Expert-based Scientific Knowledge: Communicating over Models*

Philippe Colo[1]

[1]Paris School of Economics & Université Paris I

27th Nov, 2020

JOB MARKET PAPER

This paper's latest version here

Abstract

This paper studies the transmission of complex scientific knowledge. Scientific models are formalised as probability distributions over possible scenarios. An expert is assumed to know the most likely model and seeks to communicate it to a decision maker, but cannot certify it. As a result, communication of scientific knowledge is a cheap talk game over models. The decision maker is in a situation of model-uncertainty and is ambiguity sensitive. I show that information transmission depends on both the strategic misalignment of players and the degree of consensus among scientific models. When science is divided, there is an asymmetry in information transmission when the receiver has maxmin expected utility preferences. Types below a certain threshold are necessarily pooled, regardless of the misalignment. All equilibria of the game are outcome equivalent to a partitional equilibria and, unlike similar models in the literature, the most informative one is interim Pareto dominant. These results bring new insights regarding the current COVID-19 and climate crises. They show why scientific recommendations calling for more efforts in the provision of a public good, such as social distancing or pollution reduction, may lack influence over the general public.

Keywords: Ambiguity, cheap talk, contribution to a public good

JEL classification : D81, D83, C72, Q54

*I thank Jean-Marc Tallon, Stéphane Zuber, Frédéric Koessler, Massimo Marinacci, Marco Ottaviani, Valentina Bossetti, Mark Le Quement, Francis Bloch, Larry Samuelson, Loïc Berger and Guillaume Pommey for helpful discussions. I also thank seminar and conference participants at PSE (TOM, SRE), Bocconi and FUR 2018 and Bocconi University for its 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.

1 Introduction

We are laymen on most of the knowledge we claim to possess. Most of today's science is too complex for an individual to understand it at first hand. It is often so complex that even for the experts who do understand it, it is hard to convey the evidence supporting their claims in a convincing way. In such cases, our knowledge relies much more on the word of these experts than on the evidence they can convincingly provide. The importance of our confidence in experts is even higher if one considers that, on topics such as climate change or the COVID-19 pandemic, science is highly uncertain. This is not only because data is scarce. The mechanism through which tobacco causes lung cancer is well understood and can give rise to a precise estimate of the chances of getting cancer upon smoking a packet a day. Issues like climate change give rise to various models predicting widely different probabilistic scenarios. The uncertainty over the models themselves is of a different nature than the mere randomness over outcomes that one has to face even with a single well-established model.

In this paper, I test the expert-laymen bound of trust by studying the transmission of information in the context of this complex, uncertain science. Information is about models, which I represent as probability distributions over states of the world. The transmission is strategic, as 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 among decision makers on issues such as global warming or the spread of a deadly virus. The expert reviews a set of scientific models and decides which is the most accurate. This model is the expert's type. He then communicates its findings to the decision maker who acts upon it. Given the strategic nature of the communication, the expert is typically not able to truthfully reveal which model is the most accurate. The resulting uncertainty over models creates a situation which has been extensively studied and designated as ambiguity. It calls for the use of specific ambiguity-sensitive preferences for the decision maker.

The game I study is in the tradition of Crawford and Sobel (1982)'s (hereafter CS) cheap talk game. My assumption on the utility functions are the same as in the general case of their model. The main difference lies in the fact that communication bears on probability distributions over states of the world rather than on states themselves. I mainly focus on the cases where the receiver displays maxmin expected utility preferences (MEU) and subjective expected utility (SEU) ones. At equilibrium, the sender points out sets of models as being the most accurate. In the SEU case, the equilibria of the game are similar to CS. But in the MEU case, the change in the nature of information has a major impact on the outcome of the game. Because of ambiguity aversion, the most pessimistic model is a strong point of attraction for the receiver. When the sender's preferred action leans towards this model, his influence is extremely high. When his interest is to induce an action in the opposite direction though, his influence is nonexistent. This effect can create an asymmetry regarding information transmission. Two cases arise. In the first case, science is divided because all models do not agree on which state is the worst, whatever the action. In that case, senders can be separated into two categories. For types of senders belonging to one category, whatever the difference of interest between both parties, the sender is unable to convey any information. For types in the other category, the precision of information transmission depends on the difference of interest between them. In the second case, there is no such disagreement among scientific models. All models fall into one of two previous categories and I say that science is consensual. In both cases, it is always in the sender's interest to be as precise as he can, even after he learns his type. This second result also strongly contrasts with the SEU case, where such incentives do not exist. Then, the precision of information transmission depends only on the difference of interest between both parties. When the latter is small, information transmission can be almost perfect. But once the sender learns his type, he does not always have an interest in being as informative as he can.

Restricting attention to the popular linear-quadratic example of CS, I characterise all equilibria when the receiver has both MEU and SEU preferences. In that context I show that when science is divided and the receiver is MEU, less alignment in the players' interest is required for an equilibrium with a given number of cut-off types to be possible. In addition, the sender is in a better situation ex-ante when facing a MEU receiver. When science is consensual, I show that ambiguity aversion has dramatically opposed effects on communication. If the expert's preferences lean in in direction of science's more optimistic models, no information transmission can happen at all, whatever the misalignment. But if his preferences lean towards the more pessimistic models, he can, on average, induce a better action for himself through cheap-talk communication then by fully revealing his information (if he could). This last result strongly contrasts with the case of an SEU decision maker, where the expert would always be better-off if he could certify his information. When the expert's interests parallel the effect of ambiguity aversion, his influence is extremely high. To the contrary, when his interest is to go against ambiguity aversion, his influence is nonexistent.

Finally, I extend the model to the case of α-MEU decision-making, a preference relation that allows to continuously vary the level of ambiguity aversion. I show that, whatever the misalignment, there is a degree of ambiguity aversion such that an absence of consensus in science leads to the asymmetry in information transmission I observed in the MEU case. This result suggests a form of robustness of what I observe in the MEU case. In addition, I show that when science is consensual, ambiguity aversion always eases information transmission, but that when science is divided, ambiguity aversion can have the opposite effect.

My results apply when, in any state of the world, trade-offs are to be made. Some of the greatest challenges of our time are of this nature. In a pessimistic global-warming scenario, a high level of green house gases (GHG) abatement should be chosen, at the expense of economic growth. In an optimistic one, the opposite choice should be made. When decision-making is of this nature, two situations are to be noticed. Either one state of the world is worse than the others, whatever the decision maker's action. Then, science is necessarily consensual. Or conversely, the decision maker is caught between a rock and a hard place, and science can be divided. Climate change is a good example of the former situation. Scenarios such as the melting of the Antarctic ice sheet or the collapse of the Atlantic thermohaline circulation have been called "tipping elements" Lenton et al. (2008) because they imply a radical change in the climate system. In other words, they are worse than any other. Paradoxically, it is the existence of these catastrophic threats that creates the conditions for information transmission over all models. Conversely, current decisions regarding the COVID-19 pandemic are a choice between the lesser of two evils. One possibility is that the virus lethality is limited and that sparse sanitary measures are enough to contain casualties, while safeguarding the economy. But the opposite possibility cannot be discarded, triggering decisions which would lead to a much more rigorous limitation of social life (Hollingsworth et al., 2011). Arguably, both scenarios have a comparable profile of consequences. In this situation, science can be divided and my results show that information can only be conveyed regarding the models that consider the optimistic scenario (epidemiologically speaking) as the most likely. More generally, models that call for more efforts in the provision of a public good, such as social distancing or reduction in GHG emissions, will be on the inaudible side of science.

The first specificity of my paper is to model complex science as non-certifiable information. Complex science underlies some of the greatest challenges ours societies have to face. Consider the estimation of the effects of GHG on global temperature, which relies heavily on *black box* prospective computer simulations. Firstly, the process through which these simulations provide predictions is obscure; as pointed out by Pidgeon and Fischhoff (2011), black box simulations are hardly considered as convincing supporting evidence, even for scientists whose disciplines use observational methods. It is also extremely difficult for an expert of this field to justify why a given prospective simulation was chosen, a given methodology implemented or given assumptions made. What distinguishes an expert is precisely his direct understanding of the scientific foundation supporting existing models and of their relative quality. The modelling choice made by epidemiologists in order to evaluate the impact of sanitary measures on the COVID-19 pandemic is another good example. Two main approaches exist: process-based models, that try to capture the mechanisms by which diseases spread and curve-fitting approaches that aim at mathematically approximate the shape of the growth of the epidemic (Ferguson et al., 2003). The latter class of models does not attempt to characterise the underlying transmission process. As argued by Berger et al. (2020), choosing among these models is a fine art. It requires balancing between simplicity and comprehensiveness as a function of available data and the general understanding of the underlying mechanism. A task which makes the expert who he is. A final appropriate example, is the one of economists when they represent the social world through models. Constantly we have to navigate among modelling choices for the sake of tractability, compatibility with the rest of the literature or empirical testability. Economic modelling is complex because it requires this expertise. The resulting choices can be extremely hard to justify outside of the profession, a difficulty that has and still does attract a lot criticism.

The second specificity of my paper is to identify model uncertainty as the specific uncertainty surrounding scientific knowledge. Consider again the case of the effects of GHG emissions on the global temperature. Its estimation differs widely among the numerous existing models, largely because they rely on very different modelling choices. Predicting the impact of the rise of global temperature involves, for instance, modelling the socioeconomic response of our societies. As argued by Heal and Millner (2014), this can be done in a great variety of ways, leading to model uncertainty. The same challenge is present in the management of the COVID-19 pandemic. At the beginning of the pandemic, there was uncertainty about some of the more fundamental characteristics of the virus, such as its transmission channel, assuming one or another medium did, naturally, highly impact the resulting public policy recommendation (see Hellewell et al. (2020) or Anderson et al. (2020)). Following a tradition in statistics and decision theory dating back at least to Wald (1949) (see Marinacci (2015) for a survey), I represent scientific models as probability distributions over states of the world. Scientific models are simplified representations of reality, or models, capturing the main effects at stake in a given situation. For different courses of actions, a model predicts consequences as a function of the state of the environment, and uses probabilities to estimate their likelihood. When various models exist to represent the same phenomena, and offer different predictions, we face *model uncertainty*.

Under model uncertainty, preferences generally fail to satisfy the expected-utility requirements, as famously pointed out by Ellsberg (1961). In particular, decision makers may display ambiguity aversion. An individual is exposed to ambiguity when the expected payoff to his strategy varies with the probabilities over which he is uncertain. An ambiguity-averse individual will tend to favour strategies that reduce that exposure. MEU preferences, introduced by Gilboa and Schmeidler (1989), which I will focus on throughout most of the paper, are the more popular example of preferences capturing that trend. Ghirardato et al. (2004)'s α-MEU preference is the most tractable extension of MEU, where α captures the degree of ambiguity aversion of the decision maker. When $\alpha = 1$, the two criteria coincide.

The third specificity of my paper is to take into account the consequences of the often overlooked difference of interest between experts and decision makers regarding science's relative lack of persuasive power. The case of climate change is again an example of utmost importance. Consider a receiver as a country deciding on its GHG emissions in the widely studied context of a game of contribution to a public bad (Barrett, 1994). Such a receiver benefits from her own emissions but fails to internalise the externalities she produces on others, leading to inefficiently high levels of emission. An expert communicating on the corresponding scientific knowledge might be expected to take this inefficiency into account. As a result, if he cares about all countries' welfare, there is always an asymmetry of interests between the two. The rate of vaccination in a population is a similar example. Vaccination is a public good: the higher it is, the better the population is protected from the disease. Yet, individuals face a private cost in doing so and may be assumed to care only about their own welfare. As pointed out by Geoffard and Philipson (1997), in this context, the overall rate is inefficient. Again, there is an asymmetry of interest between a public health authority acting as a social planner and the individual members of the population.

This study contributes to the recent literature on cheap talk communication with ambiguity sensitive preferences. Kellner and Le Quement (2017) were the first to study this question. In their model, communication is on states of the world. They allow for Ellsbergian communication strategies which are a kind of mixed strategy of the sender, where the mixing probability is ambiguous. They show that the use of these strategies reduces misalignment between players, creating equilibria which ex-ante Pareto dominates the corresponding ones in CS. Kellner and Le Quement (2018) explore a simple two actions two states setting, with only standard mixed strategies allowed, but an ambiguous prior over the states. They show that the optimal communication strategy of the sender is a randomisation over partitions. Because my communication is over probability distributions and only pure strategies are allowed in my model, these results differ from mine. In addition, as pointed out by Hanany et al. (2020), because communication is over states, the updating assumed in these papers violate sequential optimality. This is an issue I don't face when communication is over models.

Hansen and Sargent (2001) and Hansen et al. (2006) have explored the effect of ambiguity aversion on model uncertainty in the context of dynamic decision making, showing its connection with robust control. Millner et al. (2013) and Berger et al. (2016) have argued for the relevance of model uncertainty and ambiguity aversion in the context of climate change management, where knowledge is scarce. This paper belongs to that line of thought, highlighting the informational and decisional consequences of this type of uncertainty when the source of information is explicitly modelled.

This paper also relates to a continuing debate in epistemology regarding the role of testimony in the foundation of knowledge. In classical epistemology, beliefs qualify as knowledge only if by perception or inference one can verify their truth. This position has been called *reductionist* and has notably been defended by Hume (1740) and Chisholm et al. (1989). But then, why can we say, for instance, that we know that GHG emissions are responsible for global warming? For most of us, this comes neither from perception nor from logical inference. As argued by Burge (1993), perception and inference cannot be seen as *warrants* for most of what we collectively designate as knowledge. An alternative *anti-reductionist* approach argues in favour of adding testimony to the list of primary warrants of knowledge (Hardwig, 1985). For supporters of this view, it is the confidence in an expert's testimony which rationally entitles the layman to hold the expert's judgement for knowledge (Goldman, 2001). It is the strength of this bound of trust that epistemologically *entitles* the layman to knowledge. This paper's contribution is to formally model the relationship of trust between expert and layman as a strategic interaction. How much expert-based knowledge the layman is entitled to posses is the information he holds *at equilibrium*. For instance, in the absence of misalignment, the layman is *entitled* to the same knowledge as the expert, as the former has no strategic reason to manipulate his information. The study of this game's equilibria thus contributes to the study of the foundation of expert based-knowledge.

Section 2 introduces the base model and provides important preliminary results (Proposition 1). Section 3 establishes general results regarding the structure of equilibria. No full revelation can happen at equilibrium: the sender never discloses his private information on models. Yet, when the interest of both parties is not too distant, partial information transmission can happen. All equilibria are outcome equivalent to those where the sender credibly points out an interval of models containing the most accurate one (Proposition 2). In general, multiple equilibria may exist in which the designated intervals are more or less broad (Proposition 3). In section 4, I show that when the receiver has MEU preferences, information transmission can only be conveyed for models below a given threshold, even if misalignment is arbitrarily small (Theorem 1). I show that the sender always prefers to convey as much information as possible. That is, I show that all equilibria can be ranked by informativeness and that the sender is always interim better off (i.e. after having learned his type) playing the most informative one (Theorem 2). This does not hold when the receiver has SEU preferences. In section 5, adapting from the CS linear-quadratic example with uniform prior, I show that when science is divided, the sender is ex-ante better-off facing a MEU receiver than an SEU one (Proposition 6). When science is consensual and that the consensus does not match the sender's preferences, he is ex-ante better-off then under full revelation (Theorem 3). Section 6 extends to α-MEU preferences and shows that whatever the misalignment, there is a degree of ambiguity aversion for which no information transmission is possible for types above a given threshold (Proposition 10 and 11). As an application to those results, section 7 explicitly models the COVID-19 example - a case of conflicting science - by introducing multiple receivers playing a game of contribution to a public good under model uncertainty. The sender, who knows the state generating model, aims at maximising social welfare. Proposition 13 shows that only models which recommend less social distancing can have an influence on contributors. Then, I suggest an ex-ante measure of the impact of behavioural responses on social welfare. I use it to argue that despite their limitations, MEU contributors dominate SEU ones to that respect. The appendix contains all generalisations and proofs.

