



Public Persuasion

Marie Laclau, Ludovic Renou

► **To cite this version:**

| Marie Laclau, Ludovic Renou. Public Persuasion . 2016. <hal-01285218>

HAL Id: hal-01285218

<https://hal-pse.archives-ouvertes.fr/hal-01285218>

Submitted on 8 Mar 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Public Persuasion*

Marie Laclau[†] and Ludovic Renou[‡]

This version: March 2, 2016

First version: May 10, 2013

Abstract

This short paper studies the problem of public persuasion, that is, when a sender has to persuade multiple receivers, possibly having heterogenous beliefs, with the same information for all. We show that public persuasion constrains the sender in how he can influence the beliefs of receivers: if the sender wants to influence the beliefs of one particular receiver, he loses all controls over the beliefs of the others. This observation partially generalizes to targeted persuasion.

Keywords: Public, targeted, persuasion, multiple priors, splitting, concavification.

1 Introduction

This short paper studies the problem of *public persuasion*. The leading examples we have in mind are how advertising campaigns persuade consumers to adopt a firm's products, how public appearances on social and mass media persuade electors to vote for a politician, how national health campaigns persuade individuals to quit smoking, to immunize or to exercise, to name just a few.

A common feature of all the above examples is that a sender needs to persuade multiple receivers, possibly having diverse opinions, with the same information for all. In other words,

*We thank the hospitality of the Cowles Foundation at Yale University, where part of this project was written. This project was supported by the French National Research Agency, under grants CIGNE (ANR-15-CE38-0007-01) and AmGame (ANR 12-FRAL-0008-01) and through the program Investissements d'Avenir (ANR-10-LABX-93-0).

[†]Paris School of Economics and CNRS, 48 boulevard Jourdan, 75014 Paris, France, marie.laclau@psemail.eu

[‡]Department of Economics, University of Essex, CO4 3SQ, United Kingdom, lrenou@essex.ac.uk

the sender cannot target its information release to each receiver (or groups of receivers), but has to release the same information to all receivers. This has important implications for the distributions of posterior beliefs a sender can generate: the likelihood ratios of any two receivers, i.e., the ratios of posterior beliefs over prior beliefs, must be collinear (see Lemma 1, condition (2)(ii)). This collinearity condition encapsulates the constraints public persuasion imposes on the sender. The best the sender can do is to “split” the prior beliefs of a selected but unique receiver. The posterior beliefs of all other receivers are then uniquely pinned down as a consequence of the collinearity condition. This is the main economic insight of our analysis and underlines our main result, the complete characterization of the sender’s optimal payoff (Theorem 1). Moreover, we show that this insight generalizes without much difficulties to targeted persuasion, i.e., to situations where the sender can target its release of information to groups of individuals.

Related literature. This paper contributes to the growing literature on Bayesian persuasion, as defined by Kamenica and Gentskow. In Kamenica and Gentskow (2011), there is a single sender and a single receiver, sharing a common prior. In Alonso and Câmara (2015a), there is a single receiver, whose beliefs may differ from the sender’s beliefs. In Alonso and Câmara (2015b), there are multiple receivers, but they all share the same prior beliefs as the sender. The present paper generalizes all these contributions in assuming multiple receivers with heterogenous beliefs. Moreover, the emphasis of the paper complements the work of Alonso and Câmara (2015b). We are interested in characterizing the optimal payoff the sender can obtain in general persuasion games with multiple receivers, while Alonso and Câmara are interested in particular games, namely voting games, and in understanding how the voting procedure impacts on the information release.

Mathematically, the paper is related to the seminal work of Aumann and Maschler (1967, reprinted in 1995) on zero-sum repeated games with one-sided incomplete information. We generalize the “splitting” lemma to account for multiple receivers with heterogeneous beliefs and, as Aumann and Maschler, characterize the optimal payoff of the sender as the concavification of a certain function.¹

2 The general setup

We consider a reduced-form model, where the sender’s payoff depends directly on profiles of posterior beliefs and states of the world. As we will see later, many economic applications reduce to our model.

¹Kamenica and Gentskow refer to a splitting as Bayes plausible posteriors and to the concavification of a function as the convex hull of a function.

Let Ω be a finite set of states of the world and I an index set, e.g., $I = \{1, \dots, n\}$ or $I = [0, 1]$. Let $(p_i)_{i \in I}$ be a profile of prior beliefs, with $p_i \in \text{int}\Delta(\Omega)$ for each $i \in I$, and $p^* \in \Delta(\Omega)$ the sender's prior beliefs.²

Prior to learning the state of the world, the sender has the opportunity to commit to release some public information. Formally, the sender commits to a signaling function $\pi : \Omega \rightarrow \Delta(S)$, with S a finite set of signals.³ We write $p_i^{s,\pi}$ for the posterior of p_i conditional on the signal s when the sender commits to π , i.e.,

$$p_i^{s,\pi}(\omega) = \frac{\pi(s|\omega)p_i(\omega)}{\sum_{\omega \in \Omega} \pi(s|\omega)p_i(\omega)},$$

if $\sum_{\omega \in \Omega} \pi(s|\omega)p_i(\omega) > 0$, and is arbitrary otherwise. Finally, if the profile of posterior beliefs is $(p_i^s)_{i \in I}$ and the state of the world is ω , the sender's payoff is $u((p_i^s)_{i \in I}, \omega)$. This completes the description of our reduced-form model.

The sender's objective is therefore to choose a signaling function π so as to maximize his expected payoff, i.e., to solve the optimization problem:

$$\max_{\pi: \Omega \rightarrow \Delta(S)} \sum_{s, \omega} u((p_i^{s,\pi})_{i \in I}, \omega) \pi(s|\omega) p^*(\omega), \quad (\mathcal{P})$$

Before solving (\mathcal{P}) , we present two sets of economic applications.

Applications I: Persuading an audience. As a first set of applications, consider persuasion games between a sender and multiple receivers, with possibly *heterogenous* beliefs. There is a population of receivers $I = \{1, \dots, n\}$, with $p^i \in \text{int}\Delta(\Omega)$ the prior belief of receiver i . Receiver i has a non-empty compact set of actions A^i and a utility function $v^i : \times_{i \in I} A^i \times \Omega \rightarrow \mathbb{R}$.⁴ The sender's utility is $\tilde{u} : \times_{i \in I} A^i \times \Omega \rightarrow \mathbb{R}$. For each public signal s , let $a^*((p_i^s)_{i \in I})$ be a Nash equilibrium of the game between receivers, when the profile of posteriors is $(p_i^s)_{i \in I}$ (if two signals generate the same set of posteriors, the same Nash equilibrium is selected). Letting $u((p_i^s)_{i \in I}, \omega) := \tilde{u}(a^*((p_i^s)_{i \in I}), \omega)$, the sender's optimal choice is equivalent to the optimization problem (\mathcal{P}) . The seminal problem analyzed in Kamenica and Gentzkow (2011) is a special case of our more general model. As concrete applications, receivers may be consumers, who have to decide whether to purchase a good, or voters, who have to decide whether to vote for a candidate.

Applications II: Bayesian persuasion with ambiguity. As a second set of applica-

²For simplicity, we assume that each prior p_i has full support. A weaker assumption is to assume that priors are absolutely continuous with respect to each others, i.e., $p_i(\omega) > 0 \Rightarrow p_j(\omega) > 0$ for all (i, j, ω) . This weaker assumption does not affect our results.

³We argue later that the finiteness of S is without loss of generality.

