

Persuasion with unknown beliefs

SVETLANA KOSTERINA

Department of Economics, University of Pittsburgh

A sender designs an information structure to persuade a receiver to take an action. The sender is ignorant about the receiver's prior, and evaluates each information structure using the receiver's prior that is the worst for the sender. I characterize the optimal information structures in this environment. I show that there exists an optimal signal with two realizations, characterize the support of the signal realization recommending approval, and show that the optimal signal is a hyperbola. The lack of knowledge of the receiver's prior causes the sender to hedge her bets: the optimal signal induces the high action in more states than in the standard model, albeit with a lower probability. Increasing the sender's ignorance can hurt both the sender and the receiver.

KEYWORDS. Bayesian persuasion, robust mechanism design.

JEL CLASSIFICATION. D8.

1. INTRODUCTION

When trying to persuade someone, one finds it useful to know the beliefs the target of persuasion holds. Yet often such beliefs are unknown to the persuader. How should persuasion be designed when knowledge of prior beliefs is limited?

The following example illustrates the problem. The prosecutor decides how to collect evidence to convince a judge or a jury to convict a defendant. Kamenica and Gentzkow (2011) discuss this application under the assumption that the judge's prior belief about the defendant is known to the prosecutor. However, prosecutors may have to collect evidence and prepare their arguments before knowing which judge will be assigned to the case and before knowing the composition of the jury if one is involved. The identity of the key decision makers in trials is often unknown *ex ante*: in a variety of circumstances judges are assigned to cases randomly (see, for example, Depew, Eren, and Mocan (2017)) and citizens are randomly chosen to appear for the jury duty. Furthermore, even when the identities of these decision makers are revealed, the prosecutor may still have limited information about their beliefs.

Svetlana Kosterina: svk14@pitt.edu

I am indebted to my advisor, Wolfgang Pesendorfer, for his encouragement and support in developing this project. I am also grateful to Faruk Gul and Germán Gieczewski for insightful conversations. For helpful comments, I would like to thank David Dillenberger, Pietro Ortoleva, Kris Ramsay, Evgenii Safonov, Juan Pablo Xandri, Leeat Yariv, the audiences at the 2019 meetings of the American Economic Association (Atlanta, Georgia), the 29th International Conference on Game Theory at Stony Brook University (2018), the 2019 Midwest Economic Theory Conference (Bloomington, Indiana), Young Economists Symposium 2019 (Columbia University), Canadian Economic Theory Conference 2020, and seminar audiences at Princeton, Pittsburgh, Cornell and Vanderbilt.

In this setting, the sender (e.g., a prosecutor) designs an experiment (i.e., a way of collecting evidence). The receiver (e.g., a judge) can take one of two actions (e.g., convict or acquit). The sender wishes to convince the receiver to take the high action in all states. Thus the prosecutor aims to convince the judge to convict the defendant regardless of whether the defendant is guilty.¹ The receiver takes the action desired by the sender only if the receiver's expectation of the state given her information is above a threshold and takes the other action otherwise. We call this threshold a *threshold of doubt*. In line with this reasoning, the judge only convicts the defendants she believes to be guilty.

The main economic problem I consider that is illustrated by the example above is the problem of persuasion when the prior belief of the receiver is unknown to the sender. My model isolates the impact of the sender's ignorance about the receiver's prior by focusing on a setting where no other entities are ignorant and the sender is not ignorant about any other elements of the game. That is, the sender and the receiver (believe that they) know the distribution of the state (though their beliefs may differ) and the information structure the sender commits to. For instance, the prosecutor knows how likely the defendant is to be guilty and everyone involved understands how evidence is obtained.

In the standard Bayesian persuasion model (Kamenica and Gentzkow (2011), Kolotilin (2015)), the sender and the receiver have a common prior belief about the state. An optimal signal in that model recommends the high action with probability 1 in all states above a threshold and recommends the low action with probability 1 in all states below the threshold. We call this threshold a *threshold of action*. The threshold of action is below the threshold of doubt, so that the receiver takes the high action on a greater range of states than he would under complete information. If the sender and the receiver have commonly known heterogeneous priors (Alonso and Camara (2016a)), the optimal signal is not necessarily threshold but is still partitional: the high action is recommended either with probability 1 or with probability 0 given a state. As the results in this paper will establish, when the receiver's beliefs are unknown, the optimal signal is very different.

I model the sender's ignorance by assuming that the sender believes that the receiver's prior is chosen from a set of priors to minimize the sender's payoff. The sender has a known prior over the states and designs an experiment to maximize her payoff in the worst case scenario. This approach is in the spirit of robust mechanism design which studies mechanisms that are robust to the informed party's information and independent of the higher-order beliefs. One reason why a sender may focus on the worst case scenario is ambiguity aversion (Ellsberg (1961)), which may be especially relevant when the process by which uncertainty resolves is unknown or poorly understood. In terms of the trial example, the kind of uncertainty that a prosecutor faces about the judge's prior is a plausible candidate for the kind of uncertainty that is likely to generate ambiguity aversion.

I focus on the case where the sender knows that the receiver's prior assigns no less than a certain probability to each state but is otherwise ignorant about the receiver's

¹We can modify the model to allow the prosecutor to not want to convict the defendants who are completely innocent and our results would still go through.

prior. Formally, the set of the receiver's priors is all priors that put on each state w at least $(1 - a)g(w)$ for some $a \in (0, 1)$ and a density g . This set of priors has the advantage of allowing me to flexibly model the lack of the sender's knowledge about the receiver's prior. Because the higher a is, the smaller the mass the receiver's prior has to put on each state, a measures the sender's ignorance. If g is higher on one subset of states than on the other, then the sender has more information about the receiver's prior on the first subset.

The main contribution of this paper is a characterization of the optimal information structure in a model of persuasion when the receiver's beliefs are unknown. First, I show that there exists an optimal information structure with only two signal realizations. Second, I show that there is an *importance index* that determines which states are in the support of the signal realization recommending the high action. Third, I provide a formula for the probability that the high action is recommended on the support of the signal, showing that the optimal signal is a *hyperbola*. Fourth, I analyze comparative statics on the optimal information structure with respect to the sender's knowledge of the receiver's prior.

