Communicating about Confidence: Cheap Talk with an Ambiguity-Averse Receiver

By Philippe Colo*

Draft: September 4, 2023

An expert, who is only informed of the probability of possible states, communicates with a decision maker through cheap talk. The decision maker considers different probability distributions over states as possible and is ambiguity averse. I show that all equilibria of the game are equivalent to partitional ones and that the most informative is interim dominant for the expert. Information transmission regarding probabilities that are bad news for the decision maker is facilitated by ambiguity aversion. However, ambiguity aversion also makes information transmission impossible, whatever the preference misalignment, regarding probabilities that are good news for him.

JEL: D81, D83, C72

Keywords: Ambiguity, cheap talk

We are laymen about most of the knowledge we claim to possess. On many topics, our claimed knowledge relies much more on experts than on evidence we can directly access. Reliance in experts is particularly important if one considers the case of topics about which experts are uncertain, either because of the lack of consensus among them, or because the existing evidence is too scarce to make

* ETH Zurich, Rämistrasse 101, 8092 Zürich, Switzerland (email: pcolo@ethz.ch). I thank Frédéric Koessler, Jean-Marc Tallon, and Stéphane Zuber for valuable guidance. I also thank Mohammed Abdellaoui, Sarah Auster, Pierpaolo Battigalli, Loic Berger, Antoine Bommier, Valentina Bosetti, Adam Dominiak, Marc Fleurbaey, Alexis Ghersengorin, Ani Guerdjikova, Brian Hill, Christian Kellner, Nenad Kos, Mark Le Quement, Massimo Marinacci, David Martimort, Sujoy Mukerji, Marco Ottaviani, Guillaume Pommey, Francesco Squintani, and three extremely accurate referees for helpful discussions, as well as seminar and conference participants at PSE, Bocconi, ETH Zurich, IESEG, Duisburg-Essen, FUR 2018, SAET2019, and ESEM2019. I also thank Bocconi and ETH for their hospitality. Financial support through ANR CHOp (ANR-17-CE26-0003), ANR ADE (ANR-18-ORAR-0005-01), ANR INDUCED (ANR-17-CE03-0008), and EUR PGSE is gratefully acknowledged.

accurate predictions. For instance, this might be the case in the face of unexpected events, such as the COVID-19 pandemic, investments on very volatile markets, or complex topics such as predicting the economic consequences of climate change[1]. In those instances, even experts cannot accurately predict a most likely outcome, but only convey the degree of confidence they hold over possible options[2].

In this paper, an expert communicates to a decision maker the degree of confidence she holds over possible options, through a game of strategic information transmission. Because the expert is uncertain, her confidence is represented by a probability distribution over possible states. This probability is the expert's type. I focus on the case where the decision maker's beliefs entirely rely on his trust in the expert. Therefore, I assume information is non-certifiable. The transmission is strategic: the sender (the expert) does not necessarily have the same interests as the receiver (the decision maker). For instance, the expert can be concerned with externalities resulting from the decision maker's behaviour on issues such as the spread of a deadly virus or the limitation of greenhouse gas emissions.

The decision maker is aware that the expert can only assign a probability to the possible states. However, as he does not know which probability the expert would assign, the decision maker is exposed to ambiguity[3]; this occurs when the expected payoff to his strategy varies with the probabilities over which he is uncertain. It is then natural to assume that the decision maker may have ambiguity sensitive preferences. These generally fail to satisfy the expected utility requirements, as famously pointed out by Ellsberg (1961). An ambiguity-averse individual will, for instance, tend to favour actions that reduce his exposure to probabilistic uncertainty. I mainly focus on cases where the receiver displays Gilboa and Schmeidler (1989)'s maxmin expected utility preferences (MEU) or

[1]Expert uncertainty can be the result of limits in existing knowledge regarding phenomena fundamentally deterministic. But it can also be the case in mature sciences that are inherently stochastic, as, for instance, quantium physics.

[2]The IPCC's uncertainty language regarding confidence communication (Mastrandrea et al., 2010) has been widely discussed. See for instance Bradley, Helgeson and Hill (2017).

[3]This is typically the case for emerging sciences, where multiple theories compete to explain a given phenomenon. Eliaz and Spiegler (2020) shows that this situation can persist in the long run.

Savage (1972)'s subjective expected utility (SEU).

The game I study is in the tradition of Crawford and Sobel (1982)'s (hereafter CS) cheap talk game. Communication is about the sender's probabilistic confidence over possible states rather than about the states themselves, making the sender's type a probability distribution. Given the strategic nature of the communication, the sender is typically not able to truthfully reveal the confidence she has in each state. Instead, in equilibrium, the sender conveys an interval of probabilities which contains her real confidence. The SEU case will be kept as a benchmark, as it is identical to CS. In that case, the size of the intervals designated in equilibrium will depend only on the *misalignment*: the difference of interest between parties. However, in the MEU case, even for arbitrarily small misalignment, it can be that no information transmission happens in equilibrium. This result strongly contrasts with the SEU case, where the precision of information transmission depends only on the degree of misalignment between parties. When the misalignment is small, information transmission is almost perfect over the entire set of types.

To understand the MEU case result, assume there are only two states. Two cases may arise. In the first case, the receiver's payoff in one state is so *dominated* by the pay-off in the other, that the maximal expected payoff of the receiver is monotone in function of the probability of that dominated state. Then, the more confidence the sender has in the dominated state, the lower the receiver's maximal expected payoff. In the second case, no state is dominated in that manner, and the maximal expected payoff of the receiver in function of the probability of a given state exhibits an interior minimum and is increasing in either direction. In other words, an increase in the sender's confidence in any given state can have a non-monotone effect on the receiver's maximal expected payoff.

Assume the sender is misaligned toward a larger action relative to the receiver. The key effect of ambiguity aversion is then that if the maximal expected payoff of the receiver is decreasing in function of the sender's confidence in the higher state, ambiguity aversion helps communication. This is because the receiver makes decisions based on the worst possible type, which is in the direction of the sender's misalignment. Conversely, ambiguity aversion plays against the sender's misalignment when the maximal expected payoff of the receiver is increasing, making communication impossible in equilibrium. Thus, overall, information transmission regarding probabilities that are bad news for the decision maker is facilitated by ambiguity aversion. However, ambiguity aversion makes information transmission impossible, whatever the misalignment, for probabilities that are good news for the decision maker.

I further show that when the receiver has MEU preferences, the sender is always better off being as informative as possible, even after she learns her type. In the SEU case, no interim ordering of equilibria - from the sender's perspective - is, in general, possible: the sender does not always have an interest in being as informative as she could be.

Scientific communication is a leading application for my model and a topic receiving growing attention in economics. Recent work by Spiess (2018), Banerjee et al. (2020), Andrews and Shapiro (2021) and Schwartzstein and Sunderam (2021) focuses on optimal choice of scientific modelling, depending on the scientist's objective or audience. Unlike these papers, I do not assume that the expert's audience observes data and uses it to assess the statistical properties of the reports provided by the expert. This is because I focus on cases where assigning a degree of confidence is a pure act of expert-judgement[4]. For example, take the epidemiological models used to evaluate the impact of health measures on the COVID-19 pandemic. Two main approaches exist: process-based models that try to capture the mechanisms by which diseases spread, and curve-fitting approaches that aim to mathematically approximate the growth of the epidemic (Ferguson et al., 2003). As argued by Berger et al. (2020), because these models

[4]Li, Rosen and Suen (2001) is an early reference which relies on a similar assumption with similar applied aims, although in a fairly different model.

employ different approaches and do not reach a consensus, balancing their epidemic predictions in order to set the best possible estimation is a difficult task. It requires experience with both epidemics and formal representations, which only experts possess. During the pandemic, this uncertainty *across* models was clearly present. For a decision maker, resolving it requires more than epidemic data; what matters is information about the models themselves.

This study also relates to the literature on cheap talk communication with ambiguity-sensitive preferences. Kellner and Le Quement (2017) were the first to study this question in a simple two-action, two-state setting, with only standard mixed strategies allowed, but an ambiguous prior over the states. Kellner and Le Quement (2018) focus on the classical linear-quadratic example of CS, but allow for Ellsbergian communication strategies. They show that the use of these strategies reduces misalignment between players, creating equilibria which ex-ante Pareto dominate the corresponding ones in CS. These results differ from mine because my model is further away from CS, as communication is about probability distributions, and the maximal payoff of the receiver is sensitive to the state. In addition, as pointed out by Hanany, Klibanoff and Mukerji (2020), the updating assumed in these papers violates sequential optimality. This is an issue I do not face when studying communication about a set of probability distributions. To the best of my knowledge, this paper is the first to study strategic communication about the set of priors of an ambiguity-sensitive decision maker.

Finally, this study also provides a game-theoretical argument in favour of the persistence of multiple prior beliefs when decision making concerns complex scientific topics. This modelling assumption has been increasingly used in the context of climate change management, for instance, by Millner, Dietz and Heal (2013) and Berger, Emmerling and Tavoni (2016), and in the case of model uncertainty, notably in Hansen and Sargent (2001) and Hansen et al. (2006). If decision makers perceive ambiguity regarding complex scientific topics, even under expert advice, ambiguity will not resolve because information transmission is always imprecise.

Section I introduces the framework. Section II clarifies the structure of the value of information in our game, and Section III identifies the consequence of decision making under ambiguity. Section IV establishes the main results in the general setting of Crawford and Sobel (1982). A simple linear quadratic example is provided and used for illustration and supplementary characterisations. Section V discusses the results. Appendix A relaxes some assumptions made in the main text and generalises to a finite number of states. Appendix B contains some of the proofs of the main text.

I. Setup

A. Primitives

I consider a game of communication between an expert acting as a sender S (she), and a decision maker acting as a receiver R (he). Let $\mathcal{A} = \mathbb{R}$ be the set of actions of R and let $\Omega = \{0, 1\}$ be the set of possible states of nature[5]. For $i = S, R$, let $u_i : \mathcal{A} \times \Omega \to \mathbb{R}$ be the von Neumann-Morgenstern utility function of player i, that maps her actions and the state into her welfare. I start by making the following assumptions:

Assumption 1 (**Utilities** - Crawford and Sobel (1982))**.** *u_i is assumed twice continuously differentiable and strictly concave in a. For every $\omega \in \Omega$, there is $a \in \mathbb{R}$ such that $\frac{\partial u_i(a,\omega)}{\partial a} = 0$. For all $a \in \mathbb{R}$, $\frac{\partial u_i(a,\omega)}{\partial a}$ is strictly increasing in ω.*

This assumption implies that u_i admits a unique maximum for each state. Define $a_i(\omega) = \arg\max_{a \in \mathcal{A}} u_i(a, \omega)$ as this maximum. It is the optimal action of player i, under perfect information, that the state is ω. Assumption 1 ensures that $a_i(\omega)$ is strictly increasing in ω. I call $\omega = 1$ ($\omega = 0$) the high (low) state as it is the one where the optimal action is the highest (lowest). Assumption 1 states that for any state, there is a single optimal action. In addition, optimal

[5]Appendix A generalises to a finite number of states.

actions are strictly increasing with the state. Assumption 1 is a single-crossing assumption: it implies that $u_i(,0)$ and $u_i(,1)$ can cross only once over $\mathcal{A}$.

There is ambiguity in the sense that, ex-ante, it is not known which exact distribution the state is drawn from. Instead, R only knows that there is a family of distributions $\mathcal{D} = \{p_\theta | \theta \in [\underline{\theta}, \overline{\theta}]\}$, where $\underline{\theta}, \overline{\theta} \in [0,1]$, containing the state-generating one, where p_θ is the probability mass function of a Bernoulli distribution of parameter θ:

$$p_\theta(\omega) = \begin{cases} \theta \text{ if } \omega = 1 \\ 1 - \theta \text{ if } \omega = 0 \end{cases}$$

Because there is a bijection between the set of probability distributions $\mathcal{D}$ and the set of their parameters $\mathcal{C} = [\underline{\theta}, \overline{\theta}]$[6], I will, for simplicity, specify all communication strategies on $\mathcal{C}$ and call its elements *distributions*. Let $A_i(\theta) = argmax_{a \in \mathcal{A}} \mathbb{E}_\theta(u_i(a, \omega))$ be the optimal action in the eyes of player i under distribution θ, where $\mathbb{E}_\theta(u_i(a, \omega)) = (1 - \theta)u_i(a, 0) + \theta u_i(a, 1)$.