⁴Assume that each v_i is continuous in $(a_i)_{i \in I}$ and concave in a_i .

tions, consider persuasion games between a sender and a receiver, who is *ambiguous* about the state of the world. The receiver has maxmin preferences (Gilboa and Schmeidler, 1989), with $P := \{p_i : i \in I\}$ the set of prior beliefs of the receiver.⁵ The receiver has a non-empty compact set of actions A and his utility function is $v : A \times \Omega \rightarrow \mathbb{R}$, while the sender's utility is $\tilde{u} : A \times \Omega \rightarrow \mathbb{R}$. We assume full Bayesian updating, which implies that given a signaling function π and a signal s , the receiver updates his set of prior beliefs into the set of posteriors beliefs $\{p_i^{s,\pi} : i \in I\}$. For each signal s , the receiver chooses an action in:

$$A^*((p_i^s)_{i \in I}) = \arg \max_{a \in A} \min_{i \in I} \sum_{\omega \in \Omega} v(a, \omega) p_i^s(\omega),$$

where $\{p_i^s : i \in I\}$ is the set of posteriors. Let $a^*((p_i^s)_{i \in I})$ be an optimal action of the receiver in $A^*((p_i^s)_{i \in I})$ (if the receiver is indifferent between several actions, then choose arbitrarily). Letting $u((p_i^s)_{i \in I}, \omega) := \tilde{u}(a^*((p_i^s)_{i \in I}), \omega)$, the sender's optimal choice is yet again equivalent to the optimization problem (\mathcal{P}) . As a concrete example, consider the pharmaceutical industry. Pharmaceutical companies (the senders) need to provide information to health agencies (the receivers), e.g., the European Medicine Agency or the Food and Drug Administration, about the effectiveness of proposed medicines in order to get the authorization to sell. Pharmaceutical companies *commit* to perform and release results from all clinical trials.⁶ Health agencies have then to decide whether the medicine can be put on the market. Even if health authorities have the results of all clinical trials, they still face large uncertainties and are likely to entertain ranges of estimates about the effectiveness of the proposed medicines. A simple precautionary principle would require health agencies to take decisions based on their most pessimistic estimates, i.e., to adopt the maxmin criterion.

3 Main results

3.1 Necessary and sufficient conditions for splitting with multiple priors

Notations. For two $|\Omega|$ -dimensional vectors p and q , define the $|\Omega|$ -dimensional vector q/p as follows: $(q/p)(\omega) = q(\omega)/p(\omega)$ if $p(\omega) > 0$, $(q/p)(\omega) = 0$ if $p(\omega) = 0$. For a $|\Omega|$ -dimensional vector p , we write $\|p\|$ for the ℓ_1 -norm of p , i.e., $\sum_{\omega \in \Omega} |p(\omega)|$.

⁵The axiomatization of Gilboa and Schmeidler requires the set of prior beliefs P to be non-empty, compact and convex.

⁶Indeed, recent developments in the pharmaceutical industry have seen firms like GlaxoSmithKline to voluntarily commit to release results from all clinical trials within a reasonable time frame. E.g., the FDA requires all clinical trials to be made available within one year.

We first characterize the set of posterior beliefs the sender can induce with a commitment to a signaling function. When I is the singleton $\{i\}$, the sufficient and necessary condition to *split* the unique prior belief p_i into the posterior beliefs $(p_i^s)_{s \in S}$ is that p_i is in the convex hull of $\{p_i^s : s \in S\}$ (see Aumann and Maschler, 1995). With multiple priors, however, the condition is not sufficient; an additional condition is needed. The following lemma establishes necessary and sufficient conditions to induce the profile of posterior beliefs $(p_i^s)_{i \in I}$ for each signal $s \in S$, starting from the profile of prior beliefs $(p_i)_{i \in I}$.

Lemma 1. *Let S be a finite set and $(p_i^s)_{i \in I}$ be a set of posterior beliefs for each $s \in S$. The following statements are equivalent:*

1. *There exists a signaling function $\pi : \Omega \rightarrow \Delta(S)$ such that $p_i^s = p_i^{s, \pi}$ for all $(i, s) \in I \times S$.*
2. *There exists $(\lambda_i^s)_{i, s}$, with $\lambda_i^s \in [0, 1]$ and $\sum_s \lambda_i^s = 1$ for all $(i, s) \in I \times S$, such that*
 - (i) $(1, \dots, 1) \in \sum_s \lambda_i^s (p_i^s / p_i)$ for all $(i, s) \in I \times S$.
 - (ii) $\lambda_i^s (p_i^s / p_i) = \lambda_j^s (p_j^s / p_j)$ for all $(i, j, s) \in I \times I \times S$.

The proof can be found in Appendix A. Few comments are in order. If I is the singleton $\{i\}$, then (2)(ii) is vacuously satisfied, while (2)(i) simply means that p_i is in the convex hull of $\{p_i^s : s \in S\}$, which is the classic splitting result. More generally, a necessary condition for splitting multiple priors is that each prior p_i is included in the convex hull of the set of posteriors $\{p_i^s : s \in S\}$, for each $i \in I$. However, this is not sufficient. For instance, with two states of the world, condition (2)(ii) of Lemma 1 implies that for all (p_i, p_j) and ω , if $p_i(\omega) \geq p_j(\omega)$, then $p_i^s(\omega) \geq p_j^s(\omega)$ (and the opposite if $p_i(\omega) \leq p_j(\omega)$). Conditions (2)(i) and (2)(ii) are geometric conditions in the space of *likelihood ratios*. Condition (2)(i) states that the unit vector is in the convex hull of the set of likelihood ratios $\{p_i^s / p_i : s \in S\}$ for each i , while condition (2)(ii) states that the likelihood ratios $(p_i^s / p_i)_i$ are collinear for each s .⁷ Lastly, condition (2)(ii) implies that the two $|\Omega|$ -dimensional vectors p_i^s / p_i and p_j^s / p_j have unit cross-ratios, hence the Hilbert distance between these two vectors is zero. In Appendix B, we prove that this implies that the Hilbert distance between p_i and p_j is the same as between p_i^s and p_j^s for every (i, j, s) . We prove similarly that the Hilbert distance between p_i and p_i^s is the same as between p_j and p_j^s for every (i, j, s) . This provides a partial geometric characterization of the splittings in the space of *probabilities* (to not be confused with the space of likelihood ratios). Unfortunately, we have not been able to obtain a complete geometric characterization for the space of probabilities.

⁷If we adopt the weaker assumption mentioned in footnote 2, the required modification is to write (2)(i) as: $p_i = \sum_s \lambda_i^s p_i^s$ for all (i, s) .

Example 1. There are two states of the world ω_1 and ω_2 , two prior beliefs p_1 and p_2 , and two signals s_1 and s_2 . For concreteness, $p_1 = (1/2, 1/2)$ and $p_2 = (1/4, 3/4)$. We claim that there exists a signaling function that induces the posteriors $(p_1^{s_1}, p_2^{s_1}) = ((1/3, 2/3), (1/7, 6/7))$ and $(p_1^{s_2}, p_2^{s_2}) = ((2/3, 1/3), (2/5, 3/5))$. To see this, we apply the second statement of Lemma 1. The likelihood ratios are given by $(p_1^{s_1}/p_1) = (2/3, 4/3)$, $(p_1^{s_2}/p_1) = (4/3, 2/3)$, $(p_2^{s_1}/p_2) = (4/7, 8/7)$, $(p_2^{s_2}/p_2) = (8/5, 4/5)$, and it is easy to verify that condition (2)(i) and (2)(ii) are satisfied with $(\lambda_1^{s_1}, \lambda_1^{s_2}) = (1/2, 1/2)$ and $(\lambda_2^{s_1}, \lambda_2^{s_2}) = (7/12, 5/12)$. See Figure 1 for a geometric illustration (the likelihood ratios are represented by black disks). The likelihood ratios in state s_1 (resp., s_2) are on the same ray, i.e., they are collinear (condition (2)(ii)). Moreover, the unit vector is on the lines connecting the likelihood ratios in states s_1 and s_2 (condition (2)(i)).




