



Persuasion without Priors

Alexei Parakhonyak¹ · Anton Sobolev²

Received: 18 September 2024 / Accepted: 24 June 2025

© The Author(s) 2025

Abstract

We consider an information design problem in which the sender faces ambiguity regarding the probability distribution over the states of the world, the utility function, and the prior of the receiver. The solution concept is minimax loss (regret), that is, the sender minimizes the difference in payoffs from the Bayesian benchmark in the worst-case scenario. In the binary state and binary action setting the mechanism contains a continuum of messages, and admits a representation as a randomization over two-message mechanisms. A small level of uncertainty makes the sender more truthful, but larger uncertainty may result in the sender lying more often than the Bayesian benchmark. If the sender either knows the probability distribution over the states of the world, or knows that the receiver knows it, then the maximal loss is bounded from above by $1/e$. When admissible mechanisms are limited to cut-off strategies, this result generalizes to a multiple state model.

Keywords Persuasion · Robustness · Multiple priors · Minimax regret

JEL Classification D81 · D82 · D83

1 Introduction

The standard Bayesian persuasion framework, pioneered by Kamenica and Gentzkow (2011), imposes strong informational assumptions on the sender. In order to design the optimal mechanism, the sender must know the probability distribution over the states of the world, the receiver's prior, and the receiver's utility function. Arguably, such assumptions are questionable in many practical settings. For example, an auto seller who designs a test-drive procedure might not know the buyer's prior thoughts

✉ Alexei Parakhonyak
aleksei.parakhonyak@economics.ox.ac.uk

Anton Sobolev
anton.sobolev@uni-mannheim.de

¹ Department of Economics and Lincoln College, University of Oxford, OX1 3UQ Oxford, United Kingdom

² Department of Economics and MaCCI, University of Mannheim, 68131 Mannheim, Germany

about a car or whether the buyer has other offers at hand. A prosecutor might not be aware of the judge's prior belief about the probability of the defendant being guilty, or of the judge's preferences for acquitting a guilty person versus sentencing an innocent one.

The optimal persuasion rule relies on precise information about the environment, and thus may be rather fragile. That is, a marginal overestimation of the receiver's propensity to take the action desired by the sender may result in choosing a mechanism which fails to persuade the receiver in any state of the world. In this paper, we study a robust sender's decision rule that performs relatively well in each possible environment or under any prior distribution over these environments. Our approach is valid when the sender faces a population of receivers with an unknown distribution of types. It can also be viewed as a compromise among senders with different priors about the receiver type who have to agree on a single persuasion mechanism. Such a compromise would leave all the sender types minimally regretting their choice.

More specifically, we consider a setting in which the sender potentially faces uncertainty about all three inputs of the optimal persuasion mechanism—the distribution of the states, the receiver's prior and utility. Moreover, the sender has no prior assumptions about the distribution of these inputs. Such an uninformed sender evaluates the persuasion mechanism based on *loss* or regret, which is the difference between the payoff from the current persuasion mechanism and what the Bayesian sender obtains. We use the term 'Bayesian sender' for a sender who knows all the parameters of the model but not the realization of the state of the world. The sender is interested in minimizing loss in the *worst-case scenario*, and by doing so achieves relatively good performance in any possible actual environment.

The receiver in our game is a standard Bayesian. Specifically, the receiver observes the choice of the mechanism used by the sender, the realization of the signal produced by this mechanism, and combines this information with her prior to choose the optimal action.

The main part of our analysis focuses on the binary state (good or bad) and binary action (adopt or reject) problem. There is a clear ranking of the states, but the sender does not know the probability of the good state, the receiver's prior, and the outside option. We show that the receiver's parameters can be combined into a single integral characteristic—the receiver's optimism. We solve for the optimal persuasion mechanism by representing the original problem as a zero-sum game between the sender and adversarial nature. The sender chooses a persuasion strategy and wants to minimize loss, while nature chooses the probability of the good state and the receiver's optimism and tries to maximize loss. As is typical for zero-sum games, our game has an equilibrium in mixed strategies only.¹

The equilibrium mixed strategy of the sender in our zero-sum game is equivalent to playing the optimal mechanism, which involves infinitely many messages; one message fully revealing that the state is bad, and each of the other messages persuading a receiver with optimism above a certain threshold. Any such message is equivalent to the choice of a single two-message mechanism, designed to persuade a specific

¹ See Duersch et al. (2012) for conditions for existence of pure strategy equilibria in symmetric zero-sum games.

receiver, from the support of the sender's mixed strategy. Thus, the robust persuasion mechanism involves considerably richer message space than the optimal Bayesian mechanism, which includes only two messages.

The strategy of adversarial nature involves a negative correlation between the receiver's optimism and the true probability of being in the good state. That is, the worst-case scenarios are those when having a high probability of being in the good state is coupled with a receiver who is hard to persuade, either because of a low prior or a high outside option, and vice versa. Intuitively, a pessimistic receiver is less likely to be persuaded by a given mechanism, which is particularly bad when the state is likely to be good and a (more) truthful mechanism would often result in adoption.

The minimax loss is increasing in uncertainty over parameters and, with our normalization of the sender's utility to an interval between 0 and 1, is bounded from above by the omega constant.² When the sender faces no uncertainty over the receiver's optimism, the minimax loss equals zero: if the sender knows the receiver's optimism, he can maximize the probability of persuasion in the bad state, even though the probability of this state may be unknown to the sender. When the sender faces maximal uncertainty over the receiver's optimism but knows the probability of the good state, the maximal loss is bounded above by $1/e$ but can be lower if the good state is relatively likely. These two results highlight the fact that information asymmetry between the sender and the receiver plays a crucial role. Interestingly, if the receiver's prior coincides with the true probability distribution over the states of the world, then the loss of the sender is also bounded above by $1/e$. Therefore, the maximal loss of the sender is the same irrespective of whether he knows the probability distribution over the states or knows that the receiver knows it. The reason is that in both cases nature cannot negatively correlate the receiver's optimism with the probability of the state being good, limiting the scope for inflicting losses on the sender.

We also investigate how the probability of lying (i.e., not sending a revealing message in the bad state) depends on uncertainty about the receiver's optimism. We show that initially, this probability decreases with uncertainty, i.e., when uncertainty is small, the sender employs a more truthful strategy than the one used by the Bayesian sender. That is, from the perspective of the receiver, a small level of uncertainty improves the quality of communication. However, as the level of uncertainty increases, at some point the sender starts lying more, potentially becoming less truthful than the Bayesian sender.

Finally, we extend our analysis to multiple states of the world. We assume that the receiver's utility is monotone and that the only uncertainty the sender faces is about the receiver's prior and her outside option. We show that the sender's loss approaches its maximal possible value of one as the number of states approaches infinity. However, this negative result relies on a peculiar optimal mechanism in which the Bayesian sender would find it optimal to exclude specific states, randomly chosen by adversarial nature, from the support of adoption recommendation. If the sender (as well as the Bayesian benchmark) is restricted to cut-off strategies, i.e., when the adoption recommendation is sent in all states above a certain threshold, then the optimal

² The omega constant is a solution to $\Omega e^{\Omega} = 1$, see Mezo (2022) for details.

mechanism is equivalent to the one we derived in the binary state model, resulting in similar strategies and identical losses.

Literature review

Our paper extends robust decision-making techniques to the Bayesian persuasion problem. We employ the minimax regret or minimax loss approach, which ensures that the sender remains close to the optimum irrespective of the prior. An alternative interpretation is that the sender faces a receiver drawn from an unknown distribution and aims to perform near the optimal sender who observes the receiver's type under the most disadvantageous distribution. This approach was axiomatized by Stoye (2011) and applied to monopoly pricing (Bergemann and Schlag 2008), search (Bergemann and Schlag 2011; Parakhonyak and Sobolev 2015; Schlag and Zapechelnyuk 2021; Schlag and Sobolev 2022), monopoly regulation (Guo and Shmaya 2025), project choice (Guo and Shmaya 2023), statistical treatment choice (Manski 2004), and dynamic Bayesian games (Schlag and Zapechelnyuk 2024).

The Bayesian persuasion literature, pioneered by Kamenica and Gentzkow (2011), has grown rapidly (see Bergemann and Morris (2019); Kamenica (2019) for surveys). Our paper relates to work on multiple receiver types (see Kolotilin et al. (2017); Laclau and Renou (2017); Guo and Shmaya (2019); Parakhonyak and Vikander (2023); Best and Quigley (2024)), differing in that the sender does not know the receiver type distribution. Another key feature of our model is that the receiver's prior may differ from the Bayesian sender's, a setting first studied in Alonso and Camara (2016).

The most popular criterion for decision-making without priors, or robust decision-making, is maximin utility. Robust persuasion literature primarily employs this approach, which, unlike minimax loss, does not allow for the interpretation of an unknown receiver population. Hu and Weng (2021) analyze a setting where the sender and receiver share a common prior (unlike our model) and nature determines the receiver's private information. They show that even in binary settings, optimal mechanisms may involve infinitely many messages, persuading some receivers with specific beliefs while failing for others. Sapiro-Gheiler (2024) applies a similar maximin utility approach to a setting where the outside option of the receiver is drawn from an unknown distribution with a fixed mean. Dworczak and Pavan (2022) also assume a common prior and examine a setting where nature sends an extra signal to the receiver based on both the sender's mechanism and signal realization, granting nature more power than in our model, where nature and sender move simultaneously. Under this approach, the sender aims for the best payoff under the worst-case prior. In our setting with unknown priors, maximin utility criterion yields a trivial result: nature can choose a non-persuadable receiver, yielding zero payoff. To address this, additional constraints are introduced, such as restricting priors to distributions with a certain probability mass above or below a given point (see Kosterina (2022)). Moreover, unlike us, she considers sequential moves by the sender and adversarial nature. In her model, the optimal mechanism is independent of the sender's prior. By contrast, in our framework, even if the sender had a prior (see Corollary 2), the optimal mechanism and minimax loss would depend on it.

The closest paper to ours is Babichenko et al. (2022). They also apply the minimax loss/regret solution concept to a Bayesian persuasion problem. The main difference lies in the structure of uncertainty. In Babichenko et al. (2022), the sender and the receiver hold a common prior, and the uncertainty is with respect to the receiver's utility function. We consider two other sources of uncertainty—about the receiver's prior and the actual distribution over the states of the world.

Differences in the sources of uncertainty imply different mechanisms for generating losses between the two papers, which implies that losses are different between the two models, with two exceptions. Firstly, our models coincide when the number of states equals two and when the following extra restrictions are imposed: (i) there is no uncertainty about the state distribution in our model, and (ii) the utility function is known to be monotone in Babichenko et al. (2022). Secondly, as the number of states approaches infinity, both models generate maximal loss. Moreover, to obtain the limiting result, Babichenko et al. (2022) require that the utility function is not uniformly bounded from below, which is arguably a stronger assumption than a non-negative (but arbitrarily small) prior. Another notable difference between the papers is that our binary state space setting allows us to study questions not covered in Babichenko et al. (2022). In particular, we explicitly derive the optimal mechanism for an arbitrary level of uncertainty, show how uncertainty about the distribution of states and about the prior interact in generating the maximal loss (i.e., either having a correct prior or knowing that the receiver holds a correct but unknown prior reduces the maximal loss to $1/e$), and characterize the impact of uncertainty on the truthfulness of the sender.

The rest of the paper is organized as follows. Section 2 introduces the model in a binary setting, derives the Bayesian sender benchmark and sets up the loss function. Section 3 contains all the main results for the binary setting. Section 4 extends our results to a framework with multiple states, and Section 5 concludes. All the proofs are presented in the Appendices, Appendix A contains proofs of the key results: Lemma 2 and Theorems 1 and 2. Online Appendix B containing proofs of all other results. Online Appendix C contains technical derivations used in the proofs. Online Appendix D provides a detailed comparison with the privately informed receiver model by Kolotilin et al. (2017).

2 Binary Model and Bayesian Sender Benchmark

In this section, we introduce necessary notation, the Bayesian sender benchmark and set up the uninformed sender problem.

We consider a Bayesian persuasion problem with binary state and binary action. Suppose that the set of possible states of the world is $\Omega = \{0, 1\}$. The receiver chooses an action $a \in \{0, 1\}$. In what follows we refer to $a = 1$ as adopting (e.g. purchasing the product) and $a = 0$ as rejecting. Utility functions of the receiver and the sender are given by

$$u_R = r + a(\omega - r), \quad u_S = a,$$

where $r \in (0, 1)$ is the outside option of the receiver.³ That is, the receiver (she) prefers to adopt when $\omega = 1$ and the sender (he) prefers adoption in any state of the world.

We assume that the receiver has a prior belief that probability of $\omega = 1$ is equal to $\beta \in (0, 1)$. The true probability of $\omega = 1$ might be different from the receiver's belief and is denoted by $\alpha \in (0, 1)$. The sender commits to a persuasion mechanism. Formally, let M be a set of messages contained in topological space⁴ and \mathcal{M} be a Borel sigma-algebra generated by this set. The sender's strategy is a mapping $\mu^M : \Omega \times \mathcal{M} \rightarrow [0, 1]$, such that for any $\omega \in \Omega$ and $B \in \mathcal{M}$, $\mu_\omega(B) \equiv \mu(B|\omega)$ is a probability measure. After observing the mechanism μ^M and a message $m \in M$ generated by this mechanism, the receiver updates the prior in a Bayesian way and chooses an action $a \in \{0, 1\}$, which maximizes her expected utility.

