



Munich Personal RePEc Archive

Expert-based Knowledge: Communicating over Scientific Models

Colo, Philippe

University Duisburg-Essen

13 October 2021

Online at <https://mpra.ub.uni-muenchen.de/110488/>
MPRA Paper No. 110488, posted 03 Nov 2021 14:21 UTC

Expert-based Knowledge: Communicating about Scientific Models*

Philippe Colo¹

¹University of Duisburg-Essen

3rd Nov, 2021

Abstract

An expert communicates about scientific models with a decision maker through cheap-talk. Models are probability distributions over states. The decision maker is ambiguity sensitive. I show that all equilibria of the game are outcome equivalent to partitional equilibria and that the most informative one is interim dominant for the expert. Unlike in similar models of the literature, information transmission depends both on the strategic misalignment of players and a form of consensus among scientific models. When science is divided and the decision maker has maxmin expected utility preferences, whatever the misalignment, no information can be conveyed above a certain threshold.

Keywords: Ambiguity, cheap talk

JEL classification: D81, D83, C72

*I thank Frédéric Koessler, Jean-Marc Tallon and Stéphane Zuber for precious guidance. I also thank Massimo Marinacci, Marco Ottaviani, Ani Guerdjikova, Mark Le Quement, Christian Kellner and Guillaume Pomme for helpful discussions as well as seminar and conference participants at PSE, Bocconi, ETH Zurich, FUR 2018 and ESEM2019 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 about most of the knowledge we claim to possess. Today's science is too vast and complex for a single individual to master it by himself. We say we know about the mechanism of climate change or the COVID-19 pandemic because of what we have heard from experts in those fields. On such topics, our claimed knowledge relies much more on the confidence we have in these experts than on the evidence we can directly access. The importance of our confidence in experts is even higher if we consider the case of scientific models. Scientific models are descriptive approximations of reality. They provide structure to our perception of the world by focusing on a phenomenon's main mechanism. By nature, scientific models are fallible. Counter-examples to their predictions can be easy to find, especially in the face of complex and unexpected events such as the COVID-19 pandemic. Therefore, it can be difficult for experts to convince the general public that a given model is a phenomenon's best available reading grid.

In this paper, I study our knowledge of scientific models when it is entirely expert-based. I try to understand the foundations of the expert-layman bond of trust that will bring the latter to see reality through the reading grid suggested by the former. In order to do so, I model expert-based knowledge as a game of strategic information transmission. Information is about models, which I represent as probability distributions over possible states.¹ 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 resulting from the decision maker's behavior on issues such as the spread of a deadly virus or the limitation of greenhouse gas emissions. 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 them.

The game I study is in the tradition of Crawford and Sobel's (1982) cheap talk game (hereafter CS). The main difference lies in the fact that communication concerns models rather than states. Given the strategic nature of the communication, the expert is typically not able to truthfully reveal which model is the most accurate. At equilibrium, the sender designates an interval of models which contains his type. As in CS, the size of such an interval will depend on the *misalignment*: the difference of interest between parties. Yet, I also show that, whatever the misalignment, information transmission can be impossible over an entire set of models. In the following, I will argue that this situation is caused by a form of division among models. Because this can happen even for

¹In doing so, I follow a tradition in statistics and decision theory dating back at least as far as Wald (1949) (see Marinacci (2015) for a survey).

arbitrarily small misalignments, this result offers an alternative explanation to mistrust in experts and science. Consider the example of the COVID-19 pandemic, where the expert is concerned about the externalities caused by the decision maker's behaviour. If over- and under-restricting social interactions are perceived as comparable threats by the decision maker, experts advising in favour of the strongest limitations are inaudible at equilibrium.

Under model uncertainty, preferences generally fail to satisfy the expected-utility requirements, as famously pointed out by Ellsberg (1961). This situation calls for the use of specific ambiguity-sensitive preferences for the decision maker. An individual is exposed to ambiguity when the expected payoff of his strategy varies with the probabilities over which he is uncertain. An ambiguity-averse individual will tend to favour strategies that reduce that exposure. I mainly focus on the cases where the receiver displays Gilboa and Schmeidler (1989)'s maxmin expected utility preferences (MEU) or Savage (1972)'s subjective expected utility (SEU). 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 the recommended one in this model, his influence is extremely high. When his interest is to induce an action in the opposite direction though, his influence is nonexistent. Yet, if there is no univocally worst state in terms of utility for the receiver, the worst possible combination of probabilities over states can be an interior element of the set of types. In that case, there is a change in the monotonicity of the maximal expected welfare of the receiver according to the model considered. I say that science is divided because, on either side of the monotonicity change, putting a higher probability on one state has an opposite effect on the receiver's maximal expected welfare. Loosely speaking, if there are two states, the two sides disagree on which state is the worse threat to the receiver.

The equilibrium analysis of this game shows that for all senders on one side of this division, it is always better to be as informative as possible, even after they learn their type. For types on the other side though, information transmission is impossible, whatever the misalignment. This result strongly contrasts with the SEU case, where such asymmetry does not exist. Under SEU preferences, the precision of information transmission depends only on the difference of interest between the two parties. When the latter is small, information transmission can be almost perfect over the entire set of types. But once the sender learns his type, he does not always have an interest in being as informative as he could be.

Arguably, several of the mechanisms described in this paper were exemplified during the COVID-19 crisis. During the early stages of the pandemic, information was too scarce to exclude the possibility that, even under optimal restriction policies, the virus would be so deadly that the outcome would be worse than in any other case. Clearly, for decision makers, there was a consensus that the more likely that possibility, the worse. In line with the results of this paper, during that first stage of the crisis, confidence in experts was high. Even models recommending drastic social restriction measures were influential. Yet, later on in the pandemic, it appeared that for many decision makers, excessive restrictions in the face of limited epidemic danger could be worse than even the worst epidemic scenario under appropriate restriction levels. Finding the appropriate balance between economic activity and social limitations had become the main challenge of the crisis. This situation paved the way for a division in science: depending on the model considered, one aspect would be prioritised over the other. The results I derive for this case offer an explanation to why scepticism appeared at that time, and why models recommending the strongest levels of social restriction lost influence among decision makers.

Scientific communication is the subject of growing attention in economic research. Recent work by Spiess (2018), Banerjee et al. (2020), Andrews and Shapiro (2021) and Schwartzstein and Sunderam (2021) focuses on optimal choice of scientific modelling, according to different objectives or audiences for the scientist. Unlike these papers, I do not assume that the expert's audience observes some data and uses it to assess the statistical properties of the models provided by the expert. This is because I focus on complex topics where discriminating among different modelling approaches is a pure act of expert-judgement. For example, take the epidemiological models used to evaluate the impact of health measures on the COVID-19 pandemic. Two main approaches exist: process-based models, that try to capture the mechanisms by which diseases spread, and curve-fitting approaches that aim to mathematically approximate the growth of the epidemic (Ferguson et al., 2003). As argued by Berger et al. (2020), even with sufficient data to evaluate these models, choosing among them is a fine art. But deciding which modelling approach is the most promising to describe an ongoing pandemic is an even harder task. It requires experience of both epidemics and formal representations which, by definition, only experts possess. During the pandemic, both of these uncertainties *across* models and *about* models were present. For a decision maker, resolving them requires more than epidemic data. What matters is information about the models themselves.

