whenever this bliss point exists. We obtain the following result.

Proposition 4. *Suppose that there is some (ε, α) such that $\tilde{\pi}_{1L}(\varepsilon, \alpha) < 1$.*

1. *If $\Delta_0 > \Delta_1$ and $\alpha > 1/2$, then there is some $\hat{\alpha} \in (1/2, 1]$ such that*
   - *there is an interior ambiguity bliss point $\varepsilon(\alpha)$ for all $\alpha < \hat{\alpha}$ and*
   - *$\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is increasing in ε for any $\alpha > \hat{\alpha}$.*

   It holds that $\varepsilon'(\alpha) > 0$ on $(1/2, \hat{\alpha})$ and $\lim_{\alpha \searrow 1/2} \varepsilon(\alpha) = 0$. Further, there is some $\hat{\Delta}_0 > \Delta_1$ such that $\hat{\alpha} = 1$ if $\Delta_0 < \hat{\Delta}_0$ and $\hat{\alpha} < 1$ if $\Delta_0 > \hat{\Delta}_0$.

2. *If $\Delta_0 < \Delta_1$ and $\alpha < 1/2$, then there is some $\hat{\alpha} \in [0, 1/2)$ such that*
   - *there is an interior ambiguity bliss point $\varepsilon(\alpha)$ for all $\alpha > \hat{\alpha}$ and*
   - *$\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is increasing in ε for any $\alpha < \hat{\alpha}$.*

   It holds that $\varepsilon'(\alpha) < 0$ on $(\hat{\alpha}, 1/2)$ and $\lim_{\alpha \nearrow 1/2} \varepsilon(\alpha) = 0$. Further, there is some $\hat{\Delta}_1 > \Delta_0$ such that $\hat{\alpha} = 0$ if $\Delta_1 < \hat{\Delta}_1$ and $\hat{\alpha} > 0$ if $\Delta_1 > \hat{\Delta}_1$.

The first part of Proposition 4 states that if the receiver is pessimistic ($\alpha > 1/2$) and the sender's preferred action is the least sensitive one ($\Delta_0 > \Delta_1$), then $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ (and the sender's equilibrium payoff) is non-monotonic in the receiver's perceived ambiguity ε for some interval of ambiguity attitudes $(1/2, \hat{\alpha})$. Over this interval, the sender's perceived ambiguity bliss-point $\varepsilon(\alpha)$, is strictly increasing in α. It may be the case that $\hat{\alpha} = 1$ and $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is non-monotonic in ε for all $\alpha > 1/2$. If instead $\hat{\alpha} < 1$, then $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is increasing in ε whenever the receiver is sufficiently pessimistic. Whether $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is non-monotonic in ε for all $\alpha > 1/2$, or whether $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ becomes increasing for α sufficiently large, depends on the remaining parameters.

For example, Proposition 4 establishes that if Δ_0 is sufficiently large, then $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ becomes monotonic in ε for large values of α. The second part of Proposition 4 is simply a mirror image for the case $\Delta_0 < \Delta_1$ and $\alpha < 1/2$. We illustrate the result in Figure 2.

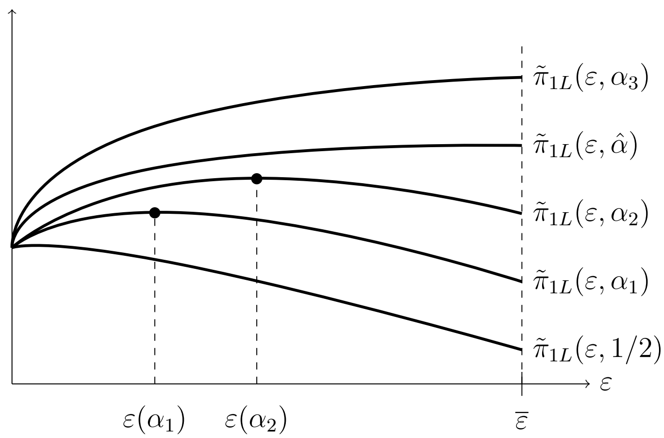

Figure 2: Illustration of part 1 of Proposition 4, with $\Delta_0 > \hat{\Delta}_0 > \Delta_1$ and $1/2 < \alpha_1 < \alpha_2 < \hat{\alpha} < \alpha_3$. Note that the slope of $\tilde{\pi}_{1L}(\varepsilon, \hat{\alpha})$ is zero at $\bar{\varepsilon}$.

Proposition 3 and 4 characterize the comparative statics with respect to the receiver's perceived ambiguity for distinct values of the remaining parameters. An intuitive way of combining them, is to observe that when $\Delta_0 < \Delta_1$, the sender's preferred level of receiver perceived ambiguity is decreasing in α, and when $\Delta_0 > \Delta_1$, the sender's preferred level of receiver ambiguity is increasing in α. This conclusion is formalized in the following corollary.

Corollary. *Suppose* $\tilde{\pi}_{1L}(\varepsilon, \alpha) < 1$ *for all* $(\varepsilon, \alpha) \in [0, \bar{\varepsilon}) \times [0, 1]$.[22] *Let* $\epsilon : [0, 1] \to [0, \bar{\varepsilon}]$ *map the receiver's ambiguity attitude to the sender's preferred level of receiver perceived ambiguity, where we set* $\epsilon(\alpha) := \bar{\varepsilon}$ *if* $\tilde{\pi}(\varepsilon, \alpha)$ *is increasing in* ε.

[22]This condition implies that we ignore the cases in which uninformative signals are optimal for some parameter values. While the result requires additional qualifications if we allow uninformative signals to be optimal, its qualitative structure remains identical.

(i) If $\Delta_0 < \Delta_1$*, then there is some* $\hat{\alpha} \in [0, 1/2)$ *such that*

$$\epsilon(\alpha) = \begin{cases} \bar{\varepsilon} & \text{if } 0 \le \alpha < \hat{\alpha} \\ \varepsilon(\alpha) & \text{if } \hat{\alpha} < \alpha < 1/2 \\ 0 & \text{if } 1/2 \le \alpha \le 1 \end{cases} .$$

(ii) If $\Delta_0 = \Delta_1$*, then* $\epsilon(\alpha) = 0$ *for all* $\alpha \in [0, 1]$.