Consider a mechanism $\mu^M = (\mu_0, \mu_1)$. We define a measure μ_R (corresponding to the posterior of a Bayesian receiver), such that for any $B \in \mathcal{M}$ we have

$$\mu_R(B) = \beta\mu_1(B) + (1 - \beta)\mu_0(B).$$

Now we define a conditional probability (of $\omega = 1$ conditional on message m) as a random variable $P_\beta(m)$, such that for all $B \in \mathcal{M}$ we have

$$\int_B P_\beta(m) d\mu_R(m) = \beta\mu_1(B). \tag{1}$$

We define the acceptance set as a set of messages that result in a posterior sufficiently high to justify adoption:

$$A = \{m \in M : P_\beta(m) \geq r\}. \tag{2}$$

Now we consider the Bayesian sender benchmark. That is, the sender (i) knows the probability of the good state α , and knows (ii) the prior of the receiver β and the outside option r . Although the Bayesian sender knows the probability distribution over the states, he does not know the realization of the state when he chooses his mechanism. The objective function of such a sender can be written as

$$\pi(\mu; \alpha, \beta, r) = \alpha \int_A d\mu_1 + (1 - \alpha) \int_A d\mu_0 = \alpha\mu_1(A) + (1 - \alpha)\mu_0(A), \tag{3}$$

where A is the acceptance set for mechanism μ^M and parameters β and r . The optimal mechanism, which we denote by $\hat{\mu}$, is described in Kamenica and Gentzkow (2011), contains two messages, one of which is sent with probability 1 in the good state.

Lemma 1 *The optimal persuasion mechanism (of the Bayesian sender) contains two messages m^+ and m^- such that:*

1. if $\beta < r$ then $\hat{\mu}_1(m^+) = 1$ and $P_\beta(m^+) = r$;

³ Our linear utility specification, naturally, is not restrictive in the binary setting. We keep this specification in the multiple states extension, but all our results for multiple states presented in Section 4 can be generalized to any utility function, provided that the sender knows the ranking of the states by the receiver.

⁴ One can think about M as an interval $[0, 1]$.

2. if $\beta \geq r$ then $\hat{\mu}_1(m^+) = 1$ and $P_\beta(m^+) = \beta$.

That is, the structure of the optimal mechanism does not depend on α . The expected profit associated with the optimal mechanism $\hat{\mu}$ is given by

$$\pi(\hat{\mu}; \alpha, \beta, r) = \alpha + (1 - \alpha) \min \left\{ \frac{\beta(1 - r)}{r(1 - \beta)}, 1 \right\}.$$

Now we proceed with the uninformed sender problem, which is the key focus of this paper. The uninformed sender knows neither the probability of the good state α , nor the receiver's prior β and her outside option r . Suppose that such a sender commits to some mechanism μ^M . We define the loss function as the difference between the payoffs of the Bayesian and the uninformed senders:

$$L(\mu^M; \alpha, \beta, r) = \pi(\hat{\mu}; \alpha, \beta, r) - \pi(\mu^M; \alpha, \beta, r). \tag{4}$$

The objective of the uninformed sender is to minimize loss in the worst-case scenario:

$$\inf_{G \in \mathcal{G}} \sup_{F \in \mathcal{F}} \int \int L(\mu^M; \alpha, \beta, r) dF dG, \tag{5}$$

where \mathcal{G} is a set of all probability measures over all possible persuasion mechanisms⁵ defined on (M, \mathcal{M}) and \mathcal{F} is a set of all probability distributions over the admissible parameters of the model, i.e., $[\underline{\alpha}, \bar{\alpha}] \times [\underline{\beta}, \bar{\beta}] \times [\underline{r}, \bar{r}] \subset (0, 1)^3$. That is, the sender may have some restriction on the range of parameters of the model. When $\underline{\alpha} = \bar{\alpha} = \alpha_0$ the sender has (subjectively) perfect knowledge of the environment but faces uncertainty about the type of the receiver. For example, a seller of a product knows the probability of the product being faulty, or a doctor knows the probability of treatment failing to be successful, but they do not know the buyer's subjective view on the likelihood of the product being good (i.e., the treatment being efficient).

3 Analysis

3.1 Optimal Persuasion Mechanism

The solution to problem (5) is obtained by solving a zero-sum game between the sender, who chooses persuasion mechanism μ^M and adversarial nature, which simultaneously chooses a probability distribution over parameters of the environment, (α, β, r) , in order to maximize loss. In our analysis we are going to focus on the equilibrium of this game.

Although the optimal strategy of the Bayesian sender has a very simple two-message structure, the strategy of the uninformed one may involve many different messages. We

⁵ For representational convenience, we derive the sender's strategy as a randomization over a set of mechanisms. Intuitively, such randomization can be represented as a single persuasion mechanism, see Corollary 1.

show that this is indeed the case (see Corollary 1). We show that despite its complexity this mechanism can be represented as a randomization over standard mechanisms (see Lemma 2), in which one message perfectly reveals the negative state.

Suppose that adversarial nature plays a distribution $\mathbb{F}(\alpha, \beta, r)$. Suppose that the sender plays mechanism $\mu^M = (\mu_0(m), \mu_1(m))$, $m \in M$. The following Lemma states that the set of best responses of the uninformed sender to any strategy of adversarial nature includes a mechanism consisting of two messages.

Lemma 2 *For any \mathbb{F} and any mechanism μ^M the expected loss of the uninformed sender $\mathbb{E}_{\mathbb{F}}L(\mu^M; \alpha, \beta, r)$ can be achieved by a randomization over mechanisms consisting of at most two messages. Moreover, in each of these mechanisms one of the messages is sent with probability 1 if the state is $\omega = 1$.*

Intuitively, the proof of Lemma 2 relies on the following observation. Consider a message m that is played in the original mechanism with probability $\mu_1(m)$. This message m persuades some types of receivers but not others. With the same probability $\mu_1(m)$, the sender could choose a corresponding binary mechanism in which message m^+ , sent with probability 1 in $\omega = 1$, persuades the same types of receivers as message m does in the original mechanism. Such randomization yields the same expected payoff as the original mechanism, implying that the payoff from any randomization over general mechanisms can be matched by an appropriate randomization over binary mechanisms. Hence, in deriving the optimal mechanism, it is sufficient to consider binary recommendation mechanisms with $\mu_1(m^+) = 1$. Thus, any mechanism is fully characterized by a single number, that is, the probability of sending the adoption recommendation in the bad state. Note, that for the binary message mechanism, the receiver adopts if her posterior is higher than r , i.e.,

$$\frac{\beta}{\beta + (1 - \beta)\mu_0(m^+)} \geq r \Leftrightarrow \lambda \geq \mu_0(m^+),$$

where

$$\lambda = \frac{\beta(1 - r)}{r(1 - \beta)} \quad (6)$$

is a characteristic of receiver's behaviour, which we refer to as the 'receiver's optimism'. We denote the range of possible values of λ as $[\underline{\lambda}, \bar{\lambda}]$.⁶ Note that the receiver's decision to adopt depends solely on the optimism parameter λ , but the actual decomposition into β and r does not play any role. In particular, uncertainty about λ can be interpreted as uncertainty about the receiver's prior beliefs, her utility, or both.⁷ Hence, in what follows, instead of writing the strategy of nature as $\mathbb{F}(\alpha, \beta, r)$, we will use corresponding distribution $F(\alpha, \lambda)$; we will denote by $F(\lambda)$ the marginal distribution of receiver's optimism. The higher the λ , the more optimistic the receiver is and the

⁶ From the restrictions on the strategy of nature we get $\underline{\lambda} = \min \left\{ 1, \frac{\beta(1-r)}{r(1-\beta)} \right\}$ and $\bar{\lambda} = \min \left\{ 1, \frac{\beta(1-r)}{r(1-\beta)} \right\}$, as for all $\lambda > 1$ the receiver does not need to be persuaded.

⁷ For most of our analysis, a reader may assume, e.g., that r is fixed and is common knowledge. Such interpretation is explicitly required for the analysis of the case of the correct prior of the receiver (Propositions 4 and 5). The multiple states analysis requires r being a free choice of adversarial nature.

more often the sender can send m^+ in the bad state. In what follows, we denote the probability of sending the adoption recommendation in the bad state as $\mu \equiv \mu_0(m^+)$. Lemma 1 implies that the optimal strategy of the Bayesian sender is $\hat{\mu}_0(m^+) = \lambda$. Using Lemma 2, we can rewrite loss function (4) as

$$L(\mu; \alpha, \lambda) = \begin{cases} (1 - \alpha)(\lambda - \mu), & \mu \leq \lambda \\ \alpha + (1 - \alpha)\lambda, & 1 \geq \mu > \lambda \end{cases} \tag{7}$$

The first line of the loss function corresponds to the case when the sender is too modest in his recommendation. Recommending adoption in the bad state with higher probability would still persuade the receiver, and thus the sender incurs a loss from too infrequent recommendation. The second line of the loss function corresponds to the case when the sender is too ambitious in his recommendation. Recommending adoption in the bad state too frequently leads to failure of message m^+ to persuade the receiver (in both states of the world). When $\mu = \lambda$, the sender chooses the optimal mechanism and loss equals zero.

As the sender and adversarial nature play a zero-sum game, there is no pure strategy equilibrium (the sender wants to choose $\mu = \lambda$ and nature wants to set $\lambda \neq \mu$ to create some loss). In Theorem 1, we characterize a mixed strategy equilibrium of this game.

Theorem 1 *There exists a mixed strategy equilibrium in which:*

1. *the sender and nature choose μ and λ from common support $[\underline{\mu}, \bar{\lambda}]$, where $\underline{\mu} \in [\underline{\lambda}, \bar{\lambda}]$;*
2. *the sender's strategy is characterized by C.D.F. $G(\mu)$ which is continuous on $(\underline{\mu}, \bar{\lambda}]$;*
3. *nature's strategy is characterized by C.D.F. $F(\lambda)$ which is continuous on $[\underline{\mu}, \bar{\lambda}]$;*
4. *nature sets the probability of the good state according to function*

$$\alpha(\lambda) = \begin{cases} \bar{\alpha}, & \lambda < \lambda_0 \\ \underline{\alpha}, & \lambda \geq \lambda_0 \end{cases} \tag{8}$$

with $\lambda_0 \in (\underline{\mu}, \bar{\lambda})$.

Moreover, this equilibrium is unique in the class of binary message mechanism strategies.

Expressions for G and F are given by equations (28) and (29) in the proof of the Theorem in the Appendix (see page 35). Now we discuss the properties of the equilibrium.

First, the sender and nature randomize over the same convex support. Moreover, if the upper bound of the support of F was below $\bar{\lambda}$, then for any strategy of the sender nature could increase the loss by choosing a higher λ (see case 1 in formula (7)). This logic does not work for the sender at the lower bound: there is no reason to choose μ below the lower bound of the support of nature's strategy, as this would only increase loss. Hence, it may be the case that $\underline{\mu} > \underline{\lambda}$.

Second, in the interior of the support their strategies are continuous. The intuition is that if the sender had an atom at some point, then nature would like to put some

probability mass just below this point, thus increasing loss by making the receiver less persuadable. Hence, the only λ that the sender can potentially play with positive probability is when it is impossible to undercut, i.e., $\underline{\lambda}$.

Third, if nature had an atom at some point, the sender would like to transfer the probability mass from just below this point slightly upwards. The only case when this is not possible is the upper bound $\bar{\lambda}$.

Fourth, conditional on nature playing some λ , its choice of α is deterministic (see equation (8)) and takes extreme values only. Intuitively, if λ is high, then the sender is more likely to end up in the first case in formula (7), implying that it is optimal to have the lowest possible α . Similarly, when λ is low, then the second case is more likely, meaning that the choice of the largest possible α maximizes loss. That is, nature makes the probability of the good state to be negatively correlated with receiver's optimism: whenever the receiver is sufficiently optimistic, the probability of the good state is minimal and vice versa.

The sender's strategy G described in Theorem 1 may be interpreted as a mechanism involving many messages. Indeed, such a strategy chooses a binary mechanism $\{\mu_0(m^+) = \mu, \mu_1(m^+) = 1\}$ with density $g(\mu)$. Note that this binary mechanism persuades any receiver with $\lambda \geq \mu$. Alternatively, the sender could commit to a large mechanism, which sends a message "adopt if your $\lambda \geq \mu$ " with density $g(\mu)$ in the good state and $\mu g(\mu)$ in the bad state. Since for each (α, λ) this mechanism induces the same adoption probability, the best response of nature is still F .⁸ That is, the sender's strategy allows for two interpretations: as a pure strategy mechanism with a continuum of messages or a mixed strategy over two-message mechanisms. We sum up this result in the following corollary.

Corollary 1 *The worst-case loss is minimized by the following mechanism:*

- send message "adopt if your $\lambda \geq \mu$ " with density $g(\mu)$ if $\omega = 1$;
- send message "adopt if your $\lambda \geq \mu$ " with density $\mu g(\mu)$ if $\omega = 0$;
- send message "reject" with probability $1 - \mathbb{E}_G \mu$ if $\omega = 0$;

where g is the density function corresponding to C.D.F. of sender's strategy G defined by (28) in the proof of Theorem 1.

Comparing the mechanism from Corollary 1 with the Bayesian sender benchmark mechanism from Lemma 1 leads to the following observations. First, in the presence of uncertainty the sender finds it optimal to use a continuum of messages instead of two. Second, under uncertainty, the sender randomizes over messages in the good state of the world, while the Bayesian sender sticks to a pure strategy in this state.⁹ Finally, both mechanisms involve a message which perfectly reveals state $\omega = 0$.

We now discuss uniqueness of the optimal mechanism. Theorem 1 established uniqueness of the equilibrium in the class of binary mechanisms, in which the sender randomizes between the messages only when $\omega = 0$. However, as it follows from

⁸ Note that this construction essentially reverse-engineers the steps in the proof of Lemma 2 where we represented an arbitrary mechanism with a randomization over binary mechanisms.

⁹ Note, that in the setting with multiple receiver types by Kolotilin et al. (2017), in the case of binary state space the optimal mechanism contains two messages. Thus, the continuum of messages arises not from the multiplicity of receiver types but from the sender's ambiguity regarding the distribution over the types.