This study also relates to the literature on cheap talk communication with ambiguity-sensitive preferences. Kellner and Le Quement (2017) were the first to study this question. In their model, communication is on states of the world, allowing for Ellsbergian communication strategies. They

show that the use of these strategies reduces misalignment between players, creating equilibria which ex-ante Pareto-dominate the corresponding ones in CS. Kellner and Le Quement (2018) explore a simple two-action two-state setting, with only standard mixed strategies allowed, but an ambiguous prior over the states. These approaches differ from mine because in the present paper, communication is over probability distributions. In addition, as pointed out by Hanany et al. (2020), the updating assumed in these papers violates sequential optimality. This is an issue I do not face when studying communication about models. To the best of my knowledge, this paper is the first to study strategic communication directly about the models that form the set of priors of an ambiguity-sensitive decision maker.

Finally, this study also contributes to the literature on model uncertainty by providing a game-theoretical foundation to multiple prior beliefs. Because knowledge about models is expert-based, information transmission is always imprecise. It is therefore justified to assume that decision makers use multiple models to analyse the world. This generally translates into multiple prior beliefs, as often assumed under model uncertainty² and experimentally measured by Abdellaoui et al. (2021).

Section 2 introduces the framework and a simple example. Section 3 establishes general results regarding the structure of equilibria and proves the main results. Section 4 provides supplementary characterisations using the example introduced above and section 5 discusses the results. Appendix A generalises by relaxing some assumptions made in the main text. Appendix B extends to α -MEU preferences (Ghirardato et al., 2004) and shows that results are, in a certain sense, robust to varying degrees of ambiguity aversion. Appendix C contains all the 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 $\Omega = \{0, 1\}$ be the set of possible states of nature. For $i = S, R$, let $u_i : \mathcal{A} \times \Omega \rightarrow \mathbb{R}$ be the von Neumann-Morgenstern utility function of player i , that maps her actions and states to a utility value. I start by making the following assumptions:

²See for instance Hansen and Sargent (2001) and Hansen et al. (2006) in the context of dynamic decision making or Millner et al. (2013) and Berger et al. (2016) in the case of climate change management.

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

This assumption implies that u_i admits a unique maximum for each state. Define $a_i(\omega) = \arg \max_{a \in \mathcal{A}} u_i(a, \omega)$ as this maximum. It is the optimal action of player i under perfect information that the state is ω . Assumption 1 ensures that $a_i(\omega)$ is strictly increasing in ω . In the example of the COVID-19 pandemic, one can think of the receiver as a political decision maker and of the sender as an epidemiologist. The virus is either extremely contagious ($\omega = 1$) or of limited spread ($\omega = 0$). \mathcal{A} may then represent the level of social restriction the decision maker has to impose. I call $\omega = 1$ ($\omega = 0$) the high (low) state as it is the one where the optimal action is the highest (lowest). The choice of a social restriction level is the result of a trade-off between economic activity and casualties due to the pandemic. Assumption 1 states that for any epidemic scenario, there is a single optimal restriction level. A lower restriction level $a < a_i(\omega)$ is not optimal for i because it might cause too much loss of life in the population. Nor is a higher restriction level $a > a_i(\omega)$ optimal for i , as it entails an over-reduction in economic activity. In addition, assumption 1 states that the optimal level of social restriction is strictly higher in the case where the virus is the most contagious. Assumption 1 is a single crossing assumption: it implies that $u_i(\cdot, 0)$ and $u_i(\cdot, 1)$ can cross only once over \mathcal{A} .

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

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

There is a bijection between the sets \mathcal{D} and $\mathcal{C} = [\underline{\theta}, \bar{\theta}]$. In the rest of the paper, for simplicity, I will specify all the communication strategies on the set \mathcal{C} which will be referred to as the set of *models*. Let $A_i(\theta) = \arg \max_{a \in \mathcal{A}} \mathbb{E}_\theta(u_i(a, \omega))$ 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.³ 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, note 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 Equilibrium concept

Ex-ante, both players are in a situation of model uncertainty. 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:

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

³In 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.

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

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

In the following, the MEU case (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 distribution⁴ $\theta_0 \in \mathcal{C}$ 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 $\sigma : \mathcal{C} \rightarrow \mathcal{M}$. An action rule for R is a strategy $y : \mathcal{M} \rightarrow \mathcal{A}$. Note that I will focus only on pure strategies. Let $\sigma^{-1}(m) \subseteq \mathcal{C}$, be the set of potential types of S, having received message m , when S follows strategy σ . An equilibrium (σ^*, y^*) is defined such that:

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

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

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

⁴In appendix A I show that this assumption can be replaced by the assumption that the sender receives a noisy signal regarding models' likelihood

2. Having received an equilibrium message $m \in \text{supp}(\sigma^*)$, 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:

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

R then evaluates action a by:

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

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

Any message m such that $m \notin \text{supp}(\sigma^*)$ is interpreted as some equilibrium message $m_* \in \text{supp}(\sigma^*)$.⁵

2.3 An example

Before directing our attention to the equilibria of this game, it is useful to take a moment to study a specificity of the receiver's pay-off structure. Consider the following parametric example:

⁵Note that this equilibrium concept corresponds to a perfect Bayesian equilibrium, where the receiver has smooth preferences (Klibanoff et al., 2005) and evaluates actions through $V_R^{KMM}(a, \sigma^{-1}(m)) = \int_{\theta \in \mathcal{C}} g(\theta|m) \phi(\mathbb{E}_\theta(u_R(a, \omega))) d\theta$ and ϕ is linear (in the SEU case) or that $-\frac{\phi''}{\phi} \rightarrow +\infty$ (in the MEU case).

Linear-quadratic example:

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

Then:

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

The example above is similar to CS's linear-quadratic one, with the difference of the $-c\omega$ term. Depending on the value c , the maximal utility in a given state is either higher or lower than in the other state. When $c = 0$, both states are comparable, in the sense that under perfect information, the receiver could achieve exactly the same pay-off in both of them. On the contrary, if for instance $c > 0$, state 0 gives a higher maximum pay-off to the receiver than state 1.

As figure 1 shows, when $c \in (-1, 1)$, R has the same utility in both states for $a = \frac{1+c}{2} \in (0, 1)$. Consider the special case where $c = 0$. As illustrated by figure 1, models on both sides of $\theta = 0.5$ increase the maximal expected utility of the receiver. For instance, $\theta = 0.2$ and $\theta = 0.8$ both improve the receiver's maximal expected utility compared with $\theta = 0.5$. There is no strict ordering over models with respect to R 's expected utility but a division of \mathcal{C} into two intervals over which R 's maximal expected utility goes in opposite directions. For models below $\theta = 0.5$, the higher the probability of the high state, the lower R 's utility. But for models above $\theta = 0.5$, the higher the probability of that same state, the higher R 's utility. In a certain sense, on both sides of $\theta = 0.5$, there is no consensus about which state is the worse threat.

Now consider the case where $c \geq 1$. As illustrated by figure 2, the maximal expected utility of the receiver is strictly decreasing in θ , the probability of the high state. In other words, for any model of \mathcal{C} , the higher the probability of the high state, the lower R 's welfare. Now, there is a consensus over the fact that the more likely the high state, the worse-off the receiver.