2 Setup

2.1 Primitives

I consider a game of communication between an expert acting as a sender S (he), and a decision maker acting as a receiver R (she). Let $\mathcal{A} = \mathbb{R}$ be the set of actions of R and let $\Omega = \{0, 1\}$ is the set of possible states of nature. 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 $a_i(\omega) = \arg\max_{a \in \mathcal{A}} u_i(a, \omega)$ 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. In the context of our climate application, think of $\mathcal{A}$ as the level of GHG abatement and of Ω as the set of climate scenarios. I call $\omega = 1$ ($\omega = 0$) the high (low) state as it is the one where the optimal action is the highest (lowest). Thus, if $\mathcal{A}$ represents a level of contribution to a public good, as for GHG abatments, the high state is the one where climate damage is the highest. Conversely, if $\mathcal{A}$ captures a level of social distancing, as in the COVID-19 example, the high state is the one where the mortality of the virus is the highest. The choice of an abatement level is the result of a trade-off between economic growth (positively correlated with abatements) and potential damages created by global warming. Assumption 1 states that for any climate scenario, there is a single optimal abatement level. A higher abatement level $a > a_i(\omega)$ is not optimal for i because it might create too much climate damage. A lower abatement level $a < a_i(\omega)$ is neither optimal for i as it implies to reduces economic growth by too much.

There is model uncertainty in the sense that, ex-ante, it is not known according to which distribution the state is drawn. Instead, there is a family of Bernoulli distributions $\mathcal{D} = \{p_\theta | \theta \in [\underline{\theta}, \overline{\theta}]\}$, where $\underline{\theta}, \overline{\theta} \in [0, 1]$, that potentially generates the true state, where p_θ is the probability mass function of a Bernoulli distribution of parameter θ defined such that:

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

There is a bijection between the sets $\mathcal{D}$ and $\mathcal{C} = [\underline{\theta}, \overline{\theta}]$. In the rest of the paper, for simplicity, I will specify all the communication strategy on the set $\mathcal{C}$ which will be referred to as the set of *models*. 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 model θ, where $\mathbb{E}_\theta(u_i(a, \omega)) = (1 - \theta)u_i(a, 0) + \theta u_i(a, 1)$.

Assumption 2 (**Model misalignment**). *For any model, the optimal actions of S and R are always misaligned:*

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

Assumption 2 states that regardless of the model, there is always a difference of interest between S and R such that optimal actions are ordered in the same way[1]. 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 a sorting condition over models.

Lemma 1. *Assumption 1 implies that:*

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

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

2.2 Strategic interaction

Ex-ante, both players are in a situation of model uncertainty, also called ambiguity. In order to model the way R acts under model uncertainty, I will consider two separate cases. First, I will consider the case where they evaluate 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 will also draw attention to 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 models $\mu \in \Delta(\mathcal{C})$. Formally, $(\mathcal{C}, \mathcal{B}(\mathcal{C}), \mu)$ is a probability space, where $\mathcal{B}(\mathcal{C})$ is the Borel σ-algebra on $\mathcal{C}$ generated by the usual topology and μ is an absolutely continuous probability measure. In order to study a case of communication over models which is similar to CS, I will assume that in this second case, R knows the objective distribution according to which the model is drawn by nature. Thus, μ is an objective distribution and I also assume that $supp(\mu) = \mathcal{C}$. R then evaluates action a under uncertainty

[1] 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 actions of both players.

through:

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

where g is the probability distribution function of μ. In the following, the MEU case (respectively SEU case) is the one where R's evaluation of action coincides with the MEU (respectively SEU) decision criteria.

The timing of the game is as follows:

1. Nature draws the state generating distribution θ_0, according to μ. S is privately informed.
2. S sends a message regarding his type.
3. R updates her beliefs and chooses an action.

Having learned the state generating distribution[2] $\theta_0 \in \mathcal{C}$ from nature, S sends a message $m \in \mathcal{M}$, where $\mathcal{M} = [0,1]$ to R. A signalling strategy for S is the strategy $\sigma : \mathcal{C} \to \mathcal{M}$. An action rule for R is a strategy $y : \mathcal{M} \to \mathcal{A}$. Notice that I will focus only on pure strategies. Let $\sigma^{-1}(m) \subseteq \mathcal{C}$, be the set of potential types of S, in the eyes of R, 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 $\max_{m \in \mathcal{M}} V_S^\theta(m)$.

2. Having received an equilibrium message $m \in supp(\sigma^*)$, an MEU receiver updates her belief such that she 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 her prior using Bayes' rule such that:

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

[2] In appendix A I show that this assumption can be replaced by the one that the sender receives a noisy signal regarding models' likelihood

R then evaluates action a by:

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

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

As usual, any message m such that $m \notin supp(\sigma^*)$ is interpreted as some equilibrium message $m_* \in supp(\sigma^*)$.

Before directing our attention to the equilibria of this game and their specificity, it is useful to take a moment to study the players' decision making under uncertainty. This will help us understanding the specificities of the SEU and MEU decision criteria in my model and provide us with the intermediate results we need for the study of the equilibria. Given the assumptions and lemma 1, when the receiver evaluates action according to SEU, the model is very similar to CS. In that case, one can identify each model as a state in CS setting, where the payoff is the expected utility under that model and μ is the prior over states in CS. This case can thus be used as a benchmark regarding CS. Yet, when the receiver evaluates action according to MEU, the game dramatically changes. In order to see why, let me define the two following key elements:

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 cautious model.*

$\widetilde{a}$ is the action that maximises the function that gives the worst possible utility to the receiver. I call it the precautionary action because it is the optimal action anticipating that the worst state will always realise. Two special cases arise.

The first case is when one state is sufficiently worst than the other so that it gives a lower utility for any action $a \in [a_R(0), a_R(1)]$: $u_R(a,0) \leq u_R(a,1)$ or $u_R(a,1) \leq u_R(a,0)$. Then the precautionary action is the optimal action in the worst state: $\widetilde{a} = a_R(0)$ or $\widetilde{a} = a_R(1)$. Consider the example illustrated by Figure 1, where the high state is the worst. Then, whatever the set of models $\mathcal{C} \subset [0,1]$, for a given action the expected utility of the receiver is strictly decreasing in θ, the probability of the high state. In other words, all models agree that the higher the probability of high state, the worst R's welfare. Thus, there is a consensus among models when it comes to identify the *worst* state. Therefore, I'll say that it is a case of *consensual* science.

The second case is when no state is sufficiently worst than the other. Because of the single crossing assumption I made on utilities, both states must give the same utility for a given action in $(a_R(0), a_R(1))$. By construction, this action must be $\widetilde{a}$, the action that maximises the function that gives the worst possible utility to the receiver. Then, $\widetilde{\theta}$ is the model - not necessarily in $\mathcal{C}$ - for which the precautionary action is the optimal action. As illustrated by Figure 2, in this case, no state fully dominates the other in terms of utility for R. When $\widetilde{\theta} \in (\underline{\theta}, \overline{\theta})$, the expected utility of the receiver is decreasing for models putting a lower weight on the high state than the cautious one ($\theta < \widetilde{\theta}$) and increasing for the others ($\theta > \widetilde{\theta}$). Figure 2 illustrates this, as $\widetilde{\theta} = 0.5$ and the optimal action when $\theta = 0.2$ and

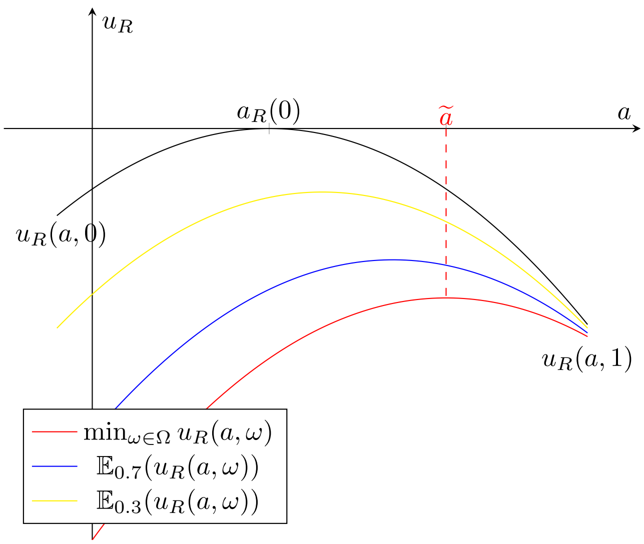

Figure 1: A case of consensual science : for all models, the receiver's welfare is decreasing with the probability of the high state.

$\theta = 0.8$ both improve the receiver's utility. Thus, there is no strict ordering over models with respect to R's expected utility but a division of $\mathcal{C}$ in two sets of models with opposite impact on R's expected utility. For models on one side of $\widetilde{\theta}$, the higher the probability of one state, the better R's welfare. But for models on the other side of $\widetilde{\theta}$, the higher the probability of that same state, the worse R's welfare. Thus, in this case, there is no consensus among models when it comes to identify the *worst* state. Therefore, I'll say it is a case of *divided science*.

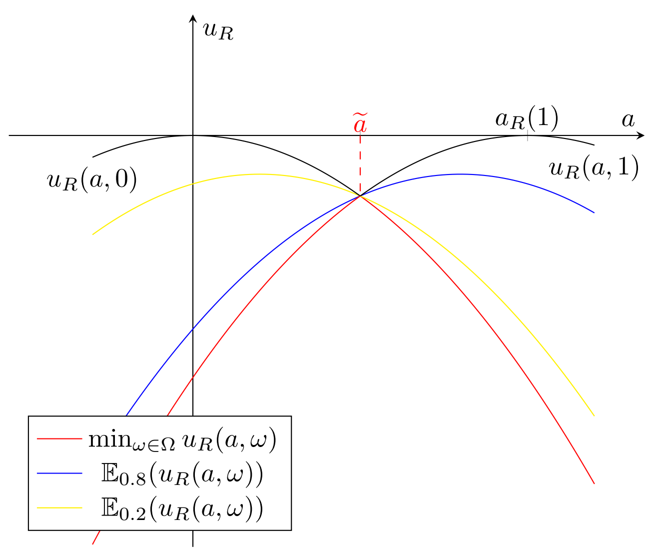

Figure 2: A case of divided science: for models above 0.5, the receiver's welfare is increasing with the probability of the high state. For models below 0.5, the opposite happens.

It is clear that the behavioural response of an MEU decision maker will be of a different nature, would science be divided or consensual. In the latter case, the precautionary action consists in anticipating the fully dominated state, thus acting as if the cautious model was the one putting the highest probability on that state. In the former case, the precautionary action consists in hedging against uncertainty, thus acting as if the cautious model was balancing odds between both states in the exact manner that leads to $\widetilde{a}$ as an optimal action. For $B \subset \mathcal{C}$, define $A_R(B) = argmax_{a\in\mathcal{A}} \min_{\theta\in B} \mathbb{E}_\theta(u_R(a,\omega))$ the optimal action of a MEU receiver given the set of priors B. Given these definitions, we can now state the following result:

Proposition 1. *Define $B = [\theta_1, \theta_2] \subset \mathcal{C}$ the set of priors of the receiver. The optimal action of a MEU receiver 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 when an MEU receiver believes that all models are below $\widetilde{\theta}$ ($\theta_0 \in [\theta_1, \theta_2]$ and $\theta_2 < \widetilde{\theta}$) she optimally acts as if the probability of the high state were maximal. When she believes that all models are above $\widetilde{\theta}$ ($\theta_1 > \widetilde{\theta}$) she optimally acts as if the probability of the high state were minimal. Notice that these two cases are the only possible ones when science is consensual. For instance, if the high state is fully dominated by the low state, $\widetilde{\theta} = 1$. Finally, when R believes that the cautious model could be the state generating model ($\widetilde{\theta} \in [\theta_1, \theta_2]$), she optimally acts as if it were the case.

Notice that in the SEU case the cautious action plays no particular part. As Proposition 5 will show, in the famous linear quadratic example of CS - a case of divided science - when R is MEU, her maximal payoff is increasing towards $\widetilde{\theta}$ and decreasing afterwards, while in the SEU case, it is constant as a function of the sender's type.

3 Equilibrium analysis

Let us now turn to the study of the equilibrium structure. First, I introduce the following definition:

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

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

A q-cut-off partition equilibrium is an equilibrium where there is a partition of the set of types in $q+1$ 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 in what cell the sender is 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 starts by showing that the number of actions induced at 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 actions 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. Notice that 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, we focus only on partition equilibria.

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

Proposition 3. *In any partition equilibria of the game (σ_q^*, y^*), the cut-off types $\theta_1^q, ..., \theta_q^q$ are defined such that for $k \in 0, ..., q-1$:*

$$V_S^{\theta_{k+1}^q}(y^*(m_k^q)) = V_S^{\theta_{k+1}^q}(y^*(m_{k+1}^q)) \tag{1}$$

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

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

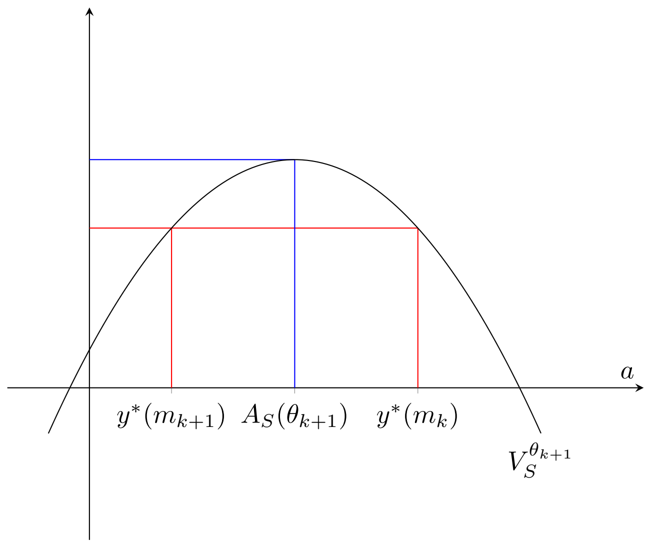

Figure 3: Identifying cut-offs

4 General results

In the following, I go one step further in the characterisation of the game's equilibria. Recall that in the MEU case, R evaluates action a by:

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

For any $m \in \mathcal{M}$, call $\sigma^{*-1}(m) = [\theta_1, \theta_2]$. Recall from Proposition 1 that:

$$A_R(\sigma^{*-1}(m)) = \begin{cases} A_R(\theta_2) \text{ if } \theta_2 < \widetilde{\theta} \\ A_R(\widetilde{\theta}) \text{ if } \widetilde{\theta} \in B \\ A_R(\theta_1) \text{ if } \theta_1 > \widetilde{\theta} \end{cases}$$

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

An upward misaligned sender is never capable of conveying information over $[\widetilde{\theta}, \overline{\theta}]$ when the receiver has MEU preferences. Recall the characterisation result of partition equilibria given by Proposition 3. For a θ_k to be a cut-off type, it must be that the message sent by types in the cell below and above θ_k induce actions that gives the same utility to S. As illustrated by Figure 4, when $\theta_k \geq \widetilde{\theta}$ and S is upwards misaligned, the action induced by any interval of types of lower bound θ_k is $A_R(\theta_k)$. It is necessarily strictly below the optimal action of the sender, $A_S(\theta_k)$. This is because only the lower bound of the designated interval of probabilities counts for R when $\theta \geq \widetilde{\theta}$.