Figure 1: The geometry of splittings in the space of likelihood ratios ($i = 1, 2$, $s = s_1, s_2$)

3.2 Characterization of the sender's optimal payoff

The second statement in Lemma 1 has one crucial implication: if we choose a splitting of an arbitrary prior, then conditions (2)(i) and (2)(ii) pin down *all* other posteriors. To see this, choose an arbitrary $i^* \in I$, and let $(\lambda_{i^*}^s, p_{i^*}^s)_s$ be a splitting of p_{i^*} , i.e., $p_{i^*} = \sum_s \lambda_{i^*}^s p_{i^*}^s$ and $\lambda_{i^*}^s \in \Delta(S)$. From conditions (i) and (ii), we have that

$$\lambda_i^s = \lambda_{i^*}^s \|p_i \cdot (p_{i^*}^s / p_{i^*})\|,$$

and

$$p_i^s = \frac{p_i \cdot (p_{i^*}^s / p_{i^*})}{\|p_i \cdot (p_{i^*}^s / p_{i^*})\|}.$$

Consequently, we can view the profile of posteriors $(p_i^s)_{i \in I}$ as the image of a map $f : \Delta(\Omega) \rightarrow \times_{i \in I} \Delta(\Omega)$, where the i -th component is given by

$$f_i(p_{i^*}^s) := \frac{p_i \cdot (p_{i^*}^s / p_{i^*})}{\|p_i \cdot (p_{i^*}^s / p_{i^*})\|},$$

for some fixed but arbitrary $i^* \in I$. In words, $f_i(p_{i^*}^s) \in \Delta(\Omega)$ is the posterior belief of p_i if the signal s is observed and $(\lambda_{i^*}^s, p_{i^*}^s)_s$ is a splitting of p_{i^*} . It is worth stressing that $f_i(p_{i^*}^s)$ only depends on $p_{i^*}^s$ and not on the entire splitting $(\lambda_{i^*}^s, p_{i^*}^s)_s$. As we shall see later, this property is specific to public persuasion; it does not extend to targeted persuasion, where different groups may receive different signals. The choice of the index $i^* \in I$ is immaterial because all prior beliefs have the same support. To ease notation, we write p for the fixed but arbitrary prior p_{i^*} in the sequel.

Let $U_{p,p^*} : \Delta(\Omega) \rightarrow \mathbb{R}$ be the function, parameterized by (p, p^*) , defined by

$$U_{p,p^*}(q) := \sum_{\omega} u(f(q), \omega) \frac{p^*(\omega)}{p(\omega)} q(\omega),$$

for all $q \in \Delta\Omega$. For any function $U : \Delta(\Omega) \rightarrow \mathbb{R}$, we write $\text{cav } U$ for the concavification of U , i.e., the smallest concave function above U . (See Aumann and Maschler, 1967, reprinted in 1995, for details.)

Theorem 1. *The sender's optimal payoff is $\text{cav } U_{p,p^*}(p)$.*

Theorem 1 is proved at the end of this section. As was already mentioned, Theorem 1 implies that the sender can restrict attention to finite sets of signals with at most $|\Omega|$ elements. Indeed, by definition of the concavification, there exists a set S such that $(\text{cav } U_{p,p^*}(p), p) = \sum_{s \in S} \lambda^s (U_{p,q^s}(p^s), p^s)$, with $\sum_{s \in S} \lambda^s = 1$ and $\lambda_s \in [0, 1]$ for all $s \in S$. Since each point $(U_{p,p^*}(p^s), p^s)$ belongs to $\{(q, u) \in \Delta\Omega \times \mathbb{R} : u = U_{p,p^*}(q)\}$, an $|\Omega| - 1$ -dimensional space, the observation follows from Carathéodory theorem.

In many economic applications of interest, there are only two states of the world, e.g., the defendant is guilty or innocent, the drug is better than a placebo or not, the quality of the good is high or low, to name just a few. With two states of the world, the characterization of the optimal payoff is simple. With a slight abuse of notation, for each $i \in I$, let p_i be the probability of the first state. Assume that there exists $i^* \in I$ such that $p_{i^*} \leq p_i$ for all $i \in I$. This assumption is satisfied when $\{p_i : i \in I\}$ is a compact set, as in our two

leading applications. (If this assumption fails, pick any arbitrary prior.) With two states of the world, the i -th component f_i of the function f is given by

$$f_i(q) := \frac{c_i q}{1 + q(c_i - 1)},$$

for all $q \in [0, 1]$, with

$$c_i := \frac{1 - p_{i^*}}{p_{i^*}} \frac{p_i}{1 - p_i}.$$

Note that f_i is strictly increasing, concave and satisfies $f_i(0) = 0$, $f_i(1) = 1$, and $f_i(q) \geq q$ for all $q \in [0, 1]$. The maximization problem (\mathcal{P}) then amounts to choose two posterior beliefs $(p_{i^*}^s, p_{i^*}^{s'})$, with $p_{i^*}^s \leq p_{i^*} \leq p_{i^*}^{s'}$, such that

$$\frac{p_{i^*}^{s'} - p_{i^*}}{p_{i^*}^{s'} - p_{i^*}^s} U_{p_{i^*}, p^*}(p_{i^*}^s) + \frac{p_{i^*} - p_{i^*}^s}{p_{i^*}^{s'} - p_{i^*}^s} U_{p_{i^*}, p^*}(p_{i^*}^{s'})$$

is maximized. (If $p_{i^*}^{s'} = p_{i^*}^s$, the value is $U_{p_{i^*}, p^*}(p_{i^*})$.) We now illustrate our results with the help of a simple example.

Example 2. This example is adapted from Aumann and Hart (2003). There is one sender and one receiver, and two states of the world ω_1 and ω_2 . The receiver has to choose an action in $\{LL, L, M, R, RR\}$ after observing the sender signal s . We assume that the receiver has maxmin preferences and prior-by-prior updating. The profile of prior beliefs of the receiver is $(\frac{1}{3} + i)_{i \in [0, 1/3]}$, while the sender's prior belief is $1/2$ (with a slight abuse of notation, we only refer to the beliefs about state ω_1).⁸ The payoffs are (the first entry in each cell is the sender's payoff):

state ω_1 :	LL	L	M	R	RR
	0, 10	1, 8	0, 5	1, 0	0, -8
state ω_2 :	LL	L	M	R	RR
	0, -8	1, 0	0, 5	1, 8	0, 10

As argued, it is without loss of generality to assume that the sender only uses two signals, which we denote s_1 and s_2 . It is easy to verify that $i^* = 0$, which corresponds to the prior $1/3$. Moreover, if the sender splits the prior $1/3$ into $(p^s, \lambda^s)_{s \in \{s_1, s_2\}}$, the set of posteriors is the closed interval $[f_0(p^s), f_{1/3}(p^s)]$ when the signal is s , $s \in \{s_1, s_2\}$. By construction, f_0 is the identity function, while $f_{1/3}$ is given by

$$f_{1/3}(q) = \frac{4q}{1 + 3q},$$

⁸Equivalently, the set of prior beliefs of the receiver is $[\frac{1}{3}, \frac{2}{3}]$.

for all $q \in [0, 1]$. For instance, if $q = 1/4$, then $f_{1/3}(1/4) = 4/7$.