I show that a state w is in the support of the signal realization recommending the high action if the importance index $I(w) = f_s(w)/(g(w)(w^* - w))$ exceeds a threshold, where f_s is the density of the sender's prior and w^* is the threshold of doubt. Thus the optimal signal is more likely to induce the high action with a strictly positive probability in a state if the sender's prior assigns a high probability to the state, the sender is more ignorant about the probability of this state, and the state is close to the threshold of doubt.

I provide results showing that full support is a robust property of the signal chosen by an ignorant sender. In particular, I establish that if either the sender is sufficiently ignorant or the sender's knowledge is detail-free (by which we mean the sender knows that the receiver's prior may put probability 0 on any subset of states in which the sender and the receiver disagree about the optimal action), then the optimal signal recommends the high action with a strictly positive probability in every state.

I provide comparative statics on the optimal information structure with respect to the sender's knowledge of the receiver's prior. I show that the more ignorant the sender is, the more she hedges her bets and spreads out on the states the probability with which the high action is recommended. Formally, if we increase a , thereby decreasing the weight $(1 - a)g(w)$ that the receiver's prior has to put on each state w , then the support of the optimal signal expands, so that the high action is recommended in more states, but the probability with which the high action is recommended decreases.

The results thus change the way we think about Bayesian persuasion: unlike the intuition in the standard model, it is not optimal to pool all sufficiently high states together and give up on persuading the receiver in the lower states. Instead, the sender must allow persuasion to fail with some probability on some of the high states and is able to persuade the receiver with a positive probability on the low states. The model thus makes clear the impact of the sender's lack of knowledge about the receiver's prior on the optimal signal: the lack of knowledge causes her to hedge her bets and spread out the probability with which the high action is recommended. Oftentimes, if the state

goes up a little, the sender is able to increase the probability with which the receiver is persuaded only by a small amount.

I next consider the welfare implications of the sender's ignorance. I show that the impact of increasing the sender's ignorance on the receiver's welfare is ambiguous: it can either benefit or hurt the receiver. Because greater ignorance always hurts the sender, this implies that the sender's ignorance about the receiver's prior can hurt *both* the sender and the receiver. I also show that the receiver strictly prefers to face an ignorant sender rather than a sender who is perfectly informed about the receiver's prior. Finally, I show that if the optimal signal recommends the high action with a strictly positive probability in every state, then greater sender ignorance benefits the receiver. Going back to the trial example, these results may have implications for the assignment of judges to cases, showing why prosecutors should be given only limited information about the judge who will be assigned to their case.

My results imply that when the receiver's beliefs are unknown, especially pernicious outcomes are possible. For instance, there are parameters under which the judge convicts even completely innocent defendants with a strictly positive probability, whereas if the receiver's prior is known (and coincides with the sender's prior), the probability that completely innocent defendants are convicted is 0. Thus a model of persuasion with unknown beliefs can rationalize the occurrence of adverse outcomes that cannot be explained by the standard model.

The final contribution of the paper lies in solving a mechanism design problem to which the revelation principle does not apply. Solving such problems tends to be challenging. I show that my model can be solved by using a fixed-point argument to define the receiver's prior that is realized after the sender chooses an information structure. Importantly, the set of possible priors of the receiver has to be sufficiently rich to admit a belief that would match the information structure chosen by the sender, and my analysis highlights the correct notion of richness for this problem.

The rest of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 introduces the model. Section 4 contains examples illustrating the main results and provides intuition for the results. Section 5 presents the characterization of the optimal information structure. Section 6 provides a sketch of the proof. Section 7 contains results on the comparative statics of the optimal signal with respect to the sender's ignorance and conducts welfare analysis. Section 8 shows that full support is a robust property of the signal chosen by an ignorant sender.

2. RELATED LITERATURE

The present paper is related to two strands of literature: Bayesian persuasion and robust mechanism design. Early Bayesian persuasion papers include Brocas and Carillo (2007), Ostrovsky and Schwarz (2010), and Rayo and Segal (2010). Kamenica and Gentzkow (2011) introduce a general model of Bayesian persuasion and characterize the sender's value. Alonso and Camara (2016b) and Chan, Seher, Fei, and Yun (2019) study persuasion of voters. Alonso and Camara (2016a) consider Bayesian persuasion where the sender and the receiver have heterogeneous priors. In my robust model, if

the receiver's prior that is the worst for the sender was independent of the signal chosen by the sender (which, for example, happens if there are two states), then the sender would effectively face a persuasion problem with heterogeneous priors, which is solved in Alonso and Camara (2016a). In general, however, there is not a worst prior that is independent of the signal, so the method of Alonso and Camara (2016a) does not apply.

Kolotilin (2017) and Kolotilin, Tymofiy, Andriy, and Ming (2017), among others, consider models of Bayesian persuasion with an infinite state space. Guo and Shmaya (2019) also consider persuasion of a privately informed receiver and show that the optimal mechanism takes the form of nested intervals. Perez-Richet (2014) and Hedlund (2017) consider a Bayesian persuasion model with a privately informed sender. Ely (2017), among others, studies a dynamic Bayesian persuasion model.

One strand of the Bayesian persuasion literature studies models where the distribution of the state is endogenous: after observing the information structure chosen by the sender, the agent can take an action affecting the distribution of the state. Variants of this problem are examined in Rodina (2017), Rosar (2017), Boleslavsky and Kim (2018), and Perez-Richet and Skreta (2018). Whereas these papers look at settings where either the sender or the receiver can manipulate the state (e.g., a student can exert effort in studying), my paper looks at settings without manipulation (e.g., a judge cannot manipulate the evidence and the prosecutor cannot present evidence that has not been obtained legally).