(iii) If $\Delta_0 > \Delta_1$*, then there is some* $\hat{\alpha} \in (1/2, 1]$ *such that*

$$\epsilon(\alpha) = \begin{cases} 0 & \text{if } 0 \le \alpha \le 1/2 \\ \varepsilon(\alpha) & \text{if } 1/2 < \alpha < \hat{\alpha} \\ \bar{\varepsilon} & \text{if } \hat{\alpha} < \alpha \le 1 \end{cases} .$$

The corollary (which is illustrated in Figure 3) summarizes two main qualitative features of our results. First, as mentioned, the sender's preferences for receiver perceived ambiguity are monotonic in the receiver's ambiguity attitude. Second, the sender's preferences for receiver perceived ambiguity are tilted toward zero receiver ambiguity. In particular, notice that regardless of which of the two actions is the more sensitive one, the sender prefers zero receiver ambiguity for (slightly more than) half of all possible values of α. The symmetric counterpart to this observation, would be for the sender to prefer maximal receiver ambiguity for the remaining values of α. Instead, there is an interior bliss-point $\varepsilon(\alpha)$ that is arbitrarily close to zero when α is close to 1/2, and only possibly a cutoff value $\hat{\alpha}$ above which the sender prefers maximal receiver ambiguity.

These results can be understood in terms of the interaction of two main forces. First, the expansion of the prior set can make the receiver more or less inclined to take the sender's preferred action, depending on which action is the more sensitive one and depending on her ambiguity attitude. For example, if the sender's preferred action is the least sensitive one, an optimistic ($\alpha \le 1/2$) receiver views the smaller uncertainty of that action unfavorably, while a pessimistic ($\alpha \ge 1/2$) receiver views it favorably. This translates to an unfavorable response to more perceived ambiguity for an optimistic receiver, and a favorable response for a pessimistic receiver.[23] Therefore, the

[23] This effect can be anticipated by simply staring at the expression for the sender's constraint in (1).

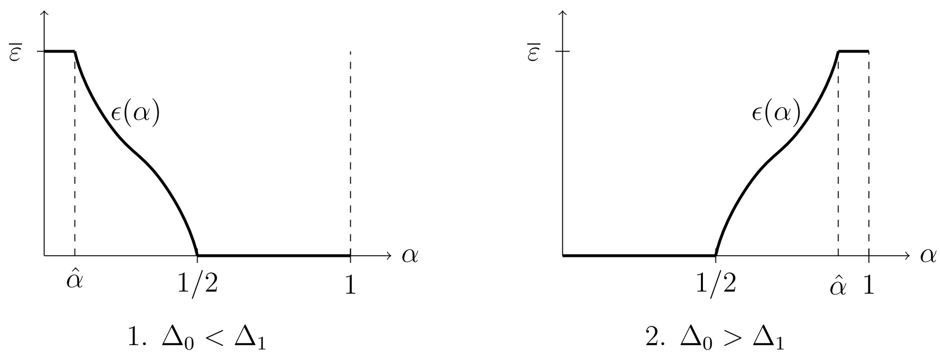

Figure 3: Sender's preferred level of receiver perceived level of ambiguity as a function of the receiver's ambiguity attitude (examples are chosen such that $\Delta_0 > \hat{\Delta}_0$ and $\Delta_1 > \hat{\Delta}_1$).

sender finds that it becomes easier (harder) to persuade a pessimistic (optimistic) receiver as the prior set expands. In fact, the effect of expanding the prior set tends to become more favorable for the sender as the receiver becomes more pessimistic. This explains the monotonicity of the sender's preferred level of receiver perceived ambiguity, present in both parts of the Corollary.

Second, as the prior set expands it becomes inherently more difficult to persuade the receiver, because of the fundamental above discussed concavity property of Bayes' rule. This property implies that as the receiver's prior set expands, the posterior set tends to expand faster at the lower bound than at the upper bound. This produces a sort of sluggishness, which requires the sender to use more informative signals, to push up the receiver's posterior set to the point where she is willing to take the sender's preferred action. That is, the receiver becomes more difficult to persuade as her prior set expands, regardless of her ambiguity attitude and the sensitivities of her actions. This explains the tilt toward sender preferences for degenerate receiver prior sets, present in both parts of the Corollary.

The comparative statics described in the Corollary reflect the resolution of these two forces. When the receiver is optimistic (pessimistic) and the sender's preferred action is the most (least) sensitive one, both forces push in the same direction, resulting in a sender preference for a collapsed prior set. When the receiver is pessimistic (optimistic) and the sender's preferred action is the least (most) sensitive one, both forces push in opposite directions. The tension then produces an interior bliss point, with the resolution leaning more towards a maximal prior set, as the receiver's uncertainty preferences become more extreme.

As a final remark, we note that the Corollary combined with Proposition 2 produces the following summary of our comparative statics results.

Remark 2. If $\Delta_0 > (<)\Delta_1$, then the sender's equilibrium expected payoff, as well as the sender's preferred degree of receiver perceived ambiguity, are both increasing (decreasing) in the receiver's degree of pessimism α.

4. Concluding remarks

In this paper we investigate the effect of receiver ambiguity on optimal persuasion. We find that the sender's preferences over receiver ambiguity attitude and perceived ambiguity depend crucially on the receiver's payoff differences across states given each action (the sensitivities of the actions). Our analysis also uncovers a certain inherent sender preference for persuading a receiver with a clearly defined prior, which arises because of curvature properties of Bayes' rule in the prior. Conversely, this means that there is a certain tendency for a receiver who perceives more ambiguity to be exposed to more informative equilibrium signals, which implies more accurate equilibrium receiver decisions (given any objective prior on the state).

We believe that the binary environment of this paper is an important special case of the Bayesian persuasion framework but it is also an obvious limitation of our analysis.[24] While it is of interest to understand how our results generalize to less restrictive environments, extensions to larger state spaces or action spaces have thus far proven challenging. One might speculate, however, that some of the mechanics underlying the tilt of the sender's preferences, will be important also when the receiver's action space is larger. For example, if the action space is an interval and the receiver's payoff function is a quadratic loss function, a fully pessimistic receiver would choose the action that is optimal at the midpoint of her posterior set. That is, the sender's optimal choice of signal would involve moving the midpoints of the receiver's posterior sets optimally. As the receiver's prior set expands, this becomes a more sluggish process because of the curvature properties of Bayes' rule that drive the tilt of the sender's preferences in the binary environment. We leave a detailed analysis for future research.

[24]Notice that several important extensions of the Bayesian persuasion framework emphasize similar binary environments, see, e.g., Bergemann et al. (2018), Felgenhauer and Loerke (2017), Orlov et al. (2018), and Perez-Richet (2014).