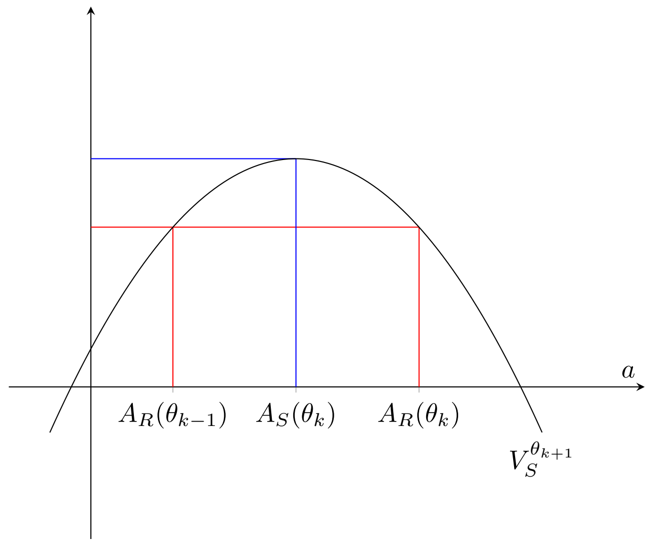

Figure 4: MEU cut-off above $\widetilde{\theta}$ for $A_S(\theta_k) > A_R(\theta_k)$

A consequence of Theorem 1 is that when science is consensual but all models are above $\widetilde{\theta}$, the only equilibrium is the babbling equilibrium. That is, whatever the sender's type, whatever the message he sends, the induced action is always the same. In other words, in this situation, the sender is inaudible. The next proposition shows that all partitional equilibria of our game can be built from the same finite set of cut-off types.

Proposition 4. *When the receiver evaluates actions following the MEU criteria and the sender is upwards misaligned, there are $M > 0$ partition equilibria. Call $\theta_M^- < ... < \theta_1^-$ the cut-offs of the equilibrium with most cut-offs. Then the cut-offs of any equilibrium are the q first elements of that sequence: $\theta_q^- < ... < \theta_1^-$, for $2 \leq q \leq M$.*

In a q cut-off equilibrium, for $k \in 0, ..., q-1$:

$$V_S^{\theta_{k+1}^-}(m_{k+1}^-) = V_S^{\theta_{k+1}^-}(m_k^-)$$

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

As illustrated by Figure 5, in the context of MEU preferences, the cut-off types of every equilibrium are the same. That is: the interior cut-off of the one cut-off equilibrium is the same as the first interior cut-off of the two cut-off equilibrium. Similarly, the second cut-off of the two cut-off equilibrium is the same as the second cut-off of the three cut-off equilibrium. The same holds for all cut-offs of the existing equilibria. In particular, there is always an equilibrium with two cut-offs - $\underline{\theta}$ and $\overline{\theta}$ - corresponding to the babbling equilibrium.

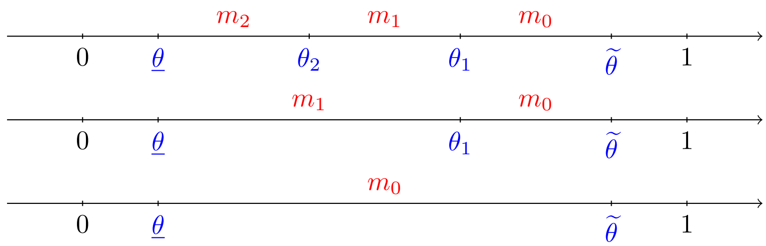

Figure 5: MEU equilibria for $\underline{\theta} < \widetilde{\theta}$

Given Theorem 1, all cut-offs are in $[\underline{\theta}, \widetilde{\theta}]$. As a result, when S points out an interval of models, R only cares about its upper bound. As a result, cut-offs types will not be determined by an indifference between pairs of expectations of intervals of models (the two adjacent cells) but by an indifference between pairs of expectations of models (the lower bounds of the two adjacent 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 set, 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. Then, all other equilibria are characterised by the k first terms ($1 \leq q \leq M$) of that sequence, assuming it has M elements.

A direct consequence of Proposition 4 is that all equilibria of the game can be ranked by informativeness. Notice that this is never possible in the SEU case. A direct corollary can be established regarding interim Pareto dominance among equilibria.

Theorem 2. *When the receiver has MEU preferences:*

1. *The sender is always interim weakly better off by playing the most informative equilibrium strategy*
2. *When the sender's type is not in a cell containing $\widetilde{\theta}$, he is strictly better off playing the most informative equilibrium strategy*

The intuition of the proof, for S downwards misaligned, is the following. Consider 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 the 3 cut-off equilibrium interim Pareto 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 Pareto dominates the 3 cut-off one.

Thus, S is always interim better-off communicating following the most informative equilibrium strategy. 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 is maximal for the equilibrium with most cut-off. Condition (M) is satisfied if for any two sequence of cut-off types the k-th cut-off of each sequence 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 suggests 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. Unlike in CS, in my case, one could adopt interim Pareto dominance as a selection criterion, which is immune to the limitations of ex-ante Pareto dominance. Yet, it brings out the same (most informative) equilibrium and provides a foundation for the attention it receives in applications.

The question of equilibrium selection in the CS framework has been recently tackled in Chen et al. (2008). In this paper, a condition on utility functions, NITS, has been proposed. 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 his type (if he could). NITS is not satisfied in my framework. To see this consider the case of the babbling equilibrium. Would R know that the sender is of the lowest type, her optimal action would be $A_R(\underline{\theta})$. Yet, at equilibrium, the optimal action of the receiver is $A_R(\widetilde{\theta})$. Were misalignment small enough for there to be at least one cut-off below $\widetilde{\theta}$, then, by definition, the sender would prefer $A_R(\underline{\theta})$ to $A_R(\widetilde{\theta})$.

5 The linear-quadratic example

In order to give a further insight of the results in the MEU case, I start by characterising all partitional equilibria in the context of the following example adapted from the widely used linear-quadratic example of CS. I also provide the same characterisation for the SEU case.

Linear-quadratic example:

- $u_S(a, \omega) = -(a - \omega - b)^2$, where $b > 0$
- $u_R(a, \omega) = -(a - \omega)^2$
- $\mathcal{C} = [0, \overline{\theta}]$, where $\overline{\theta} \geq \frac{1}{2}$ and $\mu \sim U([0, 1])$

Then, the ex-post utilities in my game are:

$$\begin{cases} \mathbb{E}_{\theta_0}(u_S(a, \omega)) = -(1 - \theta_0)(a - b)^2 - \theta_0(a - 1 - b)^2 \\ \mathbb{E}_{\theta_0}(u_R(a, \omega)) = -(1 - \theta_0)a^2 - \theta_0(a - 1)^2 \end{cases}$$

and the corresponding optimal actions are:

$$\begin{cases} A_S(\theta_0) = \theta_0 + b \\ A_R(\theta_0) = \theta_0 \end{cases}$$

Notice that players' optimal actions ex-post are similar to the ones of CS: the receiver tries to match the model and the sender does the same, plus the bias. Here, $u_R(a, 0) = u_S(a, 1)$ when $a = \frac{1}{2}$. As a result, for $\overline{\theta} \geq \frac{1}{2}$, the hedging model is $\widetilde{\theta} = \frac{1}{2}$. As a result, when $\overline{\theta} = \frac{1}{2}$ the linear-quadratic example is a case of consensual science. When $\overline{\theta} > \frac{1}{2}$ it is a case of divided science. I start by characterising all cut-offs of the game in both the SEU and MEU cases.

Proposition 5. *In the context of our linear-quadratic example:*

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

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

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

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

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

$$0 < b < \frac{1}{4n} \tag{3}$$

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

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

A corollary of Proposition 3 is that it is possible to characterise each equilibrium's cell sizes.

Corollary 1. *Consider a q-cut-off partition equilibrium. When $b > 0$ and R is SEU, 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, non terminal cells are of constant size. For all $k \in 1,...,q-2$:

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

In the MEU case non terminal cells have always the same size ($2b$), whatever the considered equilibrium. In the SEU case, it depends on the considered equilibrium. This illustrates why the general result proved in Theorem 4 holds. If all non terminal cells have the same size in any given equilibrium, and that in addition the first cut-off is always the same (as proven in Proposition 5), it is straightforward that those equilibria can be ranked by informativeness in the Blackwell sense. Corollary 1 also states that in the SEU case, wells are at least of size $4b$ and are thus always strictly larger.

Overall, the structure of equilibria can be summarised by Figure 6. Every shaded zone corresponds to the set of distributions the receiver will consider at equilibrium. The blue shaded area corresponds to the expected utility of R, when models are in $[\frac{1}{2}, 1]$. The green area corresponds to R's expected utility when models are in $[0.3, 0.5]$, the yellow one to $[0.1, 0.3]$ and the brown one to $[0, 0.1]$.

5.1 An example of divided science

In this subsection I now assume that $\overline{\theta} = 1$ making my example a case of divided science. In the SEU case, this case is the exact reproduction of CS's linear quadratic example when communication is over states.

First, notice that because science divided, in the MEU case all cut-offs are concentrated on $[0, \frac{1}{2}]$. In the SEU case are distributed on the entire $[0, 1]$ segment. When b converges to 0, the paritional equilibrium with most cut-offs converges to full revelation in the SEU case, whereas in the MEU one this only happens on $[0, \frac{1}{2}]$. In our example,

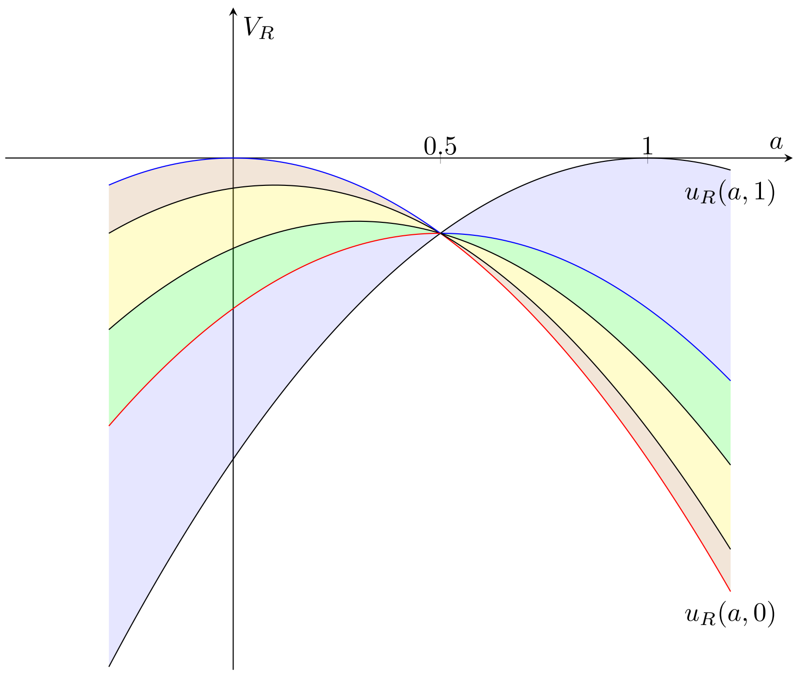

Figure 6: MEU equilibrium valuations for $b = 0.2$

whatever the degree of positive misalignment between S and R, information transmission can happen only regarding models below $\frac{1}{2}$.

Second, one can see that whatever the bias, the sender is able ta partition the set of types much more when the receiver is MEU. Consider a given positive bias such that it is possible to get a n cut-off equilibrium with a 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 a n-cut-off equilibrium is possible in the MEU case $b_M(n) = \frac{1}{4n}$. 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 $n \geq 2$, $b_S(n) = b_M(\frac{n(n+1)}{2})$. Thus, there is a n cut-off equilibrium between an SEU receiver of bias b and the sender if and only if there is a $\frac{n(n+1)}{2}$ cut-off equilibrium between an MEU receiver of bias b.

In the following, I will compare the ex-ante welfare of the sender in two versions of the game where in one the receiver is SEU and one where the receiver is MEU. It will be natural to consider the $\frac{n(n+1)}{2}$-cut-off equilibrium for the MEU version (as the ex-ante and interim Pareto dominant and most informative equilibrium) and the n-cut-off equilibrium for the SEU one (as the ex-ante Pareto dominant).

Proposition 6. *In the context of our running example, for a given bias $b > 0$, the sender is ex-ante better-off in the partitional equilibrium with the maximum number of cut-offs of the MEU version of the game ($\frac{n(n+1)}{2}$) compared to the partitional equilibrium with the maximum number of cut-offs of the SEU version of the game (n) for $n \geq 2$.*

When comparing the ex-ante welfare of the sender in both games, two effects go in opposite directions. On one hand, equilibria with much more cut-offs can be sustained, for a given bias, in the MEU version of the game. On the other hand, these cut-offs are all on one side of $\widetilde{\theta}$ while, ex-ante, all $\theta \in [0,1]$ are equally likely. Proposition 6 shows that the former effect dominates even the case of divided science where S can't communicate to an MEU receiver on $(\frac{1}{2}, 1]$.

5.2 An example of consensual science

I now turn to a case of consensual science by assuming that $\overline{\theta} = \frac{1}{2}$. Given that $b > 0$ it is a case where communication is possible on the entire set $\mathcal{C}$ when R is MEU.

Theorem 3. *In the context of our linear-quadratic example with $\mathcal{C} = [0, \frac{1}{2}]$, when R has MEU preferences the sender can induce an action on average closer to his preferred action than under full revelation. When he does, he is ex-ante better-off than under full revelation*

This result follows from the fact that when R is MEU, all non terminal cells are of length $2b$. As the equilibrium action of R is one bound of a cell and the expected preferred action of the sender is average between both bounds, the distance between them is always b. The terminal cell is, by construction, smaller then the others. As a result, the distance of actions is also smaller and the average ex-ante distance is at most b. The ex-ante welfare of the sender is decreasing with this distance and computation shows that it is higher in the most informative equilibrium than under full revelation.

This result shows all the strength of cheap-talk communication when R has MEU preferences and that science is consensual. In that case, the sender can, on average, induce an action closer to his preferred one then under full revelation. This can never happen when R is SEU or in CS's model of communication over states : the induced action is always further away of the sender's preferred action then under full revelation. This result changes the interpretation of what cheap-talk communication is in our model. Under SEU preferences, the model describes a situation where the expert would always be better-off if he could certify his information to the decision maker. Here, cheap-talk is not a way for him to manipulate the decison maker to his advantage, but a stopgap to the burden that strategic misalignment represents for him. But when the consensus of science is favourable to the sender, the receiver MEU preferences makes cheap-talk communication a powerful tool of persuasion. This contrasts dramatically with the case where $\mathcal{C} \subset [\frac{1}{2}, 1]$, where Theorem 1 implies that no communication can happen at all.