Corollary 1 this randomization can be formulated as an alternative mechanism with many messages. It is possible to construct other mechanisms, for example, by simply relabelling the messages. However, according to Lemma 2, each such mechanism can be replicated by an appropriate mixture of binary mechanisms. Proposition 1 verifies that this replication is unique, i.e., any equilibrium mechanism is pay-off equivalent to the mixed strategy described in Theorem 1. Although pay-off equivalence follows from the Minimax Theorem, Proposition 1 establishes a stronger result: the strategy of nature is *identical* for each sender’s mechanism, and each sender’s mechanism passes the same amount of information to the receiver and induces the same distribution of actions.

Proposition 1 *Let G and F be the equilibrium strategies of the sender and nature described in Theorem 1. Then, the set of equilibrium mechanisms and strategies of nature (G', F') is such that $F' = F$, and G' and G result in the same loss, i.e.,*

$$\int \int L(\mu^M; \alpha, \beta, r) dF dG' = \int \int L(\mu^M; \alpha, \beta, r) dF dG.$$

3.2 Properties of Minimax Loss

Now we proceed with characterization of the minimax loss. We start with deriving global properties and then focus on the impact of uncertainty on the minimax loss.

Lemma 3 *The minimax loss of the sender is given by*

$$\bar{L} = \ln \left(\frac{\underline{\alpha} + (1 - \underline{\alpha})\bar{\lambda}}{\underline{\alpha} + (1 - \underline{\alpha})\lambda_0} \right). \tag{9}$$

This result is obtained directly from the equilibrium strategies derived in Theorem 1. Note that λ_0 (i.e., the point where nature switches from the highest to the lowest α) depends on $\bar{\alpha}$ and $\bar{\lambda}$, meaning that the minimax loss depends on all four boundaries of the support of the nature strategy. We now consider the properties of the minimax loss.

Proposition 2 *The minimax loss is characterized by the following properties:*

1. \bar{L} strictly increases in $\bar{\alpha}$ and $\bar{\lambda}$, strictly decreases in $\underline{\alpha}$ and weakly decreases in $\underline{\lambda}$;
2. $\sup_{\{\underline{\alpha}, \bar{\alpha}, \underline{\lambda}, \bar{\lambda}\}} \bar{L} = \bar{\Omega}$, where $\bar{\Omega}$ is the solution to $\bar{\Omega}e^{\bar{\Omega}} = 1$;¹⁰
3. if $\underline{\lambda} = \bar{\lambda}$ then $\bar{L} = 0$;
4. if $\alpha = \underline{\alpha} = \bar{\alpha}$ then there exists $\hat{\alpha}_0$ such that \bar{L} is weakly increasing for $\alpha \leq \hat{\alpha}_0$ and strictly decreasing otherwise; moreover, $\bar{L} \leq \frac{1}{e}$.

The first part of Proposition 2 is very intuitive: more uncertainty always harms the sender, as it makes it harder to target the optimal mechanism, while the Bayesian sender benchmark is always optimal given the parameters.

¹⁰ Alternatively, $\bar{\Omega} = W_0(1)$, where W is Lambert’s W function, see https://en.wikipedia.org/wiki/Omega_constant or Mezo (2022) for details.

The second part states that if uncertainty is maximal, then the minimax loss approaches what is called the ‘omega constant’, which equals approximately 0.567.

The third part of Proposition 2 says that when there is no uncertainty about the receiver’s type then the minimax loss equals zero. This is intuitive: the optimal mechanism employed by the Bayesian sender does not depend on α , thus the uninformed sender can perfectly match this mechanism. The only impact of uncertainty is that the uninformed sender does not know the exact value of the equilibrium payoff but is certain that the maximal possible payoff will be achieved, whatever its value is.

Finally, the fourth part of Proposition 2 illustrates the role of uncertainty in λ . Suppose that the probability of the good state is known and takes value α . For generic $\underline{\lambda}$ and $\bar{\lambda}$ the minimax loss first increases and then decreases in α . If we allow for the largest possible uncertainty in λ , i.e., $\underline{\lambda} = 0$ and $\bar{\lambda} = 1$, the minimax loss is maximal and equals to $1/e$ for all $\alpha < 1/e$ and then strictly decreases in α , becoming vanishingly small as $\alpha \rightarrow 1$.

This case of known α can be of particular interest for many applications. For example, a seller might know the true quality of the product, modelled as a probability to fit a particular buyer. What the seller does not know, however, is which product the buyer is currently using and what the buyer’s beliefs about how likely it is that the product will suit her. The seller designs an advertising strategy or a testing procedure, i.e., a persuasion mechanism, to attain the maximal probability of selling the product. Note that the strategy of the sender in this case is obtained directly from Theorem 1 by plugging in $\underline{\alpha} = \bar{\alpha} = \alpha$.¹¹ We present these much simpler expressions in the following Corollary.¹²

Corollary 2 *Suppose that $\underline{\alpha} = \bar{\alpha} \equiv \alpha$ and $\lambda \in [\underline{\lambda}, \bar{\lambda}]$. Then the equilibrium strategy of the sender is*

$$G(\mu) = 1 + \ln \left(\frac{\alpha + (1 - \alpha)\mu}{\alpha + (1 - \alpha)\bar{\lambda}} \right), \quad (10)$$

with support $[\underline{\mu}, \bar{\lambda}]$, where

$$\underline{\mu} = \max \left\{ \underline{\lambda}, \frac{1}{1 - \alpha} \left(\frac{\alpha + (1 - \alpha)\bar{\lambda}}{e} - \alpha \right) \right\}.$$

The associated minimax loss is given by

$$\bar{L} = [\alpha + (1 - \alpha)\underline{\mu}] \ln \left(\frac{\alpha + (1 - \alpha)\bar{\lambda}}{\alpha + (1 - \alpha)\underline{\mu}} \right) \leq 1/e.$$

Interestingly, when nature is not restricted in its strategies, the value of knowing the true state of the world reduces the highest minimax loss by approximately one third.

¹¹ See formulae (28) and (29) in the proof of the Theorem in the Appendix.

¹² Our problem for known α has a similar structure to Bergemann and Schlag (2008), and therefore the functional forms of the optimal distribution (over persuasion probability or price) and loss are the same in both papers.

That is,

$$\Delta \bar{L} = \bar{\Omega} - \frac{1}{e}.$$

This difference can be interpreted in the following way. Suppose, that nature is free in its choice of λ and α but is now required to reveal its choice of α to the sender prior to the choice of the persuasion mechanism. In this case, nature would be indifferent between all $\alpha < 1/e$ and the sender would enjoy a decrease in the expected loss of $\Delta \bar{L}$. This shape of the maximal loss as a function of α is similar to the one in Babichenko et al. (2022) when nature chooses a monotone utility function over the states.

Another observation is that under full information, the choice of the optimal mechanism does not depend on α . However, in the case of uncertainty about λ , the actual (known) value of α affects both the choice of the mechanism and the strategy of nature. The reason is that nature plays a more pessimistic distribution over λ when α is higher, which makes the sender more truthful; that is, the distribution function $G(\mu)$ shifts to the right.

Proposition 2 considers global properties of the minimax loss. Now we are going to explore the marginal impact of a small uncertainty in parameters on the minimax loss. Suppose that either α or λ is fixed, with a complimentary variable taking values from a permissible range. That is, we allow nature to vary its corresponding parameter, for example, the probability of the good state, in a neighbourhood $[\alpha - \varepsilon/2, \alpha + \varepsilon/2]$. Part 1 of Proposition 2 establishes that the minimax loss should increase. Now we are going to establish by *how much* it increases in response to a small increase in uncertainty.

Proposition 3 *Consider small parameter uncertainty.*

1. *Suppose that $\alpha \in [\underline{\alpha}, \bar{\alpha}]$ and let $\lambda \in [\hat{\lambda} - \varepsilon/2, \hat{\lambda} + \varepsilon/2]$. Then, when $\varepsilon \rightarrow 0$, $\bar{L} = (1 - \alpha)\varepsilon + o(\varepsilon)$.*
2. *Suppose that $\lambda \in [\underline{\lambda}, \bar{\lambda}]$ and let $\alpha \in [\hat{\alpha} - \varepsilon/2, \hat{\alpha} + \varepsilon/2]$. Then, when $\varepsilon \rightarrow 0$, $\bar{L} = \bar{L}_0 + k\varepsilon + o(\varepsilon)$, where $k < 1/2$ and $\bar{L}_0 = [\hat{\alpha} + (1 - \hat{\alpha})\underline{\lambda}] \ln \left(\frac{\hat{\alpha} + (1 - \hat{\alpha})\bar{\lambda}}{\hat{\alpha} + (1 - \hat{\alpha})\underline{\lambda}} \right)$.*

The first part of Proposition 3 says that the impact of the uncertainty in λ crucially depends on $\underline{\alpha}$, but does not depend on $\bar{\alpha}$. The intuition is as follows. Consider equation (9). We know that $\bar{\lambda}$ increases by $\varepsilon/2$, and in the worst-case λ_0 can decrease by $\varepsilon/2$ (i.e., not by more than the decrease in $\underline{\lambda}$) which, for small ε , results in a total impact of $(1 - \underline{\alpha})\varepsilon$ (due to the coefficient $(1 - \underline{\alpha})$ in front of both λ 's). The upper bound $\bar{\alpha}$ affects the minimax loss only indirectly via λ_0 , but does not play any role for sufficiently small ε . Thus, the overall impact of uncertainty in λ is generically less than one-to-one. For $\varepsilon = 0$ the minimax loss equals zero, as follows from part 3 of Proposition 2.

The second part of Proposition 3 says that uncertainty in α of size ε generates at most a $\frac{1}{2}\varepsilon$ increase in the minimax loss. The precise value of k depends on the parameter values, explicit formulae for each of the cases are given by equations (37) and (39) in the Online Appendix. For $\varepsilon = 0$ the minimax loss is positive, unlike in the first part of the proposition. This is because uncertainty in λ alone, unlike only uncertainty in α , is sufficient to generate a positive loss. Moreover, from Proposition 2, we have that $\bar{L}_0 \leq 1/e$. Intuitively, for any $\varepsilon \geq 0$ the highest minimax loss is obtained when nature is not restricted in its choice of λ .

3.3 Probability of Lying

We now consider the probability of the sender lying in the bad state, that is, recommending adoption when $\omega = 0$. This probability is given by the expected value of μ given the optimal strategy of the sender. The Lemma 4 characterizes the average probability of lying.

Lemma 4 *The probability of lying, i.e., recommending adoption in the bad state, can be represented as*

$$\mathbb{E}_G \mu = \bar{\lambda} - \frac{\bar{L}}{1 - \underline{\alpha}}.$$

It strictly increases in $\bar{\lambda}$, weakly increases in $\underline{\lambda}$, and strictly decreases in $\bar{\alpha}$ and $\underline{\alpha}$.

If either $\bar{\lambda}$ or $\underline{\lambda}$ goes up, then the receiver is on average easier to persuade. The sender naturally responds to a more optimistic distribution of receivers with a choice of a more ambitious strategy, that is, he is more likely to recommend adoption in the bad state.

Whenever $\bar{\alpha}$ or $\underline{\alpha}$ increases, the probability of being in the good state goes up. This increases the cost of not persuading the receiver, as choosing too high μ would result in the receiver rejecting regardless of the state of the world, while using a more modest strategy with a lower μ harms the sender only in the bad state, which is relatively less likely than before. Thus, the higher costs of failing to persuade the receiver result in the choice of a more truthful mechanism.

Overall, increasing uncertainty affects the probability of lying through two channels: either directly through $\bar{\lambda}$ and $\underline{\alpha}$ or indirectly through \bar{L} ; ceteris paribus a higher value of loss is associated with a more truthful sender. Since \bar{L} is weakly increasing in uncertainty, the indirect effect on the probability of lying is always negative. In particular, this implies that an increase in uncertainty through $\underline{\lambda}$ and $\bar{\alpha}$ increases the minimax loss (see Proposition 2) and disciplines the sender. Therefore, a policy maker who is interested in the receiver's welfare never wants to reduce the set of admissible parameters by excluding the most pessimistic receiver types and most favourable distributions over the states, i.e., prefers to maximize sender's minimax loss \bar{L} given parameters $\bar{\lambda}$ and $\underline{\alpha}$.


The overall impact of increasing uncertainty on the probability of lying is ambiguous. First, consider the impact of an increase in uncertainty in λ , where $\bar{\lambda}$ increases and $\underline{\lambda}$ decreases by $\varepsilon/2$. In this case, if uncertainty is small, we know that the minimax loss increases proportionally to $(1 - \underline{\alpha})\varepsilon$, as shown in Proposition 3. Then, from Lemma 4, we can derive that

$$\Delta \mathbb{E}_G \mu \approx \frac{\varepsilon}{2} - \varepsilon = -\frac{\varepsilon}{2}.$$

That is, the sender who is only marginally unsure about the receiver's prior is more truthful than the Bayesian one, because the indirect channel dominates the direct one. Intuitively, such sender responds to the risk of failure to persuade, which results in zero payoff, by choosing a more cautious strategy.

Second, consider larger levels of uncertainty in λ . At some point it may be the case that the lower bound $\underline{\mu}$ becomes larger than $\underline{\lambda}$. In this case \bar{L} does not depend

Fig. 1 $\mathbb{E}_G(\mu)$ as a function of ε , $\underline{\alpha} = 0.1$, $\bar{\alpha} = 0.3$, $\hat{\lambda} = 0.5$



on $\underline{\lambda}$, thus we can apply Lemma 4, which states that $\mathbb{E}_G\mu$ increases in $\bar{\lambda}$ and thus, increases in uncertainty. Intuitively, when uncertainty is large reducing risk to persuade is extremely costly and the sender prefers to gamble with a more aggressive strategy.

A combination of these two results is presented in Figure 1. In this figure, we impose $\lambda \in [\hat{\lambda} - \varepsilon/2, \hat{\lambda} + \varepsilon/2]$ and vary ε between zero and one. If there is no uncertainty, i.e., $\varepsilon = 0$, then the sender uses the optimal strategy, that is, recommends adopting with probability $\mu = \hat{\lambda} = 1/2$. For small ε the probability of lying decreases in uncertainty, as explained above. For ε very large, when the lower bound of the support is larger than $\underline{\lambda}$ (all points to the right of the vertical dashed line), the probability of lying strictly increases in ε . As can be seen, there is an interior minimum of the probability of lying, which is attained at such levels of uncertainty that the lower bound is still determined by $\underline{\lambda}$. Interestingly, the largest level of uncertainty may correspond to the least truthful strategy of the sender.