Several papers, written simultaneously with and independently from the present paper, consider models related to ambiguity in Bayesian persuasion. Laclau and Renou (2017) consider a model of publicly persuading receivers with heterogeneous priors under the unanimity rule. Their model is equivalent to a model of persuading a single receiver who has multiple priors where, after the sender commits to an information structure and *after* a signal is realized, the receiver's prior is chosen to minimize the sender's payoff. In contrast, in the present paper the receiver's prior is chosen after the sender commits to an information structure but *before* a signal is realized. The difference has several important implications. First, unlike my model, the model of Laclau and Renou (2017) has a concave closure characterization. Second, my model has the interpretation of the sender not knowing the receiver's beliefs, while the model of Laclau and Renou (2017) does not have this interpretation and is instead suitable for settings where a committee is voting on an issue and unanimous support of all members is needed.

A paper by Hu and Weng (2018) features a sender persuading a receiver who will see private information unknown by the sender. The sender believes that the private information will be chosen to minimize the sender's payoff. One difference from the present paper is that in Hu and Weng (2018), the sender is ignorant about the receiver's private information, whereas in my paper she is ignorant about the receiver's prior. This means that the applications the papers address are different: my paper addresses applications where the information environment of the receiver is controlled by the sender (e.g., the jury cannot consider any information not given to it in the trial), while the paper by Hu and Weng addresses applications where the receiver can see external information (e.g., a consumer can obtain information about a product not just via advertisement, but also by talking to friends). Moreover, Hu and Weng (2018) solve for the optimal signal in the

example with two states (and two actions), whereas I characterize the optimal signal with a continuum of states.² Dworczak and Pavan (2020), written subsequent to my paper, considers a sender who is ignorant about information that receivers may get, as well as about the equilibrium the receivers will play, and characterizes properties of robust solutions.

The literature on robust mechanism design studies the design of optimal mechanisms when the designer does not know the distribution of agents' types and designs a mechanism to maximize his utility in the worst case scenario (see, for example, Chassang (2013), Carroll (2015, 2017), and Carrasco, Vitor, Nenad, Matthias, Paulo, and Humberto (2017)). To the best of my knowledge, the present paper is the first one to consider a model of robust Bayesian persuasion in which the prior belief of the receiver is unknown to the sender.

3. MODEL

3.1 Payoffs

The state space is an interval $[l, h]$ such that $l > 0$.³ The sender's preferences are state-independent. The sender gets a utility of $u(e)$ if the receiver's expected value of the state given the receiver's prior and the signal realization is e . The sender's utility function u is

$$u(e) = \begin{cases} 0 & \text{if } e \in [l, w^*), \\ 1 & \text{if } e \in [w^*, h]. \end{cases}$$

The model can be interpreted as one where the receiver can take one of the two actions, 0 or 1, the payoff to taking action 0 is 0, and the payoff to taking action 1 is linear in the state. Here the receiver takes action 1 if and only if his expectation of the state is weakly greater than w^* . I will call w^* the threshold of doubt.

3.2 Priors

The sender has a known prior over the states, while the receiver has a set of priors. This is in contrast to the standard model of Bayesian persuasion, where the sender and the receiver have a known common prior. Let φ denote the set of all cumulative distribution functions (CDFs) of probability measures on $[l, h]$. The sender's prior is a probability measure on $[l, h]$ with a CDF F_s . I assume that F_s admits a density f_s that is C^1 .⁴ For the sake of convenience, I also assume that f_s has full support.⁵

²See also Beauchêne, Li, and Li (2019), who consider a model in which the sender and the receiver share a common prior, the sender can commit to ambiguous information structures, and both the sender and the receiver are ambiguity averse. Their model cannot be interpreted as one where the sender is ignorant.

³The assumption that $l > 0$ is without loss of generality and is made for convenience.

⁴Some of the results in this paper, including the result that there exists an optimal signal with only two realizations, do not require this assumption.

⁵All results generalize to the case when the sender's prior does not have full support.

The set of receiver's priors is indexed by a reference prior G and an ignorance index $a \in [0, 1]$. The reference prior G admits a C^1 density g . For convenience, in most of the paper I also assume that g has full support. The receiver's set of priors is⁶

$$C_{ag} = \{F \in \varphi : \mu_F(A) \geq (1-a)\mu_G(A) \text{ for all } A \in \mathcal{B}([l, h])\}.$$

That is, the receiver's set of priors is all priors that put on each (measurable) set A a mass at least as large as the mass the reference prior G puts on A , scaled down by $1-a$. One way in which we can interpret C_{ag} is that the receiver has access to an estimate of the distribution of the state but does not fully trust it and believes that the state is drawn from this distribution with probability $1-a$ and from some other distribution with probability a . Moreover, the sender knows the estimate but not the other distribution the receiver uses. For a related way to interpret C_{ag} , imagine that a large sample of realizations of the state has been observed. Proportion $1-a$ of the sample was observed by both the sender and the receiver, and proportion a was observed by the receiver only.

To understand the assumption on the set of priors, consider a version of the model in which the state space is discrete. Then the receiver's set of priors consists of all priors that assign a probability of at least $(1-a)g(w)$ to each state w . Thus the sender knows that the receiver believes that the probability of each state w is at least $(1-a)g(w)$, but does not know what exactly this probability is. The fact that $a \in (0, 1)$ implies that $\int_{[l, h]} (1-a)g(w)dw < 1$, which ensures that the receiver's prior is not completely pinned down by this requirement.

Observe that $a = 1 - (1-a) \int_{[l, h]} g(w)dw$ is the difference between 1 and the mass put on the state space $[l, h]$ by the measure $(1-a)\mu_G$. Thus the ignorance index a measures the sender's ignorance: the larger a is, the more ignorant the sender is. In particular, if $a = 0$, so that there is no ignorance, then the set of the receiver's priors collapses to just one prior—and this is the reference prior G —while if $a = 1$, so that there is complete ignorance, then the set of the receiver's priors includes all priors.

I assume that $\int_{[w^*, h]} g(w)dw > 0$ and $(1-a) \int_{[l, h]} wg(w)dw + al < w^*$. These assumptions ensure that the sender's problem is nontrivial. In particular, the assumption that $\int_{[w^*, h]} g(w)dw > 0$ ensures that there exists a feasible information structure that induces the receiver to take action 1 with a strictly positive probability. The assumption that $(1-a) \int_{[l, h]} g(w)wdw + al < w^*$ ensures that if no information is provided, then the receiver will take action 0.