Another interesting direction for future research would be to allow the receiver to choose between both actions probabilistically. As is well known, decision makers who perceive ambiguity and are ambiguity averse sometimes have strict preferences for mixed strategies – sometimes referred to as hedging. In our setting, allowing the receiver to mix would indeed have an important impact on the optimal behavior of the sender. In particular, it is possible to show that the sender's optimal strategy of persuasion is sometimes to induce receiver hedging. The analysis, however, is rather orthogonal to the one in the present paper. We therefore also leave this for future research.

Acknowledgements. We thank two anonymous referees and Jürgen Eichberger for very helpful feedback.

Appendix A. Explicit Form of Optimal Signal

Here we provide explicit formulae for the sender's optimal signal for cases (A) and (B) in Proposition 1. These are given by the solutions to equations (*) and (**).

$$\pi_{1L} = \frac{(m-\varepsilon)((1-\alpha)\Delta_0 + \alpha\Delta_1 - \Delta_L)}{(1-(m-\varepsilon))2\Delta_L} + \frac{(m+\varepsilon)(\alpha\Delta_0 + (1-\alpha)\Delta_1 - \Delta_L)}{(1-(m+\varepsilon))2\Delta_L} + \sqrt{\left[\frac{(m-\varepsilon)((1-\alpha)\Delta_0+\alpha\Delta_1-\Delta_L)}{(1-(m-\varepsilon))2\Delta_L} + \frac{(m+\varepsilon)(\alpha\Delta_0+(1-\alpha)\Delta_1-\Delta_L)}{(1-(m+\varepsilon))2\Delta_L}\right]^2 + \frac{(m-\varepsilon)(m+\varepsilon)\Delta_H}{(1-(m-\varepsilon))(1-(m+\varepsilon))\Delta_L}} \tag{A}$$

$$\pi_{1L} = \frac{(m-\varepsilon)((1-\alpha)\Delta_H - \alpha\Delta_L)}{(1-(m-\varepsilon))2\Delta_L} + \frac{(m+\varepsilon)(\alpha\Delta_H - (1-\alpha)\Delta_L)}{(1-(m+\varepsilon))2\Delta_L} + \sqrt{\left[\frac{(m-\varepsilon)((1-\alpha)\Delta_H-\alpha\Delta_L)}{(1-(m-\varepsilon))2\Delta_L} + \frac{(m+\varepsilon)(\alpha\Delta_H-(1-\alpha)\Delta_L)}{(1-(m+\varepsilon))2\Delta_L}\right]^2 + \frac{(m-\varepsilon)(m+\varepsilon)\Delta_H}{(1-(m-\varepsilon))(1-(m+\varepsilon))\Delta_L}} \tag{B}$$

The expression for π_{1L} for case (B) confirms the fact that the optimal signal is characterized by the ratio Δ_H/Δ_L in the case of comonotonic actions (as discussed in the main text). This, however, is not the case when actions are not comonotonic, as reflected in in the expression for π_{1L} for case (A) above.

Appendix B. Proofs

Proof of Proposition 1. Suppose $G(m-\varepsilon, m+\varepsilon) < 0$. The formulas in the statement follow from the arguments in the main text. Existence of an optimal signal follows from the intermediate value theorem, because $G(\beta_1(\pi, m-\varepsilon), \beta_1(\pi, m+\varepsilon))$ is continuous in π_{1L}, and because $G(\beta_1(\pi, m-\varepsilon), \beta_1(\pi, m+\varepsilon)) = G(1,1) = \Delta_H > 0$ for $\pi_{1H} = 1$ and $\pi_{1L} = 0$, and $G(\beta_1(\pi, m-\varepsilon), \beta_1(\pi, m+\varepsilon)) = G(m-\varepsilon, m+\varepsilon) < 0$ for $\pi_{1H} = 1$ and $\pi_{1L} = 1$. Uniqueness follows because $G(\beta_1(\pi, m-\varepsilon), \beta_1(\pi, m+\varepsilon))$ is strictly decreasing in π_{1L}. $\square$

Proof of Proposition 2. Continuity of $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ in α follows from the maximum theorem, as $\tilde{\pi}_{1L}(\varepsilon, \alpha) = \max\{\pi_{1L} \in [0,1] : G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon)) \geq 0\}$ and G is continuous.

Suppose $\tilde{\pi}_{1L}(\varepsilon, \alpha) < 1$. Since $\beta(\pi_{1L}, m-\varepsilon)$ and $\beta(\pi_{1L}, m+\varepsilon)$ are strictly decreasing in π_{1L}, and since G is strictly increasing in both its arguments, $\partial G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))/\partial \pi_{1L}|_{\pi_{1L}=\tilde{\pi}_{1L}(\varepsilon,\alpha)} < 0$. The implicit function theorem then yields that $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ is differentiable in α with

$$\frac{\partial \tilde{\pi}_{1L}(\varepsilon, \alpha)}{\partial \alpha} = -\left.\frac{\partial G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))/\partial \alpha}{\partial G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))/\partial \pi_{1L}}\right|_{\pi_{1L}=\tilde{\pi}_{1L}(\varepsilon,\alpha)}.$$

As the denominator is negative, $\partial \tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial \alpha$ shares sign with $\partial G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))/\partial \alpha$.

In case (B), we have $\Delta_0 \geq 0 \geq \Delta_1$ and, therefore, $\Delta_0 - \Delta_1 > 0$. In case (C), we have $\Delta_0 \leq 0 \leq \Delta_1$ and, therefore, $\Delta_0 - \Delta_1 < 0$. The explicit formula for $\partial G(\mu_1, \mu_2)/\partial \alpha$ in the main text then implies that $\partial G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))/\partial \alpha$ shares sign with $\Delta_0 - \Delta_1$. $\square$

Proof of Lemma 2. Suppose $I = \{\varepsilon \in [0, \bar{\varepsilon}) : \tilde{\pi}_{1L}(\varepsilon, \alpha) < 1\} \neq \varnothing$. Continuity of $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ follows from the maximum theorem, as $\tilde{\pi}_{1L}(\varepsilon, \alpha) = \max\{\pi_{1L} \in [0,1] : G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon)) \geq 0\}$ and G is continuous. Differentiability of $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ in ε on I follows by an argument similar to the one for the differentiability of $\tilde{\pi}_{1L}(\varepsilon, \alpha)$ in α in the proof of Proposition 2. The remainder of the proof consists of three steps.

Step 1. The set I is a right-open interval.