The impact of uncertainty in α (i.e., when $\underline{\alpha} = \hat{\alpha} - \varepsilon/2$ and $\bar{\alpha} = \hat{\alpha} + \varepsilon/2$) on the sender’s strategy is ambiguous. Our numerical analysis shows that the probability of lying can both increase or decrease in uncertainty even in the neighbourhood of $\varepsilon = 0$, depending on the initial choice of parameters.

3.4 Informed Receiver

In Corollary 2, we showed that the sender can reduce the minimax loss from $\bar{\Omega}$ to $\frac{1}{e}$ by being informed about the distribution over the states of the world. We now consider an alternative setting where the sender knows the outside option r and knows that the receiver holds a *correct prior*, i.e., $\beta = \alpha$. This restricts nature’s strategy, preventing independent choices for λ and α .¹³ We show that, in this setting, the upper bound on the sender’s minimax loss remains $1/e$ as in Corollary 2. We focus on the case where nature’s choice of α is not restricted, i.e., $\alpha \in (0, 1)$. We define

$$\lambda_r(\alpha) \equiv \min \left\{ \frac{\alpha}{1-\alpha} \frac{1-r}{r}, 1 \right\}. \tag{11}$$

¹³ The assumption that the receiver knows r is crucial. Without it, nature could adjust r to choose λ , making the model identical to the baseline.

Thus, $\lambda_r(\alpha)$ is the receiver's optimism, similar to (6), but here it depends on nature's choice of α . For a given mechanism μ and α , the loss function is:

$$L_r(\mu; \alpha) = \begin{cases} (1 - \alpha)(\lambda_r(\alpha) - \mu), & \mu \leq \lambda_r(\alpha) \leq 1, \\ \alpha + (1 - \alpha)\lambda_r(\alpha), & \mu > \lambda_r(\alpha). \end{cases} \quad (12)$$

The second case can be rewritten as $\alpha + (1 - \alpha)\lambda_r(\alpha) = \alpha/r$.

As in the baseline model, both the sender and nature play mixed strategies. We denote the lower and upper bounds of the sender's strategy support as $\underline{\mu}$ and $\bar{\mu}$, and the bounds of nature's strategy as $\underline{\alpha}_F$ and $\bar{\alpha}_F$. Now, we characterize the sender's optimal mechanism.

Proposition 4 *Suppose $r \in (0, 1)$. Then, there exists a mixed strategy equilibrium where:*

1. *The sender chooses the mechanism according to the continuous distribution function*

$$G(\mu) = \frac{1}{1 - \underline{\mu}^r} \left[1 - \left(\frac{\mu}{\underline{\mu}} \right)^r \right],$$

with support $[\underline{\mu}, 1]$, where $\underline{\mu}$ satisfies

$$\underline{\mu}^r + \underline{\mu}r^2/(1 - r) = (1 - r); \quad (13)$$

2. *Nature plays a continuous distribution function F on support $[\underline{\alpha}_F, \bar{\alpha}_F]$, where $\underline{\alpha}_F = \lambda_r^{-1}(\underline{\mu}) = r\underline{\mu}/(1 - r + r\underline{\mu})$ and $\bar{\alpha}_F = \lambda_r^{-1}(1) = r$.*

The supports of λ and μ coincide, and the sender's equilibrium strategy is continuous, with no atom at the lower bound. Unlike in the baseline model, where α and λ are negatively correlated, here they are positively related via $\lambda_r(\alpha)$. Nature may choose large α , making $\lambda_r(\alpha) > 1$, but this is suboptimal: in case $\mu \leq \lambda_r(\alpha)$, loss (12) decreases with α when $\lambda_r(\alpha) > 1$. However, nature still pushes $\bar{\alpha}_F$ to r , so the largest λ equals 1, as for $\mu \leq \lambda_r(\alpha)$ loss (12) increases in α .

The receiver's correct prior belief replaces the negative correlation between α and λ with a positive one, reducing the sender's equilibrium loss that we characterize in the following proposition.

Proposition 5 *The minimax loss in the informed receiver model is:*

$$\bar{L}_r = \frac{\underline{\alpha}_F}{\bar{\alpha}_F} = \frac{\underline{\mu}}{1 - r + r\underline{\mu}}. \quad (14)$$

\bar{L}_r decreases in r and is bounded above by $\lim_{r \rightarrow +0} \bar{L}_r = 1/e$.

Equation (14) may seem counter-intuitive, as a higher ratio between the lower and upper bounds increases equilibrium loss. However, note that both $\underline{\alpha}_F$ and $\bar{\alpha}_F$ depend on r . As r increases, the receiver becomes harder to persuade, leading the sender to use

more truthful strategies, making the case $\mu \leq \lambda_r(\alpha)$ more likely. Also, as r increases, the distribution $F(\alpha)$ shifts up, reducing the loss in the case $\mu \leq \lambda_r(\alpha)$.

Proposition 5 shows that the maximum loss equals $1/e$, the same as when the sender knows the true state, as in Corollary 2. This implies that, without restrictions on nature’s strategy, the sender is equally well-off whether they know α or know that the receiver knows α . In both cases, nature’s strategy is limited to choosing a single variable, leading to the same minimax loss.

4 Model with Large Number of States

In this section, we generalize our model to a richer state space. We assume that the state space consists of $n + 1$ states:¹⁴

$$\Omega = \left\{ 0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n}, 1 \right\}.$$

We also assume that, similar to the setting in Corollary 2, the uninformed sender knows the true probability distribution over the states of world, that is now represented as a vector $(\alpha_\omega)_{\omega \in \Omega}$ with $\sum_{\Omega} \alpha_\omega = 1$. In all other features, the model we analyze here is identical to our baseline. That is, adversarial nature chooses a distribution over the receiver’s prior β_ω and the outside option r , while the sender chooses a persuasion mechanism which maps Ω to the message space M . Upon learning about the mechanism and a signal it generated, the receiver updates her beliefs in a Bayesian way and decides whether to adopt or reject.

We start with a negative result for large number of states, i.e., $n \rightarrow \infty$. We assume that for any two states $\omega', \omega'' \in \Omega$ we have $\alpha_{\omega'} / \alpha_{\omega''} < \bar{A}$, where $\bar{A} \in \mathbb{R}$. That is, as the number of states grows larger, none of the states becomes infinitely more likely than some other state. Note that $\alpha_\omega < \bar{A}/n$ for any $\omega \in \Omega$. This implies that the distribution of states α_ω converges to a continuous distribution function. The minimax loss is denoted as \bar{L} .

Theorem 2 *Suppose that the distribution of states $(\alpha_\omega)_{\omega \in \Omega}$ is known to the sender and nature is unrestricted in its choice of the receiver’s prior and the outside option. Furthermore, suppose that for any n , $\alpha_\omega < \bar{A}/n$ for some $\bar{A} \in \mathbb{R}_{++}$. Then,*

$$\lim_{n \rightarrow \infty} \bar{L} = 1.$$

The mechanism behind Theorem 2 is as follows. For illustration purposes, suppose that all states are equally likely. Then adversarial nature can choose the following combination of the receiver’s prior and the outside option. Nature assigns the prior of the receiver in the following way:

1. puts a large probability mass on state $\omega = 1$;

¹⁴ Note, that we maintain the specification of the receiver’s utility function as $u_R = r + a(\omega - r)$. Our analysis remains valid for an arbitrary utility function, provided that the utility the receiver obtains in each state is known to the sender. In this case, the states can be re-labeled to make the function monotone.

2. randomly selects $\lfloor \sqrt{n} \rfloor$ states among the remaining ones,¹⁵ which we call ‘cursed states’, and assigns the same small probability mass to each of them; we denote this mass as δ , requiring that the total mass of ‘cursed states’ approaches zero as n approaches infinity;
3. assigns probability mass ε , such that $\sqrt{n}\varepsilon/\delta \rightarrow 0$, to each of the $n - \lfloor \sqrt{n} \rfloor$ remaining states;
4. chooses $r \in (\frac{n-1}{n}, 1)$, such that pooling state $\omega = 1$ with state $\omega = (n-1)/n$ happening with probability δ leads to rejection, but if state $\omega = 1$ is pooled even with the lowest $n - \lfloor \sqrt{n} \rfloor$ states, each happening with probability ε , then the receiver finds it optimal to adopt.

Intuitively, nature puts $\lfloor \sqrt{n} \rfloor$ ‘cursed states’ in n states below $\omega = 1$. These states are assigned with high enough prior δ , that pooling them with the top state drags the expected utility of the receiver down sufficiently to result in rejection. The Bayesian sender can avoid these ‘cursed states’, and pool all the remaining states, as the receiver’s prior belief of being in these states, ε , is so low that the expected utility does not decrease enough to justify rejection. Thus, the Bayesian sender can persuade the receiver in $n - \lfloor \sqrt{n} \rfloor$ states, or almost always in large state spaces since

$$\lim_{n \rightarrow \infty} \frac{n - \lfloor \sqrt{n} \rfloor}{n} = 1.$$

In order to persuade the receiver in a number of states that does not become vanishingly small as $n \rightarrow \infty$, the uninformed sender has to include $\lfloor \gamma n \rfloor$ out of n states in the receiver’s posterior, $\gamma \in (0, 1)$. However, for any fixed γ the probability of having a ‘cursed state’ among any $\lfloor \gamma n \rfloor$ states approaches one as n approaches infinity. As having a ‘cursed state’ in a posterior leads to rejection, the uninformed sender can only persuade in a vanishingly small fraction of states, meaning that the minimax loss approaches one.

Theorem 2 has an interesting connection to the results of Babichenko et al. (2022). They consider a setting with a common prior, but the receiver’s utility is unknown. In their case, if the sender knows the ranking of the states (which is equivalent to the monotone utility assumption), the loss is bounded above by $1/e$ for any n (Theorem 3.6). In contrast, in our case, it approaches 1 as the number of states tends to infinity. This difference arises from nature’s freedom in choosing the receiver’s prior. In their setting, the loss approaches 1 if the utility that nature can assign to cursed states is not bounded from below as $n \rightarrow \infty$. Although the tool used by adversarial nature in their case is stronger compared to the choice of the prior in our paper, the mechanism is ultimately similar: some states become ‘cursed’, and pooling each of these states with the rest results in a failure to persuade. The difference is in how aggressively nature could assign probability mass to cursed states (one can show that the loss in their setting is weakly higher than in ours).

Although the results of Theorem 2 might seem rather pessimistic for the prospects of the uninformed sender in large state environments, it is important to understand that the result relies on having rather peculiar optimal persuasion mechanisms. In

¹⁵ $\lfloor x \rfloor$ is the largest integer smaller or equal to x .

such mechanisms, the Bayesian sender hand-picks some small number of states in which rejection is recommended. Most real-life mechanisms, such as certification procedures, are unlikely to exhibit this feature: they take the form of a simple pass-fail rule (that is, the utility level is above a certain threshold in our context).¹⁶ Moreover, it is straightforward to show that the optimal mechanism indeed exhibits the cut-off property as long as $[\beta(\omega)/\alpha(\omega)](\omega - r)$ is monotone in ω , which holds when the receiver’s beliefs do not differ substantially from the sender’s.

We now look at a restricted set of mechanisms, which arguably are more likely to be implemented in reality, namely cut-off mechanisms. In a cut-off mechanism, there is a state ω_0 , such that message m^+ is sent in all $\omega > \omega_0$, sends message m^- in all states $\omega < \omega_0$ and potentially randomizes between both messages at ω_0 . We assume that both the Bayesian and the uninformed senders are restricted to cut-off strategies. We will now establish the equivalence between the problem restricted to cut-off strategies with the binary problem studied in Section 3.

We start by showing that any cut-off rule can be described by a real number in $[0, 1]$. Suppose that the sender plays a mechanism in which message m^+ is sent with probability one in all $\omega > \omega_0$, where $\omega_0 < 1$,¹⁷ and with probability $\eta \in [0, 1]$ in state ω_0 .

We define $\lambda = 1 - \omega_0 - (1 - \eta)\frac{1}{n}$. Then, since

$$\omega_0 = \frac{\lfloor(1 - \lambda)n\rfloor}{n} \quad \text{and} \quad \eta = 1 - ((1 - \lambda)n - \lfloor(1 - \lambda)n\rfloor),$$

we obtain that there is one-to-one mapping between any cutoff mechanism and $\lambda \in [0, 1]$. Analogously to our main model, λ can be interpreted as receiver optimism: in all states $\omega > \lambda$ the Bayesian sender can send m^+ with probability 1, and in state just below λ with probability η , and still persuade the receiver.

We define

$$y(\lambda) = \frac{1}{1 - \alpha_n} \left[-\alpha_n + \sum_{\omega > \lfloor(1-\lambda)n\rfloor/n} \alpha_\omega + \alpha_{\lfloor(1-\lambda)n\rfloor/n} (1 - (1 - \lambda)n + \lfloor(1 - \lambda)n\rfloor) \right],$$

that is, the total probability of sending message m^+ conditional on $\omega \neq 1$ if a mechanism characterized by λ is being played. Note that $y(\lambda) : [0, 1] \rightarrow [0, 1]$ is a strictly increasing function. The term in square brackets equals to the total probability of adoption minus α_n .

Now, suppose that for some realization of the strategy of nature the Bayesian sender uses the optimal mechanism which persuades the receiver with optimism λ , i.e., a mechanism that recommends adoption in all states except the top with probability $y(\lambda)$ and with probability one in $\omega = 1$. Suppose that the uninformed sender uses a

¹⁶ Zapechelnyuk (2020) shows that such mechanisms are optimally chosen by a regulator.