3.3 Information structures and evaluation of payoffs

The order of moves is as follows. First, the sender commits to an information structure π .⁷ Next, the receiver's prior $F \in C_{ag}$ is chosen to minimize the sender's payoff.

⁶Here $\mathcal{B}([l, h])$ denotes the Borel sigma algebra on $[l, h]$ and μ_F denotes the measure corresponding to the CDF F .

⁷An information structure is a Markov kernel π . Here $\pi(\sigma|\omega)$ is the probability of signal realization σ given that the state is ω . Formally, letting $\mathcal{B}(M)$ and $\mathcal{B}([l, h])$ denote the Borel sigma algebras on the message space M and the state space $[l, h]$, respectively, a Markov kernel π is defined as a mapping $\pi : [l, h] \times \mathcal{B}(M) \rightarrow [0, 1]$ such that for every $\omega \in [l, h]$, $B \mapsto \pi(B|\omega)$ is a probability measure on M and for every $B \in \mathcal{B}(M)$, $\omega \mapsto \pi(B|\omega)$ is $\mathcal{B}([l, h])$ -measurable. See Pollard (2002) for more details.

Then the state is realized (from the sender's perspective, the state is drawn from the distribution F_s). After this, the signal is realized according to the information structure π . Then, having seen a signal realization σ , the receiver forms an expectation of the state given that the receiver's prior is F and that the information structure is π .⁸

I restrict my attention to information structures that have a finite number of signal realizations given each state. I conjecture that my results hold for all information structures.

If the sender chooses an information structure π and a receiver with a prior F sees signal realization σ , then the receiver's expectation of the state is $E_{F\pi}[\omega|\sigma]$. Then, if the sender chooses an information structure π and the receiver's prior is $F \in C_{ag}$, the sender's payoff is

$$\int_{[l, h]} \sum_{\sigma} \mathbb{1}_{E_{F\pi}[\omega|\sigma] \geq w^*} \pi(\sigma|w) dF_s(w).$$

Recall that the prior of the receiver is chosen from the set C_{ag} to minimize the sender's payoff. Thus the sender's payoff from choosing an information structure π is

$$U(\pi) = \min_{F \in C_{ag}} \int_{[l, h]} \sum_{\sigma} \mathbb{1}_{E_{F\pi}[\omega|\sigma] \geq w^*} \pi(\sigma|w) dF_s(w)$$

and the sender's equilibrium payoff is

$$\sup_{\pi} \min_{F \in C_{ag}} \int_{[l, h]} \sum_{\sigma} \mathbb{1}_{E_{F\pi}[\omega|\sigma] \geq w^*} \pi(\sigma|w) dF_s(w). \quad (1)$$

4. EXAMPLES AND INTUITION

In this section, I present simple examples that illustrate my general results and provide some intuition for why the results hold.

4.1 Preliminary examples

First consider two extreme benchmarks: a completely ignorant sender and a sender who is fully informed. If the sender is completely ignorant, any prior of the receiver is possible. In particular, priors that put probability 1 on states below the threshold w^* are possible. Receivers with such priors will take action 0 no matter what information is provided, so the sender gets her lowest possible payoff of 0. If the sender is fully informed (and the prior for the receiver is the same as the sender's prior), it is known that the solution to the sender's problem has a threshold structure: there is an optimal signal p with two signal realizations, σ_1 and σ_0 , such that $p(\sigma_1|w) = 1$ for $w \in [w', h]$, $p(\sigma_1|w) = 0$ for $w \in [l, w')$ for some threshold $w' \in [l, w^*)$, and the receiver takes action 1 if and only if the realized signal is σ_1 . When a prior common to the sender and the receiver is fixed,

⁸If g does not have full support, then we need to specify how the receiver updates his beliefs after observing signal realizations that have zero probability under the receiver's prior. In this case, reasonable updating rules such as the receiver not changing his prior belief or putting mass 1 on the lowest state in the support of the signal realization ensure that the results in the present paper hold.

recommending action 1 in higher states below the threshold of doubt w^* yields a strictly greater benefit to the sender than recommending action 1 in lower states, so the sender recommends action 1 in all sufficiently high states such that the receiver's expectation upon seeing the signal realization σ_1 is exactly w^* .

Next consider an example with only two states. Suppose that the two states are 0 and 1, the sender wants the receiver to take action 1 no matter what the state is, and the receiver only takes action 1 if she believes that the probability that the state is 1 is higher than some threshold p^* . Because the state space is binary, we can identify a prior with the probability that the state is 1. Note that the priors can be ranked such that, given priors p_1 and p_2 with $p_1 < p_2$ and a signal realization from a fixed information structure, the posterior under p_1 is lower than the posterior under p_2 . Because higher posterior beliefs yield a higher payoff to the sender, this means that prior p_1 is worse for the sender than prior p_2 , so the receiver's priors can be ranked according to their value to the sender independently of the information structure. Because the lowest prior is the worst for the sender, the information structure optimal for an ignorant sender is the one that maximizes the sender's payoff given the lowest feasible prior of the receiver. Because this information structure is a solution to the standard Bayesian persuasion problem with known priors of the sender and the receiver, we know that the solution has a concave closure characterization.

4.2 Three-state example

Finally, consider an example with three states. Suppose that the states are 0, 1, and 2, the sender wants the receiver to take action 1 no matter what the state is, and the receiver only takes action 1 if her expectation is above $w^* = 1.6$. Suppose also that there are two possible prior beliefs for the receiver, assigning the probabilities $(0, 1/2, 1/2)$ and $(1/2, 0, 1/2)$, respectively, to each state. To see that the receiver's prior that minimizes the sender's payoff depends on the information structure the sender chooses, consider two information structures. Under the first one, signal σ_1 is sent with probability 1 in states 1 and 2, and signal σ_0 is sent with probability 1 in state 0. Under the second one, σ_1 is sent with probability 1 in states 0 and 2, and σ_0 is sent with probability 1 in state 1. Because action 0 must be taken after signal σ_0 , the receiver's prior that minimizes the expected state conditional on signal σ_1 minimizes the sender's payoff. Under the first information structure, this is the prior that maximizes the probability of state 1, and under the second one, this is the prior that maximizes the probability of state 0. Thus the worst prior is $(0, 1/2, 1/2)$ under the first information structure and $(1/2, 0, 1/2)$ under the second one.