Proof. The result is obvious if $I = [0, \bar{\varepsilon})$, so suppose $\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha) = 1$ for some $\tilde{\varepsilon} \in [0, \bar{\varepsilon})$. Notice that

$$G(m-\varepsilon, m+\varepsilon) = \begin{cases} (m-\varepsilon)(\alpha\Delta_1 + (1-\alpha)\Delta_0) + (m+\varepsilon)(\alpha\Delta_0 + (1-\alpha)\Delta_1) - \Delta_L & \text{(A)} \\ (1-\alpha)(m-\varepsilon)(\Delta_0 + \Delta_1) + \alpha(m+\varepsilon)(\Delta_0 + \Delta_1) - \Delta_L & \text{(B)} \\ \alpha(m-\varepsilon)(\Delta_0 + \Delta_1) + (1-\alpha)(m+\varepsilon)(\Delta_0 + \Delta_1) - \Delta_L & \text{(C)} \end{cases}$$

is a linear function of ε. Therefore, either $G(m-\varepsilon, m+\varepsilon) \geq 0$ (and $\tilde{\pi}_{1L}(\varepsilon, \alpha) = 1$) on some segment $[0, x]$ and $G(m-\varepsilon, m+\varepsilon) < 0$ (and $\tilde{\pi}_{1L}(\varepsilon, \alpha) < 1$) on $(x, \bar{\varepsilon})$, or $G(m-\varepsilon, m+\varepsilon) < 0$ on some segment $[0, x)$ and $G(m-\varepsilon, m+\varepsilon) \geq 0$ on $[x, \bar{\varepsilon})$. □

Step 2. If $\partial\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha)/\partial\varepsilon = 0$ for some $\tilde{\varepsilon} \in I$, then $\partial^2\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha)/\partial\varepsilon^2 < 0$.

Proof. Suppose $\partial\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha)/\partial\varepsilon = 0$ and $\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha) < 1$. The implicit function theorem implies that $\tilde{\pi}_{1L}$ is differentiable at $\tilde{\varepsilon}$ with derivative

$$\frac{\partial\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha)}{\partial\varepsilon} = -\frac{\partial G(\beta(\pi_{1L}, m-\tilde{\varepsilon}), \beta(\pi_{1L}, m+\tilde{\varepsilon}))/\partial\varepsilon}{\partial G(\beta(\pi_{1L}, m-\tilde{\varepsilon}), \beta(\pi_{1L}, m+\tilde{\varepsilon}))/\partial\pi_{1L}}\bigg|_{\pi_{1L}=\tilde{\pi}_{1L}(\tilde{\varepsilon},\alpha)}.$$

Since the denominator is strictly negative, $\partial\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha)/\partial\varepsilon$ shares sign with the numerator. It follows that $\partial G(\beta(\pi_{1L}, m-\tilde{\varepsilon}), \beta(\pi_{1L}, m+\tilde{\varepsilon}))/\partial\varepsilon = 0$.

With a slight abuse of notation, abbreviate $\partial G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))/\partial x$ as $\partial G/\partial x$. The second (implicit) derivative at $\varepsilon = \tilde{\varepsilon}$ is given by

$$\begin{aligned}\frac{\partial^2\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha)}{\partial\varepsilon^2} &= -\frac{\frac{\partial^2 G}{\partial\varepsilon^2}\left(\frac{\partial G}{\partial\pi_{1L}}\right)^2 - 2\frac{\partial^2 G}{\partial\varepsilon\partial\pi_{1L}}\frac{\partial G}{\partial\pi_{1L}}\frac{\partial G}{\partial\varepsilon} + \frac{\partial^2 G}{\partial\pi_{1L}^2}\left(\frac{\partial G}{\partial\varepsilon}\right)^2}{\left(\frac{\partial G}{\partial\pi_{1L}}\right)^3}\Bigg|_{\pi_{1L}=\tilde{\pi}_{1L}(\tilde{\varepsilon},\alpha),\varepsilon=\tilde{\varepsilon}} \\ &= -\frac{\partial^2 G(\beta(\pi_{1L}, m-\tilde{\varepsilon}), \beta(\pi_{1L}, m+\tilde{\varepsilon}))/\partial\varepsilon^2}{\partial G(\beta(\pi_{1L}, m-\tilde{\varepsilon}), \beta(\pi_{1L}, m+\tilde{\varepsilon}))/\partial\pi_{1L}}\bigg|_{\pi_{1L}=\tilde{\pi}_{1L}(\tilde{\varepsilon},\alpha)},\end{aligned}$$

which shares sign with $\partial^2 G(\beta(\pi_{1L}, m-\tilde{\varepsilon}), \beta(\pi_{1L}, m+\tilde{\varepsilon}))/\partial\varepsilon^2$. Straightforward differentiation gives

$$\frac{\partial^2 G}{\partial\varepsilon^2} = \begin{cases} -2\pi_{1L}(1-\pi_{1L})\left(\frac{\alpha\Delta_1+(1-\alpha)\Delta_0}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^3} + \frac{\alpha\Delta_0+(1-\alpha)\Delta_1}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^3}\right) & \text{(A)} \\ -2\pi_{1L}(1-\pi_{1L})\left(\frac{(1-\alpha)(\Delta_0+\Delta_1)}{(m-\varepsilon+(1-m+\varepsilon)\pi_{1L})^3} + \frac{\alpha(\Delta_0+\Delta_1)}{(m+\varepsilon+(1-m-\varepsilon)\pi_{1L})^3}\right) & \text{(B)} \\ -2\pi_{1L}(1-\pi_{1L})\left(\frac{\alpha(\Delta_0+\Delta_1)}{(m-\varepsilon+(1-m+\varepsilon)\pi_{1L})^3} + \frac{(1-\alpha)(\Delta_0+\Delta_1)}{(m+\varepsilon+(1-m-\varepsilon)\pi_{1L})^3}\right) & \text{(C)} \end{cases}.$$

Since $\Delta_0 \geq 0$ and $\Delta_1 \geq 0$ in case (A), and since $\Delta_0 + \Delta_1 > 0$, we have that $\partial^2 G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon))/\partial\varepsilon^2 < 0$, which proves the claim. □

Step 3. Suppose $\varepsilon, \varepsilon' \in I$ with $\varepsilon < \varepsilon'$. If $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon \leq 0$, then $\partial\tilde{\pi}_{1L}(\varepsilon', \alpha)/\partial\varepsilon < 0$. If $\partial\tilde{\pi}_{1L}(\varepsilon', \alpha)/\partial\varepsilon \geq 0$, then $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon > 0$.