We now compute the function $q \mapsto U_{1/3,1/2}(q)$. Whenever the set of posteriors is $[q, f_{1/3}(q)]$, the receiver chooses an action $a^*([q, f_{1/3}(q)])$ that maximizes his worst expected payoff (if there are several, choose one that favors the sender), thus giving the sender's payoff in state ω : $u(a^*([q, f_{1/3}(q)], \omega))$. For instance, if the set of posteriors is $[1/4, 4/7]$, the receiver's optimal action is M and the sender's payoff is 0. Indeed, if the receiver chooses M , he guarantees himself a payoff of 5, while if he chooses LL (resp., L , R and RR), the worst payoff is $-7/2$ (resp., 2, $24/7$, and $-2/7$).

Since the sender's payoff is independent of the state of the world, the function $U_{1/3,1/2}$ is then computed as

$$u(a^*([q, f_{1/3}(q)])) \left(\frac{1/2}{1/3}q + \frac{1 - (1/2)}{1 - (1/3)}(1 - q) \right) = \frac{3}{4}u(a^*([q, f_{1/3}(q)])) (1 + q).$$

Figure 2 plots the function $U_{1/3,1/2}$ and its concavification.




Figure 2: The payoff function (bold) and its concavification (dots)

Notice that the concavification is piecewise linear, with the equation of the upper piece given by $(3/4)q + (3/4)$. It follows that the optimal payoff to the sender is 1.

It is important to bear in mind that we *cannot* read from the figure what is the best payoff to the sender if his prior was different, since the function $U_{1/3,1/2}$ is parameterized by it. Nor can we read from the figure the sender's payoff following the splitting. Yet, we can learn about the optimal splittings. For instance, an optimal splitting is $(1/13, 3/4)$, with induced sets of posteriors $[\frac{1}{13}, \frac{1}{4}]$ and $[\frac{3}{4}, \frac{12}{13}]$. Intuitively, if the set of posteriors is $[\frac{1}{13}, \frac{1}{4}]$ (resp., $[\frac{3}{4}, \frac{12}{13}]$), the receiver's optimal action is R (resp., L), which guarantees a payoff of 1 to the sender. Partial revelation of information is therefore optimal.

Example 3. This second example is nearly identical to the previous one. The only differences are the payoff matrices:

$\omega_1:$	<table style="border-collapse: collapse; text-align: center;"> <tr><td style="border: 1px solid black; padding: 2px 10px;">L</td><td style="border: 1px solid black; padding: 2px 10px;">R</td></tr> <tr><td style="border: 1px solid black; padding: 2px 10px;">$1, -1$</td><td style="border: 1px solid black; padding: 2px 10px;">$0, 1$</td></tr> </table>	L	R	$1, -1$	$0, 1$	$\omega_2:$	<table style="border-collapse: collapse; text-align: center;"> <tr><td style="border: 1px solid black; padding: 2px 10px;">L</td><td style="border: 1px solid black; padding: 2px 10px;">R</td></tr> <tr><td style="border: 1px solid black; padding: 2px 10px;">$1, 1$</td><td style="border: 1px solid black; padding: 2px 10px;">$0, -1$</td></tr> </table>	L	R	$1, 1$	$0, -1$
L	R										
$1, -1$	$0, 1$										
L	R										
$1, 1$	$0, -1$										

Let $q^* = \frac{1}{3} \left(\sqrt{\frac{5}{2}} - 1 \right) < \frac{1}{3}$ be the solution to the equation $1 - 2\frac{4q}{1+3q} = 2q - 1$, i.e., q^* is the belief that makes the receiver indifferent between playing L and R . The sender's optimal pure action is L if $q < q^*$ and R , otherwise. (We do not consider mixed actions in this example.)




Figure 3: The payoff function (bold) and its concavification (dots)

The sender's optimal payoff is $\frac{1}{2} \frac{2 + \sqrt{\frac{5}{2}}}{4 - \sqrt{\frac{5}{2}}} < 1$. Moreover, this requires the sender to reveal partial information about the state of the world, since the optimal splitting is $(q^*, 1)$. We can contrast this result with the unambiguous case (unique priors). It is clear that if the agent thinks that the state of the world is ω_1 with probability $1/3$, then the best for the sender is to remain silent. This would guarantee him a payoff of 1. We can easily modify the example to obtain the converse, i.e., no information is released with multiple priors, but some information is released with unique priors.

We conclude this section with the proof of Theorem 1.

Proof. Choose some arbitrary index $i^* \in I$. If the sender commits to π , his expected payoff is:

$$\begin{aligned}
\sum_{s,\omega} u((p_i^{s,\pi})_{i \in I}, \omega) p^*(\omega) \pi(s|\omega) &= \sum_{s,\omega} u((p_i^{s,\pi})_{i \in I}, \omega) \frac{p^*(\omega)}{p_{i^*}(\omega)} p_{i^*}(\omega) \pi(s|\omega) \\
&= \sum_{s,\omega} u((p_i^{s,\pi})_{i \in I}, \omega) \frac{p^*(\omega)}{p_{i^*}(\omega)} p_{i^*}^{s,\pi}(\omega) \left(\sum_{\omega} \pi(s|\omega) p_{i^*}(\omega) \right) \\
&= \sum_s \left(\sum_{\omega} u(f(p_{i^*}^s), \omega) \frac{p^*(\omega)}{p_{i^*}(\omega)} p_{i^*}^s(\omega) \right) \lambda_{i^*}^s \\
&= \sum_s U_{p_{i^*}, p^*}(p_{i^*}^s) \lambda_{i^*}^s,
\end{aligned}$$

with $\lambda_{i^*}^s \in [0, 1]$ for all s , $\sum_s \lambda_{i^*}^s = 1$, and $\sum_s \lambda_{i^*}^s p_{i^*}^s = p_{i^*}$; the second equality follows from the definition of total probabilities and the third equality from Lemma 1. Thus, the maximization problem is equivalent to

$$\max_{(\lambda^s, p^s)_s} \sum_s U_{p, p^*}(p^s) \lambda^s,$$

subject to $\sum \lambda^s p^s = p$, where p is the arbitrarily chosen prior belief p_{i^*} . It is well-known that the solution corresponds to the concavification of U_{p, p^*} evaluated at p . \square

3.3 Beyond public persuasion: targeted persuasion

A central feature of our model is that persuasion is public. While it is a natural assumption in many economic applications (and certainly the only meaningful assumption in sender-receiver games with ambiguity), “targeted” communication best models other applications. For instance, in electoral campaigns, politicians and their teams target different groups of voters with specific messages. Door-to-door canvassing even makes it possible to individually persuade voters. Similarly, social media like Facebook or Twitter make it possible to tailor advertisement campaigns at the individual level. We now explain how our analysis generalizes to targeted persuasion.