Assumption 2 (**Misalignment under risk**). *For any distribution, the optimal actions of S and R are always misaligned:*

$$A_S(\theta) > A_R(\theta) \textit{ for all } \theta \in \mathcal{C}$$

Assumption 2 states that regardless of the distribution, there is always a difference of interest between S and R, such that optimal actions are ordered in the same way[7]. Note that excluding the case where $A_S(\theta) < A_R(\theta)$ for all $\theta \in \mathcal{C}$ is without loss of generality, as all results are symmetrical.

Finally, notice that the sorting condition over states of Assumption 1 implies the following sorting condition over distributions:

[6] Appendix A extends to parametric families of distributions with a single parameter distributed over a finite number of states

[7] In Appendix A, I show that Assumption 2 is implied by the equivalent assumption made on optimal actions as a function of the state (as in CS), plus an assumption on the ordering of the marginal utility of the actions of both players.

$$\frac{\partial^2 \mathbb{E}_\theta(u_i(a,\omega))}{\partial\theta\partial a} > 0 \tag{1}$$

as :

$$\frac{\partial^2 \mathbb{E}_\theta(u_i(a,\omega))}{\partial\theta\partial a} = \frac{\partial u_i(a,1)}{\partial a} - \frac{\partial u_i(a,0)}{\partial a}$$

and Assumption 1 gives that the latter is strictly positive.

Inequality (1) states that the marginal utility of actions is increasing with θ. As, for a given distribution, the expected utility of actions is single-peaked, it implies that the optimal action of players, $A_i(\theta)$, is a strictly increasing function of θ.

B. Equilibrium concept

Ex-ante, both players are in a situation of ambiguity. In order to model the way R acts under ambiguity, I will consider two separate cases. First, I will consider the case where R evaluates actions under uncertainty through the maxmin decision criteria (MEU) proposed by Gilboa and Schmeidler (1989). According to Gilboa and Schmeidler (1989), in addition to their utility function, players are characterised by a set of priors over Ω, which I will assume to be $\mathcal{C}$. R evaluates action $a \in \mathcal{A}$ by:

$$V_R^{MEU}(a) = \min_{\theta\in\mathcal{C}} \mathbb{E}_\theta(u_R(a,\omega))$$

Second, I consider the case where the receiver's decision making coincides with Savage (1972)'s subjective expected utility (SEU), often identified as a case of ambiguity neutrality. In that case, R's preferences are represented by a utility function and a subjective prior over distributions $\mu \in \Delta(\mathcal{C})$, admitting a probability distribution function g. In order to study a case of communication about distributions which is similar to CS, I will assume that, in this case, R knows the objective distribution from which the distribution is drawn. Thus, μ is an objective distribution, and I also assume that $supp(\mu) = \mathcal{C}$. R then evaluates action a under uncertainty through:

$$V_R^{SEU}(a) = \int_{\theta \in \mathcal{C}} g(\theta)\mathbb{E}_\theta(u_R(a, \omega)))d\theta$$

In the following, the MEU case (the SEU case) is the one where R's evaluation of an action coincides with the MEU (the SEU) decision criteria.

The timing of the game is as follows:

1. Nature draws the state-generating distribution p_{θ_0}, according to μ. S privately observes θ_0[8].
2. S sends a message regarding her type.
3. R updates his beliefs and chooses an action.

Having learned the distribution $\theta_0 \in \mathcal{C}$ corresponding to the state-generating distribution, S sends a message $m \in \mathcal{M}$, where $\mathcal{M} = [0, 1]$ to R. A signalling strategy for S is the strategy $\sigma : \mathcal{C} \rightarrow \mathcal{M}$. An action rule for R is a strategy $y : \mathcal{M} \rightarrow \mathcal{A}$. Notice that I will focus only on pure strategies. Let $\sigma^{-1}(m) \subseteq \mathcal{C}$, be the set of potential types of S, having received message m, when S follows strategy σ. An equilibrium (σ^*, y^*) is defined such that:

1. A sender of type θ evaluates message m by:

$$V_S^\theta(m) = \mathbb{E}_\theta(u_S(y^*(m), \omega))$$

   $\forall \theta \in \mathcal{C}$, any $\sigma^*(\theta) \in \mathcal{M}$ solves $\max_{m \in \mathcal{M}} V_S^\theta(m)$.

[8]This is a simplifying assumption, limiting the scope of the paper. In general the sender could also be assumed to perceive some uncertainty about the data-generating process, while here I assume that her expertise is sufficient to nail down this process.

2. Having received an equilibrium message $m \in supp(\sigma^*)$, an MEU receiver updates his belief such that he evaluates action a by:

$$V_R^{MEU}(a, \sigma^{-1}(m)) = \min_{\theta \in \sigma^{-1}(m)} \mathbb{E}_\theta(u_R(a, \omega)))$$

An SEU receiver is able to update his prior using Bayes' rule such that:

$$g(\theta|m) = \begin{cases} \frac{g(\theta)}{g(\sigma^{*-1}(m))} \text{ if } \theta \in \sigma^{*-1}(m) \\ 0 \text{ if not} \end{cases}$$

R then evaluates action a by:

$$V_R^{SEU}(a, \sigma^{-1}(m)) = \int_{\theta \in \mathcal{C}} g(\theta|m) \mathbb{E}_\theta(u_R(a, \omega)) d\theta$$

In both cases, R chooses action $y^*(m)$ which solves $\max_{a \in \mathcal{A}} V_R^{SEU}(a, \sigma(m))$ (respectively $\max_{a \in \mathcal{A}} V_R^{MEU}(a, \sigma(m))$).

Any message m such that $m \notin supp(\sigma^*)$ is interpreted as some equilibrium message $m_* \in supp(\sigma^*)$[9].

II. Value of information

An important particularity of the game we study here is the part played by the value of information[10]. In order to see how, let us consider the following parametric example:

[9]Notice that this equilibrium concept corresponds to a perfect Bayesian equilibrium (PBE), where the receiver has smooth preferences (Klibanoff, Marinacci and Mukerji, 2005) and a linear ϕ (in the SEU case), or to the limit case of a PBE where $-\frac{\phi''}{\phi'} \to +\infty$ (in the MEU case).

[10]The specific meaning of value of information is clarified in the following section.

Main example:

- $u_S(a, \omega) = -(a - \omega - b)^2 - c\omega$ where $b > 0$ and $c \in \mathbb{R}$
- $u_R(a, \omega) = -(a - \omega)^2 - c\omega$
- $\mathcal{C} = [0, 1]$ and $\mu \sim U(\mathcal{C})$

Then:

$$\begin{cases} A_S(\theta) = \theta + b \\ A_R(\theta) = \theta \end{cases}$$

Although inspired by CS's linear-quadratic example, my main example is fundamentally different. Here, the receiver's interim utility when receiving message m, $\mathbb{E}_\mu(\theta u_R(., \omega = 1) + (1 - \theta)u_R(., \omega = 0)|m)$, is a convex combination of the *expected utility* he gets for each distribution, while, in CS's example, the receiver's valuation function is a convex combination of the *utility* he gets for each state directly. An important consequence is that, in CS's example, the maximal utility R gets in each state is the same, making it independent of the sender's type. In my example though, R's maximal expected utility is never constant in the sender's type. Depending on the value of c, it can either be monotone in θ, or *V shaped*: decreasing then increasing. When $c = 0$, both states are equivalent, in the sense that the receiver could achieve exactly the same pay-off in both of them. Then, as illustrated by Figure 1, the maximal expected welfare of the receiver is at its lowest when $\theta = 0.5$, because it does not allow the receiver to skew the decision in favour of either state. Any information that moves the prior to the right or to the left of $\theta = 0.5$ would increase R's maximal expected pay-off. For instance, $\theta = 0.2$ and $\theta = 0.8$ both improve the receiver's maximal expected utility compared with $\theta = 0.5$. Any additional information thus has positive value because it helps the receiver to skew the decision towards the more likely state. In other words, starting from prior $\theta = 0.5$, any information is *good news*.

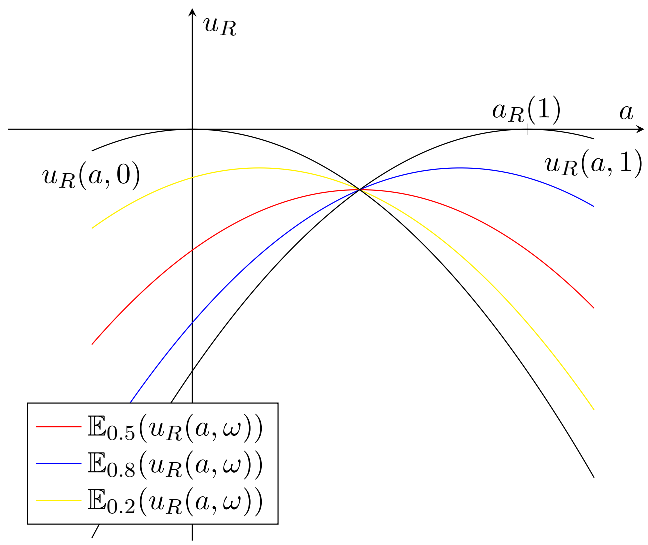

Figure 1. : $c = 0$ a case of comparable states. For distributions above 0.5, the receiver's maximal welfare is increasing with the probability of the high state. For distributions below 0.5, the opposite happens.

In fact, as long as $c \in (-1, 1)$, no state is dominated, in the sense that no state gives a higher utility, whatever the receiver's action. Then, at least for some prior, any information has positive value, as, by increasing the probability of one or the other state, it increases R's maximal expected utility. To see this, notice that when $c \in (-1, 1)$, R's maximal expected welfare is at its worst when $\theta = \frac{1+c}{2} \in (0, 1)$-that is, when odds and stakes between states perfectly balance each other out in terms of maximal expected utility for R. It follows that under the prior $\theta = \frac{1+c}{2}$, any information has positive value, as it would allow R to skew the decision in favour of a state. But when $c \geq 1$ or $c \leq -1$ this is not the case any more. As illustrated by figure 2, one state is now dominated by the other. Any information that increases its probability has negative value and is thus *bad news*, whatever the prior. It follows that the maximal expected utility of the receiver is strictly monotone in θ.

The recent COVID-19 crisis exemplifies a case of dominated states. Consider the example of a political decision-maker in charge of choosing a level of social restriction. During the early stages of the pandemic, information was too scarce to exclude the possibility that, even under optimal restriction policies, the virus would be so deadly that the outcome would be worse than in any other case. That scenario was, clearly, a dominated state. Other situations are cases of comparable states. Consider a farmer envisioning the possibility of a transition from conventional agriculture to agroecology. Prior to the transition, he has to choose a selling price for his products. There is uncertainty regarding the cost of transitioning from one method to the other and, as a result, about the optimal price he should choose under the agroecological method[11]. For the farmer, a fully functional agroecological farming system can generate comparable levels of profit to those provided by conventional methods, if he sells his product at the right (higher) price. Thus, both systems are comparable states for him.

The properties of the value of information in the main example will transfer to the general case. Let us define the following elements to state this.

Definition 1. *Define* $\widetilde{a} = argmax_{a\in\mathcal{A}} \min_{\omega\in\Omega} u_R(a,\omega)$ *as the precautionary action, and* $\widetilde{\theta} \in [0,1]$ *such that* $A_R(\widetilde{\theta}) = \widetilde{a}$ *as the most pessimistic distribution.*

I call $\widetilde{a}$ the precautionary action because it is the optimal action anticipating the worst possible state. $\widetilde{\theta}$ is the distribution for which the precautionary action is the optimal action[12]. From these definitions, it is straightforward that for $\widetilde{\theta} \in \mathcal{C}$ it is necessary that states are comparable[13]. As the main example suggests, and as Proposition 1 will prove, the maximal expected utility of the receiver decreases for distributions putting lower weight on the high state than the most pessimistic one ($\theta < \widetilde{\theta}$) and increases for the others ($\theta > \widetilde{\theta}$). Therefore, any information that

[11]Here each potential state is defined by potential transition costs and the decision maker's action variable is the selling price.