Proof. Suppose $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon \leq 0$ and, to contradiction, that

$$\partial\tilde{\pi}_{1L}(\varepsilon', \alpha)/\partial\varepsilon > 0$$

for some $\varepsilon' > \varepsilon$. Let $\tilde{\varepsilon} = \max\{x \in [\varepsilon, \varepsilon'] : \partial\tilde{\pi}_{1L}(x, \alpha)/\partial\varepsilon = 0\}$, where $\tilde{\varepsilon}$ is well defined by the hypothesis, and because $\partial\tilde{\pi}_{1L}(x, \alpha)/\partial\varepsilon$ is differentiable on I and, therefore, continuous in its first argument on I. Step 2 implies $\partial^2\tilde{\pi}_{1L}(\tilde{\varepsilon}, \alpha)/\partial\varepsilon^2 < 0$, and, combining with the definition of $\tilde{\varepsilon}$, we obtain $\partial\tilde{\pi}_{1L}(x, \alpha)/\partial\varepsilon < 0$ on $(\tilde{\varepsilon}, \varepsilon']$, a contradiction. Therefore, $\partial\tilde{\pi}_{1L}(\varepsilon', \alpha)/\partial\varepsilon \leq 0$.

Suppose $\partial\tilde{\pi}_{1L}(\varepsilon', \alpha)/\partial\varepsilon = 0$. Step 2 implies $\partial^2\tilde{\pi}_{1L}(\varepsilon', \alpha)/\partial\varepsilon^2 < 0$, and there is then an interval (z, ε') on which $\partial\tilde{\pi}_{1L}(x, \alpha)/\partial\varepsilon > 0$, which produces the same contradiction as in the previous paragraph. Hence, $\partial\tilde{\pi}_{1L}(\varepsilon', \alpha)/\partial\varepsilon < 0$. □

The statement that we must have either case (i), (ii), or (iii), follows by combining Step 1 and Step 3. More precisely, suppose $\partial\tilde{\pi}_{1L}(\varepsilon^*, \alpha)/\partial\varepsilon = 0$ for some $\varepsilon^* \in I$. If ε^* is an element of the interior of I, then Step 3 implies that we are in case (iii), and otherwise $\varepsilon^* = \min I$ and we are in case (ii). If there is no such ε^*, then the continuity of $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon$ in ε implies that we are either in case (i) or (ii). □

Proof of Proposition 3. Suppose $I = \{\varepsilon \in [0, \bar{\varepsilon}) : \tilde{\pi}_{1L}(\varepsilon, \alpha) < 1\} \neq \varnothing$. We prove each direction of the statement separately.

Step 1. If either (a) $\Delta_0 = \Delta_1$, (b) $\alpha = 1/2$, (c) $\Delta_0 > \Delta_1$ and $\alpha < 1/2$, or (d) $\Delta_0 < \Delta_1$ and $\alpha > 1/2$, then $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon \leq 0$ for all $\varepsilon \in I$.

Proof. Suppose, to contradiction, $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon > 0$ for some $\varepsilon \in I$. Lemma 2 then implies $\tilde{\pi}_{1L}(0, \alpha) < 1$ and $\partial\tilde{\pi}_{1L}(0, \alpha)/\partial\varepsilon > 0$. The implicit function theorem implies that $\partial\tilde{\pi}_{1L}(0, \alpha)/\partial\varepsilon$ shares sign with

$$\left.\frac{\partial G(\beta(\pi, m-\varepsilon), \beta(\pi, m+\varepsilon))}{\partial\varepsilon}\right|_{\varepsilon=0}$$
$$= \begin{cases} \frac{\tilde{\pi}_{1L}(0,\alpha)(2\alpha-1)(\Delta_0-\Delta_1)}{(m+\tilde{\pi}_{1L}(0,\alpha)(1-m))^2} & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \geq 0 \text{ (A)} \\ \frac{\tilde{\pi}_{1L}(0,\alpha)(2\alpha-1)(\Delta_0+\Delta_1)}{(m+\tilde{\pi}_{1L}(0,\alpha)(1-m))^2} & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \leq 0 \text{ (B)} \\ \frac{\tilde{\pi}_{1L}(0,\alpha)(1-2\alpha)(\Delta_0+\Delta_1)}{(m+\tilde{\pi}_{1L}(0,\alpha)(1-m))^2} & \text{if } \Delta_0 \leq 0 \text{ and } \Delta_1 \geq 0 \text{ (C)} \end{cases}$$

Therefore, we must have either $\Delta_0 > \Delta_1$ and $\alpha > 1/2$ or $\Delta_0 < \Delta_1$ and $\alpha < 1/2$, which is the negation of the hypothesis in the claim. □

Step 2. If $\partial\tilde{\pi}_{1L}(\varepsilon, \alpha)/\partial\varepsilon \leq 0$ for all $\varepsilon \in I$, then either (a) $\Delta_0 = \Delta_1$, (b) $\alpha = 1/2$, (c) $\Delta_0 > \Delta_1$ and $\alpha < 1/2$, or (d) $\Delta_0 < \Delta_1$ and $\alpha > 1/2$.

Proof. Suppose $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon \leq 0$ for all $\varepsilon \in I$. If $\tilde{\pi}_{1L}(0,\alpha) < 1$, then $\partial\tilde{\pi}_{1L}(0,\alpha)/\partial\varepsilon \leq 0$, and the formula derived in Step 1 immediately implies the result. Suppose, therefore, that $\tilde{\pi}_{1L}(0,\alpha) = 1$, which implies $G(m,m) \geq 0$. We have that

$$\frac{\partial G(m-\varepsilon, m+\varepsilon)}{\partial\varepsilon} = \begin{cases} (2\alpha-1)(\Delta_0-\Delta_1) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \geq 0 \quad \text{(A)} \\ (2\alpha-1)(\Delta_0+\Delta_1) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \leq 0 \quad \text{(B)} \\ (1-2\alpha)(\Delta_0+\Delta_1) & \text{if } \Delta_0 \leq 0 \text{ and } \Delta_1 \geq 0 \quad \text{(C)} \end{cases}.$$

If $\Delta_0 > \Delta_1$ and $\alpha > 1/2$ we are either in case (A) or (B) and $\partial G(m-\varepsilon, m+\varepsilon)/\partial\varepsilon > 0$. If $\Delta_0 < \Delta_1$ and $\alpha < 1/2$, we are either in case (A) or (C) and $\partial G(m-\varepsilon, m+\varepsilon)/\partial\varepsilon > 0$. But then $G(m-\varepsilon, m+\varepsilon) \geq 0$ for all $\varepsilon \in [0,\bar{\varepsilon})$, which contradicts the assumption that $\tilde{\pi}_{1L}(\varepsilon,\alpha) < 1$ for some $\varepsilon \in [0,\bar{\varepsilon})$. Therefore, one of conditions (i)-(iv) in the statement of the claim must hold. □

Proof of Proposition 4. Let $G(\mu_1,\mu_2,\alpha)$ be the sender's constraint function G such that α is explicitly included in the list of arguments. We first prove the following claim.