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




Figure 2: $c = 1$, a case of consensual science: for all models, the receiver's maximal welfare is decreasing with the probability of the high state.

2.4 Consensus and division in science

As the above example illustrates, the pay-off structure of the receiver can be of two kinds. Either the maximal expected welfare of the receiver is always increasing (or decreasing) with the model $\theta \in \mathcal{C}$ considered, or this is not the case. I say that in the former case science is *consensual* for the receiver and that in the latter case, it is *divided*.

Definition 1. *Science is consensual for the receiver if any model that puts a higher probability on the state giving the lowest maximal utility to the receiver decreases her maximal expected welfare.*

$$\forall \theta, \theta' \in \mathcal{C}, \theta < \theta', \begin{cases} u_R(a_R(0), 0) \geq u_R(a_R(1), 1) \Rightarrow \mathbb{E}_\theta(u_R(A_R(\theta), \omega)) > \mathbb{E}_{\theta'}(u_R(A_R(\theta'), \omega)) \\ u_R(a_R(0), 0) < u_R(a_R(1), 1) \Rightarrow \mathbb{E}_\theta(u_R(A_R(\theta), \omega)) < \mathbb{E}_{\theta'}(u_R(A_R(\theta'), \omega)) \end{cases}$$

Science is divided for the receiver if it is not consensual.

Consensus in science is a monotonicity condition on the maximal expected welfare of the receiver $\mathbb{E}_\theta(u_R(A_R(\theta), \omega))$. It thus depends on both the receiver's ex-post preferences u_R and the set of models \mathcal{C} . There is consensus when, the more models in \mathcal{C} put weight on a given state, the more they reduce the maximal expected welfare of the receiver. When this is not the case, science is divided.

For science to be divided, two things must happen. First, it must be that no state fully dominates the other in terms of utility for the receiver. Because of the single crossing assumption I made on utilities, both states can give the same utility for a single given action in $(a_R(0), a_R(1))$. This action, \tilde{a} , which I define below, maximises the function that gives the worst possible utility to the receiver.

Definition 2. *Define $\tilde{a} = \operatorname{argmax}_{a \in \mathcal{A}} \min_{\omega \in \Omega} u_R(a, \omega)$ as the precautionary action and $\tilde{\theta} \in [0, 1]$ such that $A_R(\tilde{\theta}) = \tilde{a}$ as the cautious model.*

I call \tilde{a} the precautionary action because it is the optimal action anticipating that the worst state will always occur. $\tilde{\theta}$ is the model for which the precautionary action is the optimal action.⁶ For science to be divided, the second condition is that $\tilde{\theta}$ is an interior element of \mathcal{C} . The maximal expected utility of the receiver will be decreasing for models putting lower weight on the high state than on the cautious one ($\theta < \tilde{\theta}$), and increasing for the others ($\theta > \tilde{\theta}$). Thus, when $\tilde{\theta} \in (\underline{\theta}, \bar{\theta})$, there is a change in the monotonicity of the maximal expected welfare of the receiver over \mathcal{C} .

It is clear that the behavioural response of an MEU decision maker will be of a different nature, whether science is divided or consensual. In the latter case, the precautionary action

⁶The fact that $\tilde{\theta}$ exists and is unique is proven in Lemma 4.

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 \tilde{a} as an optimal action. For $B \subset \mathcal{C}$, let $A_R(B) \subset \operatorname{argmax}_{a \in \mathcal{A}} \min_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega))$ be the set of optimal actions of a MEU receiver given the set of priors B .

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

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

Proposition 1 states that an MEU receiver has a unique optimal action for any belief $\theta_0 \in B$ where B is an interval of \mathcal{C} . When she further believes that all models are below $\tilde{\theta}$ ($\theta_0 \in [\theta_1, \theta_2]$ and $\theta_2 < \tilde{\theta}$), she optimally acts as if the probability of the high state were maximal. When she believes that all models are above $\tilde{\theta}$ ($\theta_1 > \tilde{\theta}$), she optimally acts as if the probability of the high state were minimal. But when science is divided, $\tilde{\theta}$ is in the interior of B . R will always act as if the probability of the high state were maximal for beliefs below $\tilde{\theta}$ ($\theta_2 < \tilde{\theta}$) and minimal for beliefs above $\tilde{\theta}$ ($\theta_1 > \tilde{\theta}$). When R believes that the cautious model could be the state-generating model ($\tilde{\theta} \in [\theta_1, \theta_2]$), she optimally acts as if it were the case.

Note that in the SEU case, R 's actions are not sensitive to the change in monotonicity of her maximal expected welfare. Consensus and division in science play no particular role. In that case, the game is in fact very similar to CS. One can indeed equate each model with a state in CS's setting, where the corresponding payoff is the expected utility under that model and μ is the prior over states. This case can thus be used as a benchmark.

3 Equilibrium analysis

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

Definition 3. Set $\{\theta_0, \dots, \theta_q\} \subseteq \mathcal{C}$ such that:

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

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

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

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

The proof of Proposition 2 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. Note 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. Note that there is always at least one partition equilibrium: the 2 cut-off equilibrium, often called the babbling equilibrium, where all types send the same message.

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

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

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

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

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




Figure 3: Identifying cut-offs

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

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

When the receiver has MEU preferences, there is no cut-off type in $[\tilde{\theta}, \bar{\theta}]$. To see why, assume that there is a $q + 1$ cut-off equilibrium. Recall the characterisation result of partition equilibria given by Proposition 3. For θ_q to be a cut-off type, the message sent by types in the cell below and above θ_q must induce actions that give the same utility to a sender of type θ_q . If θ_q was a cut-off type, it would follow from proposition 1 that:

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

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

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

and as $V_S^{\theta_q}$ is strictly increasing for $a \leq A_S(\theta_q)$, it is impossible for types in the cell below and above θ_q to induce actions that give the same utility to a sender of type θ_q . As a result, the indifference between actions induced by messages m_{q-1} and m_q needed for θ_q to be a cut-off type (as displayed in figure 3) is impossible.

A consequence of Theorem 1 is that when science is consensual such that $\tilde{\theta} \leq \underline{\theta}$, 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.

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

Lemma 2. *When the receiver evaluates actions following the MEU criteria there are $M > 0$ partition equilibria. Call $\theta_0 < \dots < \theta_M$ the cut-offs of the equilibrium with most cut-offs. Then the q cut-off partition equilibrium is defined by cut-offs $\theta_0 < \theta_{M-q} < \dots < \theta_M$, for $0 \leq q \leq M$.*




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

As illustrated by Figure 5, lemma 2 states that all equilibria of our game can be built from the same set of cut-off types. More specifically, it states that if one considers the equilibrium with most cut-offs, one can describe all other equilibria by successively removing cut-offs starting from the left.