Suppose for now that the sender can only choose information structures with two realizations such that after the realization σ_1 receivers with all possible priors take action 1. Then after σ_1 , both the receiver with prior $(0, 1/2, 1/2)$ and the receiver with prior $(1/2, 0, 1/2)$ have to take action 1. If the prior is $(0, 1/2, 1/2)$, the receiver's expectation after σ_1 is $(\pi(\sigma_1|1) + 2)/(\pi(\sigma_1|1) + 1)$, and if the prior is $(1/2, 0, 1/2)$, the receiver's expectation after σ_1 is $2/(\pi(\sigma_1|0) + 1)$. Setting both expectations equal to $w^* = 1.6$ yields $\pi(\sigma_1|1) = 2/3$ and $\pi(\sigma_1|0) = 1/4$.

We can make several observations about this solution. First, in contrast to the standard model of Bayesian persuasion, action 1 is induced in states 0 and 1 with a strictly interior probability. Second, since $\pi(\sigma_1|0) < \pi(\sigma_1|1)$, the sender is able to induce action 1 with a lower probability in lower states, where the receiver's loss from the wrong action is larger. Third, we can show that the points $\pi(\sigma_1|0)$ and $\pi(\sigma_1|1)$ lie on a hyperbola (here hyperbola refers to the probability of action 1 as a function of the state).

To explore the effects of increasing sender's ignorance, now suppose that, in addition to the priors above, receiver's priors $(2/3, 0, 1/3)$ and $(1/4, 1/4, 1/2)$ are also possible. Note that if a receiver with prior $(2/3, 0, 1/3)$ takes action 1, then so does a receiver with prior $(1/2, 0, 1/2)$. Moreover, if receivers with priors $(0, 1/2, 1/2)$ and $(1/2, 0, 1/2)$ take action 1, then so does a receiver with prior $(1/4, 1/4, 1/2)$. Then we just need to ensure that the expectation of a receiver with prior $(2/3, 0, 1/3)$ after σ_1 , which is $2/(2\pi(\sigma_1|0) + 1)$, is equal to 1.6. This yields $\pi(\sigma_1|0) = 1/8$. We thus see that increasing sender's ignorance leads the sender to induce action 1 with lower probability in states where it was recommended with a positive probability. I refer to this as the *probability effect* below.

Suppose next that we decrease sender's ignorance by removing all receiver's priors but $(1/4, 1/4, 1/2)$. To ensure that the receiver takes action 1 after σ_1 , we need $(\pi(\sigma_1|1) + 4)/(\pi(\sigma_1|0) + \pi(\sigma_1|1) + 2) \geq 1.6$, which is $\pi(\sigma_1|0) \leq 1/2 - 3/8\pi(\sigma_1|1)$. If $\pi(\sigma_1|1) = 1$, the constraint is $\pi(\sigma_1|0) \leq 1/8$, while if $\pi(\sigma_1|1) = 0$, the constraint is $\pi(\sigma_1|0) \leq 1/2$. The sender will want to set $\pi(\sigma_1|1) = 0$ and $\pi(\sigma_1|0) = 1/2$ if the probability that the sender's prior assigns to state 0 is high enough. Thus decreasing sender's ignorance can cause the sender to induce action 1 in fewer states. I refer to this as the *support effect* below.

4.3 Intuition for the hyperbola

I next generalize the three-state example I discussed above to provide further intuition for the hyperbolic functional form of the optimal signal. Suppose that, as in the general model, the state space is an interval. Suppose also that any receiver's prior has to put a mass of at least $1 - a$ on some state $\alpha > w^*$, and, subject to this constraint, any prior is allowed.

Consider an information structure with two realizations, σ_1 and σ_0 , satisfying $\pi(\sigma_1|\alpha) = 1$, and the receiver's prior that puts a mass of $1 - a$ on α and a mass of a on some state ω below the threshold w^* . Then the receiver's expectation conditional on seeing σ_1 is $E[w|\sigma_1] = (\omega a \pi(\sigma_1|\omega) + \alpha(1 - a))/(a \pi(\sigma_1|\omega) + 1 - a)$. In order for the receiver to take action 1 after seeing σ_1 , we need $E[w|\sigma_1] \geq w^*$, which turns out to be equivalent to $\pi(\sigma_1|\omega) \leq (1 - a)(\alpha - w^*)/(a(w^* - \omega))$. Note that the bound $(1 - a)(\alpha - w^*)/(a(w^* - \omega))$ is a hyperbola in ω . This reveals that the hyperbolic functional form of the optimal signal arises from the way conditional expectations are calculated.

Note that the receiver's prior is chosen after the sender chose the information structure and a prior putting a mass of a on any state ω below the threshold w^* is possible. Thus if the probability of σ_1 exceeds the bound $(1 - a)(\alpha - w^*)/(a(w^* - \omega))$ at any state below the threshold, then a receiver's prior putting mass a on this state would ensure

that the receiver never takes action 1. Therefore, $\pi(\sigma_1|\omega)$ must be below the bound in all states. On the other hand, the sender's payoff is increasing in the probability that action 1 is taken, implying that it is best for the sender to maximize $\pi(\sigma_1|\omega)$ subject to the constraint that it be below the bound. Thus setting $\pi(\sigma_1|\omega)$ equal to the bound in all states yields the signal that is optimal (in the class of all signals with two realizations such that after one realization receivers with all possible priors take action 1).