Claim. Suppose $\tilde{\pi}_{1L}(\varepsilon,\alpha) < 1$ and $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon = 0$. If $\Delta_0 > \Delta_1$, then $\partial^2\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon\partial\alpha > 0$. If $\Delta_0 < \Delta_1$, then $\partial^2\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon\partial\alpha > 0$.

Proof. Suppose $\tilde{\pi}_{1L}(\varepsilon,\alpha) < 1$, $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon = 0$ and $\Delta_0 > \Delta_1$. With a slight abuse of notation, abbreviate $\partial G(\beta(\pi_{1L}, m-\varepsilon), \beta(\pi_{1L}, m+\varepsilon), \alpha)/\partial x$ as $\partial G/\partial x$. Notice that

$$\frac{\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)}{\partial\varepsilon} = -\left.\frac{\frac{\partial G(\beta(\pi_{1L},m-\varepsilon),\beta(\pi_{1L},m+\varepsilon),\alpha)}{\partial\varepsilon}}{\frac{\partial G(\beta(\pi_{1L},m-\varepsilon),\beta(\pi_{1L},m+\varepsilon),\alpha)}{\partial\pi_{1L}}}\right|_{\pi_{1L}=\tilde{\pi}_{1L}(\varepsilon,\alpha)} \implies$$

$$\frac{\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)}{\partial\varepsilon\partial\alpha} = \left.\frac{\left(\frac{\partial^2 G}{\partial\pi_{1L}^2}\frac{\partial\tilde{\pi}_{1L}}{\partial\alpha} + \frac{\partial^2 G}{\partial\pi_{1L}\partial\alpha}\right)\frac{\partial G}{\partial\varepsilon} - \left(\frac{\partial^2 G}{\partial\varepsilon\partial\pi_{1L}}\frac{\partial\tilde{\pi}_{1L}}{\partial\alpha} + \frac{\partial^2 G}{\partial\varepsilon\partial\alpha}\right)\frac{\partial G}{\partial\pi_{1L}}}{\left(\frac{\partial G}{\partial\pi_{1L}}\right)^2}\right|_{\pi_{1L}=\tilde{\pi}_{1L}(\varepsilon,\alpha)}$$

$$= -\left.\frac{\frac{\partial^2 G}{\partial\varepsilon\partial\pi_{1L}}\frac{\partial\tilde{\pi}_{1L}}{\partial\alpha} + \frac{\partial^2 G}{\partial\varepsilon\partial\alpha}}{\frac{\partial G}{\partial\pi_{1L}}}\right|_{\pi_{1L}=\tilde{\pi}_{1L}(\varepsilon,\alpha)},$$

where the second line uses the fact that $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon = 0 \Leftrightarrow \partial G/\partial\varepsilon = 0$. We have that $\partial G/\partial\pi_{1L} < 0$, so the sign of $\partial^2\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon\partial\alpha$ coincides with that of the numerator. Proposition 2 implies that $\partial\tilde{\pi}_{1L}/\partial\alpha > 0$. The remaining terms need to be computed.

We compute[25]

$$\frac{\partial G}{\partial \varepsilon} = \begin{cases} \pi_{1L}\left(-\frac{(\alpha\Delta_1+(1-\alpha)\Delta_0)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2} + \frac{(\alpha\Delta_0+(1-\alpha)\Delta_1)}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^2}\right) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \geq 0 \text{ (A)} \\ \pi_{1L}\left(-\frac{1-\alpha}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2} + \frac{\alpha}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^2}\right) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \leq 0 \text{ (B)} \end{cases}$$

Therefore,

$$\frac{\partial G}{\partial \varepsilon} = 0 \Leftrightarrow \begin{cases} \frac{(\alpha\Delta_1+(1-\alpha)\Delta_0)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2} = \frac{(\alpha\Delta_0+(1-\alpha)\Delta_1)}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^2} & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \geq 0 \text{ (A)} \\ \frac{1-\alpha}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2} = \frac{\alpha}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^2} & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \leq 0 \text{ (B)} \end{cases} \tag{B.1}$$

We compute

$$\begin{aligned} \frac{\partial^2 G}{\partial \varepsilon \partial \pi_{1L}} &= \begin{cases} 2\pi_{1L}\left(\frac{(1-m+\varepsilon)(\alpha\Delta_1+(1-\alpha)\Delta_0)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^3} - \frac{(1-m-\varepsilon)(\alpha\Delta_0+(1-\alpha)\Delta_1)}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^3}\right) & \text{(A)} \\ 2\pi_{1L}\left(\frac{\alpha(1-m+\varepsilon)(\Delta_1+\Delta_0)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^3} - \frac{(1-\alpha)(1-m-\varepsilon)(\Delta_0+\Delta_1)}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^3}\right) & \text{(B)} \end{cases} \\ &= \begin{cases} 2\pi_{1L}\frac{(\alpha\Delta_1+(1-\alpha)\Delta_0)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2}\left(\frac{(1-m+\varepsilon)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))} - \frac{(1-m-\varepsilon)}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))}\right) & \text{(A)} \\ 2\pi_{1L}\frac{\alpha(\Delta_1+\Delta_0)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2}\left(\frac{(1-m+\varepsilon)}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))} - \frac{(1-m-\varepsilon)}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))}\right) & \text{(B)} \end{cases} \end{aligned} \tag{B.2}$$

where the first line follows by applying the product rule and (B.1), and the second line follows by again applying (B.1). Hence, $\partial^2 G/\partial\varepsilon\partial\pi_{1L} > 0$.