[12]The fact that $\widetilde{\theta}$ exists and is unique is proven in Lemma 3.

[13]When $\mathcal{C} = [0,1]$ the fact that states are comparable is necessary and sufficient for $\widetilde{\theta} \in \mathcal{C}$.

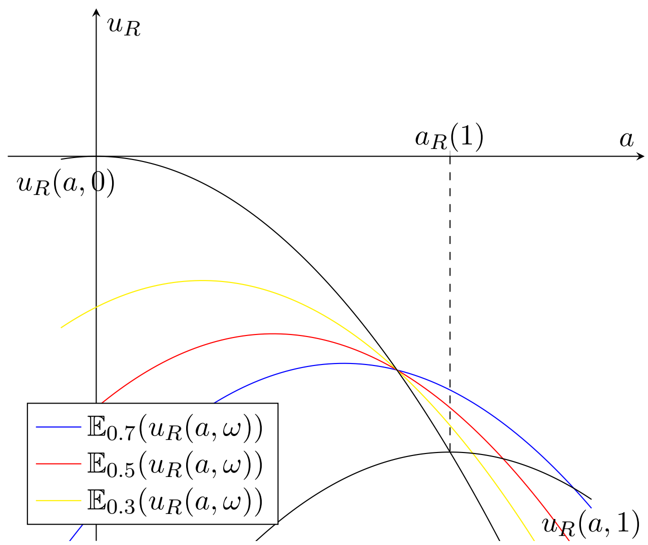

Figure 2. : $c = 1$, a case of dominated state. For all distributions, the receiver's maximal welfare is decreasing with the probability of the high state.

increases the probability of the high state, but leads to a posterior below $\widetilde{\theta}$, is bad news. In the special case where all elements of $\mathcal{C}$ are below $\widetilde{\theta}$, any information that increases the probability of the high state is bad news. Conversely, if all elements of $\mathcal{C}$ are above $\widetilde{\theta}$, any information that increases the probability of the high state is good news.

III. Implications for decision making

Now that we have a clear view of the value of information in our game, let us see how it affects decision making.

In the SEU case, the receiver will choose an action that is based on his average expected utility for the interval of probabilities communicated in equilibrium by the sender[14]. Some of these distributions may be good news for R, relative to the prior, and others may be bad news. However, because of the linearity of the ex-

[14]That is: $\mathbb{E}_{\hat{\theta}}(u_R(.,\omega))$ where $\hat{\theta} = \int_{\theta\in\sigma^{-1}(m)} \mu(\theta)\theta d\theta$.

pectation, this will play no role in R's choice of action. In other words, R's actions are not sensitive to the change in monotonicity of his maximal expected welfare, that is, to the value of information. Under SEU preferences, equilibria are in fact very similar to CS. One can equate each distribution with a state in CS's setting, where the corresponding payoff is the expected utility under that distribution and μ is the prior over states. All equilibria are partitional equilibria,[15] and the tightness of the partition depends only on the level of misalignment between parties. When the latter is small, information transmission can be almost perfect over the entire set of types. Multiple equilibria can exist and, once the sender learns her type, she does not always have an interest in being as informative as she could be.

In contrast, the behavioural response of an MEU receiver will be of a different nature, depending on whether the information is good or bad news. This is because his action will be based on the worst-case scenario within the interval of distributions communicated. Proposition 1 will formalise R's optimal response under MEU preferences. For $B \subset \mathcal{C}$, let $A_R(B) \subset argmax_{a \in \mathcal{A}} \min_{\theta \in B} \mathbb{E}_\theta(u_R(a,\omega))$ be the set of optimal actions of a MEU receiver given the set of priors B.

Proposition 1. *Define $B = [\theta_1, \theta_2] \subset \mathcal{C}$ as the set of priors of the receiver. Given that $\theta_0 \in B$, an MEU receiver has a unique optimal action which is given by:*

$$A_R(B) = \begin{cases} A_R(\theta_2) \ \textit{if} \ \theta_2 < \widetilde{\theta} \\ A_R(\widetilde{\theta}) \ \textit{if} \ \widetilde{\theta} \in B \\ A_R(\theta_1) \ \textit{if} \ \theta_1 > \widetilde{\theta} \end{cases}$$

Proposition 1 states that an MEU receiver has a unique optimal action for any belief $\theta_0 \in B$ where B is an interval of $\mathcal{C}$[16]. When he further believes that

[15] A formal definition is given below.
[16] Notice that this result also implies that the single peakedness assumption on utility functions trans-

all distributions are below $\widetilde{\theta}$ ($\theta_0 \in [\theta_1, \theta_2]$ and $\theta_2 < \widetilde{\theta}$), he optimally acts as if the probability of the high state were maximal. When he believes that all distributions are above $\widetilde{\theta}$ ($\theta_1 > \widetilde{\theta}$), he optimally acts as if the probability of the high state were minimal. But when $\widetilde{\theta}$ is in the interior of B, R will always act as if the probability of the high state were maximal for beliefs below $\widetilde{\theta}$ ($\theta_2 < \widetilde{\theta}$) and minimal for beliefs above $\widetilde{\theta}$ ($\theta_1 > \widetilde{\theta}$). When R believes that the worst possible distribution could be the state-generating distribution ($\widetilde{\theta} \in [\theta_1, \theta_2]$), he optimally acts as if it were the case.

It follows that, given the sender's desire for larger actions, it becomes impossible for her to communicate good news, that is to point out distributions which would give R a higher payoff than the most pessimistic distribution. The reason is that when we are on the upward-sloping part of the V, R will interpret any interval on that part as the left-most belief ($\widetilde{\theta}$), and, because of this extreme non-responsiveness to messages sent, conveying good news becomes impossible (as opposed to SEU, where the expectation is always interior, and so R responds more). At the same time, though, it becomes easier to communicate bad news when we are on the downward-sloping part of the V, because, now, a receiver with MEU preferences will respond more than under SEU. Instead of an interior belief, the MEU decisio maker will respond as if the distribution were the right-most element of the interval, which is closer to what S wants.

Consider the COVID-19 example mentioned above, and assume an epidemiology expert communicating about the virus deadliness to a political decision maker in charge of choosing a level of social restriction. Assume that the decision maker is known to be more reluctant to impose social restrictions than the expert. Given their misalignment, under SEU preferences, an expert advocating in favour of strong social restrictions would have had limited influence. Yet, in practice, policy makers agreed to impose unprecedented social restrictions almost entirely under expert advice. This paper's model offers a rationale for this observation by

fers to the sender's valuation functions, both in the SEU and MEU cases. As a result, mixed strategies are always dominated and restricting attention to pure strategies is without loss of generality.

assuming that the decision makers had MEU preferences which led them to react very strongly to experts announcing bad news regarding the pandemic.

In the farmer's example, assume an environmental agency counsels the farmer on his method transition. There is an interest misalignment because the agency cares mostly about the environment, while the farmer accounts more for his own economic profitability. In addition, how fast and costly the transition will be is uncertain, even to the agency, because it depends on the ecological state of the farmer's soil, the biological response of the environment to the new agricultural structure, and the weather variability[17]. The effects of ambiguity aversion we discussed above explain why it is very difficult for the agency to convince the farmer to pursue an ambitious transition, even if the latter cares almost as much as the agency about his environmental impact.

The following section will formalise the intuitions given in this section more generally.

IV. Equilibrium analysis

A. *Results in the general setting*

Let us now turn to the study of the game's equilibria. First, I introduce the following definition:

Definition 2. *Set* $\{\theta_0, ..., \theta_q\} \subseteq \mathcal{C}$ *such that:*

- $\underline{\theta} = \theta_0 < ... < \theta_q = \overline{\theta}$ *where* θ_k*, for* $0 \leq k \leq q$*, is called the k-th cut-off.*
- $\cup_{k=1}^{q} [\theta_{k-1}, \theta_k] = [\underline{\theta}, \overline{\theta}]$*, where* $[\theta_{k-1}, \theta_k)$*, for* $1 \leq k < q-1$*, is called the k-th cell and* $[\theta_{q-1}, \overline{\theta}]$ *the q-th cell.*

A q-cut-off partition equilibrium is an equilibrium of the game where the signaling strategy of S is uniform on every cell. That is, for $\theta \in [\theta_{k-1}, \theta_k)$*,* $\sigma^*(\theta) = m_{k-1}$ *and for* $1 \leq k \leq q-1$ *and* $\theta \in [\theta_{q-1}, \overline{\theta}]$*,* $\sigma^*(\theta) = m_{q-1}$*.*

[17] I thank my friend Arielle Zoellin for giving me the idea of this example.

A q-cut-off partition equilibrium is an equilibrium where there is a partition of the set of types in q cells. For any cell of this partition, any sender who is in that cell credibly sends the same message to the receiver. Having received that message, the receiver learns which cell the sender is in, and acts optimally.

Proposition 2. *In every equilibrium of the game, there is a partitioning of $\mathcal{C}$ in a finite number of cells, where every cell induces a distinct action. Thus, any equilibrium is outcome equivalent to a partition equilibrium.*

The proof of Proposition 2 is relegated to the Appendix. It starts by showing that the number of actions induced in equilibrium is finite. The argument is similar to the one given in CS's Lemma 1 and follows from both the concavity of S's evaluation of actions and the fact that the optimal action of R for a given belief $B \subset \mathcal{C}$ is in the convex hull of the optimal actions for every element of B. Then I show that types that induce a given action must form an interval. This is a consequence of the concavity of S's evaluation of actions.

Proposition 2 shows that there is a finite partition of $\mathcal{C}$, where types in every cell induce a given action from the receiver. This does not imply that types in every cell send the same message, as it is possible that different messages induce the same action. As a result, every equilibrium is not necessarily a partition equilibrium, but must be outcome equivalent to one. In the following, I focus only on partition equilibria. Notice that there is always at least one partition equilibrium: the babbling equilibrium, where all types send the same message.

In the following, I give a characterisation of all partition equilibria of the game.

Proposition 3. *In any partition equilibrium of the game (σ_q^*, y^*), the cut-off types $\theta_0^q, ..., \theta_q^q$ are defined such that for $k \in 1, ..., q$:*

$$V_S^{\theta_k^q}(y^*(m_{k-1}^q)) = V_S^{\theta_k^q}(y^*(m_k^q)) \tag{2}$$

where m_k^q *is the equilibrium message of types* $\theta \in [\theta_k^q, \theta_{k+1}^q]$.

The proof of Proposition 3 is relegated to the Appendix. Figure 3 represents the interim utility of S when her type is θ_k. As a convex combination of concave and single-peaked functions, it is concave and maximal at $A_S(\theta_k)$. Figure 3 illustrates that m_{k-1} and m_k are equilibrium messages because they induce actions that give the same level of welfare to S. As a result, θ_k is a cut-off type.

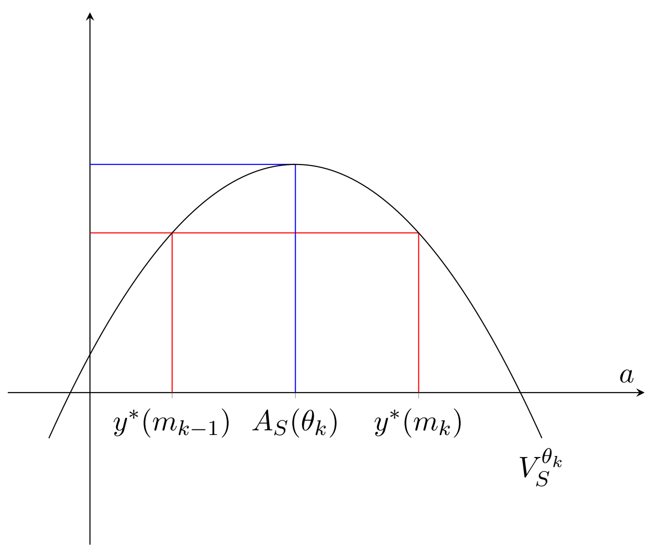

Figure 3. : Identifying cut-offs

I now state the first main result of the paper: no information transmission is possible for types above the worst possible distribution.