6 α-MEU receiver

6.1 Optimal actions and structure of equilibria

In this section, I consider the case where R evaluate actions under uncertainty through the α-maxmin decision criteria proposed by Ghirardato et al. (2004)(α-MEU). According to Ghirardato et al. (2004), in addition to their utility function, players are characterised by two more elements. First, a set of priors over Ω, which I will assume to be $\mathcal{C}$. Second, a parameter $\alpha_i \in [0, 1]$ which captures their attitude towards ambiguity. As all the analysis will be conducted at the interim stage, α_S is irrelevant. Thus, in the following, I will erase the subscript. R evaluates action $a \in \mathcal{A}$ by :

$$V_R^{\alpha}(a) = \alpha \min_{\theta \in \mathcal{C}} \mathbb{E}_{\theta}(u_R(a, \omega)) + (1 - \alpha) \max_{\theta \in \mathcal{C}} \mathbb{E}_{\theta}(u_R(a, \omega)) \tag{4}$$

Here, the behavioural consequences of ambiguity aversion are captured by α. It translates the decision maker's weighting between optimistic and pessimistic models regarding his expected utility. When $\alpha = 1$, the α-MEU decision criteria coincides with MEU[3]. Adapting Ghirardato et al. (2004)'s proposition 20 to our model one can state the following :

Definition 3 (Ghirardato et al. (2004)). *Receiver i, evaluating actions through $V_R^{\alpha_i}$, is said to be more ambiguity averse*[4] *than receiver j, evaluating actions through $V_R^{\alpha_j}$, if and only if*

$$\alpha_i > \alpha_j$$

Thus, for a fixed set of priors and utility function, increasing ambiguity aversion leads the receiver to anticipate an increasingly worst model in terms of expected utility.

As for the MEU case, having received an equilibrium message $m \in supp(\sigma^*)$, an α-MEU receiver updates her belief such that she evaluates action a by:

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

By a natural extension of the notations introduced above, for $B \subset \mathcal{C}$, define $A_R(B) = argmax_{a \in \mathcal{A}} V_R^{\alpha}(a, B)$ the optimal action of the α-MEU receiver when his belief is $\theta_0 \in B$.

[3]Notice that α-MEU does not have SEU as a special case here. Both criteria coincide only if the set of models is a singleton. Yet, this set depends of the information conveyed by the sender, which is never a singleton.

[4]In the sense of Ghirardato and Marinacci (2002)

Figure 7 illustrates the ex-ante evaluation of actions of the receiver in the context of the linear-quadratic example. All valuation functions are located in the blue area and are a convex combination between $\min_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a,\omega))$ (in red) and $\max_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a,\omega))$ (in black). Then, notice that for a given $\alpha \neq 1$ $V_R^\alpha(a)$ is not necessarily single-peaked. For instance, for $\alpha = 0.3$, $V_R^{0.3}(a)$ is maximal at 0.3 and 0.7[5].

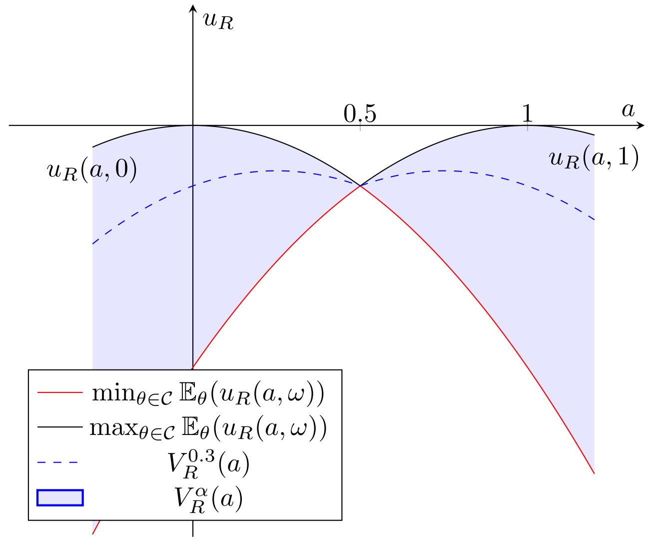

Figure 7: α-MEU ex-ante valuation

I now characterise the optimal action of R for a given set of priors.

Proposition 7. *Define $B \subset \mathcal{C}$ the set of priors of the receiver with minimal element θ_1 and maximal element θ_2. Given this belief, her optimal action $A_R(B) \in [A_R(\theta_1), A_R(\theta_2)]$. In the context of the linear-quadratic example, the optimal action of a α-MEU receiver is given by:*

$$A_R(B) = \begin{cases} \alpha A_R(\theta_2) + (1-\alpha) A_R(\theta_1) \ \textit{if } \theta_2 < \widetilde{\theta} \\ \alpha A_R(\widetilde{\theta}) + (1-\alpha) A_R(\theta_M) \ \textit{if } \widetilde{\theta} \in B \\ \alpha A_R(\theta_1) + (1-\alpha) A_R(\theta_2) \ \textit{if } \theta_1 > \widetilde{\theta} \end{cases}$$

where $\theta_M = argmax_{\theta \in \{\theta_1, \theta_2\}} \mathbb{E}_\theta(u_R(a,\omega))$

A direct consequence of Proposition 7 is that ex-ante, for any $a \in \mathcal{A}$, $\min_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a,\omega)) = \mathbb{E}_{\widetilde{\theta}}(u_R(a,\omega))$ and that $\max_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a,\omega)) = \max(\mathbb{E}_{\underline{\theta}}(u_R(a,\omega), \mathbb{E}_{\overline{\theta}}(u_R(a,\omega))$. Thus, $A_R^\alpha(\mathcal{C}) = \max\Big(A_R(\alpha\widetilde{\theta} + (1-\alpha)\overline{\theta}), A_R(\alpha\widetilde{\theta} + (1-\alpha)\underline{\theta})\Big)$, which explains the fact that in the example we considered before, optimal actions where not unique. Notice also that, when α increases, $A_R^\alpha(\mathcal{C})$ gets closer to $\widetilde{a}$, in the euclidean sense. Thus, an increase in ambiguity aversion gets R's ex-ante optimal action closer to the precautionary action.

[5]In our example there are two optimal actions for any $\alpha < 0.5$. The fact that this threshold is the one separing love and aversion for ambiguity is non-generic. For sharper utility functions - in the sens of a lower second order derivative - this threshold would be above 0.5

I now prove that under α-MEU preferences, all equilibria are still outcome equivalent to a partition equilibria.

Proposition 8. *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.*

As for the proof of Proposition 2, I start by showing that the number of actions induced at 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 actions of R for a given belief $B \subset \mathcal{C}$ is in the convex hull of the optimal actions for every element of B. This is also true when R has α-MEU preferences, as one can deduce from Proposition 7. 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.

6.2 Comparative ambiguity aversion

In the following, I'm interested in the effect that ambiguity aversion has on the structure of partitional equilibria. I compare the equilibria of two versions of the game, where the only difference is the degree of ambiguity aversion of the receivers identified to their degree of ambiguity aversion α_1 and α_2. Notice that the ex-post optimal action $A_{R_i}(\theta)$ is unaffected by ambiguity aversion, thus, I will erase the subscript.

I start by considering a case of consensual science. Consider the linear-quadratic example. For it to be a case of consensual science, it must be that $\mathcal{C}$ does not contain $\widetilde{\theta}$ as an interior point. I consider the two following examples : the set of models is either $\underline{\mathcal{C}} = [0, \frac{1}{2}]$ or $\overline{\mathcal{C}} = [\frac{1}{2}, 1]$. In the following, I characterise all the cut-offs of the corresponding partition equilibrium.

Proposition 9. *In the linear quadratic example, when R is α-MEU, for $\alpha \notin \{0, \frac{1}{2}, 1\}$:*

- *If the set of models is $\underline{\mathcal{C}} = [0, \frac{1}{2}]$, there are $N > 0$ cut-off equilibria, one by number of cut-offs, and the k-th cut-off of the $1 \leq n \leq N$ cut-off equilibrium is given by :*

$$\underline{\theta}_k^n(\alpha) = (\frac{1}{2} - \frac{2bn}{2\alpha - 1})\Big(\frac{1 - (\frac{1-\alpha}{\alpha})^k}{1 - (\frac{1-\alpha}{\alpha})^n}\Big) + \frac{2bk}{2\alpha - 1}$$

- *If the set of models is $\overline{\mathcal{C}} = [\frac{1}{2}, 1]$, there are $M > 0$ cut-off equilibria, one by number of cut-offs, and the k-th cut-off of the $1 \leq n \leq M$ cut-off equilibrium is given by :*

$$\overline{\theta}_k^n(\alpha) = (\frac{1}{2} - \frac{2bn}{2\alpha - 1})\Big(\frac{1 - (\frac{\alpha}{1-\alpha})^k}{1 - (\frac{\alpha}{1-\alpha})^n}\Big) - \frac{2bk}{2\alpha - 1} + \frac{1}{2}$$

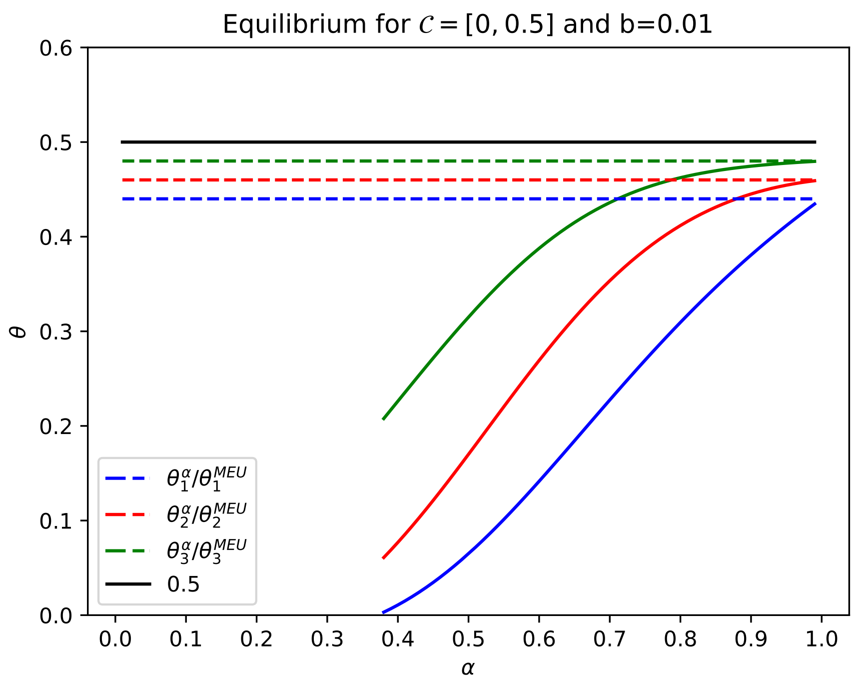

Figure 8: 3-cut-offs equilibria in $\underline{\mathcal{C}}$

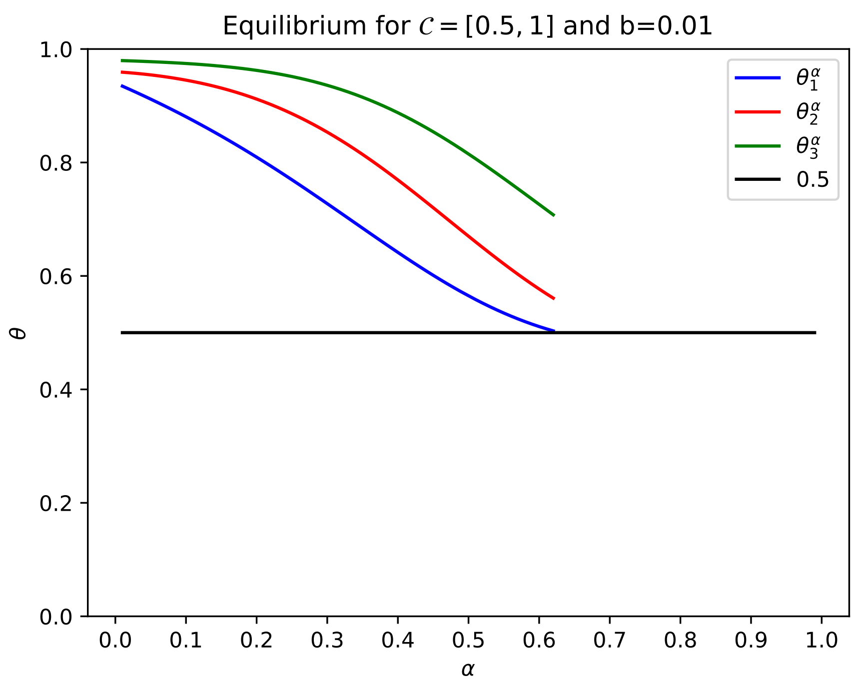

Figure 9: 3-cut-offs equilibria in $\overline{\mathcal{C}}$

Recall that in the MEU case, when the set of models is $\overline{\mathcal{C}}$ and the bias is positive, no information transmission is possible. when the set of models is $\underline{\mathcal{C}}$ and the bias is positive, information transmission is possible on the entire set. Figure 8 and 9 compute the cut-offs of the 3-cut-offs equilibrium as a function of α for a fixed positive bias $b = 0.01$. Notice that for a given level of misalignment, information transmission is possible in both cases, for given levels of ambiguity aversion. Thus, the asymmetry of the MEU case does not survive at any level of ambiguity aversion. Yet, simulations suggest that, when the set of models is $\overline{\mathcal{C}}$, cut-off values decrease with α towards $\widetilde{\theta}$. Thus for a given bias, there is a level of ambiguity aversion from which all types in $\overline{\mathcal{C}}$ must pool. Conversely, simulations suggest that, when the set of models is $\underline{\mathcal{C}}$, cut-off values continuously increase with α towards their MEU values.

I know formally prove that in the case of consensual science, no communication is possible in $\overline{\mathcal{C}}$ when α is above a given threshold. In addition, I show that for a given bias, ambiguity aversion eases the existence of a n cut-off equilibrium, for $n \geq 2$.

Proposition 10. *In the context of the linear-quadratic example, when* $\mathcal{C} = [\frac{1}{2}, 1]$*,* $b > 0$ *and* $\alpha \in (\frac{1}{2}, 1)$ *:*

1. *There is* $\alpha(b) \in (1/2, 1)$ *such that for* $\alpha \geq \alpha(b)$*, no information transmission is possible in* $[\frac{1}{2}, 1]$*. Moreover,* $\alpha(b)$ *is a decreasing function.*
2. *For two receivers* α_1 *and* α_2 *such that* $\alpha_1 < \alpha_2$*, if there is a* $n \geq 2$ *cut-off equilibrium between* S *and* α_1*, there is a* n *cut-off equilibrium between* S *and* α_2

Thus, as suggested by the simulations, for a given bias, there is a level of ambiguity aversion from which all types in $\overline{\mathcal{C}}$ must pool. This follows from the fact that for any $n \geq 2$, $\overline{\theta}^n_{n-1}(\alpha)$ is a strictly decreasing and continuous function and that $\lim_{\alpha \to +\infty} \overline{\theta}^n_{n-1}(\alpha) < \frac{1}{2}$. As a result, there must be $\alpha \in (\frac{1}{2}, 1)$ such that no partitionning of $\overline{C}$ is possible.

In addition, when there is an equilibrium with at least 2 cut-offs, ambiguity aversion eases the existence of a n cut-off equilibrium. Recall bounds of $\mathcal{C}$ are included in the count, which means that we are looking at every equilibrium which is a non-babbling one. In other words, for a given bias, increasing ambiguity aversion might enable the existence of a k-cut-off equilibrium which was not sustainable for a lower level of ambiguity aversion. In that sense, ambiguity aversion eases information transmission, when science is consensual. This second result follows from the fact that for any $n \geq 2$, $\underline{\theta}^n_{n-1}(\alpha)$ is a strictly increasing function and that $\lim_{\alpha \to +\infty} \underline{\theta}^n_{n-1}(\alpha) < \frac{1}{2}$.