¹⁷ It is always optimal to send m^+ in $\omega = 1$.

cut-off mechanism characterized by μ . Then, the loss function can be represented by

$$L(\mu; \alpha, \lambda) = \begin{cases} (1 - \alpha_n)(y(\lambda) - y(\mu)), & \mu \leq \lambda \\ \alpha_n + (1 - \alpha_n)y(\lambda), & 1 \geq \mu > \lambda \end{cases}. \quad (15)$$

That is, loss depends on how much the sender lies in states below $\omega = 1$ relative to the Bayesian benchmark. Note that this loss function (15) is equivalent to loss in the binary problem (7) with α replaced with α_n , λ and μ replaced with $y(\lambda)$ and $y(\mu)$ respectively. Thus, we can directly apply Corollary 2 with $\underline{\lambda} = 0$ and $\bar{\lambda} = 1$. This allows us to formulate the following result.

Proposition 6 *Suppose that the sender knows the probability distribution over the states of the world α and nature is free in its choice of the receiver's prior and the outside option. Moreover, suppose that the sender is restricted to cut-off strategies. Then the equilibrium strategy of the sender is*

$$G(\mu) = 1 + \ln(\alpha_n + (1 - \alpha_n)y(\mu)),$$

with support $[\underline{\mu}, 1]$, where $\underline{\mu}$ satisfies $y(\underline{\mu}) = \max\left\{0, \frac{1}{1-\alpha_n}\left(\frac{1}{e} - \alpha_n\right)\right\}$. The associated minimax loss is given by

$$\bar{L} = -[\alpha_n + (1 - \alpha_n)y(\underline{\mu})] \ln(\alpha_n + (1 - \alpha_n)y(\underline{\mu})) \leq 1/e.$$

Proposition 6 implies that the probability of the best state is the only important parameter for determining the minimax loss. If $\alpha_n < 1/e$ then the minimax loss equals $1/e$, otherwise it equals $-\alpha_n \ln \alpha_n$ and converges to 0 as α_n approaches one. This result is related to Babichenko et al. (2022), see Theorem 3.6, where the cut-off property of the optimal mechanism is delivered by a combination of monotonic utility and common prior. Proposition 6 implies that it is the cut-off property, rather than the common prior assumption, that is crucial in obtaining this result.

Note that, as equation (15) and Proposition 6 essentially establish the equivalence between the multiple state problem restricted to the cut-off mechanism and the binary state problem with a known distribution of states derived in Corollary 2, all the properties of the optimal binary-state mechanism are inherited in this setting. In particular, the optimal mechanism consists of a continuum of messages (see Corollary 1), loss is increasing in uncertainty (see Proposition 2), and small uncertainty makes the sender more truthful (see Figure 1). Moreover, the comparison with the privately informed receiver setting by Kolotilin et al. (2017) presented in Section 2 is qualitatively similar.

5 Conclusion

In this paper we characterize the optimal persuasion mechanism when the sender does not have priors about the environment he faces: probability distribution over possible states, the receiver's prior and outside option. The maximin utility approach, which is prevalent in the literature on robust persuasion, has little 'bite' in our setting without any

external constraints. Moreover, it restricts the sender’s attention to a specific worst-case prior of the receiver. Instead, we use a minimax loss approach, in which the sender is striving for the best *irrespective* of the receiver’s prior. With this approach, our setting can be interpreted as a problem of a sender facing a population of receivers with unknown characteristics.¹⁸

In a model with binary state space, the optimal mechanism involves a continuum of messages, but admits a simple representation as a randomization over two-message mechanisms. We characterize properties of the minimax loss and demonstrate that having perfect information about the receiver’s characteristics reduces loss from $\bar{\Omega}$ to zero, while having perfect information about the state space reduces it to $1/e$. Surprisingly, to achieve this loss the sender does not need to know probabilities of the states but can be equally well-off if he is sure that the receiver has an accurate knowledge of them. We show that small uncertainty makes the sender more truthful than in the perfect information, Bayesian, benchmark, while larger uncertainty may lead to more lying. This indicates that introducing some small noise into the environment might improve the quality of communication.

We show that as the number of states approaches infinity, the sender’s performance drops to zero. This result, however, arises in the environments when the optimal mechanism excludes specific states, randomly chosen by nature, from the support of adoption recommendation. As such mechanisms do not look reasonable, we focus on cut-off mechanisms, for which the results from the binary model carry over to an arbitrary number of states, producing a similar persuasion strategy and an identical minimax loss.

Appendix A: Proofs of the main results

A1: Notation

In this section we introduce notation, which we will use throughout the proofs. Denote

$$\underline{\phi}(\lambda) = \underline{\alpha} + (1 - \underline{\alpha})\lambda, \quad \bar{\phi}(\lambda) = \bar{\alpha} + (1 - \bar{\alpha})\lambda \tag{16}$$

and

$$C(\lambda) = 1 - \ln \underline{\phi}(\bar{\lambda}) - \ln \left(\frac{\bar{\phi}(\lambda)}{\underline{\phi}(\lambda)} \right). \tag{17}$$

Moreover, we define functions

¹⁸ In Appendix D, we compare our results to a setting with a sender who knows the population of receivers, but does not know the actual receiver type he faces, which is analyzed in Kolotilin et al. (2017). If such sender is used as a benchmark for the loss function, then none of our results changes. However, when played against fixed populations, robust mechanism performs well. For example, for the uniform distribution of receiver types, the difference in payoffs between the robust mechanism and the optimal mechanism is less than 1/9. At the same time, when facing adversarial nature, mechanism from Kolotilin et al. (2017) may result in the maximal possible loss of one.

$$\underline{\mu}(\lambda) \equiv \max \left\{ \lambda, \frac{1}{1 - \bar{\alpha}} \left(e^{-C(\lambda)} - \bar{\alpha} \right) \right\} \tag{18}$$

and

$$H(\lambda) \equiv \ln \left(\frac{\phi(\bar{\lambda})}{\phi(\lambda)} \right) - \max \{ e^{-C(\lambda)}, \bar{\phi}(\lambda) \} \left(1 + \ln \left(\frac{1}{\bar{\phi}(\lambda)} \min \{ e^{-C(\lambda)}, \bar{\phi}(\lambda) \} \right) \right). \tag{19}$$

A2: Proof of Lemma 2

Suppose that the sender plays mechanism $\mu^M = \{\mu_0, \mu_1\}$, on message space M . Nature randomizes over parameters (α, β, r) . We view the randomization over (β, r) as a randomization over the acceptance sets A defined in (2). Let \mathcal{A} be a collection of all acceptance sets generated by different choices of (β, r) for a given μ^M . Let \mathbb{F} be a strategy of nature defined on all (α, β, r) and $F(A)$ be a corresponding ‘marginal distribution’ over $A \in \mathcal{A}$.

Let $M' = \cup_{\mathcal{A}} A$ —i.e., M' is a set of messages which could potentially persuade at least some receiver. Let \mathcal{M}' be a sigma-algebra on M' . We claim that $\mu_0 \ll \mu_1$ (μ_0 is absolutely continuous with respect to μ_1 on \mathcal{M}'). Indeed, if there exists a set $B \in \mathcal{M}'$ such that $\mu_0(B) > 0$ and $\mu_1(B) = 0$ then $\forall A \in \mathcal{A}$ we have $B \not\subset A$ as any $m \in B$ reveals that $\omega = 0$.

We define

$$q(m) = \int_{\mathcal{A}} \mathbb{I}(m \in A) dF(A)$$

and

$$\alpha^e(m) = \mathbb{E}(\alpha | m \in M') = \frac{1}{q(m)} \int_{\mathcal{A} \times [\underline{\alpha}, \bar{\alpha}]} \alpha \mathbb{I}(m \in A) d\mathbb{F}.$$

We can write down the expected payoff of the sender as

$$\begin{aligned} \mathbb{E}_{\mathbb{F}} \pi(\mu^M; \alpha, \beta, r) &= \mathbb{E}_{\mathbb{F}} \left[\alpha \int_M \mathbb{I}(m \in A) d\mu_1 + (1 - \alpha) \int_M \mathbb{I}(m \in A) d\mu_0 \right] \\ &= \mathbb{E}_{\mathbb{F}} \left[\alpha \int_{M'} \mathbb{I}(m \in A) d\mu_1 + (1 - \alpha) \int_{M'} \mathbb{I}(m \in A) d\mu_0 \right] \\ &= \int_{M'} \alpha^e(m) q(m) d\mu_1 + \int_{M'} [1 - \alpha^e(m)] q(m) d\mu_0 \\ &= \int_{M'} \left(\alpha^e(m) + [1 - \alpha^e(m)] \frac{d\mu_0}{d\mu_1}(m) \right) q(m) d\mu_1, \end{aligned}$$

where $\frac{d\mu_0}{d\mu_1}(m)$ is a Radon-Nikodym derivative of μ_0 with respect to μ_1 . We denote $t(m) = \frac{d\mu_0}{d\mu_1}(m)$ for all $m \in M'$ and set it equal to an arbitrary $t > 1$ for all $m \in M \setminus M'$. It is straightforward to verify that $t(m) \leq 1$ for any $m \in M'$.

Consider a binary mechanism $\mu^{t(m)}$ with two messages m^+ and m^- , where message m^+ is sent with probability 1 if $\omega = 1$ and with probability $t(m)$ if $\omega = 0$. Note that if

m was in the acceptance set of some receiver under old mechanism, then m^+ is in the acceptance set under new mechanism. Indeed, the posterior of the receiver following message m^+ is

$$\tilde{p}(m) = \frac{\beta}{\beta + (1 - \beta)t(m)}.$$

Consider a receiver with a corresponding acceptance set A . Now, for any set $C \subseteq A$ we have

$$\begin{aligned} \int_C (\tilde{p}(m) - r)d\mu_R &= \int_C \tilde{p}(m)d\mu_R - r\mu_R(C) \\ &= \int_C \frac{\beta}{\beta + (1 - \beta)t(m)} \left(\beta + (1 - \beta) \frac{d\mu_0}{d\mu_1}(m) \right) d\mu_1 - r\mu_R(C) \\ &= \beta\mu_1(C) - r\mu_R(C) = \int_C P_\beta(m)d\mu_R - r\mu_R(C) \\ &= \int_C (P_\beta(m) - r)d\mu_R \geq 0. \end{aligned}$$

As this must hold for any set C we conclude that $\tilde{p}(m) \geq r$ almost everywhere. Then, the expected payoff of such binary mechanism is

$$\begin{aligned} \mathbb{E}_{\mathbb{F}}\pi(\mu^{t(m)}; \alpha, \beta, r) &= \mathbb{E}_{\mathbb{F}} [(\alpha + (1 - \alpha)t(m))\mathbb{I}(m \in A)] \\ &= [\alpha^e(m) + [1 - \alpha^e(m)]t(m)]q(m). \end{aligned}$$

Thus, we can rewrite the expected payoff from an arbitrary mechanism as

$$\begin{aligned} \mathbb{E}_{\mathbb{F}}\pi(\mu^M; \alpha, \beta, r) &= \int_{M'} \mathbb{E}_{\mathbb{F}}\pi(\mu^{t(m)}; \alpha, \beta, r)d\mu_1 \\ &= \int_M \mathbb{E}_{\mathbb{F}}\pi(\mu^{t(m)}; \alpha, \beta, r)d\mu_1. \end{aligned} \tag{20}$$

Therefore, a payoff from each mechanism can be represented as a payoff from a randomization over binary mechanisms. Let \mathcal{G} be the set of all probability measures over persuasion mechanisms, and \mathcal{G}_2 be the set of all probability measures over binary message mechanisms. From (20) and the fact that $\mathcal{G}_2 \subset \mathcal{G}$ we obtain that

$$\sup_{\mu^M \in \mathcal{G}} \mathbb{E}_{\mathbb{F}}\pi(\mu^M; \alpha, \beta, r) = \sup_{\mu^M \in \mathcal{G}_2} \mathbb{E}_{\mathbb{F}}\pi(\mu^M; \alpha, \beta, r).$$

Thus, $\inf_{\mu^M \in \mathcal{G}} \mathbb{E}_{\mathbb{F}}L(\mu^M; \alpha, \beta, r) = \inf_{\mu^M \in \mathcal{G}_2} \mathbb{E}_{\mathbb{F}}L(\mu^M; \alpha, \beta, r)$.

A3: Proof of Theorem 1

We are going to prove several Lemmata and finally proceed with the proof of the main result. We define the lower and upper bounds of the support of the strategy of nature as $(\underline{\alpha}_F, \bar{\alpha}_F)$ and $(\underline{\lambda}_F, \bar{\lambda}_F)$. Note that our restrictions imply that $\underline{\alpha} \leq \underline{\alpha}_F \leq \bar{\alpha}_F \leq \bar{\alpha}$ and

$\underline{\lambda} \leq \underline{\lambda}_F \leq \bar{\lambda}_F \leq \bar{\lambda}$. Correspondingly, the lower and the upper bounds of the support of the strategy of the sender are $\underline{\mu}_G$ and $\bar{\mu}_G$. For a fixed strategy of the sender the expected loss of nature is given by

$$L_F(\alpha, \lambda) = \int_{\underline{\mu}_G}^{\bar{\mu}_G} L(\mu; \alpha, \lambda) dG(\mu).$$

The following Lemma states that for any λ nature chooses either $\alpha = \underline{\alpha}$ or $\alpha = \bar{\alpha}$. That is, with respect to the probability of the good state nature's strategy has the widest support consisting of two points.

Lemma 5 *There is a function $\hat{\alpha}(\lambda) : [\underline{\lambda}_F, \bar{\lambda}_F] \rightarrow \{\underline{\alpha}, \bar{\alpha}\}$ such that $L_F(\hat{\alpha}(\lambda), \lambda) \geq L_F(\alpha, \lambda)$ for all $\alpha \in [\underline{\alpha}, \bar{\alpha}]$.*

Moreover, there is $\lambda_0 \in [\underline{\lambda}_F, \bar{\lambda}_F]$ such that $\hat{\alpha}(\lambda) = \bar{\alpha}$ for all $\lambda < \lambda_0$ and $\hat{\alpha}(\lambda) = \underline{\alpha}$ for all $\lambda > \lambda_0$. Moreover, $\lambda_0 > \underline{\mu}_G$.