Next,

$$\frac{\partial^2 G}{\partial \varepsilon \partial \alpha} = \begin{cases} \pi_{1L}\left(\frac{\Delta_0-\Delta_1}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2} - \frac{\Delta_1-\Delta_0}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^2}\right) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \geq 0 \text{ (A)} \\ \pi_{1L}\left(\frac{\Delta_0+\Delta_1}{(m-\varepsilon+\pi_{1L}(1-m+\varepsilon))^2} + \frac{\Delta_0+\Delta_1}{(m+\varepsilon+\pi_{1L}(1-m-\varepsilon))^2}\right) & \text{if } \Delta_0 \geq 0 \text{ and } \Delta_1 \leq 0 \text{ (B)} \end{cases}$$

[25] Case (C), where $\Delta_0 \leq 0$ and $\Delta_1 \geq 0$, is not included because we assumed $\Delta_0 > \Delta_1$.

Since $\Delta_0 > \Delta_1$, we have, $\partial^2 G/\partial\varepsilon\partial\alpha > 0$.

Summing up, we have $\partial^2 G/\partial\varepsilon\partial\pi_{1L} > 0$, $\partial\tilde{\pi}_{1L}/\partial\alpha > 0$ and $\partial^2 G/\partial\varepsilon\partial\alpha > 0$ and, therefore,

$$\frac{\partial^2 G}{\partial\varepsilon\partial\pi_{1L}}\frac{\partial\tilde{\pi}_{1L}}{\partial\alpha} + \frac{\partial^2 G}{\partial\varepsilon\partial\alpha} > 0.$$

Hence, $\partial^2\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon\partial\alpha > 0$. The proof for the case $\Delta_1 > \Delta_0$ is virtually identical, and is omitted. □

The remainder of the proof establishes statement 1 in Proposition 4 (the proof of statement 2 is analogous). Therefore, suppose $\Delta_0 > \Delta_1$ and $\alpha > 1/2$, and that $\tilde{\pi}_{1L}(\varepsilon,\alpha) < 1$ for some $(\varepsilon,\alpha) \in [0,\bar{\varepsilon}) \times (1/2,1]$. Notice that we have $G(m,m,1/2) < 0$. To see this, suppose, to contradiction, $G(m,m,1/2) \geq 0$. But

$G(m-\varepsilon, m+\varepsilon, \alpha) =$

$$\begin{cases} (m-\varepsilon)\left(\alpha\Delta_1 + (1-\alpha)\Delta_0\right) + (m+\varepsilon)\left(\alpha\Delta_0 + (1-\alpha)\Delta_1\right) - \Delta_L & \text{(A)} \\ (1-\alpha)(m-\varepsilon)(\Delta_0+\Delta_1) + \alpha(m+\varepsilon)(\Delta_0+\Delta_1) - \Delta_L & \text{(B)} \end{cases}$$

is increasing in ε and α if $\Delta_0 > \Delta_1$ and $\alpha > 1/2$. Therefore $G(m-\varepsilon, m+\varepsilon,\alpha) \geq 0$ for all admissible values of (ε,α), which contradicts $\tilde{\pi}_{1L}(\varepsilon,\alpha) < 1$ for some (ε,α).

We now prove the first part of statement 1 in Proposition 4.

Step 1. There is some $\hat{\alpha} \in (1/2,1]$ such that there is an interior ambiguity bliss point $\varepsilon(\alpha)$ for all $\alpha < \hat{\alpha}$ and such that $\tilde{\pi}_{1L}(\varepsilon,\alpha)$ is increasing in ε for any $\alpha > \hat{\alpha}$. It holds that $\varepsilon'(\alpha) > 0$ on $(1/2,\hat{\alpha})$ and $\lim_{\alpha\downarrow 1/2}\varepsilon(\alpha) = 0$.

Proof. First, the implicit function theorem implies that if $\tilde{\pi}_{1L}(\varepsilon,\alpha) < 1$ and $\partial\tilde{\pi}_{1L}(\varepsilon,\alpha)/\partial\varepsilon = 0$, then this equality defines an implicit function $\varepsilon(\alpha)$ on some neighborhood around α, which satisfies

$$\varepsilon'(\alpha) = -\frac{\frac{\partial^2\tilde{\pi}_{1L}(\varepsilon,\alpha)}{\partial\varepsilon\partial\alpha}}{\frac{\partial^2\tilde{\pi}_{1L}(\varepsilon,\alpha)}{\partial\varepsilon^2}}.$$

The Claim and Step 2 in the proof of Lemma 2 imply $\varepsilon'(\alpha) > 0$. That is, if an interior bliss point $\varepsilon(\alpha)$ exists, then $\varepsilon'(\alpha) > 0$.

Since $G(m,m,1/2) < 0$, we can confirm that $\tilde{\pi}_{1L}(0,1/2) < 1$, and consequently, by the proof of Proposition 3, $\partial\tilde{\pi}_{1L}(0,1/2)/\partial\varepsilon = 0$. Combining with the arguments in the previous paragraph, the implicit function theorem implies existence of an interior bliss point $\varepsilon(\alpha)$ on some interval $(1/2, \hat{\alpha})$.

The proof of Step 1 concludes by demonstrating that there is some $\hat{\alpha} \in (1/2, 1]$ such that an interior bliss point exists for all $\alpha < \hat{\alpha}$ but not for any $\alpha > \hat{\alpha}$. To this end, suppose $\alpha_2 > \alpha_1$ and that $\partial\tilde{\pi}_{1L}(\varepsilon(\alpha_2), \alpha_2)/\partial\varepsilon = 0$. Notice that

$$G(\beta(\pi_{1L}, m - \varepsilon(\alpha_2)), \beta(\pi_{1L}, m + \varepsilon(\alpha_2)), \alpha) = 0$$

defines $\tilde{\pi}_{1L}(\varepsilon(\alpha_2), \alpha)$ as an implicit function of α globally over the interval $(1/2, \alpha_2]$. This follows because $G(m-\varepsilon(\alpha_2), m+\varepsilon(\alpha_2), \alpha_2) < 0$ implies $G(m-\varepsilon(\alpha_2), m + \varepsilon(\alpha_2), \alpha) < 0$ for all $\alpha < \alpha_2$. Since $\partial\tilde{\pi}_{1L}(\varepsilon(\alpha_2), \alpha)/\partial\varepsilon = 0$ implies $\partial^2\tilde{\pi}_{1L}(\varepsilon(\alpha_2), \alpha)/\partial\varepsilon\partial\alpha > 0$, and $\alpha_2 > \alpha_1$, it must be that $\partial\tilde{\pi}_{1L}(\varepsilon(\alpha_2), \alpha_1)/\partial\varepsilon < 0$. Therefore, $\tilde{\pi}_{1L}(\varepsilon, \alpha_1)$ is not increasing in ε. Since we know that either $\tilde{\pi}_{1L}(\varepsilon, \alpha_1)$ is increasing in ε, or there is an interior ambiguity bliss point $\varepsilon(\alpha_1)$, we conclude the latter. Since α_2 and α_1 were arbitrary, whenever $\varepsilon(\alpha_2)$ exists, $\varepsilon(\alpha_1)$ exists for any $\alpha_1 < \alpha_2$. This implies that there is some $\hat{\alpha} \in (1/2, 1]$ such that there is an interior bliss point for all $\alpha < \hat{\alpha}$, but not for any $\alpha > \hat{\alpha}$. □