I know further prove that the first result of Proposition 10 extends to a case of divided science :

Proposition 11. *In the context of the linear-quadratic example, when* $\mathcal{C} = [0, 1]$*,* $b > 0$ *and* $\alpha \in (\frac{1}{2}, 1)$*, there is* $\alpha(b) \in (1/2, 1)$ *such that for* $\alpha \geq \alpha(b)$*, only one action can be induced by types in* $[\frac{1}{2}, 1]$*. Moreover,* $\alpha(b)$ *is a decreasing function.*

As for the consensual science case, for a given bias, there is a level of ambiguity aversion from which all types in $\overline{\mathcal{C}}$ must pool. This suggest that there is a form of continuity in the division of the set of types - on both sides of the hedging model - that we have observed in the MEU case. For any level of misalignment of S, there is degree of ambiguity aversion of R in $(\frac{1}{2}, 1)$ such that all models above $\widetilde{\theta}$ must pool. The proof of Proposition 11 builds on the one of proposition Proposition 10. I show that for any $n \geq 3$, $\overline{\theta}^n_{n-2}(\alpha)$ is a strictly decreasing and continuous function and that $\lim_{\alpha \to +\infty} \overline{\theta}^n_{n-2}(\alpha) < \frac{1}{2}$. As a result, there must be $\alpha \in (\frac{1}{2}, 1)$ such that no partitionning of $\overline{C}$ is possible. Then I prove that in the case $n = 2$, the only interior cut-off is always in $(0, \frac{1}{2})$.

However, contrary to the previous case, when science is divided, one cannot say that ambiguity aversion always eases information transmission. The following proposition characterises equilibrium with a single interior cut-off in $\mathcal{C} = [0, 1]$.

Proposition 12. *In the context of our linear-quadratic example, when $\mathcal{C} = [0, 1]$, $b > 0$ and $\alpha \in (\frac{1}{2}, 1)$, a 3-cut-off equilibrium exists if and only if :*

$$0 < b < \frac{1}{2} - \frac{\alpha}{4}$$

and the cut-off is :

$$\theta(\alpha) = \frac{2 - \alpha - 4b}{4 - 2\alpha}$$

Notice that, as illustrated by Figure 10, $\theta(\alpha)$ is a decreasing function of α, bounded above by $\frac{1}{2}$ and below by $\theta^{MEU} \equiv \lim_{\alpha \to 1} \theta(\alpha)$ the interior cut-off of the MEU 2-cutoff equilibrium. The existence constraint of that equilibrium is strictly decreasing with α. As a result, if R_1 is a receiver with given ambiguity aversion and R_2 the same receiver with increased ambiguity aversion, there can be a 2-cut-off equilibrium between R_1 and S but not between R_2 and S. Thus, ambiguity aversion has an opposite effect in this case compared to the one of consensual science we examined in Proposition 10.

7 An application : multiple receivers contributing to a public good

In the following section, I will extend the game considered so far by introducing two receivers whose actions are strategic complements. In the first stage of the game, as before, an expert acting as a sender S publicly communicates about the state generating model to the receivers. Receivers then play a game of contribution to a public good. This can be seen as a good representation of scientific authority communicating over existing knowledge on the COVID-19 pandemic. As a result, receivers will choose a level of contribution to the public good that social

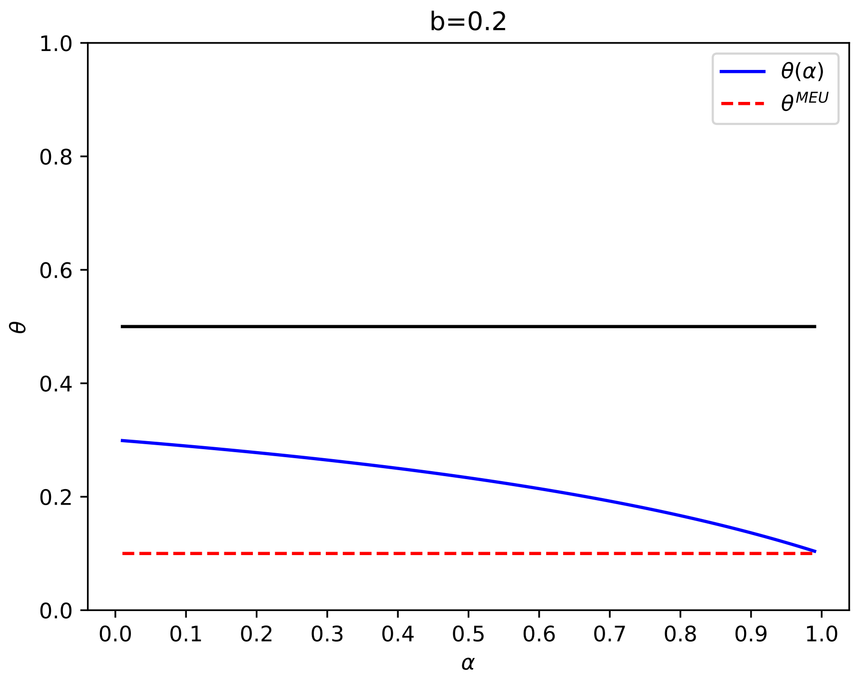

Figure 10: One cut-off equilibria

distancing constitutes in that context. In the following, actions will be called distancing and will be assumed positive $a_1, a_2 \in \mathcal{A} = \mathbb{R}^+$. As before, two epidemic scenarios are possible $\Omega = \{0, 1\}$. The timing of the game is as follows:

1. Nature selects the type of the sender, θ_0, which is privately informed.
2. **Communication stage:** The sender sends a message to the receivers regarding his type.
3. **Contribution stage:** Receivers simultaneously choose a level of distancing.

Receiver i's utility function will be:

$$u_i(a_i, a_{-i}, \omega) = (a_i + a_{-i} - \omega + 1)^2 - ra_i$$

The linear part of u_i represents the fact that distancing negatively correlates with economic growth. The $r > 0$ parameter captures the degree of correlation between social distancing and growth. The quadratic part of u_i represents the benefits of social distancing for player i. Benefits are increasing with the level of total social distancing. Player i does not take into account the impact of his actions on the benefits of social distancing for player $-i$. As a result, social distancing is a public good and its collective consequences are not internalised by players.

As before, R_i ex-ante evaluates (a_i, a_{-i}) through:

$$V_{R_i}^{MEU}(a_i, a_{-i}) = \min_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a_i, a_{-i}, \omega))$$

in the MEU case and through

$$V_{R_i}^{SEU}(a_i, a_{-i}) = \int_{\theta \in \mathcal{C}} g(\theta)\mathbb{E}_\theta(u_R(a_i, a_{-i}, \omega))d\theta$$

in the SEU one.

I call $A_i(a_{-i}, \theta) = argmax_{a_i \in \mathcal{A}} V_{R_i}(a_i, a_{-i})$ the optimal distancing level of R_i under model θ when R_{-i} distancing level is a_{-i}.

The sender seeks to maximise social welfare. That is:

$$u_S(a_1, a_2, \omega) = \sum_{i=1}^{2} \left[(a_i + a_{-i} - \omega + 1)^2 - ra_i\right]$$

S ex-ante evaluates the total distancing level $a_i + a_{-i}$ through $\mathbb{E}_{\theta_0}(u_S(a_1, a_2, \omega))$

I start by focusing on the contribution stage. Given that communication is public, at equilibrium, both receivers will have the same belief regarding θ_0. In the following proposition, for a given common belief, I characterise the equilibrium level of distancing of the receivers. I restrict my attention to intervals of $\mathcal{C}$, anticipating on the partitional nature of the sender's equilibrium strategy in the communication stage.

Proposition 13. *For a given $\theta \in \mathcal{C}$, the overall level of social distancing of the receivers $a_1^*(\theta), a_2^*(\theta) \in \mathcal{A}$ must be such that:*

$$a_i^*(\theta) + a_{-i}^*(\theta) = \frac{r}{2} + \theta$$

the socially optimal levels of distancing $a_i^P \in \mathcal{A}$ must be such that:

$$a_i^P(\theta_0) + a_{-i}^P(\theta_0) = \frac{r}{4} + \theta_0$$

Regarding the overall level of distancing, The game is outcome equivalent to the parallel one sender, one receiver communication game where the action variable is the total level of distancing, $t = a_1 + a_2$ and the utility of the sender is:

$$U_S(t, \omega) = -(t - \frac{r}{4} - \omega + 1)^2$$

and the utility of the receiver is:

$$U_R(t, \omega) = -(t - \frac{r}{2} - \omega + 1)^2$$

A consequence of Proposition 13 is that all equilibria are outcome equivalent to a partitional equilibrium. Notice that misalignment is increasing with r. This means that the more the economies are affected by social distancing, the more a social welfare maximising sender is misaligned with the receivers. As the hedging model of the parallel game is $\widetilde{\theta} = \frac{1}{2}$ and S is upwards misaligned, Proposition 1 implies that whatever the level of dependence of the receivers to social distancing, information transmission can only happen regarding models below $\widetilde{\theta}$. That is, for the models which recommend a low optimal level of social distancing, below $\frac{1}{2}$

In the spirit of Proposition 6, one can show that the sender is always ex-ante better-off in the MEU version of the game compared with the SEU one, when the prior over models is uniform (if there is a non-babbling equilibrium in the SEU version of the game). Here, ex-ante analysis is an interesting tool regarding welfare analysis. One can see it as comparing the efficiency of MEU receivers' strategic reaction to cheap talk communication to the one of SEU receivers, in the eyes of the expected ex-post social welfare. This measure can be seen as objective in the sense that it measures the total ex-post utility of receivers, for given equilibrium actions, weighted by μ, the objective probability of each model. Formally, call $W(a_1, a_2)$ the expected social welfare of the receivers when they play actions a_1 and a_2:

$$W(a_1, a_2) = \int_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_1(a_1, a_2, \omega) + u_2(a_2, a_1, \omega))d\mu$$

The ex-ante welfare of the sender in the MEU version of the game is $W(a_1^{M*}, a_2^{M*})$, where $a_i^{M*} = y_i^{M*}(\sigma^*(\theta))$ is the equilibrium distancing level of of receiver i in the $\frac{n(n+1)}{2}$-cut-off equilibrium of the MEU version of the game. Similarly, the ex-ante welfare of the sender in the SEU version of the game is $W(a_1^{S*}, a_2^{S*})$, where $a_i^{S*} = y_i^{S*}(\sigma^*(\theta))$ is the equilibrium distancing level of receiver i in the n-cut-off equilibrium strategies of the SEU version of the game. When there is a non-babbling equilibrium in the SEU version of the game, we thus have that $W(a_1^{M*}, a_2^{M*}) > W(a_1^{S*}, a_2^{S*})$.

Regarding the normative choice of a decision maker, this result brings a rather unexpected result. The fact that receivers display ambiguity aversion leads to a higher welfare, for receivers, then if they would have SEU preferences. This might specially feel unexpected if one sees ambiguity aversion as a behavioural bias which can only harm those who are subject to it. Having leaders who display MEU preferences can lead for a better social outcome then while facing a SEU ones. In this regard, this result contributes to a literature interested in the comparative robustness of ambiguity neutral and averse preferences. For instance, Condie (2008) and Guerdjikova and Sciubba (2015) study the survival of ambiguity neutral and averse decision makers on financial markets.

8 Conclusion

This paper models the transmission of expert-based scientific knowledge as cheap-talk communication over models, in a framework similar to Crawford and Sobel (1982). Because models can be represented as probability distributions, a receiver of this game can naturally be assumed to be ambiguity sensitive. For every preferences I considered, I showed that all equilibria are outcome equivalent to a partition equilibria. When the receiver is MEU, information transmission can only happen for models below a given threshold, even if misalignment is arbitrarily small. In addition, the sender always prefers to convey as much information as possible as the most informative equilibrium is interim Pareto dominant. This is not true when the receiver has SEU preferences, a case which is equivalent to the model of communication over states proposed in Crawford and Sobel (1982). In the linear-quadratic example with a uniform prior over models, this structure of equilibria leads to a situation which is also ex-ante dominant for the sender, when the receiver is MEU. This special case also shows that when the expert's preferred action is aligned with the effect of ambiguity aversion, his influence is extremely high; but in the opposite case, it is nonexistent. Assuming that the receiver has α-MEU preferences allows to show the robustness of these results : whatever the misalignment, there is a degree of ambiguity aversion for which no information transmission is possible for types above a given threshold. Increasing the ambiguity aversion parameter can have different effects. When science is consensual, ambiguity aversion always eases information transmission. But when science is divided, ambiguity aversion can harden it. Finally, by explicitly modelling the COVID-19 pandemic as an iconic case of conflicting science, I showed that only models which recommend less social distancing can have an influence on MEU contributors. Thus, my results show that, because of the epistemic nature of expert-based knowledge, when faced with a choice between the lesser of two evils, decision makers will tend to be unconvinced by scientific results which recommend more collective efforts. To the contrary, when the state which demands more collective efforts is also always worst then all others - as in the case of climate change - ambiguity aversion is a powerful ally for the transmission of expert-based scientific knowledge.

There are, of course, limitation to this work. In my view, the main avenue for future research is the extension of the current framework to multiple experts. Many situations where expert-based scientific knowledge plays an important part involve multiple senders, communicating both simultaneously or sequentially. Existing work in the context of communication over states, such as Battaglini (2002) or Krishna and Morgan (2001) show that significant differences can appear.

Appendix

A Supplementary Assumptions

A.1 Assumptions on states

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

Assumption 3 (**Misalignment** - 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 11. I call that property sharpness, in the sense that it translates a sharper preference for the optimal action.

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

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

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

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 $A_i(\theta)$. We 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 $h^\theta : a \to f_R^\theta(a) - f_S^\theta(a)$.

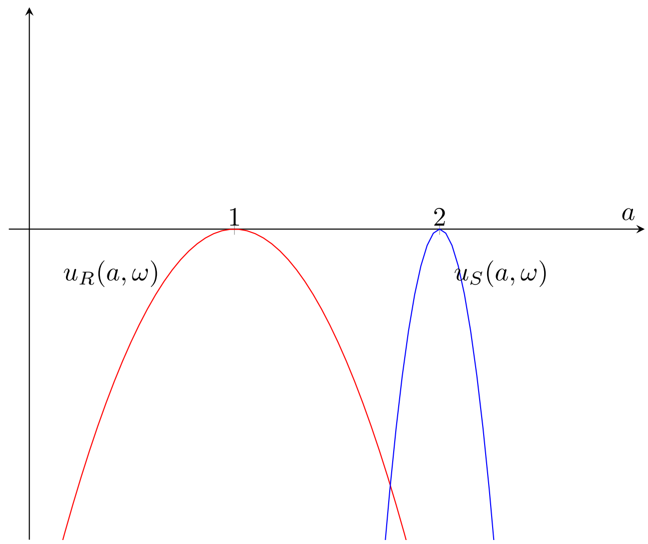

Figure 11: Sharpness Assumption

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

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

□

Lemma 2 states that whatever the realised model, R and 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 $A_S(\theta) = A_R(\theta)$.

A.2 Imperfect knowledge of the model

In the following I show that the assumption that the sender observes the state generating distribution - the *true* model - can be replaced without significant change in the result. Instead, I will assume that S observes a noisy signal regarding the state generating distribution. I focus on the case where both players have MEU preferences. Yet, results regarding the linear-quadratic example differ. The noise decreases the precision of information transmission (cell sizes), acting as an additional bias.