Proof Fix the strategy of the sender G . Then, the expected loss of nature is given by

$$L_F(\alpha, \lambda) = [1 - G(\lambda)](\alpha + (1 - \alpha)\lambda) + (1 - \alpha) \int_{\underline{\mu}_G}^{\lambda} (\lambda - \mu) dG(\mu),$$

which is linear in α for any λ . Thus, for given G and λ the expected loss of nature $L_F(\alpha, \lambda)$ is maximized by $\alpha \in \{\underline{\alpha}, \bar{\alpha}\}$. Moreover,

$$\frac{\partial L_F(\alpha, \lambda)}{\partial \alpha} = (1 - \lambda)[1 - G(\lambda)] - \int_{\underline{\mu}_G}^{\lambda} (\lambda - \mu) dG(\mu) = (1 - \lambda) - \int_{\underline{\mu}_G}^{\lambda} (1 - \mu) dG(\mu),$$

which is decreasing in λ . Thus, either α jumps downward from $\bar{\alpha}$ to $\underline{\alpha}$ at some interior point λ_0 or it equals to either $\underline{\alpha}$ or $\bar{\alpha}$ over the whole support (in which case $\lambda_0 = \underline{\lambda}_F$ or $\lambda_0 = \bar{\lambda}_F$ respectively).

Note that as at the point $\lambda = \underline{\mu}_G$ we have $\frac{\partial L_F}{\partial \alpha} > 0$ it must be the case that $\lambda_0 > \underline{\mu}_G$. Moreover, at $\lambda = \bar{\lambda}_F$

$$\frac{\partial L_F(\alpha, \lambda)}{\partial \alpha} = (1 - \bar{\lambda}_F) - \int_{\underline{\mu}_G}^{\bar{\lambda}_F} (1 - \mu) dG(\mu) < 0$$

as playing $\mu > \bar{\lambda}_F$ is dominated by playing $\mu = \bar{\lambda}_F$ for the sender, so $G(\bar{\lambda}_F) = 1$. \square

We show that nature does not play λ_0 with a positive probability.

Lemma 6 *If $\lambda_0 < \bar{\lambda}_F$, then in equilibrium $F(\lambda)$ is continuous at $\lambda = \lambda_0$.*

Proof Suppose that nature plays λ_0 with positive probability. We show that the sender prefers $\mu = \lambda_0$ to $\mu = \lambda_0 + \varepsilon$. For any function $a(x)$ we define $a(x^-) = \sup\{a(t) : t < x\}$. Consider loss of the sender from choosing $\mu = \lambda_0$:

$$\int_{\underline{\lambda}_F}^{\lambda_0^-} [\bar{\alpha} + (1 - \bar{\alpha})\lambda]dF(\lambda) + [F(\lambda_0) - F(\lambda_0^-)]\alpha(\lambda_0) \times 0$$

$$+ \int_{\lambda_0}^{\bar{\lambda}_F} (1 - \underline{\alpha})(\lambda - \lambda_0)dF(\lambda).$$

The expected loss from playing $\mu = \lambda_0 + \varepsilon$ equals to

$$\int_{\underline{\lambda}_F}^{\lambda_0^-} [\bar{\alpha} + (1 - \bar{\alpha})\lambda]dF(\lambda) + [F(\lambda_0) - F(\lambda_0^-)][\alpha(\lambda_0) + (1 - \alpha(\lambda_0))\lambda_0] +$$

$$+ \int_{\lambda_0}^{\lambda_0 + \varepsilon} [\bar{\alpha} + (1 - \bar{\alpha})\lambda]dF(\lambda) + \int_{\lambda_0 + \varepsilon}^{\bar{\lambda}_F} (1 - \underline{\alpha})(\lambda - \lambda_0 - \varepsilon)dF(\lambda).$$

The difference between the two losses is

$$[F(\lambda_0) - F(\lambda_0^-)][\alpha(\lambda_0) + (1 - \alpha(\lambda_0))\lambda_0] +$$

$$\int_{\lambda_0}^{\lambda_0 + \varepsilon} [\bar{\alpha} + (1 - \bar{\alpha})\lambda - (1 - \underline{\alpha})(\lambda - \lambda_0)]dF(\lambda) - \int_{\lambda_0 + \varepsilon}^{\bar{\lambda}_F} (1 - \underline{\alpha})\varepsilon dF(\lambda),$$

which is positive for ε small enough. Thus, the sender prefers not to play μ in some neighbourhood above λ_0 . Then, any point in that neighbourhood gives nature a higher payoff than λ_0 , arriving to contradiction. \square

Lemma 5 defines $\hat{\alpha}(\cdot)$ at all points except λ_0 . Lemma 6 states that the measure of this point is zero, so the minimax loss is unaffected by the choice of $\hat{\alpha}(\lambda_0)$. We define function at this point as $\hat{\alpha}(\lambda_0) = \underline{\alpha}$, so that $\hat{\alpha}$ is now uniquely defined on the full domain $[\underline{\alpha}, \bar{\alpha}]$. This allows us to redefine nature’s loss as

$$L_F(\lambda) = \int_{\underline{\mu}_G}^{\bar{\mu}_G} L(\mu, \hat{\alpha}(\lambda), \lambda)dG(\mu).$$

For a given strategy of nature we define the expected loss of the sender as

$$L_G(\mu) = \int_{\underline{\lambda}_F}^{\bar{\lambda}_F} L(\mu, \hat{\alpha}(\lambda), \lambda)dF(\lambda).$$

The following Lemma characterises supports of the equilibrium strategies.

Lemma 7 *In equilibrium,*

1. $F(\lambda)$ is continuous on $[\underline{\lambda}_F, \bar{\lambda}_F]$;
2. $G(\mu)$ is continuous on $(\underline{\mu}_G, \bar{\mu}_G]$;
3. Distributions $G(\mu)$ and $F(\lambda)$ have the same compact support with $\underline{\lambda}_F = \underline{\mu}_G \equiv \underline{\mu}$ and $\bar{\lambda}_F = \bar{\mu}_G \equiv \bar{\mu}$. Moreover, $\bar{\mu} = \bar{\lambda}$.

Proof Claim 1. Suppose, that the marginal distribution of the strategy of nature has an atom at some λ_1 , which is not the upper bound of its support. Suppose, that $\hat{\alpha}$ is constant (either $\underline{\alpha}$ or $\bar{\alpha}$) in the neighbourhood of λ_1 . Then there exists $\varepsilon_1 > 0$ such that for all $\varepsilon \leq \varepsilon_1$ $L_G(\lambda_1) < L_G(\lambda_1 + \varepsilon)$. Thus the sender would not play μ just above λ_1 and $G(\lambda_1 + \varepsilon) - G(\lambda_1) = 0$. Then we have that $L_F(\lambda_1) < L_F(\lambda_1 + \varepsilon/2)$, as loss, defined by either case of equation (7), is increasing in λ and probability of having each case does not depend on λ if $\lambda \in (\lambda_1, \lambda_1 + \varepsilon)$. Contradiction. Now consider the case when $\hat{\alpha}$ changes from $\bar{\alpha}$ to $\underline{\alpha}$ at λ_1 . Again, as $(1 - \bar{\alpha})(\lambda - \mu) < \underline{\alpha} + (1 - \underline{\alpha})\lambda$ we have that $L_G(\lambda_1) < L_G(\lambda_1 + \varepsilon)$ and the proof follows the same steps as in the previous case.

Claim 2. Suppose that G has an atom at some point μ_1 . Then, there exists $\varepsilon_1 > 0$ such that for all $\varepsilon \leq \varepsilon_1$ we have that

$$L_F(\mu_1) < L_F(\mu_1 - \varepsilon). \quad (21)$$

Now consider two cases. First, suppose that $\mu_1 = \bar{\lambda}_F$. From (21) we get that nature must play $\bar{\lambda}_F$ with probability zero, hence by playing μ_1 the sender never persuades the receiver, thus μ_1 cannot be played by the sender with positive probability. Second, suppose that $\mu_1 < \lambda_F$. From (21) it follows that points at and just above μ_1 are dominated from nature's perspective with points above μ_1 , hence there exists $\varepsilon_2 > 0$ such that $F(\mu_1 + \varepsilon_2) - F(\mu_1) = 0$. Then, we get that $L_G(\mu_1) > L_G(\mu_1 + \varepsilon_2/2)$, regardless of whether $\hat{\alpha}$ is constant or jumps down from $\bar{\alpha}$ to $\underline{\alpha}$ at μ_1 . Thus, the sender cannot play μ_1 with positive probability, a contradiction.

Claim 3. Suppose, that for some $a < b$ we have $G(b) - G(a) = 0$ but $F(b) - F(a) > 0$. Then note that for any $c \in [a, b)$ we have that $L_F(b) > L_F(c)$ (as $\hat{\alpha}(b) \leq \hat{\alpha}(c)$), and thus no values in $[a, b)$ can be played with positive probability. Similarly, if $F(b) - F(a) = 0$, we have that the sender has lower losses at a than in any other point in the gap. Thus, supports must coincide. Now, if both supports have the same gap $[a, b)$ we have that $L_F((a + b)/2) > L_F(a)$, so there is a profitable deviation. Thus, both supports coincide and convex. Now, suppose that the upper bound of both supports $\bar{\mu} < \bar{\lambda}$. Note that $L_F(\bar{\lambda}) > L_F(\bar{\mu})$, thus there is a profitable deviation, hence $\bar{\mu} = \bar{\lambda}$. \square

Now we are ready to prove the main theorem.

Proof of Theorem 1. First, note that from Lemma 7 we have that $\bar{\mu} = \bar{\lambda}$, as well as the continuity of strategies in the interior of common support. Second, in the mixed strategy equilibrium both the sender and nature should be indifferent between all actions in the support of their distributions.

Denote the lower bound of the equilibrium support as $\underline{\mu} \geq \underline{\lambda}$. From Lemma 5 there exists $\lambda_0 \in (\underline{\mu}, \bar{\lambda}]$ such that $\hat{\alpha}(\lambda) = \bar{\alpha}$ for all $\lambda < \lambda_0$ and $\hat{\alpha}(\lambda) = \underline{\alpha}$ for all $\lambda \geq \lambda_0$. At the end of the proof we show that λ_0 is determined by equation $H(\lambda_0) = 0$.

Strategy of nature. We start by characterizing the strategy of nature for given $\underline{\mu}$ and $\lambda_0 > \underline{\mu}$ (this inequality will be verified later using Lemma C.1). Consider the equilibrium loss of the sender. Using the indifference condition for $\mu < \lambda_0$ we obtain

that

$$L_G(\mu) = \int_{\lambda_0}^{\bar{\lambda}} (1 - \underline{\alpha})(\lambda - \mu)dF(\lambda) + \int_{\underline{\mu}}^{\lambda_0} (1 - \bar{\alpha})\lambda dF(\lambda) - \mu(1 - \bar{\alpha})[F(\lambda_0) - F(\mu)] + \bar{\alpha}F(\mu).$$

By setting $\mu = \underline{\mu}$ and using the fact that the sender must be indifferent across all μ we get

$$\int_{\lambda_0}^{\bar{\lambda}} (1 - \underline{\alpha})(\lambda - \mu)dF(\lambda) - \mu(1 - \bar{\alpha})[F(\lambda_0) - F(\mu)] + \bar{\alpha}F(\mu) = \int_{\lambda_0}^{\bar{\lambda}} (1 - \underline{\alpha})(\lambda - \underline{\mu})dF(\lambda) - \underline{\mu}(1 - \bar{\alpha})F(\lambda_0),$$

which simplifies to

$$[1 - F(\lambda_0)](1 - \underline{\alpha})(\mu - \underline{\mu}) + F(\lambda_0)(1 - \bar{\alpha})(\mu - \underline{\mu}) = [\bar{\alpha} + (1 - \bar{\alpha})\mu]F(\mu).$$

Thus,

$$F(\mu) = \frac{B(\mu - \underline{\mu})}{\bar{\alpha} + (1 - \bar{\alpha})\mu},$$

where $B = [1 - F(\lambda_0)](1 - \underline{\alpha}) + F(\lambda_0)(1 - \bar{\alpha})$.