5. MAIN RESULT

This section characterizes the optimal signal. I show that there is an optimal signal with two realizations such that after one of the signal realizations receivers with all possible priors take action 1. I will use σ_1 to denote this signal realization and I will refer to this signal realization as the realization that recommends approval. Under this signal, the probability of the signal realization recommending the high action is 1 above the threshold w^* and is a hyperbola on the support below w^* . The support of this signal realization below w^* is the set of all states such that an importance index exceeds a threshold t . There is a trade-off between adding more states to the support (by decreasing the threshold t) and recommending the high action with a greater probability (by increasing the constant c scaling the hyperbola), and the optimal signal balances these considerations. I start by defining a class of distributions over the receiver's actions that have the above form.

Letting $I(w) = f_s(w)/(g(w)(w^* - w))$ and $\underline{t} = \min_{w \in [l, w^*]} I(w)$, I define a class of distributions S^{tc} over the receiver's actions as follows: given constants $t \geq \underline{t}$, $c \geq 0$, the probability of action 1 in state ω is given by

$$S^{tc}(\omega) = \begin{cases} 1 & \text{for } \omega \in [w^*, h], \\ \min\{c/(w^* - \omega), 1\} & \text{for } \omega \in \Omega(t), \\ 0 & \text{for } \omega \in [l, w^*) \setminus \Omega(t), \end{cases}$$

where

$$\Omega(t) = \{w \in [l, w^*) : I(w) \geq t\}.$$

Theorem 1 will show that, given an optimal information structure with two realizations, the support below the threshold w^* of the signal realization recommending approval is $\Omega(t)$. $\Omega(t)$ consists of all states w below the threshold w^* such that the importance index $I(w) = f_s(w)/(g(w)(w^* - w))$ is greater than some threshold t . Thus the sender is more likely to induce the receiver to approve in state w if w is more important to the sender because the sender believes that this state is very likely. The sender is also more likely to induce approval in w if the conflict of interest between the sender and the receiver is not too large because the distance $w^* - w$ between the state and the threshold of doubt w^* is small, and if the sender is more ignorant about the probability the receiver assigns to the state w —because $g(w)$ is small.

To ensure the essential uniqueness of the distribution of the high action induced by an optimal information structure, I make use of the following assumption.⁹

⁹Here μ is the Lebesgue measure.

ASSUMPTION 1. $\mu(\{w \in [l, w^*] : I(w) = c_0\}) = 0$ for all $c_0 > 0$.

Assumption 1 says that the set of states for which the importance index $I(w)$ is equal to a constant c_0 is of measure zero for all constants $c_0 > 0$.¹⁰

A signal is said to be optimal if it solves the sender's problem (1). Theorem 1 describes an optimal signal and the distribution over the receiver's actions induced by it.

THEOREM 1. *There exist unique $t \geq \underline{t}$, $c \geq 0$ such that any optimal information structure induces a distribution s over the receiver's actions satisfying $s = S^{tc}$ almost everywhere under μ_{F_s} .*

An optimal information structure inducing the distribution s is given by $\pi(\sigma_1|\omega) = s(\omega)$, $\pi(\sigma_0|\omega) = 1 - s(\omega)$ for all $\omega \in [l, h]$.

Theorem 1 says that the distribution of the receiver's actions induced by an optimal information structure is unique. Moreover, there is an optimal information structure with two realizations, σ_1 and σ_0 . The receiver takes action 1 after seeing signal σ_1 and action 0 after seeing signal σ_0 . If the state is in $[w^*, h]$, the receiver takes action 1 with probability 1. If the state is in $\Omega(t)$, the receiver takes action 1 with probability $\min\{c/(w^* - \omega), 1\}$ for some constant c . Note that on the support $\Omega(t)$, the probability of action 1 follows a hyperbola. Finally, if the state is below w^* and not in $\Omega(t)$, then the receiver takes action 1 with probability 0.

6. SKETCH OF THE PROOF

In this section, I sketch the proof of Theorem 1. The proof has two parts. The first part establishes that there exists an optimal signal with only two realizations such that after one of the realizations, receivers with all possible priors take action 1. I explain the first part of the proof in several steps. I first present a counterexample showing that if the set of the receiver's priors does not satisfy my assumptions, there may not exist a sender-optimal signal with the property described above. After this, I show via an example how allowing for a larger set of priors fixes this problem. I then generalize this example to a setting where the sender is restricted to choosing "simple" signals with "non-overlapping support" below the threshold w^* and show that for any signal of this form there is a feasible receiver's prior that renders this signal worse than the binary signal with the property described above. In this setting, I obtain a closed-form solution for this prior of the receiver. Finally, I show that when the sender can choose arbitrary signals, there still exists a prior of the receiver that ensures the sender cannot benefit from complicated signals, only now there may not be a closed-form expression for this prior. Instead, the prior is given by a fixed point.

The second part of the proof characterizes the optimal signal. There I show that the optimal signal has a hyperbolic functional form and reduce the problem of characterizing the optimal signal to the problem of computing two numbers: the threshold t and the constant of proportionality $c(t)$ for the hyperbola.

¹⁰It can be shown that a sufficient condition for Assumption 1 to be satisfied is that f_s and g are real-analytic functions on $[l, h]$, and the importance index $I(w)$ is not a constant function on $[l, w^*]$.

6.1 Binary optimal signal: Counterexample

I first show that if the set of the receiver's priors does not satisfy the assumptions I imposed on it when I introduced the model, there may not exist a binary sender-optimal information structure with the property described above. To see this, recall the example with three states I introduced earlier. There the states are 0, 1, and 2, and there are two possible priors for the receiver, assigning probabilities $(0, 1/2, 1/2)$ and $(1/2, 0, 1/2)$, respectively, to the states. Suppose also that the sender's prior assigns probabilities $(1/2, 1/2, 0)$ to each state.

Consider an information structure π' with two signal realizations, σ'_1 and σ'_2 , sent with probabilities $(1, 0, 1/2)$ and $(0, 1, 1/2)$, respectively, in each state. Note that if the receiver has the prior $(0, 1/2, 1/2)$, then she takes action 1 after σ'_1 , and if the receiver has the prior $(1/2, 0, 1/2)$, she takes action 1 after σ'_2 . Thus, regardless of which prior the receiver has, the lowest payoff the sender can get under this information structure is $1/2$.