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

To see why lemma 2 is true, first note that, given Theorem 1, all interior cut-offs are in $[\underline{\theta}, \tilde{\theta}]$. As a result, when S points out an interval of models, R only cares about its upper bound. Thus, cut-offs types will not be determined by an indifference between two adjacent cells of models as in the SEU case, but by an indifference between the models at the upper bound of these cells. In the former case, each indifference condition depends on three distinct types and the prior. Thus, in order to determine the cut-off types, the entire sequence of indifference conditions is needed. In the latter case, each indifference condition depends on two distinct types only. Given that $[\underline{\theta}, \tilde{\theta}]$ is a closed interval, it is then possible to find the first cut-off starting from $\underline{\theta}$ and then to iterate the process to find the following ones. In doing so, I derive the cut-off types of the equilibrium that has the most cut-offs. Call the corresponding number of cut-offs M . Any signalling strategy of the

sender characterised by the q first terms ($1 \leq q \leq M$) of that sequence induces exactly the same incentive constraints for the receiver, which implies that they form part of an equilibrium.

A direct consequence of Proposition 2 is that all equilibria of the game can be ranked by informativeness, something which is never possible in the SEU case.⁷ The following result can thus be established regarding interim Pareto dominance among equilibria.

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

The intuition of the proof is the following. Consider for instance the equilibria described in figure 5. Whatever the equilibrium considered, types in $[\theta_1, \tilde{\theta}]$ will induce the same action $\tilde{\theta}$. But types in $[\underline{\theta}, \theta_1]$ will induce action $\tilde{\theta}$ 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 $\tilde{\theta}$. 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.

4 Characterisations on the linear-quadratic example

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

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

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

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

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

⁷The informativeness ranking comes from the fact that when receiving $m \in \mathcal{M}$ from a type in $[\theta_1, \theta_2]$ with $\theta_2 < \bar{\theta}$, an MEU receiver acts exactly the same as when receiving $m' \in \mathcal{M}$ from a type in $[\theta'_1, \theta_2]$, for any $\theta'_1 < \theta_1$. For an SEU receiver, this behavioural pattern is impossible, the optimal action would necessarily shift to the left.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5 Discussion

This paper models the transmission of expert-based scientific knowledge as cheap talk communication about 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 preference I considered, I showed that all equilibria are outcome equivalent to a partition equilibrium. 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, since 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 about states proposed in Crawford and Sobel (1982). In the linear-quadratic example I introduced, more cut-offs can exist in the MEU case than in the SEU one, for a given bias. This shows that when the expert's preferred action is aligned with the effect of ambiguity aversion, his influence is extremely high; but in the opposite case, it is nonexistent.

In this final section, I discuss further implications of my framework and other possible applications.

Equilibrium selection. Theorem 2 gives that S is always interim better-off by adopting the most informative equilibrium strategy in his communication. This result differs significantly from those obtained in CS’s framework. Under their monotonicity condition (M), CS show that the ex-ante expected payoffs for both Sender and Receiver are maximal for the equilibrium with the most cut-offs. Condition (M) is satisfied if for any two sequences 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. An alternative approach regarding equilibrium selection has been presented by Chen et al. (2008), who propose a condition on utility functions called NITS. Under this condition, combined with Assumption (M), only the equilibrium with the 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). In my case, one could adopt interim Pareto dominance as a selection criterion, which is immune to the limitations of ex-ante Pareto dominance and does not require supplementary assumptions as NITS does. Nevertheless, it brings out the same (most informative) equilibrium and provides a foundation for the attention it receives in applications.

Objective imprecision. Let us assume we stick to the interpretation of \mathcal{C} as the set of objective possible models. The size of \mathcal{C} captures the degree of objective imprecision in scientific knowledge. For instance, if $\mathcal{C} = [0, 1]$, the objective probability of the high state is between 0 and 1. It is then possible to analyse the effects of a change in this objective imprecision. In particular, an increase in precision can move science from a state of division (where $\tilde{\theta}$ is an interior point of \mathcal{C}) to one of consensus (where $\tilde{\theta}$ is at the boundary of \mathcal{C}). As exposed until now, such an increase in precision could result in a major change in communication possibilities, making R fully influential over \mathcal{C} or, on the contrary, completely inaudible. Take our linear-quadratic example in the case when $c = 0$. Then science is divided because $\tilde{\theta} = 0.5$ and all interior cut-offs are in $[0, 0.5]$, whatever the misalignment $b > 0$ of the sender. Assume the objective imprecision \mathcal{C} shifts from $[0, 1]$ to $\mathcal{C}' = [0, 0.5]$. Then science becomes consensual and the partitioning of the set of types is possible

over the entire \mathcal{C}' . Conversely, objective imprecision could shift to $\mathcal{C}'' = [0.5, 1]$. Science would then be consensual as well, but partitioning impossible.

Other applications While communication around the COVID-19 pandemic is a good example of the phenomena described in this paper, there are other interesting examples of communication about scientific models. The estimation of the effects of greenhouse gas (GHG) emissions on global temperature is one. It relies heavily on *black box* prospective computer simulations. 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 extremely difficult for an expert in this field to justify why a given prospective simulation was chosen, a given methodology implemented or given assumptions made. Predicting the impact of the rise of global temperatures also involves 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. 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. They show that under ambiguity-averse preferences, model disagreement is the main driver of GHG abatements. 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.

Another interesting application of this study is the modelling of the social world by economists. Constantly, economists 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 great deal of criticism.

The epistemological status of expert assessments 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 one can verify their truth by perception or inference. This position has been called *reductionist* and has notably been defended by Hume (1740) and Chisholm et al. (1989). But in that case, how 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 consider the expert's judgement as knowledge (Goldman, 2001). It is the strength of this bond 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 possess is the information he holds *at equilibrium*. Because this approach is formal, it is possible to finely describe the knowledge of scientific models a layman is, solely through experts, entitled to. In addition, it emerged from the present analysis that the degree of consensus among models also plays a major part in the foundation of expert-based knowledge.

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 state, 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 about 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 6. 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 the two players' optimal actions are never aligned, whatever the model.

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

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




Figure 6: Sharpness Assumption

Proof of lemma 3:

For player i and any $\theta \in \mathcal{C}$, define $f_i^\theta : a \rightarrow (1 - \theta)\frac{\partial u_i(a,0)}{\partial a} + \theta\frac{\partial u_i(a,1)}{\partial a}$. 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 $h^\theta : a \rightarrow f_R^\theta(a) - f_S^\theta(a)$.

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

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

□

Lemma 3 states that whatever the model realised, R 's and S 's optimal actions are always ordered in the same direction. Note 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. However, results regarding the linear-quadratic example differ. The noise decreases the precision of information transmission (cell sizes), acting as an additional bias.

Following Gajdos et al. (2008), I assume that S does not know the true model θ_0 but only observes an interval of models $[\theta_0 - \epsilon, \theta_0 + \epsilon]$ of size $2\epsilon > 0$.

Assume that S's preferences under uncertainty are MEU. Then, having observed $[\theta_0 - \epsilon, \theta_0 + \epsilon]$, S evaluates action a through:

$$V_S^{MEU}(a) = \min_{\theta \in [\theta_0 - \epsilon, \theta_0 + \epsilon]} \mathbb{E}_\theta(u_S(a, \omega))$$

Then, note that the structure of equilibria is unaffected by those changes. Proposition 2, which guarantees that all equilibria 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 $\tilde{\theta}$ (Theorem 1) is also unaffected under my assumptions. Recall that there cannot be a cut-off θ_k above $\tilde{\theta}$ 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 players (assumption 2), it is not possible for $A_R(\theta_k - \epsilon) > A_S(\theta_k - \epsilon)$.

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 2 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 utility, as the noise and the sender's ambiguity aversion reduce the precision of communication.