Next, we consider the case $\mu \geq \lambda_0$. The expected loss of the sender $L_G(\mu)$ is given by

$$\int_{\mu}^{\bar{\lambda}} (1 - \underline{\alpha})(\lambda - \mu)dF(\lambda) + \int_{\lambda_0}^{\mu} [\underline{\alpha} + (1 - \underline{\alpha})\lambda]dF(\lambda) + \int_{\underline{\mu}}^{\lambda_0} [\bar{\alpha} + (1 - \bar{\alpha})\lambda]dF(\lambda).$$

By plugging in $\mu = \lambda_0$ and using the indifference we obtain:

$$\int_{\mu}^{\bar{\lambda}} (1 - \underline{\alpha})(\lambda - \mu)dF(\lambda) + \int_{\lambda_0}^{\mu} [\underline{\alpha} + (1 - \underline{\alpha})\lambda]dF(\lambda) = \int_{\lambda_0}^{\bar{\lambda}} (1 - \underline{\alpha})(\lambda - \lambda_0)dF(\lambda),$$

which simplifies to

$$[1 - F(\lambda_0)](1 - \underline{\alpha})(\mu - \lambda_0) = [\underline{\alpha} + (1 - \underline{\alpha})\mu]F(\mu),$$

which can be rewritten as

$$F(\mu) = 1 - \frac{A}{\underline{\alpha} + (1 - \underline{\alpha})\mu}.$$

where $A = [1 - F(\lambda_0)][\underline{\alpha} + (1 - \underline{\alpha})\lambda_0]$. By plugging in $\mu = \bar{\lambda}$ and using the result from Lemma 7 that the distribution function must be continuous at λ_0 we obtain

$$F(\lambda_0) = \frac{B(\lambda_0 - \underline{\mu})}{\bar{\alpha} + (1 - \bar{\alpha})\lambda_0} = 1 - \frac{A}{\underline{\alpha} + (1 - \underline{\alpha})\lambda_0}.$$

Recall that $B = (1 - \underline{\alpha}) - (\bar{\alpha} - \underline{\alpha})F(\lambda_0)$. Then, by solving for $F(\lambda_0)$ and A and using our notation $\underline{\phi}$ and $\bar{\phi}$ we obtain

$$A = \frac{\bar{\phi}(\underline{\mu})\underline{\phi}(\lambda_0)}{\underline{\phi}(\lambda_0) - \underline{\phi}(\underline{\mu}) + \bar{\phi}(\underline{\mu})}, \quad B = (1 - \underline{\alpha}) \frac{\bar{\phi}(\lambda_0)}{\underline{\phi}(\lambda_0) - \underline{\phi}(\underline{\mu}) + \bar{\phi}(\underline{\mu})} \tag{22}$$

and

$$F(\lambda_0) = \frac{\underline{\phi}(\lambda_0) - \underline{\phi}(\underline{\mu})}{\underline{\phi}(\lambda_0) - \underline{\phi}(\underline{\mu}) + \bar{\phi}(\underline{\mu})}.$$

It remains to show that $F(\lambda_0) \in [0, 1]$. If λ_0 solves $H(\lambda) = 0$, which we are going to verify later, then following the result of Lemma C.1 we have that $\lambda_0 > \underline{\mu}$. Then, using $\underline{\phi}' > 0$ we obtain

$$0 < A < \frac{\bar{\phi}(\underline{\mu})\underline{\phi}(\lambda_0)}{\bar{\phi}(\underline{\mu})} = \underline{\phi}(\lambda_0).$$

This implies that $F(\lambda_0) = 1 - \frac{A}{\underline{\phi}(\lambda_0)} \in (0, 1)$.

Strategy of the sender. Next we derive the equilibrium strategy of the sender G . From Lemma 7 we know that G is continuous on $(\underline{\mu}, \bar{\lambda}]$. Consider the equilibrium loss of nature that plays $\lambda \geq \lambda_0$:¹⁹

$$L_F(\lambda) = (1 - G(\lambda))(\bar{\alpha} + (1 - \bar{\alpha})\lambda) + (1 - \bar{\alpha}) \int_{\underline{\mu}}^{\lambda} G(\mu)d\mu.$$

Similarly we can rewrite the loss for the case of $\lambda < \lambda_0$, so we get that

$$L_F(\lambda) = \begin{cases} (1 - G(\lambda))(\bar{\alpha} + (1 - \bar{\alpha})\lambda) + (1 - \bar{\alpha}) \int_{\underline{\mu}}^{\lambda} G(\mu)d\mu, & \lambda < \lambda_0 \\ (1 - G(\lambda))(\underline{\alpha} + (1 - \underline{\alpha})\lambda) + (1 - \underline{\alpha}) \int_{\underline{\mu}}^{\lambda} G(\mu)d\mu, & \lambda \geq \lambda_0 \end{cases}. \tag{23}$$

Consider the case of $\lambda < \lambda_0$. Define a differentiable function $J(\lambda) = \int_{\underline{\mu}}^{\lambda} G(\mu)d\mu$. Then, we can rewrite equation (23) as

$$[1 - J'(\lambda)]\bar{\phi}(\lambda) + (1 - \bar{\alpha})J(\lambda) = \bar{L}. \tag{24}$$

¹⁹ Here we apply integration by parts to Riemann-Stieltjes integral, using the fact that $(\lambda - \mu)$ is continuous and G is non-decreasing, see Theorem 21.67 in Hewitt and Stromberg (2013).

Solving homogeneous equation $J'(\lambda)\bar{\phi}(\lambda) = (1 - \bar{\alpha})J(\lambda)$ gives $J(\lambda) = c(\lambda)\bar{\phi}(\lambda)$. Plugging in back to the original equation gives $[1 - c'(\lambda)\bar{\phi}(\lambda)]\bar{\phi}(\lambda) = \bar{L}$. Thus,

$$c'(\lambda) = \frac{\bar{\phi}(\lambda) - \bar{L}}{[\bar{\phi}(\lambda)]^2},$$

which gives a solution (by replacing argument λ with μ): $G(\mu) = J'(\mu) = \ln \bar{\phi}(\mu) + C_0$. Similarly, for the case $\lambda \geq \lambda_0$ we obtain $G(\mu) = \ln \underline{\phi}(\mu) + C_1$.

Using the continuity of G at the upper bound of the support, $G(\bar{\lambda}) = 1$, we get that $C_1 = 1 - \underline{\phi}(\bar{\lambda})$. To determine C_0 we use the fact that $G(\mu)$ is continuous at $\mu = \lambda_0$ as $\lambda_0 > \underline{\mu}$, so

$$\ln \bar{\phi}(\lambda_0) + C_0 = \ln \underline{\phi}(\lambda_0) + 1 - \ln \underline{\phi}(\bar{\lambda})$$

and therefore $C_0 = C(\lambda_0)$, where $C(\cdot)$ is defined by equation (17). Thus, we obtain

$$G(\mu) = \begin{cases} \ln \bar{\phi}(\mu) + C(\lambda_0), & \mu < \lambda_0 \\ \ln \underline{\phi}(\mu) + 1 - \ln \underline{\phi}(\bar{\lambda}), & \mu \geq \lambda_0 \end{cases}.$$

The equilibrium lower bound. Next we are ready to characterize $\underline{\mu}$ as a function of λ_0 . By substituting $\underline{\lambda}$ into the obtained function $G(\mu)$ we have that $\bar{G}(\underline{\lambda}) > 0$ if and only if $\ln \bar{\phi}(\underline{\lambda}) + C(\lambda_0) > 0$ or equivalently when $\bar{\phi}(\underline{\lambda}) > e^{-C(\lambda_0)}$. In this case $\underline{\mu} = \underline{\lambda}$. Otherwise $\underline{\mu}$ is determined by the solution of $\bar{\phi}(\underline{\mu}) = e^{-C(\lambda_0)}$ and since $C' > 0$ it is easy to see that $\underline{\mu} \geq \underline{\lambda}$ where the equality occurs if and only if $\bar{\phi}(\underline{\lambda}) = e^{-C(\lambda_0)}$.

Summing up we established that the lower bound is determined by equation (18):

$$\underline{\mu} = \underline{\mu}(\lambda_0) = \max \left\{ \underline{\lambda}, \frac{1}{1 - \bar{\alpha}} \left(e^{-C(\lambda_0)} - \bar{\alpha} \right) \right\}.$$

The lower bound is defined as a function of λ_0 which we are going to characterise using Lemma 5. This Lemma also guarantees that $\underline{\mu}(\lambda_0) < \lambda_0$.

Characterization of λ_0 . It remains to show that λ_0 , i.e., the point where α jumps from the highest to the lowest possible value is defined by $H(\lambda_0) = 0$. Then, by Lemma 5 such λ_0 always exists and is uniquely defined, and moreover, the lower bound is also correctly defined, so that $\underline{\mu}(\lambda_0) < \lambda_0 \leq \bar{\lambda}$.

Note that the loss of nature is defined by

$$L_F(\alpha, \lambda) = (1 - G(\lambda))(\alpha + (1 - \alpha)\lambda) + (1 - \alpha) \int_{\underline{\mu}}^{\lambda} G(\mu) d\mu.$$

Lemma 5 states that for $\lambda = \lambda_0$ nature is indifferent between all possible values of α (and prefers $\bar{\alpha}$ above and $\underline{\alpha}$ below this point, see (23)). Thus, $\partial L_F / \partial \alpha = 0$ at $\lambda = \lambda_0$ and

$$(1 - G(\lambda_0))(1 - \lambda_0) = \int_{\underline{\mu}(\lambda_0)}^{\lambda_0} G(\mu) d\mu. \tag{25}$$

Let us first compute the left-hand side of (25)

$$(1 - G(\lambda_0))(1 - \lambda_0) = (\ln \underline{\phi}(\bar{\lambda}) - \ln \underline{\phi}(\lambda_0))(1 - \lambda_0).$$

The right-hand side of (25) can be rewritten as

$$\int_{\underline{\mu}(\lambda_0)}^{\lambda_0} G(\mu) d\mu = \frac{1}{1 - \bar{\alpha}} \left[\bar{\phi}(\mu) \ln \bar{\phi}(\mu) \right] \Big|_{\underline{\mu}(\lambda_0)}^{\lambda_0} - (1 - C(\lambda_0))(\lambda_0 - \underline{\mu}(\lambda_0)).$$

To characterize λ_0 we consider two cases determining the lower bound $\underline{\mu}(\lambda_0)$.

Case 1: $\bar{\phi}(\underline{\lambda}) > e^{-C(\lambda_0)}$. First, suppose that λ_0 satisfies $\bar{\phi}(\underline{\lambda}) > e^{-C(\lambda_0)}$ and $\underline{\mu}(\lambda_0) = \underline{\lambda}$. Then, equation (25) results in

$$(1 - C(\lambda_0) - \ln \bar{\phi}(\lambda_0))(1 - \lambda_0) = \frac{1}{1 - \bar{\alpha}} \ln \bar{\phi}(\lambda_0) - \frac{\bar{\phi}(\underline{\lambda}) \ln \bar{\phi}(\underline{\lambda})}{1 - \bar{\alpha}}.$$

Notice that $(1 - \bar{\alpha})(1 - \underline{\lambda}) = (1 - \bar{\phi}(\underline{\lambda}))$. Using the formula for $C(\lambda_0)$ we come obtain

$$\ln \left(\frac{\bar{\phi}(\bar{\lambda})}{\bar{\phi}(\lambda_0)} \right) - \bar{\phi}(\underline{\lambda})(1 - C(\lambda_0) - \ln \bar{\phi}(\underline{\lambda})) = 0. \tag{26}$$

Notice that if $\bar{\phi}(\underline{\lambda}) > e^{-C(\lambda_0)}$, then equations $H(\lambda_0) = 0$ and (26) coincide.

Case 2: $\bar{\phi}(\underline{\lambda}) \leq e^{-C(\lambda_0)}$. Suppose that λ_0 satisfies $\bar{\phi}(\underline{\lambda}) \leq e^{-C(\lambda_0)}$. Moreover, suppose that $\underline{\phi}(\underline{\mu}(\lambda_0)) = e^{-C(\lambda_0)}$. By plugging $\underline{\mu}(\lambda_0)$ into (25) we obtain:

$$\ln \left(\frac{\bar{\phi}(\bar{\lambda})}{\bar{\phi}(\lambda_0)} \right) - \bar{\phi}(\underline{\mu}(\lambda_0))(1 - C(\lambda_0) - \ln \bar{\phi}(\underline{\mu}(\lambda_0))) = 0.$$

Using that $\underline{\phi}(\underline{\mu}(\lambda_0)) = e^{-C(\lambda_0)}$ we have the final equation

$$\ln \left(\frac{\bar{\phi}(\bar{\lambda})}{\bar{\phi}(\lambda_0)} \right) - e^{-C(\lambda_0)} = 0. \tag{27}$$

Notice that if $\bar{\phi}(\underline{\lambda}) \leq e^{-C(\lambda_0)}$, then equations $H(\lambda_0) = 0$ and (27) coincide. So we showed that λ_0 is indeed determined by equation $H(\lambda_0) = 0$ and therefore according to Lemma C.1 there exists a unique λ_0 such that $\underline{\mu}(\lambda_0) < \lambda_0 \leq \bar{\lambda}$. \square

At the end of this section we explicitly present the equilibrium strategies derived in the proof of Theorem 1. For all values of parameters the sender’s strategy which minimizes loss in the worst-case scenario is given by

$$G(\mu) = \begin{cases} \ln \bar{\phi}(\mu) + C(\lambda_0), & \mu < \lambda_0 \\ \ln \underline{\phi}(\mu) + 1 - \ln \underline{\phi}(\bar{\lambda}), & \mu \geq \lambda_0 \end{cases}, \tag{28}$$

with support $[\underline{\mu}(\lambda_0), \bar{\lambda}]$, where expressions for $\underline{\phi}(\cdot)$, $\bar{\phi}(\cdot)$, $C(\cdot)$, $\underline{\mu}(\cdot)$ are given by equations (16) - (18) and λ_0 solves $H(\lambda) = 0$.

Corresponding strategy of nature, which maximizes expected loss of the sender has the same support and is defined by

$$F(\lambda) = \begin{cases} \frac{B(\lambda - \underline{\mu}(\lambda_0))}{\bar{\phi}(\lambda)}, & \lambda < \lambda_0 \\ 1 - \frac{A}{\bar{\phi}(\lambda)}, & \lambda_0 \leq \lambda < \bar{\lambda} \\ 1 & \lambda \geq \bar{\lambda} \end{cases}, \tag{29}$$

where A and B are given by (22).

Uniqueness. Note, that the mixed strategy equilibrium is uniquely defined for any support by constant expected payoff condition. Lemma 7 states that the upper bound of the support is uniquely defined. Then, for a given values of the upper bound and λ_0 the lower bound is uniquely defined by equation (18). Finally, the uniqueness of λ_0 is guaranteed by Lemma (C.1).