Let K be an index set and $\kappa : I \rightarrow K$ an onto map, with the interpretation that $\kappa(i)$ is the target group individual i belongs to.⁹ Individual i in group $\kappa(i)$ receives the private signal $s_{\kappa(i)} \in S_{\kappa(i)}$. We stress that all individuals in group k receive the same signal. Public and individual persuasion are thus two polar cases. Public persuasion corresponds to the case where $\kappa(i) = k$ for all i , while individual persuasion corresponds to the case, where $K = I$

⁹An alternative interpretation is that $\kappa^{-1}(k)$ indexes the set of prior beliefs of individual k , in a world with maxmin preferences.

and κ is the identity mapping. Throughout, all sets of signals are finite.

As before, prior to learning the state of the world, the sender commits to a signaling function $\pi : \Omega \rightarrow \Delta(S)$, where $S := \times_{k \in K} S_k$. Denote $IS := \{(i, s_k) \in I \times S_k : k = \kappa(i)\}$. We assume that the sender's payoff is a function u of the profile of posteriors $(p_i^{s_{\kappa(i)}})_{i \in I}$, with $p_i^{s_{\kappa(i)}} \in \Delta(\Omega \times S_{-\kappa(i)})$ for all $(i, s_{\kappa(i)}) \in IS$, and the state of the world ω . Indeed, in the (unmodeled) game that follows the release of information, a strategy for individual i specifies an action for each private signal $s_{k(i)}$ he might receive, and individual i 's payoff depends on his beliefs about ω and $s_{-\kappa(i)}$ (through the strategies of others). This leads us to characterize the distribution of posteriors over states of the world and private signals of others the sender can achieve with targeted persuasion.

Lemma 2. *Let S be a finite set and $(p_i^{s_{\kappa(i)}})_{i \in I}$ a profile of posterior beliefs for each $s \in S$ ($p_i^{s_{\kappa(i)}} \in \Delta(\Omega \times S_{-\kappa(i)})$ for each $(i, s_{\kappa(i)}) \in IS$). The following statements are equivalent:*

1. *There exists a signaling function $\pi : \Omega \rightarrow \Delta(S)$ such that $p_i^{s_{\kappa(i)}} = p_i^{s_{\kappa(i)}, \pi}$ for all $(i, s_{\kappa(i)}) \in IS$.*
2. *There exists $(\lambda_i^{s_{\kappa(i)}})_{i, s_{\kappa(i)}}$, with $\lambda_i^{s_{\kappa(i)}} \in [0, 1]$ and $\sum_{s_{\kappa(i)} \in S_{\kappa(i)}} \lambda_i^{s_{\kappa(i)}} = 1$ for all $(i, s_{\kappa(i)}) \in IS$, such that*

- (i) $(1, \dots, 1) \in \sum_{(s_{\kappa(i)}, s_{-\kappa(i)})} \lambda_i^{s_{\kappa(i)}} (p_i^{s_{\kappa(i)}}(\cdot, s_{-\kappa(i)})/p_i)$ for all $i \in I$.
- (ii) $\lambda_i^{s_{\kappa(i)}} (p_i^{s_{\kappa(i)}}(\cdot, s_{-\kappa(i)})/p_i) = \lambda_j^{s_{\kappa(j)}} (p_j^{s_{\kappa(j)}}(\cdot, s_{-\kappa(j)})/p_j)$ for all $(i, j, s) \in I \times I \times S$.

The proof of Lemma 2 can be found in Appendix C. Lemma 2 is a direct generalization of Lemma 1. In particular, if we choose an arbitrary index i^* and a splitting $(p_{i^*}^{s_{\kappa(i^*)}}, \lambda_{i^*}^{s_{\kappa(i^*)}})_{s_{\kappa(i^*)}}$ of p_{i^*} , then the posterior $p_i^{s_{\kappa(i)}}$ of p_i conditional on $s_{\kappa(i)}$ is uniquely determined, with:

$$p_i^{s_{\kappa(i)}}(\omega, s_{-\kappa(i)}) = \frac{\lambda_{i^*}^{s_{\kappa(i^*)}} (p_{i^*}^{s_{\kappa(i^*)}}(\omega, s_{-\kappa(i^*)})/p_{i^*}(\omega)) p_i(\omega)}{\sum_{(\omega, s_{-\kappa(i)})} \lambda_{i^*}^{s_{\kappa(i^*)}} (p_{i^*}^{s_{\kappa(i^*)}}(\omega, s_{-\kappa(i^*)})/p_{i^*}(\omega)) p_i(\omega)},$$

for all $(\omega, s_{-\kappa(i)})$, if the denominator is positive. It follows that $p_i^{s_{\kappa(i)}}$ is a function $f_i^{s_{\kappa(i)}}$ of the splitting $(p_{i^*}^{s_{\kappa(i^*)}}, \lambda_{i^*}^{s_{\kappa(i^*)}})_{s_{\kappa(i^*)}}$. It is worth stressing one important difference with public persuasion. In general, the posterior $p_i^{s_{\kappa(i)}}$ of p_i depends on the entire splitting of p_{i^*} , i.e., on the chosen distribution over posteriors. When $i \in \kappa^{-1}(\kappa(i^*))$, however, the posterior $p_i^{s_{\kappa(i)}}$ of p_i is only a function of $p_{i^*}^{s_{\kappa(i^*)}}$. This case is in fact the only relevant case with public persuasion, explaining the characterization we have obtained in Lemma 1.

Following the exact same steps as in the proof of Theorem 1, we can express the sender's

expected payoff as:

$$\sum_{s_{\kappa(i^*)}} \left(\underbrace{\sum_{(\omega, s_{-\kappa(i^*)})} u((f_i^{s_{\kappa(i^*)}}((p_{i^*}^{s_{\kappa(i^*)}}, \lambda_{i^*}^{s_{\kappa(i^*)}))_{s_{\kappa(i^*)})}_{i \in I}, \omega) p_{i^*}^{s_{\kappa(i^*)}}(\omega, s_{-\kappa(i^*)}) \frac{p^*(\omega)}{p_{i^*}^*(\omega)})}_{=: U_{p_{i^*}, p^*}((p_{i^*}^{s_{\kappa(i^*)}}, \lambda_{i^*}^{s_{\kappa(i^*)}))_{s_{\kappa(i^*)})} \right) \lambda_{i^*}^{s_{\kappa(i^*)}}.$$

Theorem 2. *The sender's optimal payoff is the value of the maximization problem:*

$$\max_{(p_{i^*}^{s_{\kappa(i^*)}}, \lambda_{i^*}^{s_{\kappa(i^*)}))_{s_{\kappa(i^*)}}} \sum_{s_i} U_{p_{i^*}, p^*}((p_{i^*}^{s_{\kappa(i^*)}}, \lambda_{i^*}^{s_{\kappa(i^*)}))_{s_{\kappa(i^*)})} \lambda_{i^*}^{s_{\kappa(i^*)}},$$

subject to $\sum_{(s_{\kappa(i^*)}, s_{-\kappa(i^*)})} \lambda_{i^*}^{s_{\kappa(i^*)}} p_{i^*}^{s_{\kappa(i^*)}}(\cdot, s_{-\kappa(i^*)}) = p_{i^*}$, if it exists.

Few remarks are in order. First, unlike the case of public persuasion, the solution does not correspond to the concavification of the function $U_{p^*, p_{i^*}}$, as the function itself depends on the choice of splittings. Second, if $S_k = S_{k'}$ for all (k, k') , then public persuasion is admissible and, thus, the sender is always at least better off by targeting his persuasion.¹⁰ Third, if a solution to the maximization problem is such that $\lambda_{i^*}^{s_{\kappa(i^*)}} = 1$ and $p_{i^*}^{s_{\kappa(i^*)}}(\cdot, s_{-\kappa(i^*)}^*) = p_{i^*}$ for some $(s_{\kappa(i^*)}^*, s_{-\kappa(i^*)}^*)$, then the maximum is also achievable with public persuasion. This is indeed trivial since the later solution corresponds to no information being released, which is always possible. Unfortunately, we have not been able to obtain sharper results, without imposing additional structure on the problem.