B α -MEU receiver

B.1 Optimal actions and structure of equilibria

In this section, I consider the case where R evaluates 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 :

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

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 coincide with MEU.⁸ Adapting Ghirardato et al. (2004)'s proposition 20 to our model, we can state the following:

Definition 4 (Ghirardato et al. (2004)). *Receiver i , evaluating actions through $V_R^{\alpha_i}$, is said to be more ambiguity-averse⁹ than receiver j , evaluating actions through $V_R^{\alpha_j}$, if and only if*

$$\alpha_i > \alpha_j$$

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

As for the MEU case, having received an equilibrium message $m \in \text{supp}(\sigma^*)$, 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 $A_R(B) = \text{argmax}_{a \in \mathcal{A}} V_R^\alpha(a, B)$ the set of optimal actions of the α -MEU receiver when his belief is B .

Figure 7 illustrates the ex-ante evaluation of actions by 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, note that for a given α $V_R^\alpha(a)$ is not necessarily single-peaked. For instance, for $\alpha = 0.3$, $V_R^{0.3}(a)$ is maximal at 0.3

⁸In general, for non-symmetric utility functions, α -MEU does not have SEU as a special case here. For instance, both criteria coincide if the set of models is a singleton. Yet, this set depends on the information conveyed by the sender, which, at equilibrium, is never a singleton.

⁹In the sense of Ghirardato and Marinacci (2002)

and 0.7.¹⁰




Figure 7: α -MEU ex-ante valuation

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

Proposition 5. *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 set of actions $A_R(B) \subset [A_R(\theta_1), A_R(\theta_2)]$. In the context of the linear-quadratic example, the set of optimal actions of an α -MEU receiver is given by:*

$$A_R(B) = \begin{cases} \alpha A_R(\theta_2) + (1 - \alpha) A_R(\theta_1) & \text{if } \theta_2 < \tilde{\theta} \\ \{\alpha A_R(\tilde{\theta}) + (1 - \alpha) A_R(\theta_M) \mid \theta_M \in \operatorname{argmax}_{\theta \in \{\theta_1, \theta_2\}} \mathbb{E}_\theta(u_R(a, \omega))\} & \text{if } \tilde{\theta} \in B \\ \alpha A_R(\theta_1) + (1 - \alpha) A_R(\theta_2) & \text{if } \theta_1 > \tilde{\theta} \end{cases}$$

A direct consequence of Proposition 5 is that ex-ante, for any $a \in \mathcal{A}$, $\min_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a, \omega)) = \mathbb{E}_{\bar{\theta}}(u_R(a, \omega))$ and that $\max_{\theta \in \mathcal{C}} \mathbb{E}_\theta(u_R(a, \omega)) = \max(\mathbb{E}_{\underline{\theta}}(u_R(a, \omega)), \mathbb{E}_{\bar{\theta}}(u_R(a, \omega)))$. Thus, $A_R(\mathcal{C}) = A_R(\alpha \tilde{\theta} + (1 - \alpha) \bar{\theta})$ when $\mathbb{E}_{\underline{\theta}}(u_R(a, \omega)) < \mathbb{E}_{\bar{\theta}}(u_R(a, \omega))$, $A_R(\mathcal{C}) = A_R(\alpha \tilde{\theta} + (1 - \alpha) \underline{\theta})$ when $\mathbb{E}_{\underline{\theta}}(u_R(a, \omega)) > \mathbb{E}_{\bar{\theta}}(u_R(a, \omega))$ and $A_R(\mathcal{C}) = \{A_R(\alpha \tilde{\theta} + (1 - \alpha) \underline{\theta}), A_R(\alpha \tilde{\theta} + (1 - \alpha) \bar{\theta})\}$ when $\mathbb{E}_{\underline{\theta}}(u_R(a, \omega)) = \mathbb{E}_{\bar{\theta}}(u_R(a, \omega))$. This explains the fact that in the example we considered before, where $\mathbb{E}_0(u_R(a, \omega)) = \mathbb{E}_1(u_R(a, \omega))$, optimal actions were not unique. Note also that, when α increases, $A_R(\mathcal{C})$ gets closer to \tilde{a} , in the Euclidean sense. Thus, an increase in ambiguity aversion moves R's ex-ante optimal action closer

¹⁰In the linear-quadratic example for $c = 0$, there are two optimal actions for any $\alpha \leq 0.5$. The fact that this threshold is the one separating love of and aversion to ambiguity is non-generic. For sharper utility functions, this threshold would be above 0.5. A formal definition is given in assumption 4 of the appendix

to the precautionary action.

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

Proposition 6. *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 5. 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.

B.2 Comparative ambiguity aversion

In the following, I examine the effect that ambiguity aversion has on the structure of partition 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 by their degree of ambiguity aversion α_1 and α_2 . Note that the ex-post optimal action $A_{R_i}(\theta)$ is unaffected by ambiguity aversion, thus, I will erase the subscript. I will only consider the linear-quadratic example introduced before. Recall that in the MEU case, when $c = 1$, science is consensual and communication is possible over the entire \mathcal{C} ; when $c = 0$, science is divided and communication is only possible over $(0, \frac{1}{2})$, and finally when $c = -1$, science is consensual but no information transmission is possible.

I start by considering the consensual science cases: $c = 1$ or -1 . In the following, I characterise all the cut-offs of the corresponding partition equilibrium.

Proposition 7. *In the linear quadratic example, when R is α -MEU, for $\alpha \notin \{0, \frac{1}{2}, 1\}$:*

- *When $c = 1$, there are $N > 0$ cut-off equilibria, one for each cut-off, and the k -th cut-off of the $1 \leq n \leq N$ cut-off equilibrium is given by:*

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

- When $c = -1$, there are $M > 0$ cut-off equilibria, one for each cut-off, and the k -th cut-off of the $1 \leq n \leq M$ cut-off equilibrium is given by:

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




Figure 8: 3-cut-offs equilibria for $c = 1$

Figures 8 and 9 compute the cut-offs of the 3-cut-off equilibrium as a function of α for a fixed positive bias $b = 0.01$. Note 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 $c = -1$, cut-off values decrease with α towards $\tilde{\theta}$. Thus for a given bias, there is a level of ambiguity aversion from which all types in \mathcal{C} must pool. Conversely, simulations suggest that, when $c = 1$, cut-off values continuously increase with α towards their MEU values.




Figure 9: 3-cut-offs equilibria for $c = -1$

I now formally prove that in the case of consensual science, no communication is possible in $c = -1$ when α is above a given threshold. In addition, I show that for a given bias, ambiguity aversion eases the existence of an n cut-off equilibrium, for $n \geq 2$.

Proposition 8. *In the context of the linear-quadratic example, when $c = -1$, $b > 0$ 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 $[0, 1]$. 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 an $n \geq 2$ cut-off equilibrium between S and α_1 , there is an n cut-off equilibrium between S and α_2*

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

In addition, when there is an equilibrium with at least 3 cut-offs, ambiguity aversion eases the

existence of an n cut-off equilibrium. Recall that 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 $N \geq 2$, $\underline{\theta}_{N-1}^N(\alpha)$ is a strictly increasing function and that $\lim_{\alpha \rightarrow +\infty} \underline{\theta}_{N-1}^N(\alpha) < \frac{1}{2}$.