Theorem 1. *When the receiver has MEU preferences, all cut-offs in* $(\underline{\theta}, \overline{\theta})$ *are below* $\widetilde{\theta}$.

Proof of Theorem 1

Assume there is a $q+1$ cut-off equilibrium and that $\theta_{q-1} < \widetilde{\theta} \leq \theta_q$. Recall the characterisation result of partition equilibria given by Proposition 3. For θ_q to be a cut-off type, the messages sent by types in the cell below and above θ_q must induce actions that give the same utility to a sender of type θ_q. If θ_q was a cut-off type, it would follow from Proposition 1 that:

$$\begin{cases} y^*(m_{q-1}) = A_R(\sigma^{*-1}(m_{q-1})) = A_R([\theta_{q-1},\theta_q)) = A_R(\widetilde{\theta}) = \widetilde{a} \\ y^*(m_q) = A_R(\sigma^{*-1}(m_q)) = A_R([\theta_q,\theta_{q+1})) = A_R(\theta_q) \end{cases}$$

As A_R is a strictly increasing function, and because S is upwards misaligned, we have that $y^*(m_{q-1}) < y^*(m_q) < A_S(\theta_q)$. As, by definition, $a \to \mathbb{E}_\theta(u_S(a,\omega))$ is strictly increasing on $[0, A_S(\theta_q)]$, we have that:

$$\mathbb{E}_{\theta_q}(u_S(y^*(m_{q-1}),\omega)) < \mathbb{E}_{\theta_q}(u_S(y^*(m_q),\omega)) \iff V_S^{\theta_q}(m_{q-1}) < V_S^{\theta_q}(m_q)$$

which is in contradiction to the assumption that θ_q is a cut-off type.

□

As illustrated by Figure 4, the utility of the sender induced by m_{q-1} is always lower than that induced by m_q. This is a direct consequence of the change in the monotonicity of R's maximal expected welfare at $\widetilde{\theta}$. When R believes that the worst possible distribution could be the state-generating one, he optimally acts as if it were the case. When he believes that $\theta_0 \in [\theta_q, \theta_{q+1})$ and $\theta_q > \widetilde{\theta}$, he will act as if the distribution were θ_q. As a result, because S is misaligned upwards, we have that:

$$\widetilde{a} < A_R(\theta_q) < A_S(\theta_q)$$

and as $V_S^{\theta_q}$ is strictly increasing for $a \leq A_S(\theta_q)$, it is impossible for messages sent by types in the cell below and above θ_q to induce actions that give the same utility to a sender of type θ_q. As a result, the indifference between actions induced by messages m_{q-1} and m_q needed for θ_q to be a cut-off type (as displayed in Figure

3) is impossible.

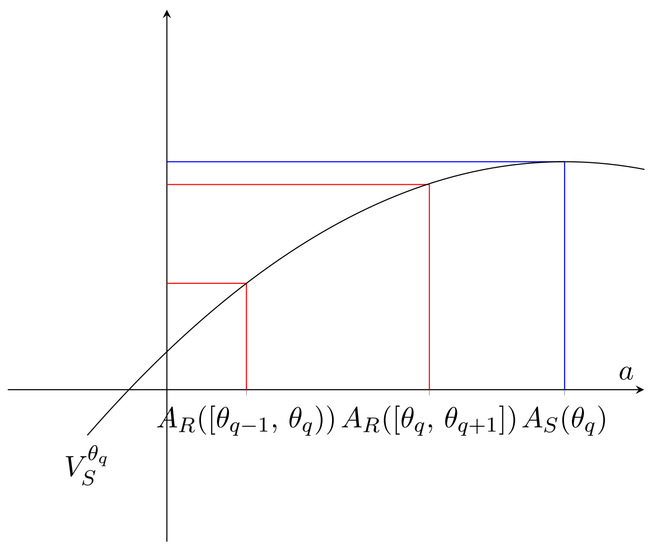

Figure 4. : MEU best responses for $\theta_{q-1} < \widetilde{\theta} < \theta_q$

A consequence of Theorem 1 is that when one state is dominated and when $\widetilde{\theta} \leq \underline{\theta}$, the only equilibrium is the babbling equilibrium. That is, whatever the sender's type, whatever the message she sends, the induced action is always the same. That this result holds even when misalignment is arbitrarily small is a consequence of the somewhat radical ambiguity aversion of an MEU decision maker. In a companion paper (Colo, 2023), I show that this result extends to α-MEU preferences, and prove a more *continuous* result: for any level of misalignment $b > 0$, there is a corresponding minimal level of ambiguity aversion $\frac{1}{2} < \alpha(b) < 1$ above which no cut-off type exists in $(\widetilde{\theta}, \overline{\theta})$.

Before moving to my second main result, I first need to state an important intermediate result.

Lemma 1. *When the receiver evaluates actions following the MEU criteria, there are $M > 0$ partition equilibria. Call $\theta_0 < ... < \theta_M$ the cut-offs of the equilibrium with the most cut-offs. Then the q cut-off partition equilibrium is defined by cut-*

offs $\theta_0 < \theta_{M-q} < ... < \theta_M$*, for* $0 \leq q \leq M$.

The proof of Lemma 1 is relegated to the Appendix. As illustrated by Figure 5, Lemma 1 states that all equilibria of our game can be built from the same set of cut-off types. More specifically, it states that if one considers the equilibrium with the most cut-offs, one can describe all other equilibria by successively removing cut-offs starting from the left.

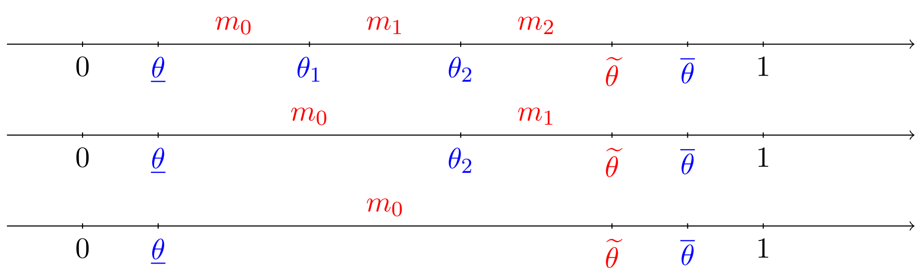

Figure 5. : MEU equilibria for $\underline{\theta} < \widetilde{\theta} < \overline{\theta}$

To see why Lemma 1 is true, first note that, given Theorem 1, all interior cut-offs are in $[\underline{\theta}, \widetilde{\theta}]$. As a result, when S points out an interval of distributions, R only cares about its upper bound. Thus, cut-offs types will not be determined by an indifference between two adjacent cells of distributions as in the SEU case, but by an indifference between the distributions at the upper bound of these cells. In the former case, each indifference condition depends on three distinct types and the prior. Thus, in order to determine the cut-off types, the entire sequence of indifference conditions is needed. In the latter case, each indifference condition depends on two distinct types only. Given that $[\underline{\theta}, \widetilde{\theta}]$ is a closed interval, it is then possible to find the first cut-off starting from $\underline{\theta}$ and then to iterate the process to find the following ones. In doing so, I derive the cut-off types of the equilibrium that has the most cut-offs. Call the corresponding number of cut-offs M. Any signalling strategy of the sender characterised by the q first terms $(1 \leq q \leq M)$ of that sequence induces exactly the same incentive constraints for the receiver.

This implies that they form part of an equilibrium.

A direct consequence of Proposition 1 is that all equilibria of the game can be ranked by informativeness, in the Blackwell sense-something which is never possible in the SEU case.[18] The following result can thus be established regarding interim dominance for the sender among equilibria.

Theorem 2. *When the receiver has MEU preferences, the sender is always interim weakly better off by playing the most informative*[19] *equilibrium strategy.*

Proof of Theorem 2:

Assume the equilibrium with the most cut-offs has M elements. For any $1 \leq q \leq M$ let a q cut-off equilibrium be characterised by S's strategy σ_q^* and elements $\theta_0, \theta_{M-q}, ..., \theta_M$.

First I will show that S is interim better off in the $q+1$ cut-off equilibrium than in the q cut-off equilibrium. Then a simple iteration gives that S is better off in the M cut-off equilibrium than in the q cut-off equilibrium, for any $q < M$.

- Assume $\theta_0 \in [\theta_q, \overline{\theta}]$.

  Then, S's interim utility in the $q+1$ cut-off equilibrium and in the q cut-off equilibrium is $\mathbb{E}_{\theta_0}(u_S((A_R(\widetilde{\theta}))$. Thus S is indifferent between both equilibria.

- Assume $\theta_0 \in [\theta_k, \theta_{k+1}]$, for $M - q \leq k \leq M$.

  Then, S's interim utility in the $q+1$ cut-off equilibrium and in the q cut-off equilibrium is $\mathbb{E}_{\theta_0}(u_S((A_R(\theta_{M-q-1}))$. Thus S is indifferent between both equilibria.

- Assume $\theta_0 \in [\underline{\theta}, \theta_{M-q-1}]$, for $M - q \leq k \leq M$.

[18]The informativeness ranking comes from the fact that when receiving $m \in \mathcal{M}$ from a type in $[\theta_1, \theta_2]$ with $\theta_2 < \overline{\theta}$ an MEU receiver acts exactly as when receiving $m' \in \mathcal{M}$ from a type in $[\theta_1', \theta_2]$, for any $\theta_1' < \theta_1$. For an SEU receiver, this behavioural pattern is impossible; the optimal action would necessarily shift to the left.

[19]In the sense of Blackwell.

Then, S's interim utility in the q+1 cut-off equilibrium is $\mathbb{E}_{\theta_0}(u_S((A_R(\theta_{M-q-1}))$ and S's interim utility in the q cut-off equilibrium is $\mathbb{E}_{\theta_0}(u_S((A_R(\theta_{M-q})).$

Yet, because θ_{M-q-1} is a cut-off type in the $q+1$ cut-off equilibrium, for any $\theta \in [\underline{\theta}, \theta_{M-q-1})$,

$$\mathbb{E}_{\theta_0}(u_S((A_R(\theta_{M-q-1})) > \mathbb{E}_{\theta_0}(u_S((A_R(\theta_{M-q}))$$

Thus, any type of sender in $[\underline{\theta}, \theta_{M-q-1})$ is interim better off in the $q+1$ cut-off equilibrium than in the q cut-off equilibrium.

□

The intuition of the proof is the following. Consider for instance the equilibria described in figure 5. Whatever the equilibrium considered, types in $[\theta_1, \widetilde{\theta}]$ will induce the same action $\widetilde{\theta}$. But types in $[\underline{\theta}, \theta_1]$ will induce action $\widetilde{\theta}$ in the babbling equilibrium, and θ_1 in the 3-cut-off equilibrium. Yet, by construction of the latter equilibrium, all types in $[\underline{\theta}, \theta_1]$ prefer to induce action θ_1 than $\widetilde{\theta}$. It follows that for the sender, the 3-cut-off equilibrium interim dominates the babbling equilibrium. The same reasoning can be applied regarding types in $[\underline{\theta}, \theta_2]$ to show that the 4-cut-off equilibrium interim dominates the 3-cut-off one.

B. Characterisations on the main example

In order to give a further insight into the results in the MEU case, I characterise all partitional equilibria in the context of the parametric example introduced above. I also provide the same characterisation for the SEU case.