4 An economic application: persuading a committee

An expert would like to persuade a committee, composed of an odd number n of members, to adopt a project. The project is adopted if a majority of members votes for its adoption. The project is either profitable (state ω_1) or unprofitable (state ω_2). If the project is profitable (resp., unprofitable), a committee member derives a utility of one from adopting (resp., rejecting) the project and of zero from rejecting (resp., adopting) it. Hence, a committee member would like to adopt the project if and only if it is profitable.

The expert derives utility from the decision chosen by the committee as well as from his *reputation*. Regardless of the project's profitability, the expert gets a payoff of one if his project is adopted and of zero, otherwise.

Intuitively, committee members expect a reputable expert to not surprise them. After

¹⁰Public persuasion corresponds to $p_{i^*}^{s_{\kappa(i^*)}}(\omega, s_{-\kappa(i^*)}) = 0$ whenever $s_{-\kappa(i^*)}$ differs from the profile $(s_{\kappa(i^*)})_{k \neq \kappa(i^*)}$.

all, the expert can commit to fully disclose the state of the world. Formally, suppose that committee member i is led to believe that the project is profitable with probability p_i^s . The quantity $-\log(p_i^s)$, known as the self-information or surprisal of the probability $(p_i^s, 1 - p_i^s)$ in state ω_1 , captures the “surprise” in observing a positive profit when one believes that this occurs with probability p_i^s . Put it differently, the self-information of $(p_i^s, 1 - p_i^s)$ in state ω_1 is equivalent to the Kullback-Leiber divergence of the probability $(p_i^s, 1 - p_i^s)$ from the probability $(1, 0)$, i.e., the probability that assigns probability one to the project being profitable (which would be the member’s posterior if the expert had committed to fully disclose the state). We thus model the reputation cost in state ω as the average surprisal, i.e., as

$$\delta \frac{1}{n} \left(\sum_i -\log(p_i^s(\omega)) \right),$$

when $(p_i^s)_i$ is the profile of beliefs of the committee members. The parameter $\delta \geq 0$ captures the disutility of having a bad reputation. In other words, this simple model captures the trade-off between current reward and future reward (through an expert’s reputation) an expert faces in recommending decisions. As a concrete illustration, pharmaceutical companies design clinical trials with the aim at getting their drugs approved by the competent agencies, e.g., the Food and Drug Administration in the US or the European Medicines Agency. In circumstances where the design of clinical trials was fraught and led to the approval of ineffective or even unsafe drugs, pharmaceutical companies are not only sanctioned with fines, but also suffers reputation’s losses, which lead to lower approval’s probabilities in the future.

The prior beliefs of the committee members are $(p_i)_i$, while the expert’s prior belief is p^* . Without loss of generality, we assume that $p_1 \leq p_2 \leq \dots \leq p_n$. Focusing on equilibria in weakly dominant strategies, the project is adopted whenever $p_{\lceil n/2 \rceil}^s(\omega_1) \geq 1/2$, i.e., whenever the median member prefers to adopt the project. (If indifferent between accepting and rejecting the project, a member accepts.)

It follows that the expert’s utility $u((p_i^s)_i, \omega)$ is

$$\mathbb{1}_{\{p_{\lceil n/2 \rceil}^s(\omega_1) \geq 1/2\}} - \delta \frac{1}{n} \left(\sum_i -\log(p_i^s(\omega)) \right).$$

Note that when $\delta = 0$, we have a model of pure persuasion: the expert wants the project to be adopted, regardless of the state.

To simplify the exposition, we choose to split the beliefs of the median member, i.e., we choose the index $i^* = \lceil n/2 \rceil$ in calculating $U_{p_{i^*}, p^*}$. If the median member has posterior belief

q , the expert's payoff function is:

$$U_{p_{\lceil n/2 \rceil}, p^*}(q) = \left(\frac{p^*}{p_{\lceil n/2 \rceil}} q + \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} (1-q) \right) \mathbb{1}_{\{q \geq 1/2\}} + \delta \frac{1}{n} \sum_i \left(\log \left(\frac{c_i q}{1 + (c_i - 1)q} \right) \frac{p^*}{p_{\lceil n/2 \rceil}} q + \log \left(1 - \frac{c_i q}{1 + (c_i - 1)q} \right) \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} (1-q) \right),$$

with

$$c_i = \frac{1-p_{\lceil n/2 \rceil}}{p_{\lceil n/2 \rceil}} \frac{p_i}{1-p_i}.$$

Note that the first term corresponds to the payoff the expert derives from the adoption of the project (which occurs if the median member believes that the project is profitable with probability at least one-half). The second term corresponds to the reputation cost (more accurately, it is the negative of the reputation cost). It is always negative, takes value zero at $q = 0$ or $q = 1$, i.e., when the expert designs a fully informative signaling function, and crucially convex in q .

Proposition 1. *If $\frac{p^*}{p_{\lceil n/2 \rceil}} \frac{1}{2} \geq U_{p_{\lceil n/2 \rceil}, p^*}(1/2)$, the concavification of $U_{p_{\lceil n/2 \rceil}, p^*}$ is given by*

$$\frac{p^*}{p_{\lceil n/2 \rceil}} q,$$

for all $q \in [0, 1]$. If $\frac{p^}{p_{\lceil n/2 \rceil}} \frac{1}{2} < U_{p_{\lceil n/2 \rceil}, p^*}(1/2)$, the concavification of $U_{p_{\lceil n/2 \rceil}, p^*}$ is given by*

$$\begin{cases} 2U_{p_{\lceil n/2 \rceil}, p^*}(1/2)q & \text{if } q \leq 1/2 \\ 2 \left[U_{p_{\lceil n/2 \rceil}, p^*}(1) - U_{p_{\lceil n/2 \rceil}, p^*}(1/2) \right] q + 2U_{p_{\lceil n/2 \rceil}, p^*}(1/2) - U_{p_{\lceil n/2 \rceil}, p^*}(1) & \text{if } q > 1/2. \end{cases}$$

The formal proof of Proposition 1 is relegated to Appendix D. Yet, the intuition is relatively clear. To start with, suppose that the reputation cost is not too large (i.e., δ close to zero) so that $\frac{p^*}{p_{\lceil n/2 \rceil}} \frac{1}{2} < U_{p_{\lceil n/2 \rceil}, p^*}(1/2)$. When the median member believes the project to be profitable with probability at least one-half, the expert can secure the adoption of the project by remaining silent. However, this wouldn't minimize the reputation cost (unless $\delta = 0$). In order to minimize the reputation cost while securing the adoption of the project, the best the expert can do is to disclose that the project is profitable only when it is and to make the median member indifferent, otherwise.¹¹ In other words, the expert generates the two most extreme beliefs 1 and 1/2 that guarantee the project's approval, and this is the best he can do by convexity of the reputation cost. A similar reasoning applies when

¹¹More formally, this corresponds to the splitting of $p_{\lceil n/2 \rceil} > 1/2$ into 1 and 1/2. The signaling function λ , with $\lambda(s_1|\omega_1) = 2 - (1/p_{\lceil n/2 \rceil})$ and $\lambda(s_1|\omega_2) = 0$, achieves that splitting.

the median member believes the project to be profitable with probability less than one-half. Alternatively, when the reputation cost is sufficiently large (i.e., δ large enough), the best the expert can do is to fully disclose whether the project is profitable, as the expected cost of surprising the committee is disproportionately larger than the gain of having the project adopted.