Next consider an information structure π with two signal realizations, σ_0 and σ_1 , such that after σ_1 receivers with both priors take action 1. Then our previous discussion of this example suggests that the probabilities of σ_1 in each state are $(1/4, 2/3, 1)$. The sender's payoff under this information structure then is $11/24 < 1/2$. Therefore, the binary signal with action 1 taken by all types of receivers after one realization is strictly worse for the sender than the binary signal where receivers with different priors take action 1 after different realizations.

6.2 Binary optimal signal: Proof restricted to simple signals

The reason for the above result is that the set of receiver's priors is not "rich enough," that is, it does not contain enough priors of the right kind. To understand this claim, consider a larger set of priors that contains all convex combinations of the two original priors (here we take convex combinations of the probability mass functions). Thus a feasible prior has the form $\alpha(1/2, 0, 1/2) + (1 - \alpha)(0, 1/2, 1/2)$ for some $\alpha \in [0, 1]$.

I will show that we can find a feasible prior such that both σ'_1 and σ'_2 lead the receiver to take action 0 (which would imply that π' is inferior to π). Such a prior needs to satisfy $E_{F, \pi'}[w|\sigma'_1] < w^*$ and $E_{F, \pi'}[w|\sigma'_2] < w^*$ for $w^* = 1.6$. Consider $\alpha = 1/2$. Note that $E_{F, \pi'}[w|\sigma'_1] = 2/(2\alpha + 1) = 1 < 1.6$ and $E_{F, \pi'}[w|\sigma'_2] = (2(1 - \alpha) + 2)/(2(1 - \alpha) + 1) = 1.5 < 1.6$. Thus the prior corresponding to $\alpha = 1/2$ indeed leads the receiver to take action 0 after both signal realizations.

The above example reflects a general principle that tells us how to find a prior making π' inferior to π . Suppose now that the sender is restricted to choosing information structures with "non-overlapping support" below the threshold w^* . That is, if a signal realization has positive probability in some state below w^* , no other signal realization can have positive probability in this state.

Then $E_{F, \pi'}[w|\sigma'_1] = (2P[\sigma'_1|2]1/2)/(P[\sigma'_1|0]P[0] + P[\sigma'_1|2]1/2)$ and $E_{F, \pi'}[w|\sigma'_2] = (P[\sigma'_2|1]P[1] + 2P[\sigma'_2|2]1/2)/(P[\sigma'_2|1]P[1] + P[\sigma'_2|2]1/2)$. Compare this to the receiver's expectation after σ_1 under π given the priors $f_1 = (1/2, 0, 1/2)$ and $f_2 = (0, 1/2, 1/2)$,

where $E_{F_1, \pi}[w|\sigma_1] = (21/2)/(P[\sigma_1|0]1/2 + 1/2)$ and $E_{F_2, \pi}[w|\sigma_1] = (P[\sigma_1|1]1/2 + 21/2)/(P[\sigma_1|1]1/2 + 1/2)$.

Now suppose that $P[\sigma'_1|0] = P[\sigma_1|0]$ and $P[\sigma'_2|1] = P[\sigma_1|1]$, and let us choose a prior such that these expectations are equal, i.e., $E_{F_1, \pi}[w|\sigma_1] = E_{F, \pi'}[w|\sigma'_1]$ and $E_{F_2, \pi}[w|\sigma_1] = E_{F, \pi'}[w|\sigma'_2]$. Setting $P[0] = P[\sigma'_1|2]1/2$ and $P[1] = P[\sigma'_2|2]1/2$ accomplishes this goal. Observe that we have constructed this prior by redistributing the mass $1/2$ that the receiver's prior must put on states below w^* according to the probabilities $P[\sigma'_1|2]$ and $P[\sigma'_2|2] = 1 - P[\sigma'_1|2]$ with which signal realizations inducing action 1 with support on states 0 and 1, respectively, are sent in state 2 above w^* . Observe also that this is a feasible prior with $\alpha = P[\sigma'_1|2]$.

The above implies that if the receiver's expectation is w^* under π , we can find a feasible prior such that, whenever the sender chooses a signal with $P[\sigma'_1|0] > P[\sigma_1|0]$ or $P[\sigma'_2|1] > P[\sigma_1|1]$ (i.e., recommending action 1 with a larger probability in some state), the receiver with this prior will take action 0. This means that the sender cannot induce action 1 under π' with a larger probability than under π .

6.3 Binary optimal signal: Fixed-point proof

The previous section shows that having receivers with different priors take action 1 after different signal realizations cannot benefit the sender if she can only choose signals with “non-overlapping support” below w^* . To complete the proof, we have to show that the sender cannot benefit even if she can choose arbitrary information structures.

To show this, I use the same logic as in the examples above. In particular, I first observe that a sender-optimal signal π in the class of signals with two realizations (such that after one realization receivers with all possible priors take action 1) results in action 1 taken with a certain probability in each state ω . We can use $p(\omega; \pi)$ to denote this probability. I next show that, for each information structure, there exists a feasible receiver's prior such that the probability of approval in each state ω is bounded above by $p(\omega; \pi)$. In greater detail, if a signal recommended action 1 with probability greater than $p(\omega; \pi)$ in ω , the receiver with this prior would not follow the recommendation after some signal realizations.

In the example in the previous section, I was able to obtain a closed-form expression for the receiver's prior. In general, it is not possible to solve for the receiver's prior in closed form. Nevertheless, I show that we can define the receiver's prior implicitly as a fixed point. We then know that a receiver's prior with the required property exists, because a fixed-point theorem ensures existence.

To simplify the sketch of the proof, suppose that $g(\omega) = 0$ for all $\omega \leq w^*$, so that the receiver's prior can put the mass a on any state below w^* . Suppose also that there is only one state 2 above w^* .