Following Lu (2017), I assume that S observes a noisy signal $s \in \mathcal{C}$ whose density is measurable on $\mathcal{C} \times \mathcal{C}$. I introduce the following definition :

Definition 4. *(Lu (2017)) A noise has local size less than* $\epsilon > 0$ *if for any* $\theta \in \mathcal{C}$, $|s - \theta_0| < \epsilon$ *surely*

I assume that the signal is noisy and has local size less than ϵ. Assume that S's preference under uncertainty are MEU. Then, having observed s, S evaluates action a through:

$$V_S^{MEU}(a) = \min_{\theta \in [s-\epsilon, s+\epsilon]} \mathbb{E}_\theta(u_S(a, \omega)))$$

Then, notice that the structure of equilibria is unaffected by those changes. Proposition 2 which guarantees that all equilibrium are outcome equivalent to a partition equilibrium only depends on the sender's type, and not the state generating distribution.

The fact that information transmission can only take place below $\widetilde{\theta}$ (Theorem 1) is also unaffected under my assumptions. Recall that there can not be a cut-off θ_k above $\widetilde{\theta}$ because $A_S(\theta_k) > A_R(s \in [\theta_k, \theta_{k+1}])$. Yet, the optimal action when the sender's signal is in $[\theta_k, \theta_{k+1}]$ is $A_R(\theta_k - \epsilon)$ and the optimal action of S when his type is θ_k is $A_S(\theta_k - \epsilon)$. Because of the misalignment of playser (assumption 2), it can not be that $A_R(\theta_k - \epsilon) > A_R(\theta_k - \epsilon)$.

The evaluation of actions by R changes. Take $B = [\theta_1, \theta_2] \subset \mathcal{C}$, if R learns that $s \in B$ it implies that $\theta_0 \in [\theta_1 - \epsilon, \theta_2 + \epsilon]$. As a result, given that the sender's type is in B, R evaluates action a through:

$$\begin{aligned} V_R^{MEU}(a, B) &= \min_{\theta \in [\theta_1 - \epsilon, \theta_2 + \epsilon]} \mathbb{E}_\theta(u_R(a, \omega)) \\ &= \mathbb{E}_{\theta_2 + \epsilon}(u_R(a, \omega)) \end{aligned}$$

Thus, R's evaluation of actions, for a given interval of parameters, still depends only on the upper bound of that interval. As a result, Theorem 2 still holds as well.

However, the characterisation in the linear quadratic will differ. The arbitrage condition of proposition 4 gives that :

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

Thus, it is as if the bias of the sender was $b + \frac{\epsilon}{2}$. The cells' length will change to a size of $2b + \epsilon$. This will have an effect on ex-ante evaluation of welfare, as the noise and the sender's ambiguity aversion decreases the precision of communication.

B Proofs of the results in the main text

Proof of lemma 1:

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

Assumption 1 gives that the latter is strictly positive.

□

Proof of Proposition 1:

In order to prove our 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:

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

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 $A_R(B) \in [A_R(\theta_1), A_R(\theta_2)]$.

$\square$

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 at 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 closes 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 the set of S types who induce u and Θ_v 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 $\frac{\partial^2 \mathbb{E}_{\theta}(u_S(a,\omega))}{\partial a \partial \theta} > 0$ (Lemma 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:

$$\begin{cases} W^{\theta}(a_2) \leq W^{\theta}(a_1) \\ W^{\theta^{'}}(a_1) \leq W^{\theta^{'}}(a_2) \\ W^{\theta^{''}}(a_2) \leq W^{\theta^{''}}(a_1) \end{cases} \tag{5}$$

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 models. Two special cases are to be noticed. Either the bad state is sufficiently worst 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 bad 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 $a_1 < \widetilde{\widetilde{a_S}}$ and $a_2 \geq \widetilde{a}_S$ can cross at most once and system (5) 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 (5) 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 (5) 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 monotony 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 at 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 (1). 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[6]. 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:

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

- 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:

[6]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.

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

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

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

□

Proof of Theorem 1

Assume there is a q cut-off equilibrium ant that $\theta_{q-1} < \widetilde{\theta} < \theta_q$. As $\theta_q > \widetilde{\theta}$, $y(m_q) = a_R(\theta_q)$ and $y(m_{q-1}) = a_R(\widetilde{\theta})$. As a_R is a strictly increasing function, we have that $y(m_{q-1}) < y(m_q) < a_S(\theta_q)$. As by definition, $a \rightarrow \mathbb{E}_\theta(u_S(a,\omega))$ is strictly increasing on $[0, a_S(\theta_q)]$ and strictly decreasing on $[a_S(\theta_q), 1]$, we have that $V_S^{\theta_q}(m_{q-1}) < V_S^{\theta_q}(m_q)$, which is a contradiction.

□

Proof of Proposition 4:

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: $\theta_M^- < ... < \theta_1^-$ (step 1). Define $\mathcal{E}^- = \{(\theta_1^-, ..., \theta_k^-)|1 \leq k \leq M\}$. Then, I show that any 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 $1 \leq q < M-1$, $C_M^- = [\underline{\theta}, \theta_M^-]$ and $C_0^- = [\theta_1^-, \widetilde{\theta}]$

Step 1:

Assume there is a N cut-off equilibrium. Then the signalling strategy of the sender σ must be such that for $q \in 0, ..., m-1$, $\forall \theta \in C_q$, $\sigma(\theta) = m_k^-$

First notice that $V_R^{m_0}(a) = \mathbb{E}_{\widetilde{\theta}}(u_R(a,\omega))$. For σ to be an equilibrium strategy we need that $\forall \theta \in C_1^-$ and $m \neq m_1^-$:

$$V_S^\theta(m_1^-) \geq V_S^\theta(m)$$

In C_1^-, type θ_1^- has the most incentive to deviate from sending m_1^- to sending m_0, which would induce a smaller action, as, $V_R^{m_1^-}(a) = \mathbb{E}_{\theta_1^-}(u_R(a,\omega))$ and $A_R(\theta)$ is strictly increasing by Lemma 1.

Thus, a necessary condition for all types in C_1^- to send m_1^- is that:

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

Furthermore, it is also necessary that all types in C_0^- prefer message m_0^-. In particular it must be the case for type θ_1^-, thus: $V_S^{\theta_1^-}(m_1^-) \leq V_S^{\theta_1^-}(m_0^-)$. As a consequence, a necessary condition for σ to be an equilibrium strategy is:

$$V_S^{\theta_1^-}(m_1^+) = V_S^{\theta_1^-}(m_0^-) \tag{6}$$

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 $q \in 0, ..., M$:

$$V_S^{\theta_{q+1}^-}(m_q^-) = V_S^{\theta_{q+1}^-}(m_{q+1}^-) \tag{7}$$

Furthermore, the fact that $[\widetilde{\theta}, \overline{\theta}] \cup C_0^- \cup C_M^- = \mathcal{C}$ and the fact that for every pair of consequent cell of the partition the incentive constraints are transitive gives that conditions (7) is both necessary and sufficient. As $A_R(\theta)$ is strictly monotone, it implies that $A_R(\theta_k) \neq A_R(\theta_{k+1})$. $\widetilde{\theta}$ being known, it is possible to derive θ_1^- directly from (6). By repeating the reasoning by induction, θ_{q+1}^- can be derived from θ_q^- for $q \in 1, ..., M-1$ from (7) as long as there is $\theta_M < \overline{\theta}$.

Step 2:

I show that any partitional strategy of the sender characterised by $\theta_q^- < ... < \theta_1^-$ is an equilibrium strategy. I proceed by iteration:

- Step 1 proves that $\theta_M^- < ... < \theta_1^-$ characterise an equilibrium. Let's show that $\theta_{M-1}^- < ... < \theta_1^-$ does as well.

  Assume S's strategy is σ_{M-1} such that for $0 \leq k \leq M-2$, $\forall \theta \in C_k^-$, $\sigma_{M-1}(\theta) = m_k^-$ and $\forall \theta \in [\underline{\theta}, \theta_{M-1}^+]$, $\sigma_{M-1}(\theta) = m_{M-1}^-$. Then for $k \in 1, ..., n-1$, when learning its type $\theta \in C_k^-$, by construction of the previous equilibrium, S's preferred message is m_k^-. When $\theta \in [\theta_{M-1}^-, \theta_M^-]$, m_{M-1}^- induces the same outcome as in the M cut-off equilibrium and is preferred to all other messages.

When $\theta \in [\underline{\theta}, \theta_M^+]$ the fact that, for every pair of consequent cell of the partition, the incentive constraints are transitive implies that message m_{M-1}^- is preferred to any other message.

- Let's assume that for $q \geq 2$, σ_q defined as above is an equilibrium strategy for S. By the same reasoning as above, it is straightforward to show that σ_{q-1} is one as well. 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 $B \subset \mathcal{C}$ and $W^\theta(a) = \mathbb{E}_\theta(u_S(a, \omega))$ the evaluation of action $a \in \mathcal{A}$ by a sender of type θ.

- Proposition 2 gives that all equilibria are partitional. First I'll show that any equilibria only characterised by elements of $\theta_q^-, ..., \theta_1^-$ must be characterised by elements of $\mathcal{E}^-$. It is straightforward to see that any equilibria only characterised by elements of $\theta_q^-, ..., \theta_1^-$ which is not in $\mathcal{E}^-$ can be constructed from an element of $\mathcal{E}^-$ by removing at least one element which is not an extrema. To prove our claim, it is thus sufficient to prove that no equilibrium constructed from an element of $\mathcal{E}^-$ by removing exactly one element which is not an extrema exists.

  Consider a strategy σ_p characterised by cut-offs $\theta_1^-, ..., \theta_{p-1}^-, \theta_{p+1}^-, ..., \theta_l^-$ for $1 \leq q \leq M$ and assume it is an equilibrium strategy[7]. 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.

- Thus σ must have a cut-off type $\theta^* \notin \{\theta_1^-, ..., \theta_M^-\}$. Assume without loss of generality that $\theta_p^- < \theta^* < \theta_{p+1}^-$ for $p \in 1, ..., M-1$. Then we have that:

$$W^{\theta^*}([\theta_p^-, \theta^*]) = W^{\theta^*}([\theta^*, \theta_{p+1}^-])$$
$$\iff \mathbb{E}_{\theta^*}(u_S(a_R(\theta_p^-))) = \mathbb{E}_{\theta^*}(u_S(a_R(\theta^*)))$$

  Yet, by the construction in step 1, the above implies that $\theta^* = \theta_{p+1}^-$, which is a contradiction.

□

Proof of Theorem 2:

[7] the choice of removing an element $\theta_p^- > \widetilde{\theta}$ is without loss of generality

Let $\theta_0 \in [\theta_q, \theta_{q+1}]$.

First, assume that $\widetilde{\theta} \notin [\theta_q, \theta_{q+1}]$. Then, for $k \leq q-1$, σ_k's interim payoff is $V_S^{\theta_0}(m_k)$ which, by definition is strictly lower than the interim payoff of σ_q: $V_S^{\theta_0}(m_q)$. Similarly, for $k^{'} \geq q$, the interim payoff of $\sigma_{k'}$ is the same as the one of σ_q: $V_S^{\theta_0}(m_q)$. This proves (2) and it is straightforward that for $q \geq 1$, (2) $\Rightarrow$ (1)

Now, consider the case where $\widetilde{\theta} \in [\theta_q, \theta_{q+1}]$. Then, whatever the equilibrium strategy the sender chooses, it's interim payoff is $\mathbb{E}_{\theta_0}(u_S((a_R(\widetilde{\theta}))$. This concludes the proof of (1).

□

Proof of Proposition 5:

In the following, for simplicity, I write $\mathbb{E}(\theta) = \mathbb{E}_\theta(\omega)$, for any $\theta \in \mathcal{C}$.

1. Assume R has SEU preferences. Assume there are n equilibrium cut-offs in $(0,1)$: $\theta_0, ..., \theta_n$ and thus $\theta_0 = \underline{\theta}$ $\theta_n = \overline{\theta}$. When receiving equilibrium message m_k^n:

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

As a result, the equilibrium action of R is $y^*(m_k^n) = \frac{\mathbb{E}(\theta_k)+\mathbb{E}(\theta_{k+1})}{2}$ and the optimal action in the eyes of S is $A_S(\theta_0) = \mathbb{E}(\theta_0) + b$. The arbitrage condition gives that a sender of type θ_k must be indifferent between m_{k-1}^n and m_k^n. That is, for $k \in 2, ..., n$:

$$\begin{aligned} -\mathbb{E}_{\theta_k}(\frac{\mathbb{E}(\theta_k)+\mathbb{E}(\theta_{k+1})}{2} - \omega - b)^2 &= -\mathbb{E}_{\theta_k}(\frac{\mathbb{E}(\theta_k)+\mathbb{E}(\theta_{k-1})}{2} - \omega - b)^2 \\ \iff -(1-\theta_k)\Big(\frac{\mathbb{E}(\theta_k)+\mathbb{E}(\theta_{k+1})}{2} - b\Big)^2 - \theta_k\Big(\frac{\mathbb{E}(\theta_k)+\mathbb{E}(\theta_{k+1})}{2} - 1 - b\Big)^2 &= \\ -(1-\theta_k)\Big(\frac{\mathbb{E}(\theta_k)+\mathbb{E}(\theta_{k-1})}{2} - b\Big)^2 - \theta_k\Big(\frac{\mathbb{E}(\theta_k)+\mathbb{E}(\theta_{k-1})}{2} - 1 - b\Big)^2 & \end{aligned}$$

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

$$\mathbb{E}(\theta_{k+1}) - \mathbb{E}(\theta_k) = \mathbb{E}(\theta_k) - \mathbb{E}(\theta_{k-1}) + 4b \tag{8}$$

Equation (8) further gives that:

$$\begin{aligned}\mathbb{E}(\theta_k) &= k(\mathbb{E}(\theta_1)-\mathbb{E}(\theta_0)) + \frac{k(k-1)}{2}4b \\ &= k(\theta_1 - \underline{\theta}) + \frac{k(k-1)}{2}4b\end{aligned}$$

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

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

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

$$0 < b < \frac{\overline{\theta}}{2n(n+1)}$$

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

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

Thus, when $\theta_1 \leq \widetilde{\theta}$, $V_R(a|m^n_0) = \mathbb{E}_{\theta_1}(u_R(a))$ and the arbitrage condition giving the cut-off types gives that $A_S(\theta_1) = \mathbb{E}(\theta_1) + b$ must thus be at equal distance from $\mathbb{E}(\theta_1)$ and $\mathbb{E}(\theta_2)$. For this to be possible, it must be that $b > 0$ [8]. Thus, when there is a n-cut-off equilibrium, it must be that $\widetilde{\theta} > \theta_n$. When receiving message m^n_k, for $k \geq 1$:

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

The equilibrium action of R when receiving the equilibrium message $[\theta_k, \theta_{k+1}]$ is $y(m^n_k) = \mathbb{E}(\theta_{k+1})$. 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

$$\begin{aligned}\mathbb{E}(\theta_{k+1}) + b &= \frac{\mathbb{E}(\theta_{k+1}) + \mathbb{E}(\theta_{k+2})}{2} \\ \iff \mathbb{E}(\theta_{k+2}) &= \mathbb{E}(\theta_{k+1}) + 2b\end{aligned}$$

[8] Proposition 1 proves this result in the general context of the model

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

$$\begin{aligned}\frac{\mathbb{E}(\widetilde{\theta}) + \mathbb{E}(\theta_n)}{2} &= \mathbb{E}(\theta_n) + b \\ \iff \mathbb{E}(\theta_n) &= \frac{1}{2} - 2b\end{aligned}$$

Which implies that, for all $k \geq 1$:

$$\mathbb{E}(\theta_k) = \theta_k = \frac{1}{2} - 2b(n - k + 1)$$

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

$$\begin{aligned}\theta_1 &> 0 \\ \iff \frac{1}{2} - 2bn &> 0 \\ \iff 0 < b &< \frac{1}{4n}\end{aligned}$$

□

Proof of Corollary 1:

It is possible to derive from Proposition 5 that in the SEU case:

$$\mathbb{E}(\theta_{k+1}) - \mathbb{E}(\theta_k) = \mathbb{E}(\theta_k) - \mathbb{E}(\theta_{k-1}) + 4b$$

Yet, for $g_\theta \in \mathcal{D}_1$, $\mathbb{E}(\theta) = \theta$ thus:

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

It is also possible to derive from Proposition 5 that in the MEU case:

$$\mathbb{E}(\theta_{k+1}) - \mathbb{E}(\theta_k) = -2b$$
$$\iff \theta_{k+1} - \theta_k = -2b$$

□

Proof of Proposition 6:

1. Given that, for $g_\theta \in \mathcal{D}_1$, $\mathbb{E}(\theta) = \theta$, the ex-ante payoff in the SEU case is calculated as in CS's section 4. Thus, when R has SEU preferences, her ex-ante expected payoff of a n-cut-off equilibrium is:

$$V_R^{SEU} = -\frac{1}{12n^2} - \frac{b^2(n^2-1)}{3}$$

For the MEU case, consider the $n+1$ cut-off equilibrium with cut-offs $\theta_1, ..., \theta_{n+1}$. As the sender's pay-off is decreasing with the distance between the sender's preferred action and the receiver's equilibrium action, and that this distance is decreasing with the size of each cell, a lower bound to the expected pay-off of the sender is his expected pay-off when the when all cells are of size $2b$. Thus, I focus on the expected pay-off of the sender when all cells are of size $2b$.