I know further prove that the first result of Proposition 8 extends to the case where $c = 0$ and science is divided:

Proposition 9. *In the context of the linear-quadratic example when $c = 0$, $b > 0$ 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 $[\frac{1}{2}, 1]$ must pool. This suggests 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 $\tilde{\theta}$ must pool. The proof of Proposition 9 builds on the one of proposition Proposition 8. I show that for any $N \geq 2$, $\bar{\theta}_{N-1}^N(\alpha)$ is a strictly decreasing and continuous function and that $\lim_{\alpha \rightarrow +\infty} \bar{\theta}_{N-1}^N(\alpha) < \frac{1}{2}$. As a result, there must be $\alpha \in (\frac{1}{2}, 1)$ such that no partitioning of $[\frac{1}{2}, 1]$ is possible.

C 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 4. *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 $A_R(B) \subset [A_R(\theta_1), A_R(\theta_2)]$.*

Proof of lemma 4:

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

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

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

□

A consequence of the Lemma 4 is that when looking for optimal actions for a given B , it is sufficient to look for actions in $[A_R(\theta_1), A_R(\theta_2)]$. Note 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 $\tilde{a} = a_R(0)$ and thus $\tilde{\theta} = 0$, $\mathbb{E}_\theta(u_R(a, \omega))$ 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 $\tilde{a} = a_R(1)$ and $\tilde{\theta} = 1$. 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 \tilde{a} it must be that this point is \tilde{a} . As a result:

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

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 $\theta_2 < \tilde{\theta}$ and strictly increasing with θ when $\theta_1 > \tilde{\theta}$, which gives the corresponding result. The above system also implies that when $\tilde{\theta} \in B$, $\mathbb{E}_\theta(u_R(a, \omega))$ is always minimal for $\theta = \tilde{\theta}$. 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:

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

The above system implies that when $\tilde{\theta} \in B$, $\min_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega))$ is increasing on $(A_R(\theta_1), \tilde{a})$

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

□

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 obtain the same equilibrium outcome must form an interval. The continuity and the strict monotonicity of the sender's preferences closes the argument.

Lemma 5. *There exists $\epsilon > 0$ such that if u and v are actions induced in equilibrium, $|u - v| \geq \epsilon$. Furthermore, the set of actions induced in equilibrium is finite.*

Proof of Lemma 5

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, note that if θ induces \bar{a} , it must be that $V_S^\theta(\bar{a}) = \max_{a \in Y} V_S^\theta(a)$. 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 θ' 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 $\theta \rightarrow V_S^\theta(u) - V_S^\theta(v)$, 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, note 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:

$$\begin{aligned} \forall \theta \in \Theta_u, \theta &\leq \hat{\theta} \\ \forall \theta \in \Theta_v, \theta &\geq \hat{\theta} \end{aligned}$$

Thus, when R is MEU, Lemma 4 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 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 4 implies that for any belief $B \subset \mathcal{C}$, the optimal action of the receiver is in $[A_R(\underline{\theta}), A_R(\bar{\theta})]$. Thus, the set of actions induced in equilibrium is bounded by $A_R(\underline{\theta})$ and $A_R(\bar{\theta})$ and at least ϵ away from one another, which completes the proof. □

Lemma 6. *In every equilibrium of the game, if a is an action induced by type θ and by type θ'' for some $\theta < \theta''$, then a is also induced by $\theta' \in (\theta, \theta'')$*

Proof of Lemma 6:

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 θ'' 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 θ' prefers and that θ'' does not. Formally, this is:

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

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 $\tilde{a}_S = \operatorname{argmax}_{a \in \mathcal{A}} \min_{\omega \in \Omega} u_S(a, \omega)$. \tilde{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 noted. Either the high state is sufficiently worse than the good one for it to give a lower utility at its optimal point: $u_S(a_S(1), 1) \leq u_S(a_S(1), 0)$. Then the hedging action is the optimal action in the high state $\tilde{a}_S = a_S(1)$. Or 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, \tilde{a}_S is the action that gives the same utility in both states.

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

- When $a_1 < \tilde{a}_S$ and $a_2 \geq \tilde{a}_S$ can cross at most once and system (5) is impossible.
- Assume $\tilde{a}_S \leq a_1 < a_2$. 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 $a \geq \tilde{a}_S$, $u_S(a, 1)$ 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 $a_1 < a_2 < \tilde{a}_S$. 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 symmetrical.

□

By Lemma 5, there is a finite number of outcomes induced in equilibrium. The continuity of $A_S(\theta)$ gives that there is a type of the sender which is indifferent between any pair of outcomes induced in equilibrium and the monotonicity of $A_S(\theta)$ implies there are only a finite number of types which are indifferent between any pair of outcomes. Hence, Lemma 6 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 prefers it to any other message.¹¹ In particular, it must be preferred to message m_{k-1} ,

¹¹Any 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.

which induces the preferred equilibrium action of types in C_{k-1}^q . For all $\theta \in C_k^q$:

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

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

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

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

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

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

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

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

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

□

Proof of Theorem 1

Assume there is a $q + 1$ cut-off equilibrium and that $\theta_{q-1} < \tilde{\theta} < \theta_q$. As $\theta_q > \tilde{\theta}$, we have that:

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

As A_R is a strictly increasing function and because S is misaligned upwards, 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 $[0, A_S(\theta_q)]$, we have that $\mathbb{E}_{\theta_q}(u_S(y^*(m_{q-1}), \omega)) < \mathbb{E}_{\theta_q}(u_S(y^*(m_q), \omega)) \iff V_S^{\theta_q}(m_{q-1}) < V_S^{\theta_q}(m_q)$, which is a contradiction to the assumption that θ_q is a cut-off type.

□

Proof of Lemma 2:

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

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

Step 1:

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

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

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

In C_0 , type θ_1 has the most incentive to deviate from sending m_0 to sending m_1 , which would

induce a higher action, as, $V_R^{MEU}(a, C_1) = \mathbb{E}_{\theta_2}(u_R(a, \omega))$ and $A_R(\theta)$ is strictly increasing by Lemma 1.

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

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

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

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

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

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

Furthermore, the fact that $\cup_{k=0}^M C_k = \mathcal{C}$ and that for every pair of consecutive cells of the partition the incentive constraints are transitive gives that condition (7) is both necessary and sufficient. As $A_R(\theta)$ is strictly monotonic, it implies that $A_R(\theta_k) \neq A_R(\theta_{k+1})$. $\underline{\theta}$ being known, it is possible to derive θ_1 directly from (6). By repeating the reasoning by induction, θ_{k+1} can be derived from θ_k for $k \in 1, \dots, M - 1$ from (7) as long as there is $\theta_M < \tilde{\theta}$.

Step 2:

I show that any partitional strategy of the sender characterised by $\theta_0 < \dots < \theta_q$ is an equilibrium strategy. I proceed by iteration:

- Step 1 proves that $\theta_0 < \dots < \theta_M$ characterise an equilibrium. Let us show that $\theta_0, \theta_2 < \dots < \theta_M$ does as well.