Finally, note that

$$\begin{aligned} & \frac{p^*}{p_{\lceil n/2 \rceil}} - 2U_{p_{\lceil n/2 \rceil}, p^*}(1/2) = \\ & -\frac{p^*}{p_{\lceil n/2 \rceil}} \delta \sum_i \frac{1}{n} \log \left(\frac{c_i}{1+c_i} \right) - \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} \left(1 + \delta \sum_i \frac{1}{n} \log \left(\frac{1}{1+c_i} \right) \right), \end{aligned}$$

and is thus continuous in all parameters of the model. It follows that if we start from a situation where $\frac{p^*}{p_{\lceil n/2 \rceil}} - 2U_{p_{\lceil n/2 \rceil}, p^*}(1/2) > 0$, local changes in the parameters won't affect the optimal information disclosure: the best would remain to fully disclose the state of the world. Of course, the expert's payoff will change so as to reflect the change in the beliefs. More precisely, the expert's payoff will change from p^* to \bar{p}^* , if the expert's belief has changed from p^* to \bar{p}^* . A similar argument applies if $\frac{p^*}{p_{\lceil n/2 \rceil}} - 2U_{p_{\lceil n/2 \rceil}, p^*}(1/2) < 0$. The optimal information disclosure is therefore robust to the choice of parameters (at least generically).

Appendix

A Proof of Lemma 1

(1) \Rightarrow (2). Let $\lambda_i^s = \sum_{\omega} \pi(s|\omega) p_i(\omega)$. It immediately follows that $\sum_s \lambda_i^s p_i^s = p_i$, i.e., (i) holds. Moreover, for all $\omega \in \Omega$, we have

$$\lambda_i^s (p_i^s(\omega)/p_i(\omega)) = \pi(s|\omega) = \lambda_j^s (p_j^s(\omega)/p_j(\omega)),$$

so that (ii) holds. Note that if there exists i such that $\lambda_i^s = 0$, then $\pi(s|\omega) = 0$ for all ω . It follows that $\lambda_j^s = 0$ for all j , since p_i and p_j have full support.

(2) \Rightarrow (1). Choose any $i^* \in I$ and let

$$\pi(s|\omega) := \frac{\lambda_{i^*}^s p_{i^*}^s(\omega)}{p_{i^*}(\omega)}.$$

This is well-defined since $p_{i^*}(\omega) > 0$. Moreover, since (i) is satisfied, we have that $\sum_s \pi(s|\omega) =$

1. The posterior of p_i , conditional on s , is given by

$$\frac{p_i(\omega)\lambda_{i^*}^s(p_{i^*}^s(\omega)/p_{i^*}(\omega))}{\sum_{\omega} p_i(\omega)\lambda_{i^*}^s(p_{i^*}^s(\omega)/p_{i^*}(\omega))} = \frac{p_i(\omega)\lambda_i^s(p_i^s(\omega)/p_i(\omega))}{\sum_{\omega} p_i(\omega)\lambda_i^s(p_i^s(\omega)/p_i(\omega))} = p_i^s(\omega),$$

where the first equality follows from (ii).

B Hilbert geometry

We introduce the Hilbert distance which helps in understanding the geometric of Lemma 1 on the simplex.

Definition 1. Hilbert distance.

We first define the Hilbert distance between two vectors $x \in \mathbb{R}_{++}^K$ and $y \in \mathbb{R}_{++}^K$.¹² Let

$$M(x/y) = \begin{cases} \inf\{\pi \in \mathbb{R}_{++} : x_k \leq \pi y_k \ \forall k \in K\} & \text{if possible,} \\ +\infty & \text{otherwise,} \end{cases}$$

and

$$m(x/y) = \sup\{\mu \in \mathbb{R}_{++} : \mu y_k \leq x_k \ \forall k \in K\}.$$

The Hilbert distance $d_H(x, y)$ between x and y is then as:

$$d_H(x, y) = \log \frac{M(x/y)}{m(x/y)}.$$

There is an alternative definition of the Hilbert distance on Euclidean spaces.¹³ Notice that:

$$\begin{aligned} M(x/y) &= \begin{cases} \inf\{\pi \in \mathbb{R}_{++} : \frac{x_k}{y_k} \leq \pi \ \forall k \in K\} & \text{if possible,} \\ +\infty & \text{otherwise,} \end{cases} \\ &= \max_{k \in K} \left\{ \frac{x_k}{y_k} \right\}. \end{aligned}$$

Similarly, $m(x/y) = \sup\{\mu \in \mathbb{R}_{++} : \mu \leq \frac{x_k}{y_k} \ \forall k \in K\} = \min_{k \in K} \left\{ \frac{x_k}{y_k} \right\}$.

¹²The Hilbert distance can be defined more generally on any closed solid cone of a real Banach space (see for instance Bushell, 1973).

¹³This is not true for the Hilbert distance on more general sets (see Bushell, 1973).

It follows that the Hilbert distance between two probabilities p_i and p_j in $\Delta(\Omega)$:

$$\begin{aligned} d_H(p_i, p_j) &= \log \frac{\max_{\omega \in \Omega} \left\{ \frac{p_i(\omega)}{p_j(\omega)} \right\}}{\min_{\tilde{\omega} \in \Omega} \left\{ \frac{p_i(\tilde{\omega})}{p_j(\tilde{\omega})} \right\}} = \log \left[\max_{\omega \in \Omega} \left\{ \frac{p_i(\omega)}{p_j(\omega)} \right\} \times \max_{\tilde{\omega} \in \Omega} \left\{ \frac{1}{\frac{p_i(\tilde{\omega})}{p_j(\tilde{\omega})}} \right\} \right] \\ &= \log \max_{(\omega, \tilde{\omega}) \in \Omega \times \Omega} \left\{ \frac{p_i(\omega) p_j(\tilde{\omega})}{p_j(\omega) p_i(\tilde{\omega})} \right\}. \end{aligned}$$

Condition (2)(iii) of Lemma 1 then implies:

$$d_H(p_i^s/p_i, p_j^s/p_j) = \log \max_{(\omega, \tilde{\omega}) \in \Omega \times \Omega} \left[\frac{\frac{p_i^s(\omega)}{p_i(\omega)} \frac{p_j^s(\tilde{\omega})}{p_j(\tilde{\omega})}}{\frac{p_j^s(\omega)}{p_j(\omega)} \frac{p_i^s(\tilde{\omega})}{p_i(\tilde{\omega})}} \right] = \log 1 = 0.$$

More generally, it is easy to see that two vectors are collinear if and only if their Hilbert distance is zero. Because of this property, the Hilbert distance is sometimes referred as a pseudo-distance, since the statement $d_H(x, y) = 0 \Rightarrow x = y$ is not true, but we have instead $d_H(x, y) = 0 \Rightarrow x = \pi y$ for some $\pi > 0$.

The following lemma shows that a necessary condition for the splitting in Lemma 1 is that for every two priors p_i and p_j , the Hilbert distances between them and their respective posteriors p_i^s and p_j^s must be the same.