As cells are of constant size, when R has MEU preferences, her ex-ante expected payoff of a $(n+1)$-cut-off equilibrium is:

$$\begin{aligned} V_R^{MEU} &= \sum_{k=1}^{\frac{n(n+1)}{2}} g([\theta_k, \theta_{k+1}]) max_{a \in \mathcal{A}} V_R^{m_k^{n+1}}(a) \\ &= -\sum_{k=1}^{\frac{n(n+1)}{2}} \frac{2}{n(n+1)} max_{a \in \mathcal{A}} \mathbb{E}_{\theta_k}((a-\omega)^2) \\ &= -\frac{2}{n(n+1)} \sum_{k=1}^{\frac{n(n+1)}{2}} \mathbb{E}_{\theta_k}((\mathbb{E}_{\theta_k}(\omega) - \omega)^2) \\ &= -\frac{2}{n(n+1)} \sum_{k=1}^{\frac{n(n+1)}{2}} \frac{b^2}{3} \\ &= -\frac{b^2}{3} \end{aligned}$$

As a result:

$$
\begin{aligned}
V_R^{MEU} - V_R^{SEU} &= -\frac{b^2}{3} + \frac{1}{12n^2} + \frac{b^2(n^2-1)}{3} \\
&= \frac{b^2}{9}(n^2-4) + \frac{1}{12n^2}
\end{aligned}
$$

For $n \geq 2$ the latter is positive. For $n = 1$ it is positive if and only if $b \leq \frac{1}{2}$. As we have assumed there is an $n+1$ cut-off equilibrium, Proposition 5 implies that $b \leq \frac{1}{2}$. Thus, $V_R^{MEU} - V_R^{SEU} \geq 0$.

2. Ex-ante, the sender's expected payoff when R has MEU preferences is:

$$
\begin{aligned}
V_S^{MEU} &= -\sum_{k=1}^{\frac{n(n+1)}{2}} \int_{\theta_k}^{\theta_{k+1}} (\theta_k - \omega - b)^2 \\
&= -\sum_{k=1}^{\frac{n(n+1)}{2}} [\int_{\theta_k}^{\theta_{k+1}} (\theta_k - \omega)^2 d\omega + b^2 - 2b \int_{\theta_k}^{\theta_{k+1}} (\theta_k - \omega) d\omega] \\
&= V_R^{MEU} - b^2 - b \sum_{k=1}^{\frac{n(n+1)}{2}} (\theta_k - \theta_{k+1})^2 \\
&= V_R^{MEU} - b^2 - 2n(n+1)b^3
\end{aligned}
$$

The ex-ante payoff of S when the receiver has SEU preferences is also calculated as in Crawford and Sobel (1982)'s section 4 giving: $V_S^{SEU} = V_R^{SEU} - b^2$. As a result:

$$
\begin{aligned}
V_S^{MEU} - V_S^{SEU} &= V_R^{MEU} - V_R^{SEU} - 2n(n+1)b^3 \\
&= \frac{b^2}{9}(n^2-4) + \frac{1}{12n^2} - 2n(n+1)b^3
\end{aligned}
$$

For the equilibrium described until here to exist, it must be that $b < \frac{1}{2n(n+1)}$ and thus for the sender is ex-ante better in the partitional equilibrium with the maximum number of cut-offs of the MEU version of the game compared to the partitional equilibrium with the maximum number of cut-offs of the SEU version of the game it must be that:

$$
\begin{aligned}
& V_S^{MEU} - V_S^{SEU} > 0 \text{ and } b < \frac{1}{2n(n+1)} \\
\iff & \frac{1}{18n(n+1)}(n^2-4) + \frac{1}{12n^2} - 2n(n+1)\frac{1}{(2n(n+1))^3} > 0
\end{aligned}
$$

The latter is true for $n \geq 2$

□

Proof of Theorem 3

I start by showing that the distance between the sender's preferred action and the receiver's equilibrium action is smaller in the most informative equilibrium than under full revelation.

Assume the most informative equilibrium between players is outcome equivalent to an n cut-off equilibrium. Call y^* the corresponding equilibrium action of R and $\theta_1, ..., \theta_n$ the cut-offs. The ex-ante distance between the sender's preferred action and the receiver's equilibrium action in the n cut-off equilibrium is given by :

$$\sum_{k=0}^{n-1} \mathbb{E}_\mu(|y^*([\theta_{k+1}, \theta_k]) - A_S(\theta)|)$$

Notice that in every cell, the equilibrium action is θ_k and the expected preferred action of the sender is $\frac{\theta_k + \theta_{k+1}}{2}$. Every non-terminal cell has length $2b$. As a result the ex-ante distance between the sender's preferred action and the receiver's equilibrium action in those cells is b. Assume the terminal cell has size $0 \leq \epsilon < 2b$. Then the ex-ante distance between the sender's preferred action and the receiver's equilibrium action is $\frac{\epsilon}{2}$. Thus :

$$\sum_{k=0}^{n-1} \mathbb{E}_\mu(|y^*([\theta_{k+1}, \theta_k]) - A_S(\theta)|) = (1-\epsilon)b + \epsilon\frac{\epsilon}{2}$$

The latter is a decreasing function of ϵ for $\epsilon < b$ and increasing for $\epsilon \geq b$. At a result it is minimal at $\epsilon = b$ and is equal to $b(1 - \frac{b}{2})$. Thus :

$$b(1 - \frac{b}{2}) \leq \sum_{k=0}^{n-1} \mathbb{E}_\mu(|y^*([\theta_{k+1}, \theta_k]) - A_S(\theta)|) \leq b$$

Under full revelation, the distance between the sender's preferred action and the receiver's equilibrium action of the sender is always b.

As the distance between the sender's preferred action and the receiver's equilibrium action is always at least b, and that the sender's pay-off is decreasing with that distance, a lower bound to the expected pay-off of the sender is his expected pay-off when the when distance is exactly b. Thus, the expected pay-off of the sender in the most informative equilibrium is greater or equal to :

$$
\begin{aligned}
-\sum_{k=0}^{n-1}\int_{\theta_{k+1}}^{\theta_k}(\theta_k-\omega-b)^2 &= \sum_{k=0}^{n-1}\Big[\frac{(\theta_k-\omega-b)^3}{3}\Big]_{\theta_{k+1}}^{\theta_k} \\
&= \sum_{k=0}^{n-1}\Big(-\frac{(\theta_k-\theta_{k+1}-b)^3}{3}+\frac{(-b)^3}{3}\Big) \\
&= -\frac{2nb^3}{3}
\end{aligned}
$$

Under full revelation the pay-off of the sender is always $-b^2$, and :

$$
-\frac{2nb^3}{3} > -b^2 \iff b < \frac{3}{2n}
$$

Yet the condition for a n equilibrium to exist is $b < \frac{1}{4n}$ and $\frac{1}{4n} < \frac{3}{2n}$.

□

Proof of Proposition 13:

The first order condition on a_i gives that:

$$
BR_i(a_{-i}) = \frac{r}{2} - a_{-i} + \theta - 1
$$

Thus, at equilibrium, when the model is $\theta \in \mathcal{C}$, the receivers distancing levels must be such that:

$$
a_1^*(\theta) + a_2^*(\theta) = \frac{r}{2} + \theta - 1
$$

The socially optimal levels of distancing a_1^P, a_2^P are such that:

$$\begin{cases} \frac{\partial \mathbb{E}_{\theta_0}(u_S(a_1,a_2,\omega))}{\partial a_1} = 0 \\ \frac{\partial \mathbb{E}_{\theta_0}(u_S(a_1,a_2,\omega))}{\partial a_2} = 0 \end{cases}$$
$$\iff a_1^P + a_2^P = \frac{r}{4} + \theta_0 - 1$$

Thus, the optimal ex-post actions of the sender and receiver of the parallel game coincide with the total level of distancing given above.

Notice also that $\min_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a_i, a_{-i}, \omega)) = \min_{\theta \in \mathcal{C}} -\mathbb{E}_\theta(t-\omega)^2$ which implies that $\mathbb{E}_\theta(u_R(a_i, a_{-i}, \omega))$ is minimal for $\theta = \frac{1}{2}$. As a result, the hedging action is $\widetilde{t} = \frac{1+r}{2}$ and the corresponding hedging model is $\widetilde{\theta} = \frac{1}{2}$. In addition, U_S and U_R verify Assumptions 1 and 4. Define $t_i(\theta) = argmax_{t \in \mathcal{A}} U_i(t, \omega)$, for $i \in R, S$, the optimal total distancing level for these representative players. $t_R(\theta) - t_S(\theta) = \frac{r}{4} > 0$. That is, both player's optimal action are always misaligned. Would receivers know the true state, the outcome of their game would always be higher than what the sender would like. As a result, our parallel communication game verifies all the assumptions necessary for it to be outcome equivalent to the game studied in this section.

□

Proof of Proposition 7 :

In order to prove our 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 :

Lemma 6. *Define $B \subset \mathcal{C}$ the belief of the receiver with minimal element θ_1 and maximal element θ_2. Given this belief, her optimal action $A_R(B) \in [A_R(\theta_1), A_R(\theta_2)]$.*

Proof of lemma 6 :

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 $A_R(B) \in [A_R(\theta_1), A_R(\theta_2)]$.

$\square$

Now, we want to prove that :

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

The maximal pay-off of the receiver as a function of the sender's type is given by :

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

where $\theta_M = argmax_{\theta \in \{\theta_1, \theta_2\}} \mathbb{E}_\theta(u_R(a, \omega))$. 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, $\min_{\theta \in B} A_R(B) = A_R(\theta_1)$ and $\max_{\theta \in B} A_R(B) = A_R(\theta_2)$

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, $\min_{\theta \in B} A_R(B) = A_R(\theta_2)$ and $\max_{\theta \in B} A_R(B) = A_R(\theta_1)$

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:

$$\begin{cases} u_R(a,0) > u_R(a,1) \text{ for } a < \widetilde{a} \\ u_R(a,0) = u_R(a,1) \text{ for } a = \widetilde{a} \\ u_R(a,0) < u_R(a,1) \text{ for } a > \widetilde{a} \end{cases} \tag{9}$$

It follows from system (9) that, 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, $\min_{\theta \in B} A_R(B) = A_R(\widetilde{\theta})$.

It also follows from system (9) that, for all $a \in [A_R(\theta_1), A_R(\theta_2)]$ the maximal pay-off of the receiver as a function of the sender's type is given by:

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

where $\theta_M = argmax_{\theta \in \{\theta_1, \theta_2\}} \mathbb{E}_\theta(u_R(a,\omega))$. The above system implies that when $\widetilde{\theta} \in B$, $\max_{\theta \in B} \mathbb{E}_\theta(u_R(a,\omega))$ is decreasing on $(A_R(\theta_1), \widetilde{a})$ (as $\mathbb{E}_{\theta_1}(u_R(a,\omega))$ is maximal at $A_R(\theta_1)$) and increasing on $(\widetilde{a}, A_R(\theta_2))$ (as $\mathbb{E}_{\theta_2}(u_R(a,\omega))$ is maximal at $A_R(\theta_2)$). As a result, it is maximal at either $A_R(\theta_1)$ or $A_R(\theta_2)$. As a result, $\max_{\theta \in B} A_R(B) = A_R(\theta_M)$.

Notice that when utilities are quadratic, a simple algebra gives that for $\theta < \theta'$:

$$\begin{aligned} argmax_{a \in \mathcal{A}} \big[\alpha \mathbb{E}_\theta(u_i(a,\omega)) + (1-\alpha)\mathbb{E}_{\theta'}(u_i(a,\omega))\big] &= \alpha\big(argmax_{a \in \mathcal{A}} \mathbb{E}_\theta(u_i(a,\omega))\big) + (1-\alpha)\big(argmax_{a \in \mathcal{A}} \mathbb{E}_{\theta'}(u_i(a,\omega))\big) \\ &= A_i(\alpha\theta + (1-\alpha)\theta') \end{aligned}$$

which implies that :

$$\begin{aligned} A_R(B) &= argmax_{a\in\mathcal{A}}\big[\alpha \min_{\theta\in B}\mathbb{E}_\theta(u_R(a,\omega)) + (1-\alpha)\max_{\theta\in B}\mathbb{E}_\theta(u_R(a,\omega))\big] \\ &= \begin{cases} \alpha A_R(\theta_2) + (1-\alpha)A_R(\theta_1) \text{ if } \theta_2 < \widetilde{\theta} \\ \alpha A_R(\widetilde{\theta}) + (1-\alpha)A_R(\theta_M) \text{ if } \widetilde{\theta} \in B \\ \alpha A_R(\theta_1) + (1-\alpha)A_R(\theta_2) \text{ if } \theta_1 > \widetilde{\theta} \end{cases} \end{aligned}$$

□

Proof of Proposition 8

Lemma 7. *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 7

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 the set of S types who induce u and Θ_v 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 $\frac{\partial^2 \mathbb{E}_\theta(u_S(a,\omega))}{\partial a \partial \theta} > 0$ (Lemma 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 α-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 $\hat{\theta}$. 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 6 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.

□

By Lemma 7 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 monotony 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 at equilibrium.

□

Proof of Proposition 9:

I focus on the case $\underline{\mathcal{C}}$. The case $\overline{\mathcal{C}}$ is symmetric.