Proposition 4. *In the context of our linear-quadratic example with uniform prior, for any $c \in \mathbb{R}$:*

- *When R has SEU preferences, an n-cut-off equilibrium exists if and only if:*

$$0 < b < \frac{1}{2n(n+1)} \tag{3}$$

and, for $k \in 1, ..., n,$ *cut-offs are:*

$$\theta_k = \frac{k}{n+1} - 2kb(n-k+1)$$

- *When R has MEU preferences, an n-cut-off equilibrium exists if and only if* $c > -1$ *and*

$$0 < b < \frac{1}{2n} \tag{4}$$

and, for $k \in 1, ..., n,$ *cut-offs are:*

$$\theta_k = 1 - 2b(n-k)$$

The proof of Proposition 4 is relegated to the Appendix. Proposition 4 shows that the value of c has no influence on communication in the SEU case. Yet, in the MEU one, when $c \leq -1$, the maximal pay off in state 0 is always lower than in state 1. As a result, the attraction exerted by ambiguity aversion plays against the sender's communication possibilities and no non-babbling equilibrium is possible. Conversely, when $c \geq 1$ the attraction exerted by ambiguity aversion plays in favour of the sender's communication possibilities and cut-offs can be on the entire set $\mathcal{C}$.

A corollary of Proposition 4 is that it is possible to characterise each equilibrium's cell sizes.

Corollary 1. *Consider a q-cut-off partition equilibrium. When R is SEU, for any $c \in \mathbb{R}$, cells are increasing in size. For all $k \in 1, ..., q-1$:*

$$\theta_{k+1} - \theta_k = \theta_k - \theta_{k-1} + 4b$$

When R has MEU preferences and $c > -1$, non-terminal cells are of constant size. For all $k \in 1, ..., q-2$:

$$\theta_{k+1} - \theta_k = 2b$$

where the cell containing $\underline{\theta}$ is called the terminal cell.

Proof of Corollary 1:

It is possible to derive from Proposition 4 that in the SEU case:

$$\theta_{k+1} - \theta_k = \theta_k - \theta_{k-1} + 4b$$

It is also possible to derive from Proposition 4 that in the MEU case:

$$\theta_{k+1} - \theta_k = 2b$$

□

In the MEU case non-terminal cells always have the same size ($2b$), whatever the equilibrium considered. In the SEU case, it depends on the considered equilibrium. This illustrates why the general result proved in Proposition 1 holds. If all non-terminal cells have the same size in any given equilibrium, and if in addition the first cut-off is always the same (as proven in Proposition 4), it is straightforward that those equilibria can be ranked by informativeness in the Blackwell sense. Corollary 1 also states that in the SEU case, cells are at least of size $4b$ and are thus always strictly larger.

The sender is able to induce a finer partition of types when the receiver is MEU. Consider a given positive bias, such that it is possible to get an n-cut-off equilibrium with an MEU receiver; then it is not always possible to sustain an n-cut-off equilibrium with an SEU receiver. More precisely: call the supremum of the bias for which an n-cut-off equilibrium is possible in the MEU case $b_M(n) = \frac{1}{2n}$. Call the equivalent value of the bias in the SEU case $b_S(n) = \frac{1}{2n(n+1)}$. Both functions are increasing in n. In addition, for $n \geq 2$, $b_S(n) = b_M(n(n+1))$. Thus, there is an n-cut-off equilibrium between an SEU receiver of bias b and the sender if and only if there is an $n(n+1)$-cut-off equilibrium between an MEU receiver of bias b.

V. Discussion

This paper models the transmission of expert-based scientific knowledge as cheap talk communication about probability distributions, in a framework similar to Crawford and Sobel (1982). I make two assumptions. First, that the state is unknown, even to the sender. Second, that the receiver perceives multiple potential data-generating processes. It follows that is natural to assume that the receiver is ambiguity sensitive. For every preference considered, I showed that all equilibria are outcome equivalent to partitional equilibria. When the receiver is MEU, information transmission can only happen for distributions below a given threshold, even if misalignment is arbitrarily small. In addition, the sender always prefers to convey as much information as possible, since the most informative equilibrium is interim dominant for the sender. This is not true when the receiver has SEU preferences, a case which is equivalent to the model of communication about states proposed in Crawford and Sobel (1982). In the linear-quadratic example I introduced, more cut-offs can exist in the MEU case than in the SEU one, for a given bias. This shows that when the expert's preferred action is aligned with the effect of ambiguity aversion, her influence is extremely high; in the opposite case however, it is nonexistent.

Equilibrium selection. Theorem 2 gives that S is always interim better off by adopting the most informative equilibrium strategy in her communication. This result differs significantly from those obtained in CS's framework. Under their monotonicity condition (M), CS show that the ex-ante expected payoffs for both sender and receiver are maximal for the equilibrium with the most cut-offs. Condition (M) is satisfied if, for any two sequence of cut-off types the k-th cut-off of each sequences can be ordered in the same direction, for any $k \geq 1$. This assumption is in particular verified by the linear-quadratic example. The resulting selected equilibrium is often the one studied in applications. Yet, as already pointed out in CS, ex-ante Pareto dominance is a questionable equilibrium-selection criterion, since once having learned their type, different sender types will necessarily have opposed preferences. CS suggest that ex-ante Pareto dominance could be retained only if there is an equilibrium-selection agreement made ex-ante between players or if it can be seen as a convention maintained over repeated plays with several opponents. An alternative approach regarding equilibrium selection has been presented by Chen, Kartik and Sobel (2008), who propose a condition on utility functions called NITS. Under this condition, combined with Assumption (M), only the equilibrium with most cut-offs survives in CS's framework. An equilibrium satisfies NITS if the sender of the lowest type weakly prefers the equilibrium outcome to the outcome induced by credibly revealing her type (if she could). In my case, one could adopt interim dominance for the sender and ex-post dominance for the receiver as selection criteria, which are immune to the limitations of ex-ante Pareto dominance and do not require supplementary assumptions as NITS does. Nevertheless, doing this brings out the same (most informative) equilibrium, and provides a foundation for the attention it receives in applications.

Extension to α-MEU preferences. In a companion paper (Colo, 2023), I extend some of these results to α-MEU preferences (Ghirardato et al., 2004), and show that they are, in a certain sense, robust to a varying degree of ambiguity aversion. Building on the parametric example presented above, I show that, for any level of misalignment of the sender, there is a degree of ambiguity aversion of the receiver $\alpha \in (\frac{1}{2}, 1)$ such that all distributions above $\widetilde{\theta}$ must pool. This suggests a form of continuity in the division of the set of types - on both sides of the worst possible distribution - that we have observed in the MEU case.

REFERENCES

Andrews, Isaiah, and Jesse M Shapiro. 2021. "A model of scientific communication." *Econometrica*, 89(5): 2117—-2142.

Banerjee, Abhijit V, Sylvain Chassang, Sergio Montero, and Erik Snowberg. 2020. "A theory of experimenters: Robustness, randomization, and balance." *American Economic Review*, 110(4): 1206–30.

Berger, Loïc, Johannes Emmerling, and Massimo Tavoni. 2016. "Managing catastrophic climate risks under model uncertainty aversion." *Management Science*, 63(3): 749–765.

Berger, Loïc, Nicolas Berger, Valentina Bosetti, Itzhak Gilboa, Lars Peter Hansen, Christopher Jarvis, Massimo Marinacci, and Richard Smith. 2020. "Uncertainty and decision-making during a crisis: How to make policy decisions in the COVID-19 context?" *University of Chicago, Becker Friedman Institute for Economics Working Paper*, , (2020-95).

Bradley, Richard, Casey Helgeson, and Brian Hill. 2017. "Climate change assessments: confidence, probability, and decision." *Philosophy of Science*, 84(3): 500–522.

Chen, Ying, Navin Kartik, and Joel Sobel. 2008. "Selecting cheap-talk equilibria." *Econometrica*, 76(1): 117–136.

Colo, Philippe. 2023. "Communicating over Uncertain Science: the α-MEU case." *Forthcoming.*

Crawford, Vincent P, and Joel Sobel. 1982. "Strategic information transmission." *Econometrica*, 1431–1451.

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

Ellsberg, Daniel. 1961. "Risk, ambiguity, and the Savage axioms." *The Quarterly Journal of Economics*, 643–669.

Ferguson, Neil M, Matt J Keeling, W John Edmunds, Raymond Gani, Bryan T Grenfell, Roy M Anderson, and Steve Leach. 2003. "Planning for smallpox outbreaks." *Nature*, 425(6959): 681–685.

Ghirardato, Paolo, Fabio Maccheroni, Massimo Marinacci, et al. 2004. "Differentiating ambiguity and ambiguity attitude." *Journal of Economic Theory*, 118(2): 133–173.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin expected utility with non-unique prior." *Journal of Mathematical Economics*, 18(2): 141–153.

Hanany, Eran, Peter Klibanoff, and Sujoy Mukerji. 2020. "Incomplete information games with ambiguity averse players." *American Economic Journal: Microeconomics*, 12(2): 135–87.

Hansen, LarsPeter, and Thomas J Sargent. 2001. "Robust control and model uncertainty." *American Economic Review*, 91(2): 60–66.

Hansen, Lars Peter, Thomas J Sargent, Gauhar Turmuhambetova, and Noah Williams. 2006. "Robust control and model misspecification." *Journal of Economic Theory*, 128(1): 45–90.

Kellner, Christian, and Mark T Le Quement. 2017. "Modes of ambiguous communication." *Games and Economic Behavior*, 104: 271–292.

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

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. "A smooth model of decision making under ambiguity." *Econometrica*, 73(6): 1849–1892.

Li, Hao, Sherwin Rosen, and Wing Suen. 2001. "Conflicts and common interests in committees." *American Economic Review*, 91(5): 1478–1497.

Mastrandrea, Michael D, Christopher B Field, Thomas F Stocker, Ottmar Edenhofer, Kristie L Ebi, David J Frame, Hermann Held, Elmar Kriegler, Katharine J Mach, Patrick R Matschoss, et al. 2010. "Guidance note for lead authors of the IPCC fifth assessment report on consistent treatment of uncertainties."

Millner, Antony, Simon Dietz, and Geoffrey Heal. 2013. "Scientific ambiguity and climate policy." *Environmental and Resource Economics*, 1–26.

Savage, Leonard J. 1972. *The Foundations of Statistics.* Courier Corporation.

Schwartzstein, Joshua, and Adi Sunderam. 2021. "Using models to persuade." *American Economic Review*, 111(1): 276–323.

Spiess, Jann. 2018. "Optimal estimation when researcher and social preferences are misaligned." *Job Market Paper*.

Supplementary Assumptions

A1. Assumptions on states

In the following I show that Assumption 2 is implied by the two following assumptions.

Assumption 3 (**Misalignment under perfect information** - Crawford and Sobel (1982))**.** *The optimal actions of S and R are always misaligned:*

$$a_S(\omega) > a_R(\omega) \text{ for all } \omega \in \Omega$$

Assumption 3 states that whatever the state, there is always a difference of interest between S and R such that optimal actions are ordered the same way.

Assumption 4 (**Sharpness**)**.** *Whatever the sate, the sender has sharper preferences than the receiver, for every action $a \in \mathcal{A}$*

$\forall a \in \mathcal{A}$, $\frac{\partial u_R(a,\omega)}{\partial a} < \frac{\partial u_S(a,\omega)}{\partial a}$

Assumption 4 is a more technical assumption on the players utility function. I assume that the player with highest optimal action in a given state has a more concave utility function in that state, as illustrated by Figure A1. I call that property *sharpness*, in the sense that it translates into a sharper preference for the optimal action.