Step 2. There is some $\hat{\Delta}_0 > \Delta_1$ such that the value $\hat{\alpha}$ in the statement of Step 1 satisfies $\hat{\alpha} = 1$ if $\Delta_0 < \hat{\Delta}_0$ and $\hat{\alpha} < 1$ if $\Delta_0 > \hat{\Delta}_0$.

Proof. Let $\tilde{\varepsilon}(\Delta_0)$ denote the interior ambiguity bliss-point as a function of Δ_0 with α fixed at $\alpha = 1$. The comparative statics with respect to α in Step 1 and 2 can be repeated (virtually identically), to conclude that (i) $\tilde{\varepsilon}'(\Delta_0) > 0$ whenever $\tilde{\varepsilon}(\Delta_0)$ exists; (ii) if Δ_0 is sufficiently close to Δ_1, then there is an interior bliss point $\tilde{\varepsilon}(\Delta_0)$; (iii) if $\tilde{\varepsilon}(\Delta_0)$ exists, then $\tilde{\varepsilon}(\Delta_0')$ exists for any $\Delta_0' \in (\Delta_1, \Delta_0)$.

To conclude the argument, it suffices to notice that there is some $\hat{\Delta}_0$ such that if $\Delta_0 > \hat{\Delta}_0$, then $\tilde{\pi}_{1L}(\varepsilon, 1) = 1$ for some (large) values of ε. In this case, $\tilde{\pi}_{1L}(\varepsilon, 1)$ is increasing, i.e., there is no $\tilde{\varepsilon}(\Delta_0)$. Therefore, there must be a value $\hat{\Delta}_0$ such that $\tilde{\varepsilon}(\Delta_0)$ exists if $\Delta_0 < \hat{\Delta}_0$, and $\tilde{\pi}_{1L}(\varepsilon, 1)$ is increasing if $\Delta_0 > \hat{\Delta}_0$. Combining with Step 1, $\hat{\alpha} = 1$ if $\Delta_0 < \hat{\Delta}_0$, and $\hat{\alpha} < 1$ if $\Delta_0 > \hat{\Delta}_0$. □

References

Alonso, R. and Camara, A. (2016), ‘Persuading voters’, *American Economic Review* **106**, 3590–3605.

Alonso, R. and Camara, O. (2018), 'On the value of persuasion by experts', *Journal of Economic Theory* **174**, 103–123.

Beauchene, D., Li, J. and Li, M. (2019), 'Ambiguous persuasion', *Journal of Economic Theory* **179**, 312 – 365.

Bergemann, D., Bonatti, A. and Smolin, A. (2018), 'The design and price of information', *American Economic Review* **108**, 1–48.

Blackwell, D. (1951), Comparison of experiments, *in* 'Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability', University of California Press, Berkeley, CA, pp. 93–102.

Eichberger, J., Grant, S., Kelsey, D. and Koshevoy, G. A. (2011), 'The α-MEU model: A comment', *Journal of Economic Theory* **146**(4), 1684–1698.

Ellsberg, D. (1961), 'Risk, ambiguity, and the savage axioms', *The Quarterly Journal of Economics* **75**, 643–669.

Felgenhauer, M. and Loerke, P. (2017), 'Bayesian persuasion with private experimentation', *International Economic Review* **58**, 829–856.

Gentzkow, M. and Kamenica, E. (2016), 'Competition in persuasion', *The Review of Economic Studies* **84**, 300–322.

Gentzkow, M. and Kamenica, E. (2017), 'Bayesian persuasion with multiple senders and rich signal spaces', *Games and Economic Behavior* **104**, 411–429.

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

Ghirardato, P. and Marinacci, M. (2002), 'Ambiguity made precise: A comparative foundation', *Journal of Economic Theory* **102**, 251–289.

Gilboa, I. and Schmeidler, D. (1989), 'Maxmin expected utility with non-unique prior', *Journal of Mathematical Economics* **18**(2), 141–153.

Hedlund, J. (2017), 'Bayesian persuasion by a privately informed sender', *Journal of Economic Theory* **167**, 229–268.

Hu, J. and Weng, X. (2018), 'Robust persuasion of a privately informed receiver', *Mimeo* .

Hurwicz, L. (1951), Optimality criteria for decision making under ignorance, Discussion paper 370, Cowles Comission.

Jaffray, J. (1989), 'Linear utility theory for belief functions', *Operations Research Letters* **8**, 107–112.

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

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

Kellner, C. and Le Quement, M. (2018), 'Endogenous ambiguity in cheap talk', *Journal of Economic Theory* **173**, 1–17.

Knight, F. (1921), *Risk, Uncertainty, and Profit*, Houghton Mifflin, Boston and New York.

Kolotilin, A. (2018), 'Optimal information disclosure: A linear programming approach', *Theoretical Economics* **13**, 607–635.

Kolotilin, A., Mylovanov, T., Zapechelnyuk, A. and Li, M. (2017), 'Persuasion of a privately informed receiver', *Econometrica* **85**, 1949–1964.

Kosterina, S. (2018), 'Persuasion with unknown beliefs', *Mimeo* .

Laclau, M. and Renou, L. (2016), 'Public persuasion', *Mimeo* .

Li, F. and Norman, P. (2018), 'Sequential persuasion', *Mimeo* .

Orlov, D., Skrzypacz, A. and Zryumov, P. (2018), 'Persuading the principal to wait', *Mimeo* .

Perez-Richet, E. (2014), 'Interim bayesian persuasion: First steps', *American Economic Review* **104**(5), 469–474.

Pires, C. P. (2002), 'A rule for updating ambiguous beliefs', *Theory and Decision* **53**, 137–152.

Riedel, F. and Sass, L. (2014), 'Ellsberg games', *Theory and Decision* **76**, 469–509.