Assume S's strategy is σ_{M-1} such that for $2 \leq k \leq M$, $\forall \theta \in C_k$, $\sigma_{M-1}(\theta) = m_k$ and $\forall \theta \in [\underline{\theta}, \theta_2]$, $\sigma_{M-1}(\theta) = m_0$. Then for $k \in 1, \dots, n-1$, 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 $\theta \in [\underline{\theta}, \theta_2]$ the fact that, for every pair of consecutive cells of the partition, the incentive constraints are transitive implies that message m_0 is preferred to any other message.

- Let us 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 will show that any equilibria only characterised by elements of $\theta_0, \dots, \theta_q$ must be characterised by elements of \mathcal{E} . It is straightforward to see that any equilibria only characterised by elements of $\theta_0, \dots, \theta_q$ 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.

For $1 \leq q \leq M$, consider a strategy σ_p characterised by cut-offs $\theta_0, \theta_q, \dots, \theta_{p-1}, \theta_{p+1}, \dots, \theta_M$ for $q+1 \leq p \leq M$ and assume it is an equilibrium strategy.¹² 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, type 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 $\theta^* \notin \{\theta_1, \dots, \theta_M\}$. Assume without loss of generality that

¹²the choice of removing an element $\theta_p < \tilde{\theta}$ is without loss of generality.

$\theta_p < \theta^* < \theta_{p+1}$ for $p \in 1, \dots, M - 1$. Then we have that:

$$\begin{aligned} W^{\theta^*}([\theta_p, \theta^*]) &= W^{\theta^*}([\theta^*, \theta_{p+1}]) \\ \iff \mathbb{E}_{\theta^*}(u_S(A_R(\theta^*))) &= \mathbb{E}_{\theta^*}(u_S(A_R(\theta_{p+1}))) \end{aligned}$$

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

Proof of Theorem 2:

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

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

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

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

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

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

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

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

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

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

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

□

Proof of Proposition 4:

1. Assume R has SEU preferences. Assume there are n equilibrium cut-offs in $[0, 1]$: $\theta_0, \dots, \theta_n$ and thus $\theta_0 = 0$ $\theta_n = 1$. When receiving equilibrium message m_k sent by types $\theta \in [\theta_k, \theta_{k+1})$ S evaluates action through:

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

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

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

Note that this arbitrage condition translates into the same condition as in CS's example:

$$\theta_{k+1} - \theta_k = \theta_k - \theta_{k-1} + 4b \tag{8}$$

Equation (8) further gives that:

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

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

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

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

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

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

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

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

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

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

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

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

$$\begin{aligned} \frac{\tilde{\theta} + \theta_{n-1}}{2} &= \theta_{n-1} + b \\ \iff \theta_{n-1} &= 1 - 2b \end{aligned}$$

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

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

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

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

□

Proof of Corollary 1:

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

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

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

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

□

Proof of Proposition 5 :

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 $u_R(a, 1) - u_R(a, 0)$. We want to prove that :

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

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

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

where $\theta_M = \operatorname{argmax}_{\theta \in \{\theta_1, \theta_2\}} \mathbb{E}_\theta(u_R(a, \omega))$. A consequence of Lemma 4 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 $\tilde{a} = a_R(0)$ and thus $\tilde{\theta} = 0$, $\mathbb{E}_\theta(u_R(a, \omega))$ 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 $\tilde{a} = a_R(1)$ and $\tilde{\theta} = 1$. 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, the two utilities are equal for a unique given action and by definition of \tilde{a} it must be that this point is \tilde{a} . As a result:

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

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:

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

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

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:

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

where $\theta_M = \operatorname{argmax}_{\theta \in \{\theta_1, \theta_2\}} \mathbb{E}_\theta(u_R(a, \omega))$. The above system implies that when $\tilde{\theta} \in B$, $\max_{\theta \in B} \mathbb{E}_\theta(u_R(a, \omega))$ is decreasing on $(A_R(\theta_1), \tilde{a})$ (as $\mathbb{E}_{\theta_1}(u_R(a, \omega))$ is maximal at $A_R(\theta_1)$) and increasing on $(\tilde{a}, A_R(\theta_2))$ (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, $\max_{\theta \in B} A_R(B) = A_R(\theta_M)$.

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

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

which implies that:

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

□

Proof of Proposition 6

Lemma 7. *There exists $\epsilon > 0$ such that if u and v are actions induced in equilibrium, $|u - v| \geq \epsilon$. Furthermore, 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 \bar{a} , it must be that $V_S^\theta(\bar{a}) = \max_{a \in Y} V_S^\theta(a)$. 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 θ' 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 $\theta \rightarrow V_S^\theta(u) - V_S^\theta(v)$, 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 :

$$\begin{aligned} \forall \theta \in \Theta_u, \theta &\leq \hat{\theta} \\ \forall \theta \in \Theta_v, \theta &\geq \hat{\theta} \end{aligned}$$

Thus, when R is α -MEU, Lemma 4 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 4 implies that for any belief $B \subset \mathcal{C}$, the optimal action of the receiver is in $[A_R(\underline{\theta}), A_R(\bar{\theta})]$. Thus, the set of actions induced in equilibrium is bounded by $A_R(\underline{\theta})$ and $A_R(\bar{\theta})$ 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 6 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 7:

I focus on the case $c = 1$. The case $c = -1$ is symmetrical.

Assume there is an $n > 0$ cut-off equilibrium. It follows from the characterisation of cut-off types in the linear-quadratic example given in the proof of Proposition 4 and the characterisation of optimal actions in the α -MEU case given in Proposition 5 that, for $1 \leq k \leq n - 2$ and $\alpha > 0$:

$$\begin{aligned}\underline{\theta}_k^n + b &= \frac{\alpha \underline{\theta}_k^n + (1 - \alpha) \underline{\theta}_{k-1}^n + \alpha \underline{\theta}_{k+1}^n + (1 - \alpha) \underline{\theta}_k^n}{2} \\ &\iff \theta_{k+1} - \theta_k = \frac{1 - \alpha}{\alpha} (\underline{\theta}_k^n - \underline{\theta}_{k-1}^n + \frac{2b}{\alpha})\end{aligned}$$

Set $V_k = \underline{\theta}_{k+1}^n - \underline{\theta}_k^n$. 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 $\alpha \notin \{0, \frac{1}{2}\}$:

$$V_k = \left(\frac{1 - \alpha}{\alpha}\right)^k \left(\underline{\theta}_1^n - \frac{2b}{2\alpha - 1}\right) + \frac{2b}{2\alpha - 1}$$

By induction, it follows that:

$$\begin{aligned}\underline{\theta}_{k+1}^n &= \sum_{j=1}^k \left[\left(\frac{1 - \alpha}{\alpha}\right)^j \left(\underline{\theta}_1^n - \frac{2b}{2\alpha - 1}\right) + \frac{2b}{2\alpha - 1} \right] + \theta_1 \\ \iff \underline{\theta}_{k+1}^n &= \sum_{j=0}^k V_j \\ \iff \underline{\theta}_k^n &= \left(\theta_1 - \frac{2bn}{2\alpha - 1}\right) \left(\frac{1 - \left(\frac{1 - \alpha}{\alpha}\right)^k}{1 - \left(\frac{1 - \alpha}{\alpha}\right)}\right) + \frac{2bk}{2\alpha - 1}\end{aligned}$$