A4: Proof of Theorem 2

Suppose that malicious nature uniformly randomizes between $\binom{n}{\lfloor \sqrt{n} \rfloor}$ possible types of the receiver. Receiver j 's prior belief about the state of the world is characterized by

$$\beta_w^j = \begin{cases} 1 - \lfloor \sqrt{n} \rfloor \delta - (n - \lfloor \sqrt{n} \rfloor) \varepsilon, & \text{if } \omega = 1, \\ \delta, & \text{if } \omega \in S^j, \\ \varepsilon, & \text{otherwise,} \end{cases}$$

where $S^j \subset \Omega$ consisting of $\lfloor \sqrt{n} \rfloor$ elements from Ω . We define²⁰

$$\delta = \frac{1}{\lfloor \sqrt{n} \rfloor n^{1/2}} \quad \text{and} \quad \varepsilon = \frac{1}{(n - \lfloor \sqrt{n} \rfloor) n^{7/2}}.$$

Suppose that there are two messages m_1 and m_2 which persuade the receiver (i.e., give expected utility higher than r), sent with probabilities $\mu_\omega(m_i)$ such that $\mu_\omega(m_1) + \mu_\omega(m_2) \leq 1$. Then, m^+ is sent with probability $\mu_\omega(m^+) = \mu_\omega(m_1) + \mu_\omega(m_2)$ and also persuades the receiver. Thus, we can focus on two-message mechanisms.

To determine the nature's choice of the outside option r , consider the expected payoff of receiver j with $S^j = \left\{ \frac{n - \lfloor \sqrt{n} \rfloor}{n}, \dots, \frac{n-1}{n} \right\}$, after receiving a message m^+ from the mechanism which sends this message if and only if $\omega \in \Omega \setminus S^j$. Then,

$$\mathbb{E}(u_R | m^+, a = 1) = \frac{1 - \lfloor \sqrt{n} \rfloor \delta - (n - \lfloor \sqrt{n} \rfloor) \varepsilon + \varepsilon \sum_{\omega \in \Omega \setminus (S^j \cup 1)} \omega}{1 - \lfloor \sqrt{n} \rfloor \delta}.$$

²⁰ Note that the receiver's prior is well-defined for all $n \geq 2$ since $\beta_1^j = 1 - \frac{1}{n^{1/2}} - \frac{1}{n^{7/2}}$ increases in n and is equal to $1 - 9/(8\sqrt{2}) > 0$ for $n = 2$, which implies that $\beta_1^j \in (0, 1)$.

Define the receiver’s outside option (for all types j) as

$$r \equiv \frac{1 - \lfloor \sqrt{n} \rfloor \delta - (n - \lfloor \sqrt{n} \rfloor) \varepsilon}{1 - \lfloor \sqrt{n} \rfloor \delta}.$$

Clearly, $\mathbb{E}(u_R | m^+, a = 1) > r$, as r is equivalent to the expected payoff in the case if all the bottom $n - \lfloor \sqrt{n} \rfloor$ were replaced with zero.

Now, consider an expected payoff of a receiver who faces a mechanism which sends message m^+ with probability $p > 0$ in a ‘cursed state’ upon receiving such message:

$$\mathbb{E}(u_R | \mu(m^+ | \omega \in S^j) > 0, m^+, a = 1) < \frac{1 - \lfloor \sqrt{n} \rfloor \delta - (n - \lfloor \sqrt{n} \rfloor) \varepsilon + \frac{n-1}{n} \delta p}{1 - \lfloor \sqrt{n} \rfloor \delta - (n - \lfloor \sqrt{n} \rfloor) \varepsilon + \delta p} \equiv \rho.$$

The right-hand-side, which we denote as ρ , is the payoff of the receiver with the most optimistic prior, i.e., with $\beta_{(n-1)/n} = \delta$, upon receiving message m^+ from a mechanism which sends this message only in $\omega = 1$ with probability 1 and in $\omega = \frac{n-1}{n}$ with probability p . Alternatively, one can interpret the right-hand-side as the upper bound of the payoff of the receiver who faces a mechanism sending m^+ at least in one of ‘cursed states’ with probability p . Now we will show that for our choice of ε and δ we have $\rho < r$ for $p = 1$, i.e., sending m^+ just in one ‘cursed state’ with probability 1 is sufficient for rejection. By plugging in our expressions we obtain that

$$\begin{aligned} \rho &= \frac{1 - 1/\sqrt{n} - 1/\sqrt{n^7} + (n - 1)/(\lfloor \sqrt{n} \rfloor \sqrt{n^3})}{1 - 1/\sqrt{n} - 1/\sqrt{n^7} + 1/\lfloor \sqrt{n} \rfloor (\sqrt{n})} < \\ &= \frac{1 - 1/\sqrt{n} - 1/\sqrt{n^7} + (n - 1)/n^2}{1 - 1/\sqrt{n} - 1/\sqrt{n^7} + 1/n}. \end{aligned}$$

Comparing the left-hand-side of this expression to r gives that $\rho < r$ whenever

$$\frac{-n^{9/2} + n^3 + n^2 + n^{3/2} + n + n^{1/2} + 1}{n^3 (n^{7/2} - n^3 + n^{5/2} - 1)} < 0.$$

The denominator of the last expression is clearly positive for $n \geq 2$. Denote the numerator:

$$\xi(n) = -n^{9/2} + n^3 + n^2 + n^{3/2} + n + n^{1/2} + 1.$$

We have that $\xi(2) = 15 - 13\sqrt{2} < 0$. Moreover, for any $n \geq 2$ we have that

$$\xi'(n) = -\frac{9}{2}n^{7/2} + 3n^2 + 2n + \frac{3}{2}n^{1/2} + 1 + \frac{1}{2}n^{-1/2} < -2n^2 + 1 < 0.$$

Thus, $\xi(n)$ is decreasing, and, therefore, is always negative, which implies that for our choice of ε and δ we have $\rho < r$, meaning that message m^+ does not persuade the receiver if sent with probability 1 in a ‘cursed state’. Thus, that there exists $\bar{p} \in (0, 1)$ such that if the sender sends message m^+ in a ‘cursed state’ with probability

$p > \bar{p}$, then the receiver rejects. Since the receiver is indifferent between adopting and rejecting for $p = \bar{p}$ we have that

$$\bar{p} = \frac{\varepsilon n(n - \lfloor \sqrt{n} \rfloor)(1 - \lfloor \sqrt{n} \rfloor \delta - (n - \lfloor \sqrt{n} \rfloor) \varepsilon)}{\delta(1 - \lfloor \sqrt{n} \rfloor \delta - n(n - \lfloor \sqrt{n} \rfloor) \varepsilon)}.$$

Note that $\lim_{n \rightarrow \infty} n \bar{p} = 0$.

Next, we show that the maximal payoff that the uninformed sender can obtain approaches 0 when n tends to ∞ . Consider two cases.

First, suppose that the sender sends message m^+ in state $\omega = 1$ with probability 1 and in some other states with probabilities not higher than \bar{p} . Then, the probability of persuading the receiver is not higher than $\alpha_1 + \sum_{\omega \in (\Omega \setminus 1)} \alpha_\omega \bar{p} = \alpha_1 + (1 - \alpha_1) \bar{p}$. Since we have that $\alpha_1 \rightarrow 0$ and $\bar{p} \rightarrow 0$ when n tends to ∞ , the probability to persuade the receiver approaches 0 when $n \rightarrow \infty$.

Second, suppose that the sender sends m^+ in $k < n$ states with probability strictly larger than \bar{p} . Then, the payoff of the uninformed sender is not larger than

$$\alpha_1 + (n - k)p + \frac{\bar{A}}{n} k \binom{n - k}{\lfloor \sqrt{n} \rfloor} / \binom{n}{\lfloor \sqrt{n} \rfloor} < \alpha_1 + (n - k)p + \frac{\bar{A}}{n} k \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right)^k \equiv \Psi.$$

If $k \leq \lfloor \sqrt{n} \rfloor$, then $\Psi \leq \alpha_1 + np + \bar{A} \frac{\lfloor \sqrt{n} \rfloor}{n} \rightarrow 0$ when n tends to ∞ . If $\lfloor \sqrt{n} \rfloor < k < n - \lfloor \sqrt{n} \rfloor$, then since $k \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right)^k$ is a weakly decreasing function for all $k \geq \lfloor \sqrt{n} \rfloor + 1$, we have that

$$\begin{aligned} \lim_{n \rightarrow \infty} \Psi &\leq \lim_{n \rightarrow \infty} \left(\alpha_1 + (n - k)p + \bar{A} \frac{\lfloor \sqrt{n} \rfloor + 1}{n} \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right)^{\lfloor \sqrt{n} \rfloor + 1} \right) \\ &= \bar{A} \lim_{n \rightarrow \infty} \frac{\lfloor \sqrt{n} \rfloor + 1}{n} e^{-\frac{\lfloor \sqrt{n} \rfloor}{n} (\lfloor \sqrt{n} \rfloor + 1)} = 0. \end{aligned}$$

We conclude that the expected payoff of the uninformed sender approaches 0 as the number of states approaches infinity.

The Bayesian sender can persuade the receiver in all states except the ‘cursed states’, as pooling all non-cursed states with the top states delivers utility r to the receiver. Thus, using the fact that nature randomises uniformly across all $\binom{n}{\lfloor \sqrt{n} \rfloor}$ choices, the expected payoff of Bayesian sender can be written as

$$\alpha_1 + (1 - \alpha_1) \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right)$$

which approaches 1 as $n \rightarrow \infty$.

Thus, as the number of states approaches infinity, the payoff of the Bayesian sender approaches 1, while the payoff of the uninformed sender approaches 0, thus $\bar{L} \rightarrow 1$.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s00199-025-01673-1>.

Acknowledgements We are grateful to Ian Ball, Dirk Bergemann, Rahul Deb, Piotr Dworczak, Yingni Guo, Peter Klibanoff, Ludwig Sinander, Karl Schlag, Nicolas Schutz, Thomas Tröger, Terence Tsui, Karolina Vocke, Haifeng Xu, seminar audiences in Bonn, Copenhagen, Mannheim, Oxford, St. Petersburg, Vienna and conference audiences at ESEM 2022 (Milan), Transatlantic Theory Workshop (Northwestern University), ASSET 2022 (Rethymno), 3rd Durham Economic Theory Conference, EARIE (2024) for their useful comments. The support of the Deutsche Forschungsgemeinschaft through PE 813/2-2 and CRC TR 224 (projects B03 and B04) is gratefully acknowledged.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alonso, R., Camara, O.: Bayesian persuasion with heterogeneous priors. *J. Econ. Theory*. **165**, 672–706 (2016)
- Babichenko, Y., Talgam-Cohen, I., Xu, H., Zabarnyi, K.: Regret-minimizing Bayesian persuasion. *Game. Econ. Behav.* **136**, 226–248 (2022)
- Bergemann, D., Morris, S.: Information design: A unified perspective. *J. Econ. Lit.* **57**(1), 44–95 (2019)
- Bergemann, D., Schlag, K.: Robust monopoly pricing. *J. Econ. Theory* **146**(6), 2527–2543 (2011)
- Bergemann, D., Schlag, K.: Pricing without priors. *J. Eur. Econ. Assoc.* **6**(2–3), 560–569 (2008)
- Best, J., Quigley, D.: Persuasion for the long run. *J. Polit. Econ.* **132**(5), 1740–1791 (2024)
- Duersch, P., Oechssler, J., Schipper, B.: Pure strategy equilibria in symmetric two-player zero-sum games. *Internat. J. Game Theory* **41**, 553–564 (2012)
- Dworczak, P., Pavan, A.: Preparing for the worst but hoping for the best: Robust (Bayesian) persuasion. *Econometrica* **90**(5), 2017–2051 (2022)
- Guo, Y., Shmaya, E.: The interval structure of optimal disclosure. *Econometrica* **87**(2), 653–675 (2019)
- Guo, Y., Shmaya, E.: Regret-Minimizing Project Choice. *Econometrica* **91**(5), 1567–1593 (2023)
- Guo, Y., Shmaya, E.: Robust monopoly regulation. *Am. Econ. Rev.* **115**(2), 599–634 (2025)
- Hewitt, E., Stromberg, K.: Real and abstract analysis: a modern treatment of the theory of functions of a real variable. Springer-Verlag (2013)
- Hu, J., Weng, X.: Robust persuasion of a privately informed receiver. *Econ. Theor.* **72**(3), 909–953 (2021)
- Kamenica, E.: Bayesian persuasion and information design. *Annu. Rev. Econ.* **11**, 249–272 (2019)
- Kamenica, E., Gentzkow, M.: Bayesian Persuasion. *Am. Econ. Rev.* **101**(6), 2590–2615 (2011)
- Kolotilin, A., Mylovannov, T., Zapechelnyuk, A., Li, M.: Persuasion of a privately informed receiver. *Econometrica* **85**(6), 1949–1964 (2017)
- Kosterina, S.: Persuasion with unknown beliefs. *Theor. Econ.* **17**(3), 1075–1107 (2022)
- Laclau, M., Renou, L.: Public persuasion. Working Paper (2017)
- Manski, C.F.: Statistical treatment rules for heterogeneous populations. *Econometrica* **72**(4), 1221–1246 (2004)
- Mezo, I.: The Lambert W function: its generalizations and applications. Chapman and Hall/CRC (2022)
- Parakhonyak, A., Sobolev, A.: Non-Reservation Price Equilibrium and Search without Priors. *Econ. J.* **125**(584), 887–909 (2015)
- Parakhonyak, A., Vikander, N.: Information design through scarcity and social learning. *J. Econ. Theory* **207**, 105586 (2023)
- Sapiro-Gheiler, E.: Persuasion with ambiguous receiver preferences. *Econ. Theor.* **77**(4), 1173–1218 (2024)

- Schlag, K., Sobolev, A.: Robust search: How to learn from the past without priors? Working paper (2022)
- Schlag, K., Zapechelnyuk, A.: Compromise, Don't Optimize: Generalizing Perfect Bayesian Equilibrium to Allow for Ambiguity. *J. Polit. Econ. Microecon.S* **2**(1), 77–128 (2024)
- Schlag, K., Zapechelnyuk, A.: Robust sequential search. *Theor. Econ.* **16**(4), 1431–1470 (2021)
- Stoye, J.: Axioms for minimax regret choice correspondences. *J. Econ. Theory* **146**(6), 2226–2251 (2011)
- Zapechelnyuk, A.: Optimal quality certification. *Am. Econ. Rev.: Insights* **2**(2), 161–176 (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.