Lemma 3. *For every (i, j) , if $d_H(p_i^s/p_i, p_j^s/p_j) = 0$, then $d_H(p_i, p_i^s) = d_H(p_j, p_j^s)$.*

Proof. Since the two vectors p_i^s/p_i and p_j^s/p_j are collinear (with a coefficient $\alpha > 0$), we have

$$\begin{aligned} d_H(p_i, p_i^s) &= \log \max_{(\omega, \tilde{\omega}) \in \Omega \times \Omega} \left\{ \frac{p_i(\omega) p_i^s(\tilde{\omega})}{p_i^s(\omega) p_i(\tilde{\omega})} \right\} \\ &= \log \max_{(\omega, \tilde{\omega}) \in \Omega \times \Omega} \left\{ \frac{1}{\alpha} \frac{p_j(\omega)}{p_j^s(\omega)} \frac{p_j(\tilde{\omega})}{p_j^s(\tilde{\omega})} \right\} \\ &= \log \max_{(\omega, \tilde{\omega}) \in \Omega \times \Omega} \left\{ \frac{p_j(\omega) p_j^s(\tilde{\omega})}{p_j^s(\omega) p_j(\tilde{\omega})} \right\} \\ &= d_H(p_j, p_j^s). \end{aligned}$$

□

Similarly, we have the following proposition.

Lemma 4. *For every (i, j) , if $d_H(p_i^s/p_i, p_j^s/p_j) = 0$, then $d_H(p_i, p_j) = d_H(p_i^s, p_j^s)$.*

Therefore, as stated in the main text, it must be that the Hilbert distance between any two priors p_i and p_j is the same as the Hilbert distance between their respective posteriors p_i^s

and p_j^s for all s . Similarly, the Hilbert distance between any prior p_i and its posterior p_i^s must be the same as the Hilbert distance between another prior p_j and its posterior p_j^s . While those conditions are necessary, we have not been able to find a complete geometric characterization on the simplex. Yet, Lemma 1 provides a complete geometric characterization on the space of likelihood ratios.

C Proof of Lemma 2

(1) \Rightarrow (2). Let $\lambda_i^{s_{\kappa(i)}} = \sum_{\omega} \pi(s_{\kappa(i)}, s_{-\kappa(i)} | \omega) p_i(\omega)$. We have that

$$p_i^{s_{\kappa(i)}}(\omega, s_{-\kappa(i)}) = \frac{\pi(s_{\kappa(i)}, s_{-\kappa(i)} | \omega) p_i(\omega)}{\sum_{(\omega, s_{-\kappa(i)})} \pi(s_{\kappa(i)}, s_{-\kappa(i)} | \omega) p_i(\omega)},$$

for all (ω, s) . It immediately follows that (i) holds. Moreover, for all $(\omega, s) \in \Omega \times S$, we have

$$\lambda_i^{s_{\kappa(i)}} \frac{p_i^{s_{\kappa(i)}}(\omega, s_{-\kappa(i)})}{p_i(\omega)} = \pi(s_{\kappa(i)}, s_{-\kappa(i)} | \omega) = \lambda_j^{s_{\kappa(j)}} \frac{p_j^{s_{\kappa(j)}}(\omega, s_{-\kappa(j)})}{p_j(\omega)},$$

so that (ii) holds.

(2) \Rightarrow (1). Choose any $i^* \in I$ and let

$$\pi(s | \omega) := \lambda_{i^*}^{s_{\kappa(i^*)}} \frac{p_{i^*}^{s_{\kappa(i^*)}}(\omega, s_{-\kappa(i^*)})}{p_{i^*}(\omega)}.$$

This is well-defined since $p_{i^*}(\omega) > 0$. Moreover, since (i) is satisfied, we have that $\sum_s \pi(s | \omega) = 1$. It is routine to verify that the posterior of p_i over $\Omega \times S_{-\kappa(i)}$ conditional on $s_{\kappa(i)}$ is indeed $p_i^{s_{\kappa(i)}}$.

D Proof of Proposition 1

We first show that the function

$$q \mapsto \frac{1}{n} \left(\sum_i \log \left(\frac{c_i q}{1 + (c_i - 1)q} \right) \frac{p^*}{p_{\lceil n/2 \rceil}} q + \log \left(1 - \frac{c_i q}{1 + (c_i - 1)q} \right) \frac{1 - p^*}{1 - p_{\lceil n/2 \rceil}} (1 - q) \right)$$

is convex. The function is twice-differentiable with second-order derivative given by

$$\begin{aligned} & \frac{p^*}{p_{\lceil n/2 \rceil}} \frac{1}{q} + \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} \frac{1}{1-q} \\ & - 2 \frac{1}{n} \sum_i \frac{c_i - 1}{1 + (c_i - 1)q} \left(\frac{p^*}{p_{\lceil n/2 \rceil}} - \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} \right) + \\ & + \frac{1}{n} \sum_i \left(\frac{c_i - 1}{1 + (c_i - 1)q} \right)^2 \left(\frac{p^*}{p_{\lceil n/2 \rceil}} q + \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} (1-q) \right). \end{aligned}$$

Since $c_i < 1$ for all $i < \lceil n/2 \rceil$, $c_i > 1$ for all $i > \lceil n/2 \rceil$, and

$$\frac{c_i - 1}{1 + (c_i - 1)q} < \frac{1}{q},$$

the above expression is bounded from below by

$$\begin{aligned} & \frac{p^*}{p_{\lceil n/2 \rceil}} \frac{1}{q} + \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} \frac{1}{1-q} \\ & - 2 \frac{1}{n} \frac{n-1}{2} \frac{1}{q} \left(\frac{p^*}{p_{\lceil n/2 \rceil}} - \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} \right) + \\ & + \frac{1}{n} \sum_i \left(\frac{c_i - 1}{1 + (c_i - 1)q} \right)^2 \left(\frac{p^*}{p_{\lceil n/2 \rceil}} q + \frac{1-p^*}{1-p_{\lceil n/2 \rceil}} (1-q) \right), \end{aligned}$$

which is strictly positive.

Proposition 1 then follows directly from the observations that $U_{p_{\lceil n/2 \rceil}, p^*}(0) = 0$ and $U_{p_{\lceil n/2 \rceil}, p^*}(q) \leq 0$ for all $q < 1/2$, $U_{p_{\lceil n/2 \rceil}, p^*}(1) = p^*/p_{\lceil n/2 \rceil} > 0$, and the convexity of the restriction of $q \mapsto U_{p_{\lceil n/2 \rceil}, p^*}(q)$ to $[1/2, 1]$ (which implies that the segment between the points $(1/2, U_{p_{\lceil n/2 \rceil}, p^*}(1/2))$ and $(1, p^*/p_{\lceil n/2 \rceil})$ is above the graph of the restriction of $q \mapsto U_{p_{\lceil n/2 \rceil}, p^*}(q)$ to $[1/2, 1]$).

References

- [1] ALONSO, R., AND CÂMARA, O. Bayesian persuasion with heterogeneous beliefs. *Working Paper* (2015a).
- [2] ALONSO, R., AND CÂMARA, O. Persuading voters. *Working Paper* (2015b).
- [3] AUMANN, R., MASCHLER, M., AND STEARNS, R. *Repeated Games with Incomplete Information*. The MIT press, 1995.
- [4] AUMANN, R. J., AND HART, S. Long cheap talk. *Econometrica* 71, 6 (2003), 1619–1660.

- [5] GILBOA, I., AND SCHMEIDLER, D. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18, 2 (1989), 141–153.
- [6] KAMENICA, E., AND GENTZKOW, M. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.