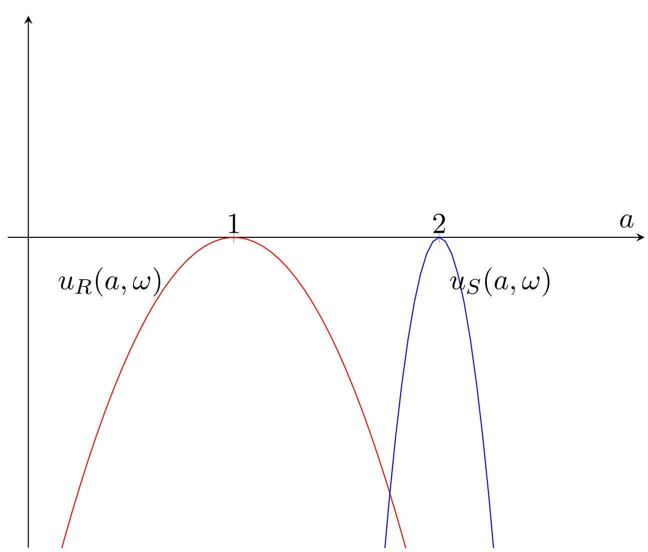

Figure A1. : Sharpness Assumption

Given Assumptions 3 and 4, I now show that both players optimal actions are never aligned, whatever the distribution.

Lemma 2. *Assumptions 3 and 4 imply that:*

$$A_S(\theta) < A_R(\theta) \textit{ for all } \theta \in \mathcal{C} \textit{ or } A_S(\theta) > A_R(\theta) \textit{ for all } \theta \in \mathcal{C}$$

Proof of Lemma 2:

For player i and any $\theta \in \mathcal{C}$, define $f_i^\theta : a \to (1-\theta)\frac{\partial u_i(a,0)}{\partial a} + \theta\frac{\partial u_i(a,1)}{\partial a}$. f_i^θ is a continuous and decreasing function crossing the x-axis only once, at $A_i(\theta)$. I want to prove that for all $\theta \in \mathcal{C}$, $A_R(\theta) < A_S(\theta)$. In order to do so, it is enough to prove that for any $\theta \in \mathcal{C}$, $f_R^\theta(a) < f_S^\theta(a)$. Set $h^\theta : a \to f_R^\theta(a) - f_S^\theta(a)$.

$$h^\theta(a) = (1-\theta)(\frac{\partial u_R(a,0)}{\partial a} - \frac{\partial u_S(a,0)}{\partial a}) + \theta(\frac{\partial u_R(a,1)}{\partial a} - \frac{\partial u_S(a,1)}{\partial a})$$

Thus, by Assumption 4, for all $a \in \mathcal{A}$, $h^\theta(a) < 0$.

□

Lemma 2 states that whatever the realised distribution, R and S's optimal actions are always ordered in the same direction. Notice that Assumption 3 isn't enough for this result. When Assumption 4 is violated, there can be $\theta \in \mathcal{C}$ such that $A_S(\theta) = A_R(\theta)$.

A2. Finite number of states

Assume that there is a finite number of state $\Omega = \{1, ..., N\}$ and that $\mathcal{D}$ is a set of probability mass functions of parametric distributions over Ω indexed by a single parameter $\theta \in \mathcal{C} \subset [0,1]$, differentiable in θ, such that there is a bijection between $\mathcal{D}$ and $\mathcal{C}$ and that:

$$\frac{\partial^2 \mathbb{E}_\theta(u_i(a,\omega))}{\partial a \partial \theta} > 0 \tag{A1}$$

where for $i \in S, R$, $\mathbb{E}_\theta(u_i(a,\omega)) = \sum_{\omega=1}^{N} p_\theta(\omega) u_i(a,\omega)$. As before, I identify models to $\mathcal{C}$. To see why the main results of the paper hold, one needs to show that Proposition 1 is still true in my new framework. First, we have that:

$$\frac{\partial \mathbb{E}_\theta(u_R(a,\omega))}{\partial \theta} = \sum_{\omega=1}^{N} \frac{\partial p_\theta(\omega)}{\partial \theta} u_R(a,\omega)$$

Because, for any $\theta \in \mathcal{C}$, p_θ is a probability mass function, it must be that:

$$\sum_{\omega=1}^{N} p_\theta(\omega) = 1 \Rightarrow \sum_{\omega=1}^{N} \frac{\partial p_\theta(\omega)}{\partial \theta} = 0$$

Denote as $\alpha_1^+, ..., \alpha_P^+$ the non-negative elements of $\{\frac{\partial p_\theta(\omega)}{\partial \theta} | \omega \in \Omega\}$ and $\alpha_1^-, ..., \alpha_Q^-$ as the negative ones. We can rewrite:

$$\frac{\partial \mathbb{E}_\theta(u_R(a,\omega))}{\partial \theta} = \sum_{j=1}^{P} \alpha_j^+ u_R(a,\omega) - \sum_{k=1}^{Q} |\alpha_k^-| u_R(a,\omega)$$

in which $\sum_{j=1}^{P} \alpha_j^+ u_R(a,\omega)$ and $\sum_{k=1}^{Q} |\alpha_k^-| u_R(a,\omega)$ are both single peaked functions as sum of functions which are. They cross at least once in $a_0 \in \mathcal{A}$. In addition, Assumption (A1) implies that this crossing is unique. It follows that there is a unique $a_0 \in \mathcal{A}$ such that:

$$\frac{\partial \mathbb{E}_\theta(u_R(a_0,\omega))}{\partial \theta} = 0$$

Thus, $\mathbb{E}_\theta(u_R(a,\omega))$ is decreasing in θ for $a < a_0$ and increasing in θ for $a \geq a_0$. It further implies that:

$$\min_{\theta\in[0,1]} \mathbb{E}_\theta(u_R(a,\omega)) = \begin{cases} \mathbb{E}_1(u_R(a,\omega)) \text{ if } a < a_0 \\ \mathbb{E}_0(u_R(a,\omega)) \text{ if } a \geq a_0 \end{cases}$$

Assumption (A1) also implies that $\mathbb{E}_1(u_R(a,\omega))$ and $\mathbb{E}_0(u_R(a,\omega))$ cross only once. It follows that $\min_{\theta\in[0,1]} \mathbb{E}_\theta(u_R(a,\omega))$ is single peaked.

Let $\widetilde{a} \equiv argmax_{a\in\mathcal{A}} \min_{\theta\in[0,1]} \mathbb{E}_\theta(u_R(a,\omega))$.[20] By definition, $\widetilde{a} = a_0$. As a result, for $B = [\theta_1, \theta_2]$, we have that:

$$\min_{\theta\in B} \mathbb{E}_\theta(u_R(a,\omega)) = \begin{cases} \mathbb{E}_{\theta_2}(u_R(a,\omega)) \text{ if } a < \widetilde{a} \\ \mathbb{E}_{\widetilde{\theta}}(u_R(a,\omega)) \text{ if } a = \widetilde{a} \\ \mathbb{E}_{\theta_1}(u_R(a,\omega)) \text{ if } a > \widetilde{a} \end{cases}$$

where $\widetilde{\theta}$ is defined as in the main text. Thus, when $\theta_2 < \widetilde{\theta}$, $\mathbb{E}_\theta(u_R(a,\omega))$ is strictly decreasing with θ for all $a \in [A_R(\theta_1), A_R(\theta_2)]$ which implies that $\min_{\theta\in B} \mathbb{E}_\theta(u_R(a,\omega)) = \mathbb{E}_{\theta_2}(u_R(a,\omega))$ and thus that $A_R(B) = A_R(\theta_2)$. Similarly, when $\theta_1 > \widetilde{\theta}$, $\mathbb{E}_\theta(u_R(a,\omega))$ is strictly increasing with θ which implies that $\min_{\theta\in B} \mathbb{E}_\theta(u_R(a,\omega)) = \mathbb{E}_{\theta_1}(u_R(a,\omega))$ and thus that $A_R(B) = A_R(\theta_1)$.

When $\widetilde{\theta} \in B$, $\min_{\theta\in B} \mathbb{E}_\theta(u_R(a,\omega))$ is increasing on $(A_R(\theta_1), \widetilde{a})$ (as $\mathbb{E}_{\theta_2}(u_R(a,\omega))$

[20]Note that this definition of $\widetilde{a}$ is equivalent to the one given in the main text when $N = 2$ and $p_\theta(1) = \theta$.

is maximal at $A_R(\theta_2) > \widetilde{a}$) and decreasing on $(\widetilde{a}, A_R(\theta_2))$ (as $\mathbb{E}_{\theta_1}(u_R(a,\omega))$ is maximal at $A_R(\theta_1) < \widetilde{a}$). As a result, it is always maximal for $\widetilde{a}$ and thus $\min_{\theta \in B} A_R(B) = A_R(\widetilde{\theta})$. Thus, Proposition 1 holds. As Theorem 1 and 2 follow directly from Proposition 1, they extend to the case of a finite number of states.

As for the two-state case, when $N > 2$, comparability and dominance across states play a central role in obtaining monotonicity in the maximal expected pay off to the receiver. Yet, interpretation will be less straight-forward than in the $N = 2$ case as it will depend not only on the distributions considered, but also on the considered family of probability distributions. If there is a complete order of dominance over states, it is easy to pick a family of distributions such that monotonicity in the maximal expected pay off of the receiver necessarily holds[21], whatever the distributions considered. Yet, if at least two states are comparable, monotonicity in the maximal expected pay off of the receiver might fail depending on the distributions considered. The above analysis shows that the reversal in the monotonicity of the maximal expected utility of the receiver can still occur at most once.

Proofs of the results in the main text

Proof of Proposition 1:

In order to prove my result, we need to study the variations of $\mathbb{E}_\theta(u_R(a,\omega))$ as a function of θ. For $a \in \mathcal{A}$,

$$\frac{\partial \mathbb{E}_\theta(u_R(a,\omega))}{\partial \theta} = u_R(a,1) - u_R(a,0)$$

Thus, we are interested in the sign of $u_R(a,1) - u_R(a,0)$. First, we need to prove the following Lemma:

[21]For instance a family of geometrical distributions.

Lemma 3. *Define $B \subset \mathcal{C}$ as the belief of the receiver with minimal element θ_1 and maximal element θ_2. Given this belief, his optimal action is $A_R(B) \subset [A_R(\theta_1), A_R(\theta_2)]$.*

Proof of lemma 3:

We prove this lemma in the more general context of α-MEU preferences. This criteria coincides with MEU when $\alpha = 1$.

First, notice that $\forall a \in \mathcal{A}$, there is $\theta_m(a) \in B$ such that $\min_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega)) = \mathbb{E}_{\theta_m(a)}(u_R(a, \omega))$. Similarly, $\forall a \in \mathcal{A}$, there is $\theta_M(a) \in B$ such that $\max_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega)) = \mathbb{E}_{\theta_M(a)}(u_R(a, \omega))$.

As a result, $\forall a \in \mathcal{A}$, $\alpha \min_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega)) + (1 - \alpha) \max_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega)) = \alpha \mathbb{E}_{\theta_m(a)}(u_R(a, \omega)) + (1 - \alpha) \mathbb{E}_{\theta_M(a)}(u_R(a, \omega)) = \mathbb{E}_{\alpha \theta_m(a) + (1-\alpha)\theta_M(a)}(u_R(a, \omega))$. As, for all $a \in \mathcal{A}$, $\theta_1 \leq \alpha \theta_m(a) + (1 - \alpha)\theta_M(a) \leq \theta_2$ and that $A_R(\theta)$ is a strictly increasing function, it must be that $A_R(B) \subset [A_R(\theta_1), A_R(\theta_2)]$.

□

A consequence of the Lemma 3 is that when looking for optimal actions for a given B, it is sufficient to look for actions in $[A_R(\theta_1), A_R(\theta_2)]$. Notice that $[A_R(\theta_1), A_R(\theta_2)] \subset [a_R(0), a_R(1)]$ and that for all $a \in [a_R(0), a_R(1)]$ either:

1. $u_R(a_R(0), 0) < u_R(a_R(0), 1)$.

   For $a > a_R(0)$, $u_R(a, 0)$ is decreasing and $u_R(a, 1)$ is increasing, utilities in both states are never equal and $u_R(a, 0) < u_R(a, 1)$ for all $a \in \mathcal{A}$. As in this case $\widetilde{a} = a_R(0)$ and thus $\widetilde{\theta} = 0$, $\mathbb{E}_\theta(u_R(a, \omega))$ is strictly increasing with θ for all $a \in [a_R(0), a_R(1)]$. As a result, $A_R(B) = A_R(\theta_1)$.