Assume there is a $n > 0$ cut-off equilibrium. It follows from Proposition 1 that, for $1 \leq k \leq n-2$ and $\alpha > 0$:

$$\theta_k + b = \frac{\alpha\theta_k + (1-\alpha)\theta_{k-1} + \alpha\theta_{k+1} + (1-\alpha)\theta_k}{2}$$
$$\iff \theta_{k+1} - \theta_k = \frac{1-\alpha}{\alpha}(\theta_k - \theta_{k-1} + \frac{2b}{\alpha})$$

Set $V_k = \theta_{k+1} - \theta_k$. It follows from the previous equality that $(V_k)_k$ is an arithmetico-geometrical sequence. As a result, for $1 \leq k \leq n-2$ and $\alpha \notin \{0, \frac{1}{2}\}$:

$$V_k = (\frac{1-\alpha}{\alpha})^k(\theta_1 - \frac{2b}{2\alpha-1}) + \frac{2b}{2\alpha-1}$$

By induction, if follows that :

$$\theta_{k+1} = \sum_{j=1}^{k}[(\frac{1-\alpha}{\alpha})^k(\theta_1 - \frac{2b}{2\alpha-1}) + \frac{2b}{2\alpha-1}] + \theta_1$$
$$\iff \quad \theta_{k+1} = \sum_{j=0}^{k} V_j$$
$$\iff \quad \theta_k = (\theta_1 - \frac{2bn}{2\alpha-1})\Big(\frac{1-(\frac{1-\alpha}{\alpha})^k}{1-(\frac{1-\alpha}{\alpha})}\Big) + \frac{2bk}{2\alpha-1}$$

In particular, it must be that $\theta_n = \frac{1}{2}$ which give that $\theta_1 = (\frac{1}{2} - \frac{2bn}{2\alpha-1})\Big(\frac{1-(\frac{1-\alpha}{\alpha})}{1-(\frac{1-\alpha}{\alpha})^n}\Big) + \frac{2b}{2\alpha-1}$. As a result, we get that :

$$\theta_k = (\frac{1}{2} - \frac{2bn}{2\alpha-1})\Big(\frac{1-(\frac{1-\alpha}{\alpha})^k}{1-(\frac{1-\alpha}{\alpha})^n}\Big) + \frac{2bk}{2\alpha-1}$$

□

Proof of Proposition 10 :

1. I start by proving that for $n \geq 2$, $\underline{\theta}^n_{n-1}(\alpha)$ is a strictly increasing function. Define $f(a) = \frac{1-a^{n-1}}{1-a^n}$, for

$a \in (0, 1/2)$. Notice that :

$$\frac{\partial f(a)}{\partial a} = \frac{a^{n-2}(a^n - na + n - 1)}{(1 - a^n)^2}$$

Thus :

$$\begin{aligned} & & \frac{\partial f(a)}{\partial a} < 0 \\ \iff & = & \frac{a^{n-2}(a^n - na + n - 1)}{(1 - a^n)^2} < 0 \\ \iff & = & a^n > n(a-1) + 1 \end{aligned}$$

Yet, $a \in (0, \frac{1}{2}) \Rightarrow a^n > 0$ and $n(a-1)+1 < 0 \iff a < 1 - \frac{1}{n}$ which is true because $a \in (0, 1/2)$ and $n \geq 2$. As a result, $\frac{\partial f(a)}{\partial a} < 0$ and f is a decreasing function. Yet :

$$\underline{\theta}^n_{n-1}(\alpha) = \frac{1}{2} f(\frac{1-\alpha}{\alpha}) + \frac{2bn}{2\alpha - 1}(1 - f(\frac{1-\alpha}{\alpha})) - \frac{2b}{2\alpha - 1}$$

$\frac{1-\alpha}{\alpha} \in (0, 1/2)$ for $\alpha \in (\frac{1}{2}, 1)$ and is decreasing in α. As a result $f(\frac{1-\alpha}{\alpha})$ is increasing in α and $\underline{\theta}^n_{n-1}(\alpha)$ as well as a sum and product of increasing functions of α. In addition, we have that :

$$\begin{aligned} & & \frac{\partial \underline{\theta}^n_{n-1}(\alpha)}{\partial b} < 0 \\ \iff & = & \frac{2n}{(2\alpha - 1}(1 - f(\frac{1-\alpha}{\alpha})) - \frac{2}{(2\alpha - 1} < 0 \\ \iff & = & f(\frac{1-\alpha}{\alpha}) > 0 \end{aligned}$$

which is true. By a symmetrical process, one can prove that $\overline{\theta}^n_{n-1}(\alpha)$ is a strictly decreasing function and that $\frac{\partial \overline{\theta}^n_{n-1}(\alpha)}{\partial b} < 0$. Yet :

$$lim_{\alpha \to 1} \overline{\theta}^n_{n-1}(\alpha) = -2b < \frac{1}{2}$$

Thus, as $\overline{\theta}^n_{n-1}(\alpha)$ is strictly decreasing and continuous, there is $\alpha(b) \in (1/2, 1)$ such that $\overline{\theta}^n_{n-1}(\alpha) = \frac{1}{2}$. As $\overline{\theta}^n_{n-1}(\alpha)$ is strictly decreasing, for $\alpha \geq \alpha(b)$, no information transmission is possible in $\overline{\mathcal{C}}$. In addition, because $\frac{\partial \overline{\theta}^n_{n-1}(\alpha)}{\partial b} < 0$, it follows that $\alpha(b)$ is a decreasing function.

2. I start by proving that for $n \geq 2$, $\underline{\theta}_1^n(\alpha)$ is a strictly increasing function. Define $f(a) = \frac{1-a}{1-a^n}$, for $a \in (0, 1/2)$. Notice that :

$$\frac{\partial f(a)}{\partial a} = \frac{n(1-a)a^{n-1}}{(1-a^n)^2} - \frac{1}{1-a^n}$$

Thus :

$$\frac{\partial f(a)}{\partial a} < 0$$
$$\iff \quad = \quad n - (n-1)a < \frac{1}{a^{n-1}}$$

Yet, $a \in (0, 1/2) \Rightarrow \frac{1}{a^{n-1}} > 2^{n-1}$ and $a \in (0, 1/2) \Rightarrow n - (n-1)a < n$. As a result, for $n \geq 2$, $n - (n-1)a < n \leq 2^{n-1} < \frac{1}{a^{n-1}}$ which implies that $\frac{\partial f(a)}{\partial a} < 0$ and f is a decreasing function. Yet :

$$\underline{\theta}_{n-1}^n(\alpha) = \frac{1}{2}f(\frac{1-\alpha}{\alpha}) + \frac{2bn}{2\alpha - 1}(1 - f(\frac{1-\alpha}{\alpha})) - \frac{2b}{2\alpha - 1}$$

$\frac{1-\alpha}{\alpha} \in (0, 1/2)$ for $\alpha \in (\frac{1}{2}, 1)$ and is decreasing in α. As a result $f(\frac{1-\alpha}{\alpha})$ is increasing in α and $\underline{\theta}_1^n(\alpha)$ as well as a sum and product of increasing functions of α. By a symmetric process, one can prove that $\overline{\theta}_1^n(\alpha)$ is a decreasing function.

Consider two receivers α_1 and α_2 such that $\alpha_1 < \alpha_2$. Assume there is a n cut-off equilibrium between S and α_1. Then $\overline{\theta}_1^n(\alpha_1) \in (0, 1)$. As $\overline{\theta}_1^n(\alpha)$ is a decreasing function, it must be that $\overline{\theta}_1^n(\alpha_2) < 1$. In addition, as $\overline{\theta}_{n-1}^n(\alpha)$ is an decreasing function, it follows that $\overline{\theta}_{n-1}^n(\alpha_2) > lim_{\alpha \to 1}\overline{\theta}_{n-1}^n(\alpha) = \frac{1}{2} - 2b > 0$ for $b < \frac{1}{4}$, which is the existence condition of the considered equilibrium. As a result, there is a n cut-off equilibrium between S and α_2

□

Proof of Proposition 11 :

First, consider the case where $n \geq 3$:

By an extension of the notations introduced in Proposition 9, I call $\overline{\theta}_k^n(\alpha)$ the k-th cut-off of the equilibrium that has n cut-offs in $[0, \frac{1}{2}]$. By a simple extension of the proof of Proposition 9 one can show that for $k \leq n-2$:

$$\overline{\theta}^n_{n-2}(\alpha) = (\overline{\theta}^n_n(\alpha) - \frac{2bn}{2\alpha - 1})\Big(\frac{1 - (\frac{\alpha}{1-\alpha})^{n-2}}{1 - (\frac{\alpha}{1-\alpha})^n}\Big) - \frac{2b(n-2)}{2\alpha - 1} + \frac{1}{2}$$

Yet, reproducing the reasoning in the proof of Proposition 10, one can show that $\overline{\theta}^n_{n-2}(\alpha)$ is a strictly decreasing function and that :

$$lim_{\alpha \to 1}\overline{\theta}^n_{n-2}(\alpha) = -2b < \frac{1}{2}$$

Thus, as $\overline{\theta}^n_{n-2}(\alpha)$ is strictly decreasing and continuous, there is $\alpha(b,n) \in (1/2, 1)$ such that $\overline{\theta}^n_{n-2}(\alpha) = \frac{1}{2}$. As $\overline{\theta}^n_{n-2}(\alpha)$ is strictly decreasing, for $\alpha \geq \alpha(b,n)$, there can't be n cut-offs in $[\frac{1}{2}, 1]$. As this holds for any number of cut-off types, it follows that there is $\alpha(b) \in (0,1)$ such that, for $\alpha \geq \alpha(b)$, only one action can be induced by types in $[\frac{1}{2}, 1]$.

Similarly, one can show that $\frac{\partial \overline{\theta}^n_{n-2}(\alpha)}{\partial b} < 0$, which implies that that $\alpha(b)$ is a decreasing function.

Second, consider the case where $n = 2$. For there to be a non babbling equilibrium between S and α, there must be $\theta(\alpha) \in (0,1)$ where such that

$$A_S(\theta(\alpha)) = \frac{z_i([0,\theta(\alpha)]) + z_i([\theta(\alpha), 1])}{2} \tag{10}$$

Then, two cases arise :

- Either $\theta(\alpha) \leq \widetilde{\theta}$. Then it follows from Proposition 7 that (10) is equivalent to :

$$\begin{aligned} \theta(\alpha) + b &= \frac{\alpha\theta(\alpha) + (1-\alpha)0 + \alpha\widetilde{\theta} + (1-\alpha)1}{2} \\ \iff \quad \theta(\alpha) &= \frac{2 - \alpha - 4b}{4 - 2\alpha} \end{aligned}$$

- Either $\theta(\alpha) \geq \widetilde{\theta}$ and it follows from Proposition 7 that (10) is equivalent to : :

$$\theta(\alpha) + b = \frac{\alpha\widetilde{\theta} + (1-\alpha)0 + \alpha_i\theta(\alpha) + (1-\alpha)1}{2}$$
$$\iff \quad \theta(\alpha) = \frac{2-\alpha-4b}{4-2\alpha}$$

It follows that for any $\alpha > \frac{1}{2}$, $b > 0 \Rightarrow \theta(\alpha) < \frac{1}{2}$

□

Proof of Proposition 12

The fact that $\theta(\alpha_i) = \frac{2-\alpha_i-4b}{4-2\alpha_i}$ has been already proven in Proposition 11. It follows that a 3-cut-off equilibrium exists if and only if $\theta(\alpha_i) \in (0,1)$ which is equivalent to :

$$0 < b < \frac{1}{2} - \frac{\alpha_i}{4}$$

□

References

Anderson, R. M., Heesterbeek, H., Klinkenberg, D., and Hollingsworth, T. D. (2020). How will country-based mitigation measures influence the course of the covid-19 epidemic? *The Lancet*, 395(10228):931–934.

Barrett, S. (1994). Self-enforcing international environmental agreements. *Oxford Economic Papers*, pages 878–894.

Battaglini, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401.

Berger, L., Berger, N., Bosetti, V., Gilboa, I., Hansen, L. P., Jarvis, C., Marinacci, M., and Smith, R. (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).

Berger, L., Emmerling, J., and Tavoni, M. (2016). Managing catastrophic climate risks under model uncertainty aversion. *Management Science*, 63(3):749–765.

Burge, T. (1993). Content preservation. *The Philosophical Review*, 102(4):457–488.

Chen, Y., Kartik, N., and Sobel, J. (2008). Selecting cheap-talk equilibria. *Econometrica*, 76(1):117–136.

Chisholm, R. M., Chisholm, R. M., Chisholm, R. M., and Chisholm, R. M. (1989). *Theory of knowledge*, volume 3. Prentice-Hall Englewood Cliffs, NJ.

Condie, S. (2008). Living with ambiguity: prices and survival when investors have heterogeneous preferences for ambiguity. *Economic Theory*, 36(1):81–108.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, pages 1431–1451.

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

Ferguson, N. M., Keeling, M. J., Edmunds, W. J., Gani, R., Grenfell, B. T., Anderson, R. M., and Leach, S. (2003). Planning for smallpox outbreaks. *Nature*, 425(6959):681–685.

Geoffard, P.-Y. and Philipson, T. (1997). Disease eradication: private versus public vaccination. *The American Economic Review*, 87(1):222–230.

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

Ghirardato, P. and Marinacci, M. (2002). Ambiguity made precise: A comparative foundation. *Journal of Economic Theory*, 102(2):251–289.

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

Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63(1):85–110.

Guerdjikova, A. and Sciubba, E. (2015). Survival with ambiguity. *Journal of Economic Theory*, 155:50–94.

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

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

Hansen, L. P., Sargent, T. J., Turmuhambetova, G., and Williams, N. (2006). Robust control and model misspecification. *Journal of Economic Theory*, 128(1):45–90.

Hardwig, J. (1985). Epistemic dependence. *The Journal of Philosophy*, 82(7):335–349.

Heal, G. and Millner, A. (2014). Reflections: Uncertainty and decision making in climate change economics. *Review of Environmental Economics and Policy*, 8(1):120–137.

Hellewell, J., Abbott, S., Gimma, A., Bosse, N. I., Jarvis, C. I., Russell, T. W., Munday, J. D., Kucharski, A. J., Edmunds, W. J., Sun, F., et al. (2020). Feasibility of controlling covid-19 outbreaks by isolation of cases and contacts. *The Lancet Global Health*.

Hollingsworth, T. D., Klinkenberg, D., Heesterbeek, H., and Anderson, R. M. (2011). Mitigation strategies for pandemic influenza a: balancing conflicting policy objectives. *PLoS Comput Biol*, 7(2):e1001076.

Hume, D. (1740). A treatise of human nature.

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

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

Krishna, V. and Morgan, J. (2001). A model of expertise. *The Quarterly Journal of Economics*, 116(2):747–775.

Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J. (2008). Tipping elements in the earth's climate system. *Proceedings of the national Academy of Sciences*, 105(6):1786–1793.

Lu, S. E. (2017). Coordination-free equilibria in cheap talk games. *Journal of Economic Theory*, 168:177–208.

Marinacci, M. (2015). Model uncertainty. *Journal of the European Economic Association*, 13(6):1022–1100.

Millner, A., Dietz, S., and Heal, G. (2013). Scientific ambiguity and climate policy. *Environmental and Resource Economics*, pages 1–26.

Pidgeon, N. and Fischhoff, B. (2011). The role of social and decision sciences in communicating uncertain climate risks. *Nature climate change*, 1(1):35–41.

Savage, L. J. (1972). *The Foundations of Statistics*. Courier Corporation.

Wald, A. (1949). Statistical decision functions. *The Annals of Mathematical Statistics*, pages 165–205.