In particular, it must be that $\underline{\theta}_n^n = \frac{1}{2}$ which give that $\underline{\theta}_1^n = \left(\frac{1}{2} - \frac{2bn}{2\alpha - 1}\right) \left(\frac{1 - \left(\frac{1 - \alpha}{\alpha}\right)}{1 - \left(\frac{1 - \alpha}{\alpha}\right)^n}\right) + \frac{2b}{2\alpha - 1}$. As a result, we obtain that:

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

□

Proof of Proposition 8:

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

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

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

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

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

$$\lim_{\alpha \rightarrow 1} \bar{\theta}_{n-1}^n(\alpha) = -2b < \frac{1}{2}$$

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

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

Thus :

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

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:

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

$\frac{1-\alpha}{\alpha} \in (0, 1/2)$ 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 symmetrical process, we can prove that $\bar{\theta}_1^n(\alpha)$ 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 $\bar{\theta}_1^n(\alpha_1) \in (0, 1)$. As $\bar{\theta}_1^n(\alpha)$ is a decreasing function, it must be that $\bar{\theta}_1^n(\alpha_2) < 1$. In addition, as $\bar{\theta}_{n-1}^n(\alpha)$ is an decreasing function, it follows that $\bar{\theta}_{n-1}^n(\alpha_2) > \lim_{\alpha \rightarrow 1} \bar{\theta}_{n-1}^n(\alpha) = \frac{1}{2} - 2b > 0$ 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 9:

We are in the case where $c = 0$. Recall that in this case $\tilde{\theta} = \frac{1}{2}$.

- Consider the case of an n cut-off equilibrium that has $n \leq N - 3$ cut-off types below $\frac{1}{2}$.

I call $\theta_k^N(\alpha)$ the k -th cut-off of our equilibrium. By Proposition 5, we get that for $n + 2 \leq k \leq N - 1$ and $\frac{1}{2} < \alpha < 1$:

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

as before, we obtain by induction that:

$$\theta_k^N(\alpha) = \left(\theta_{n+1}^N(\alpha) - \frac{2bN}{2\alpha - 1}\right) \left(\frac{1 - \left(\frac{\alpha}{1-\alpha}\right)^k}{1 - \left(\frac{\alpha}{1-\alpha}\right)}\right) - \frac{2bk}{2\alpha - 1} - \frac{1}{2}$$

In particular, it must be that $\theta_N^N = 1$ which gives that $\theta_{n+1}^N = \left(1 + \frac{2bN}{2\alpha - 1}\right) \left(\frac{1 - \left(\frac{\alpha}{1-\alpha}\right)}{1 - \left(\frac{\alpha}{1-\alpha}\right)^N}\right) + \frac{2b}{2\alpha - 1}$.

As a result, we get that :

$$\theta_{N-1}^N(\alpha) = \left(1 - \frac{2bN}{2\alpha - 1}\right) \left(\frac{1 - \left(\frac{\alpha}{1-\alpha}\right)^{N-1}}{1 - \left(\frac{\alpha}{1-\alpha}\right)^N}\right) - \frac{2b(N-1)}{2\alpha - 1} + \frac{1}{2}$$

Yet, reproducing the reasoning in the proof of Proposition 8, we can show that $\theta_{N-1}^N(\alpha)$ is a strictly decreasing function and that:

$$\lim_{\alpha \rightarrow 1} \theta_{N-1}^N(\alpha) = \frac{1}{2} - 2b < \frac{1}{2}$$

Thus, as $\theta_{N-1}^N(\alpha)$ is strictly decreasing and continuous, there is $\alpha(b) \in (1/2, 1)$ such that $\theta_{N-1}^N(\alpha) = \frac{1}{2}$. 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 $[\frac{1}{2}, 1]$.

- Second, consider the case where $n = N - 2$. Then, there is a single cut-off type in $(\frac{1}{2}, 1)$. Call that type $\theta_{N-1}^N(\alpha) \in (\frac{1}{2}, 1)$. By Proposition 5 it must be that:

$$\Leftrightarrow \begin{cases} \theta_{N-1}^N(\alpha) + b = \frac{\alpha \frac{1}{2} + (1-\alpha)\theta_n^N + \alpha\theta_{N-1}^N(\alpha) + (1-\alpha)1}{2} \\ \text{or} \\ \theta_{N-1}^N(\alpha) + b = \frac{\alpha \frac{1}{2} + (1-\alpha)\theta_{n+1}^N + \alpha\theta_{N-1}^N(\alpha) + (1-\alpha)1}{2} \\ \theta_{N-1}^N(\alpha) = \frac{1}{2} - \frac{1-\alpha}{2-\alpha}\theta_n^N - 2b \\ \text{or} \\ \theta_{N-1}^N(\alpha) = \frac{1}{2} - \frac{1-\alpha}{2-\alpha}\theta_{n+1}^N - 2b \end{cases}$$

In both cases we have that $\lim_{\alpha \rightarrow 1} \theta_{N-1}^N(\alpha) = \frac{1}{2} - 2b < \frac{1}{2}$. By the same argument as above, there must be $\alpha(b) \in (0, 1)$ such that, for $\alpha \geq \alpha(b)$, only one action can be induced by types in $[\frac{1}{2}, 1]$.

- Finally, consider the case where $N = 2$. Then, either $\theta_{N-1}^N(\alpha) \leq \frac{1}{2}$ for any $\alpha > \frac{1}{2}$, either there is $\alpha > \frac{1}{2}$ such that $\theta_{N-1}^N(\alpha) > \frac{1}{2}$. In that second case, following Proposition 5 it must be that:

$$\begin{aligned} \theta_{N-1}^N(\alpha) + b &= \frac{\alpha\tilde{\theta} + (1-\alpha)0 + \alpha\theta_{N-1}^N(\alpha) + (1-\alpha)1}{2} \\ \Leftrightarrow \theta_{N-1}^N(\alpha) &= \frac{2-\alpha-4b}{4-2\alpha} \end{aligned}$$

As $\lim_{\alpha \rightarrow 1} \frac{2-\alpha-4b}{4-2\alpha} = \frac{1}{2} - 2b < \frac{1}{2}$. By the same argument as above, there must be $\alpha(b) \in (0, 1)$ such that, for $\alpha \geq \alpha(b)$, only one action can be induced by types in $[\frac{1}{2}, 1]$.

In all three cases one can show that $\frac{\partial \theta_{N-1}^N(\alpha)}{\partial b} < 0$, which implies that that $\alpha(b)$ is a decreasing function.

□

References

- Abdellaoui, M., Colo, P., and Hill, B. (2021). Eliciting multiple prior beliefs. *HEC Paris Research Paper Forthcoming*.
- Andrews, I. and Shapiro, J. M. (2021). A model of scientific communication. *Econometrica*, 89(5):2117—2142.
- Banerjee, A. V., Chassang, S., Montero, S., and Snowberg, E. (2020). A theory of experimenters: Robustness, randomization, and balance. *American Economic Review*, 110(4):1206–30.
- 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.
- 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.
- Gajdos, T., Hayashi, T., Tallon, J.-M., and Vergnaud, J.-C. (2008). Attitude toward imprecise information. *Journal of Economic Theory*, 140(1):27–65.
- 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.
- 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.
- 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.
- Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892.
- 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.

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

Spiess, J. (2018). Optimal estimation when researcher and social preferences are misaligned.

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