2. $u_R(a_R(0), 0) > u_R(a_R(0), 1)$ and $u_R(a_R(1), 0) > u_R(a_R(1), 1)$.

   For $a > a_R(0)$, $u_R(a, 0)$ is decreasing and $u_R(a, 1)$ is increasing, but as $u_R(a_R(1), 0) > u_R(a_R(1), 1)$ it must be that utilities in both states are never equal. As a result, $u_R(a, 0) > u_R(a, 1)$ for all $a \in \mathcal{A}$. Thus, in this case $\widetilde{a} = a_R(1)$ and $\widetilde{\theta} = 1$. It follows that $\mathbb{E}_\theta(u_R(a, \omega))$ is strictly decreasing with θ for all $a \in [a_R(0), a_R(1)]$. As a result, $A_R(B) = A_R(\theta_2)$.

3. $u_R(a_R(0), 0) > u_R(a_R(0), 1)$ and $u_R(a_R(1), 0) \leq u_R(a_R(1), 1)$.

   As for $a > a_R(0)$, $u_R(a, 0)$ is strictly decreasing and $u_R(a, 1)$ is strictly increasing. Thus, both utilities are equal for a unique given action and by definition of $\widetilde{a}$ it must be that this point is $\widetilde{a}$. As a result:

$$\begin{cases} u_R(a,0) > u_R(a,1) \text{ for } a < \widetilde{a} \\ u_R(a,0) = u_R(a,1) \text{ for } a = \widetilde{a} \\ u_R(a,0) < u_R(a,1) \text{ for } a > \widetilde{a} \end{cases}$$

   Thus, for $a \in [A_R(\theta_1), A_R(\theta_2)]$, $\mathbb{E}_\theta(u_R(a, \omega))$ is strictly decreasing with θ when $\theta_2 < \widetilde{\theta}$ and strictly increasing with θ when $\theta_1 > \widetilde{\theta}$, which gives the corresponding result. The above system also implies that when $\widetilde{\theta} \in B$, $\mathbb{E}_\theta(u_R(a, \omega))$ is always minimal for $\theta = \widetilde{\theta}$. As a result, for all $a \in [A_R(\theta_1), A_R(\theta_2)]$ the minimal pay-off of the receiver as a function of the sender's type is given by:

$$\min_{\theta \in B} \mathbb{E}_\theta(u_R(a,\omega)) = \begin{cases} \mathbb{E}_{\theta_2}(u_R(a,\omega)) \text{ if } a < \widetilde{a} \\ \mathbb{E}_{\widetilde{\theta}}(u_R(a,\omega)) \text{ if } a = \widetilde{a} \\ \mathbb{E}_{\theta_1}(u_R(a,\omega)) \text{ if } a > \widetilde{a} \end{cases}$$

   The above system implies that when $\widetilde{\theta} \in B$, $\min_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega))$ is increasing on $(A_R(\theta_1), \widetilde{a})$ (as $\mathbb{E}_{\theta_2}(u_R(a, \omega))$ is maximal at $A_R(\theta_2) > \widetilde{a}$) and decreasing on $(\widetilde{a}, A_R(\theta_2))$ (as $\mathbb{E}_{\theta_1}(u_R(a, \omega))$ is maximal at $A_R(\theta_1) < \widetilde{a}$). As a result, it is always maximal for $\widetilde{a}$. As a result, $\min_{\theta \in B} A_R(B) = A_R(\widetilde{\theta})$.

□

Proof of Proposition 2

The proof is structured as follows. First, I show that the number of outcome actions induced in equilibrium is finite. Then, I prove that the set of types which get the same equilibrium outcome must form an interval. The continuity and the strict monotonicity of the sender's preferences completes the argument.

Lemma 4. *There exists $\epsilon > 0$ such that if u and v are actions induced in equilibrium, $|u - v| \geq \epsilon$. Further the set of actions induced in equilibrium is finite.*

Proof of Lemma 4

I say that action u is induced by an S-type θ if it is a best response to a given equilibrium message $m : u \in \{A_R(\theta) | \theta \in \sigma^{-1}(m)\}$. Let Y be the set of all actions induced by some S-type θ. First, notice that if θ induces $\overline{a}$, it must be that $V_S^{\theta}(\overline{a}) = \max_{a \in Y} V_S^{\theta}(a)$. Since u_S is strictly concave, $V_S^{\theta}(a)$ can take on a given value for at most two values of a. Thus, θ can induce no more than two actions in equilibrium.

Let u and v be two actions induced in equilibrium, $u < v$. Define Θ_u as the set of S types who induce u and Θ_v as the set of S types who induce v. Take $\theta \in \Theta_u$ and $\theta^{'} \in \Theta_v$. By definition, θ reveals a weak preference for u over v and $\theta^{'}$ reveals a weak preference for v over u that is:

$$\begin{cases} V_S^{\theta}(u) \geq V_S^{\theta}(v) \\ V_S^{\theta^{'}}(v) \geq V_S^{\theta^{'}}(u) \end{cases}$$

Thus, by continuity of $\theta \to V_S^{\theta}(u) - V_S^{\theta}(v)$, there is $\hat{\theta} \in [\theta, \theta^{'}]$ such that $V_S^{\hat{\theta}}(u) = V_S^{\hat{\theta}}(v)$. Since u_S is strictly concave, we have that:

$$u < A_S(\hat{\theta}) < v$$

Then, notice that since $\frac{\partial^2 \mathbb{E}_{\theta}(u_S(a,\omega))}{\partial a \partial \theta} > 0$ (Inequality (1)), it must be that all types that induce u are below $\hat{\theta}$. Similarly, it must be that all types that induce v are above $\hat{\theta}$. That is:

$$\forall \theta \in \Theta_u, \theta \leq \hat{\theta}$$
$$\forall \theta \in \Theta_v, \theta \geq \hat{\theta}$$

Thus, when R is MEU, Lemma 3 implies that the optimal action of the receiver, given that $\theta \in \Theta_u$ is below the optimal action when the type is $\hat{\theta}$. Similarly, the optimal action of the receiver, given that $\theta \in \Theta_v$ is above the optimal action when the type is $\hat{\theta}$. The same is true when when R is SEU. That is:

$$\begin{cases} A_R(\Theta_u) \leq A_R(\hat{\theta}) \\ A_R(\Theta_v) \geq A_R(\hat{\theta}) \end{cases}$$
$$\iff u \leq A_R(\hat{\theta}) \leq v$$

However, as $A_R(\theta) \neq A_S(\theta)$ for all $\theta \in \mathcal{C}$, there is $\epsilon > 0$ such that $|A_R(\theta) - A_S(\theta)| \geq \epsilon$ for all $\theta \in \mathcal{C}$. It follows that $|u - v| \geq \epsilon$.

Lemma 3 implies that for any belief $B \subset \mathcal{C}$, the optimal action of the receiver is in $[A_R(\underline{\theta}, A_R(\overline{\theta})]$. Thus, the set of actions induced in equilibrium is bounded by $A_R(\underline{\theta})$ and $A_R(\overline{\theta})$ and at least ϵ away from one another, which completes the proof.

□

Lemma 5. *In every equilibrium of the game, if a is an action induced by type* θ *and by type* $\theta^{''}$ *for some* $\theta < \theta^{''}$*, then a is also induced by* $\theta^{'} \in (\theta, \theta^{''})$

Proof of Lemma 5:

For the purpose of the proof, we introduce the notation $W^{\theta}(a) = \mathbb{E}_{\theta}(u_S(a, \omega))$, which is the evaluation of $a \in \mathcal{A}$ by a sender of type θ.

We proceed by contradiction. Suppose a_1 is induced by type θ and by type $\theta^{''}$ and that there is $\theta^{'} \in (\theta, \theta^{''})$ such that a_1 is not induced. Then there must be $a_2 \neq a_1$ that type $\theta^{'}$ prefers and that $\theta^{''}$ does not. Formally, this is:

$$\begin{cases} W^{\theta}(a_2) \leq W^{\theta}(a_1) \\ W^{\theta^{'}}(a_1) \leq W^{\theta^{'}}(a_2) \\ W^{\theta^{''}}(a_2) \leq W^{\theta^{''}}(a_1) \end{cases} \tag{B1}$$

Notice that for $a \in \mathcal{A}$:

$$\frac{\partial W^{\theta}(a)}{\partial \theta} = u_S(a,1) - u_S(a,0)$$

Similarly to S, define $\widetilde{a}_S = argmax_{a \in \mathcal{A}} \min_{\omega \in \Omega} u_S(a, \omega)$. $\widetilde{a}_S$ is the action that maximises the worst possible expected utility of the sender among the set of distributions. Two special cases are to be noticed. Either the high state is sufficiently worse than the good one for it to give a lower utility at its optimal point: $u_S(a_S(1), 1) \leq u_S(a_S(1), 0)$. Then the hedging action is the optimal action in the high state $\widetilde{a}_S = a_S(1)$. Either the former is not true $(u_S(a_S(1), 1) > u_S(a_S(1), 0))$ and both states must give the same utility for a given action in $(a_S(0), a_S(1))$. In that case $\widetilde{a}_S$ is the action that gives the same utility in both states.

As a result, $W^{\theta}(a)$ is strictly decreasing for $a < \widetilde{a}_S$, constant for $a = \widetilde{a}_S$ and strictly increasing for $a > \widetilde{a}_S$. Assume that $a_1 < a_2$:

- When $a_1 < \widetilde{a_S}$ and $a_2 \geq \widetilde{a}_S$ can cross at most once and system (B1) is impossible.
- Assume $\widetilde{a}_S \leq a_1 < a_2$. Then:

$$\frac{\partial(W^\theta(a_1) - W^\theta(a_2))}{\partial\theta} = u_S(a_1,1) - u_S(a_1,0) - (u_S(a_2,1) - u_S(a_2,0))$$

As, for $a \geq \widetilde{a}_S$, $u_S(a,1)$ is a strictly increasing function and $u_S(a,0)$ a strictly decreasing one, we have that $a_1 < a_2$ implies that $u_S(a_1,1) - u_S(a_1,0) < u_S(a_2,1) - u_S(a_2,0)$. Thus, $W^\theta(a_1) - W^\theta(a_2)$ is a strictly decreasing function of θ and $W^\theta(a_2)$ and $W^\theta(a_1)$ can cross at most once, making system (B1) impossible.

- Assume $a_1 < a_2 < \widetilde{a}_S$. Then, $W^\theta(a_1) - W^\theta(a_2)$ is a strictly increasing function of θ and $W^\theta(a_2)$ and $W^\theta(a_1)$ can cross at most once, making system (B1) impossible.

The case $a_2 > a_1$ is symmetric.

□

By Lemma 4 there is a finite number of outcomes induced in equilibrium. The continuity of $A_S(\theta)$ gives that there is a type of the sender which is indifferent between any pair of outcomes induced in equilibrium and the monotonicity of $A_S(\theta)$ implies there are only a finite number of types which are indifferent between any pair of outcomes. Hence, Lemma 5 implies that there is a partitioning of $\mathcal{C}$ in a finite number of cells where every cell induces a given action in equilibrium.

□

Proof of Proposition 3

The outline of the proof is as follows. I start by showing that the cut-off types of any equilibrium must satisfy condition (2). Any other equilibrium strategies would be outcome equivalent.

Consider a couple of strategy (σ_q^*, y_q^*) and write $C_k^q = [\theta_k^q, \theta_{k+1}^q]$.

- Assume y_q^* is the equilibrium strategy of R. Given Proposition 2, any type $\theta \in C_k^q$ induces the same action and prefers it to any other equilibrium action. Thus, for σ_q^* to be an equilibrium strategy, it is without loss of generality to assume that any type $\theta \in C_k^q$ sends the same message m_k and prefer it to any other message[22]. In particular, it must be preferred to message m_{k-1} which induces the preferred equilibrium action of types in C_{k-1}^q. For all $\theta \in C_k^q$:

$$V_S^{\theta}(y^*(m_k^q)) \geq V_S^{\theta}(y^*(m_{k-1}^q))$$

Similarly, any type $\theta \in C_{k-1}^q$ must prefer sending m_{k-1} to m_k. For all $\theta \in C_{k-1}^q$:

$$V_S^{\theta}(y^*(m_k^q)) \leq V_S^{\theta}(y^*(m_{k-1}^q))$$

Thus, for σ_q^* to be an equilibrium strategy a necessary condition is that:

$$V_S^{\theta_k^q}(y^*(m_{k-1}^q)) = V_S^{\theta_k^q}(y^*(m_k^q))$$

- Assume σ_q^* is the equilibrium strategy of S. Then, for any $\theta \in \mathcal{C}$, the best response of R in the MEU case to any equilibrium message $\sigma_q^*(\theta)$ is:

$$argmax_{a \in A} V_R^{MEU}(a, \sigma_q^{*-1}(\sigma_q^*(\theta))) = y_q^*(\sigma_q^*(\theta))$$

[22]Any other signaling strategy must induce the same action from R and will thus lead to the same pay-offs for both players, whatever the sender's type.

Similarly, in the SEU case, the best response of R to any equilibrium message $\sigma_q^*(\theta)$ is:

$$argmax_{a\in A}V_R^{SEU}(a,\sigma_q^{*-1}\sigma_q^*(\theta))) = y_q^*(\sigma_q^*(\theta))$$

□

Proof of Lemma 1:

The structure of the proof is as follows. First, I provide an algorithm that characterises the cut-off types of the equilibrium that has most cut-offs: $\theta_0 < ... < \theta_M$ (step 1). Define $\mathcal{E} = \{(\theta_0, \theta_k, ..., \theta_{M-1})|1 \leq k \leq M\}$. Then, I show that any non-babbling partitional strategy of the sender characterised by cut-offs which are elements of $\mathcal{E}$ is an equilibrium strategy (step 2). I conclude by showing that this describes every equilibrium of the game (step 3).

In the following, I call $C_q = [\theta_q, \theta_{q+1}]$, for $1 \leq q < M-1$, $C_M = [\theta_M, \overline{\theta}]$ and $C_0 = [\underline{\theta}, \theta_1]$

Step 1:

Assume there is a M cut-off equilibrium. Then the signalling strategy of the sender σ must be such that for $q \in 0, ..., M$, $\forall\theta \in C_q$, $\sigma(\theta) = m_k$

First notice that $V_R^{MEU}(a, C_0) = \mathbb{E}_{\theta_1}(u_R(a,\omega))$. For σ to be an equilibrium strategy we need that $\forall\theta \in C_0$ and $m \neq m_0$:

$$V_S^{\theta}(m_0) \geq V_S^{\theta}(m)$$

In C_0, type θ_1 has the most incentive to deviate from sending m_0 to sending m_1, which would induce a higher action, as, $V_R^{MEU}(a, C_1) = \mathbb{E}_{\theta_2}(u_R(a,\omega))$ and $A_R(\theta)$ is strictly increasing by inequality (1).

Thus, a necessary condition for all types in C_0 to send m_0 is that:

$$V_S^{\theta_1}(m_0) \geq V_S^{\theta_1}(m_1)$$

Furthermore, it is also necessary that all types in C_1 prefer message m_1. In particular it must be the case for type θ_1, thus: $V_S^{\theta_1}(m_1) \geq V_S^{\theta_1}(m_0)$. As a consequence, a necessary condition for σ to be an equilibrium strategy is:

$$V_S^{\theta_1}(m_0) = V_S^{\theta_1}(m_1) \tag{B2}$$

By repeating the argument for all C_q, $q \in 1, ..., M$, a necessary condition for σ to be an equilibrium strategy is for all $q \in 1, ..., M$:

$$V_S^{\theta_q}(m_{q-1}) = V_S^{\theta_q}(m_q) \tag{B3}$$

Furthermore, the fact that $\cup_{k=0}^{M} C_k = \mathcal{C}$ and the fact that for every pair of consequent cell of the partition the incentive constraints are transitive gives that conditions (B3) is both necessary and sufficient. As $A_R(\theta)$ is strictly monotone, it implies that $A_R(\theta_k) \neq A_R(\theta_{k+1})$. $\underline{\theta}$ being known, it is possible to derive θ_1 directly from (B2). By repeating the reasoning by induction, θ_{k+1} can be derived from θ_k for $k \in 1, ..., M-1$ from (B3) as long as there is $\theta_M < \widetilde{\theta}$.

Step 2:

I show that any partitional strategy of the sender characterised by elements of $\mathcal{E}$ is an equilibrium strategy. I proceed by iteration:

For $q \geq 1$, define σ_q such that for $q \leq k \leq M$, $\forall \theta \in C_k$, $\sigma_q(\theta) = m_k$ and $\forall \theta \in [\underline{\theta}, \theta_q]$, $\sigma_q(\theta) = m_0$

- Step 1 proves that σ_1 is an equilibrium strategy.

- Assume that for $1 \leq k \leq M-1$, σ_k is an equilibrium strategy. Let's show that σ_{k+1} is one as well. First notice that, for $\theta \geq \theta_k$ the strategy of S is the same under σ_{k+1} and σ_k. As $\sigma_k(\theta)$ for $\theta \geq \theta_k$ is a best response, $\sigma_{k+1}(\theta)$ for $\theta \geq \theta_k$ is as well. Second, note that following Theorem 1, $\theta_k \leq \widetilde{\theta}$. It follows that for any $\underline{\theta} < \theta < \theta_k$, $\sigma_{k+1}(\theta)$ induce the same action from the receiver than $\sigma_k(\theta)$, for $\theta_{k-1} < \theta < \theta_k$. It then follows from inequality (1) that $\sigma_{k+1}(\theta)$ is a best response for all types $\theta < \theta_k$. This completes the proof of step 2.

Step 3:

Assume there is an equilibrium strategy of the sender σ which is not described above. Recall $A_R(B)$ to be the optimal action of R under the belief that $\theta_0 \in B$ for $B \subset \mathcal{C}$.

Proposition 2 gives that all equilibria are partitional. First I will show that any equilibria only characterised by elements of $\theta_0, ..., \theta_q$ must be characterised by elements of $\mathcal{E}$. It is straightforward to see that any equilibria only characterised by elements of $\theta_0, ..., \theta_q$ which is not in $\mathcal{E}$ can be constructed by removing non-minimal elements to $e \in \mathcal{E}$. To prove our claim, it is thus sufficient to prove that no such equilibrium can be constructed.

For $1 \leq q \leq M$, consider a strategy σ_p characterised by cut-offs $\theta_0, \theta_q, ..., \theta_{p-1}, \theta_{p+1}, ..., \theta_M$ for $q+1 \leq p \leq M$ and assume it is an equilibrium strategy. It must be that that type θ_{p+1} prefers outcome $A_R([\theta_{p-1}, \theta_{p+1}))$ to outcome $A_R([\theta_{p+1}, \theta_{p+2}))$. Yet, by construction of the equilibrium of q cut-offs, types θ_{p+1} is exactly indifferent between outcome $A_R([\theta_p, \theta_{p+1}))$ and outcome $A_R([\theta_{p+1}, \theta_{p+2}))$. As $A_R([\theta_{p-1}, \theta_{p+1})) < A_R([\theta_p, \theta_{p+1}))$, the previous implies that type θ_{p+1} prefers outcome $A_R([\theta_{p+1}, \theta_{p+2}))$ to outcome $A_R([\theta_{p-1}, \theta_{p+1}))$, which is a contradiction.

□

Proof of Proposition 4:

1. Assume R has SEU preferences. Assume there are n equilibrium cut-offs

in $[0,1]$: $\theta_0, ..., \theta_n$ and thus $\theta_0 = 0$ $\theta_n = 1$. When receiving equilibrium message m_k sent by types $\theta \in [\theta_k, \theta_{k=1})$ S evaluates action through:

$$\begin{aligned} V_R(a|m_k) &= \int_{\theta \in [\theta_k, \theta_{k+1}]} (1-\theta) u_R(a,0) + \theta u_R(a,1) d\theta \\ &= (1 - \mathbb{E}(\theta|m_k)) u_R(a,0) + \mathbb{E}(\theta|m_k) u_R(a,1) \end{aligned}$$

where $\mathbb{E}(\theta|m_k) = \int_{\theta \in [\theta_k, \theta_{k+1}]} \theta d\theta = \frac{\theta_k + \theta_{k+1}}{2}$. A first order condition on the above gives that when evaluating actions through $V_R(a|m_k)$, the optimal action is $\mathbb{E}(\theta|m_k)$. It follows that the equilibrium action of R is $y^*(m_k) = \frac{\theta_k + \theta_{k+1}}{2}$. The optimal action in the eyes of S is $A_S(\theta_0) = \theta_0 + b$. The arbitrage condition gives that a sender of type θ_k must be indifferent between m_{k-1} and m_k. That is, for $k \in 2, ..., n$:

$$A_S(\theta_{k+1}) = \frac{y^*(m_k) + y^*(m_{k+1})}{2}$$

Notice that this arbitrage condition translates in the similar condition as in CS's example:

$$\theta_{k+1} - \theta_k = \theta_k - \theta_{k-1} + 4b \tag{B4}$$

Equation (B4) further gives that:

$$\theta_k = k(\theta_1 - \theta_0) + \frac{k(k-1)}{2} 4b$$

Specifically, $1 = \mathbb{E}(\theta_n) = n(\theta_1) + \frac{n(n-1)}{2}4b$ which gives $\theta_1 = \frac{1}{n} - 2(n-1)b$ and:

$$\mathbb{E}(\theta_k) = \theta_k = \frac{k}{n} - 2kb(n-k)$$

It follows that a n cut-off equilibrium exists if and only if:

$$0 < b < \frac{1}{2n(n-1)}$$

2. Assume R has MEU preferences and that there is a n-cut-off equilibrium. When receiving message m_k^n, for $k \geq 2$:

$$V_R(a|m_k) \quad = \quad min_{\theta \in [\theta_k, \theta_{k+1}]} \mathbb{E}_\theta(u_R(a))$$

Thus, when $\theta_1 \leq \widetilde{\theta}$, $V_R(a|m_0) = \mathbb{E}_{\theta_1}(u_R(a))$ and the arbitrage condition giving the cut-off types gives that $A_S(\theta_1) = \theta_1 + b$ must thus be at equal distance from θ_1 and θ_2. For this to be possible, it must be that $b > 0$. Thus, when there is a n-cut-off equilibrium, it must be that $\widetilde{\theta} > \theta_n$. When receiving message m_k, for $k \geq 1$:

$$V_R(a|m_k) = \mathbb{E}_{\theta_{k+1}}(u_R(a))$$

The equilibrium action of R when receiving the equilibrium message $[\theta_k, \theta_{k+1}]$ is $y(m_k^n) = \mathbb{E}(\theta_{k+1})$. The arbitrage condition giving the cut-off types gives that $A_S(\theta_{k+1})$ must thus be at equal distance from $\mathbb{E}(\theta_{k+1})$ and $\mathbb{E}(\theta_{k+2})$, giving

$$\theta_{k+1} + b = \frac{\theta_{k+1} + \theta_{k+2}}{2}$$
$$\iff \theta_{k+2} = \theta_{k+1} + 2b$$

When receiving message m_n, the equilibrium action of R is $y(m_n) = \widetilde{\theta} = \frac{1}{2}$. The arbitrage condition when S is of type θ_{n-1} gives that:

$$\frac{\widetilde{\theta} + \theta_{n-1}}{2} = \theta_{n-1} + b$$
$$\iff \theta_{n-1} = 1 - 2b$$

Which implies that, for all $1 \leq k \leq n-1$:

$$\theta_k = \theta_k = 1 - 2b(n-k)$$

It follows that a n cut-off equilibrium exists if and only if:

$$\theta_1 > 0$$
$$\iff 1 - 2bn > 0$$
$$\iff 0 < b < \frac{1}{2n}$$

□