I show that the following prior works. The prior places a mass of $(1 - a)\mu_G$ on states $[l, h]$ and, in addition, for each signal realization σ_i , places a mass of

$$v_\omega^i \pi(\sigma_i|2)a$$

on each state $\omega < w^*$. The weights ν satisfy $\nu_\omega^i \in [0, 1]$ and $\sum_{\omega < w^*} \nu_\omega^i = 1$. The weights are defined as a fixed point of the mapping given by

$$\nu_w^i = \frac{\pi(\sigma_i|w)(w^* - w) \left[\sum_k \nu_w^k \pi(\sigma_k|2) \right]}{\sum_{\omega < w^*} \pi(\sigma_i|\omega)(w^* - \omega) \left[\sum_k \nu_\omega^k \pi(\sigma_k|2) \right]}.$$

We can interpret ν_w^i as the proportion of the receiver's loss from taking the wrong action after signal realization σ_i in state w relative to the receiver's total loss after σ_i . This proportion is larger if the state w is farther below the threshold, if action 1 is recommended in w with a higher probability, and if the prior probability of w is higher.

Intuitively, we need to find a prior that redistributes the maximum probability mass a that can be put on states below w^* to make all signal realizations that are sent with excessively large probability result in action 0. Thus $\pi(\sigma_i|2)a$ identifies the portion of the mass a that “will be used” to make signal realization σ_i result in action 0. This mass $\pi(\sigma_i|2)a$ should be allocated to the states below w^* in which σ_i has positive probability. If there is only one such state for each σ_i , then $\nu_\omega^i = 1$ for this state and $\nu_\omega^i = 0$ for all other states, so we obtain the simplified case from the previous section. The weights ν_ω^i tell us exactly how the probability mass should be allocated among all states in which σ_i has positive probability.

How do we know what the weights ν should be? Analogously to what we did in the preceding examples, we are choosing ν to ensure that the probability with which action 1 is taken in each state under an arbitrary information structure is no larger than the probability of action 1 under the optimal information structure described in Theorem 1. We guess weights ν that accomplish this goal and then verify that they work. Importantly, even though to specify the prior we only need to specify as many probabilities as there are states (when the state space is finite), we have to specify the weights for each signal realization-state pair. This is needed to ensure that we are able to use a fixed-point theorem to argue that the weights are well defined.

6.4 How do we compute the optimal c and t ?

Because there is a sender-optimal signal with two realizations such that after one realization, receivers with all possible priors take action 1 and, as Section 4.3 explains, under such signal, the probability of action 1 follows a hyperbola, I have reduced the problem of computing the optimal signal to computing two numbers: the threshold t and the constant of proportionality c . Next I explain how these two numbers are computed.

Note that decreasing the threshold t or increasing the constant c decreases the expected state conditional on the approval recommendation. Then for each threshold t , there exists a maximal constant of proportionality $c(t)$ such that the receiver is willing to follow the approval recommendation. Next, for every pair $(t, c(t))$, I compute the marginal change in the sender's payoff from expanding the support. This marginal change must be zero at the unique optimum: if it is positive, then the sender prefers to

expand the support, while if it is negative, she prefers to shrink the support. Thus we set the marginal change to zero. Together with the equation providing the formula for the constant of proportionality $c(t)$, these two equations allow us to solve for the two unknowns, t and $c(t)$.¹¹

7. COMPARATIVE STATICS AND WELFARE

In this section, I ask, “As the sender becomes more ignorant, how does the optimal signal change and how is the receiver’s welfare affected?” Proposition 1 shows that if we take two sets of priors such that the sender is more ignorant under the set of priors C_{ag} than $C_{a'g}$, then the support of the optimal signal under C_{ag} contains the support under $C_{a'g}$ and the probability of approval on the support is higher for the second set. Figure 1 illustrates the comparative statics described in Proposition 1. Here I let t_a denote the optimal t at the ignorance index a .

PROPOSITION 1. *If $0 < a' < a < 1$ and the optimal signal at a does not have full support, then*

$$\Omega(t_{a'}) \subset \Omega(t_a) \quad \mu_{F_S}\text{-a.e.} \quad \text{and} \quad c(t_{a'}) > c(t_a).$$

We can interpret Proposition 1 as saying that increasing the sender’s ignorance has two effects: the support effect and the probability effect. The support effect is that the




FIGURE 1. Optimal information policy: increasing ignorance.

¹¹See the Appendix for more details.

range of states on which the sender's signal induces the receiver to take action 1 with a positive probability expands, and the probability effect is that the probability with which the receiver takes action 1 on these states decreases. Thus greater ignorance about the receiver's prior makes the sender "hedge her bets" and spread out the probability with which the high action is recommended on the states below the threshold.

Next, consider two sets of priors, C_{ag} and $C_{a'g}$ for $0 < a' < a < 1$ so that the sender is more ignorant under C_{ag} than under $C_{a'g}$. Let $V(a)$ denote the equilibrium payoff of the receiver when the set of priors is C_{ag} . Proposition 2 shows how the receiver's equilibrium payoff changes as the sender becomes more ignorant.

PROPOSITION 2. *Suppose that the receiver's prior is $F \in C_{a'g} \subset C_{ag}$ for $0 < a' < a < 1$. Then*

- (i) *If the optimal signal recommends approval in all states with a strictly positive probability under a and a' , then $V(a) \geq V(a')$. If, in addition, the reference prior g has full support, then $V(a) > V(a')$.*
- (ii) *There exist g, F_s such that $0 < V(a) < V(a')$.*

Proposition 2 first considers the case in which the optimal signal recommends the high action with a strictly positive probability in every state. The proposition shows that in this case an increase in the sender's ignorance weakly increases the receiver's payoff. Moreover, if the receiver considers all states below the threshold possible, then the increase in the receiver's payoff is strict.

The intuition for these results is as follows. An increase in the sender's ignorance has two effects: it expands the support of the signal and lowers the probability with which the high action is recommended on the support. If the signal already has full support, then only the second effect is present: more ignorant senders recommend the high action with a lower probability below the threshold. Because the receiver does not want to take the action anywhere below the threshold, this is unambiguously good for the receiver.

The second part of Proposition 2 shows that the receiver's equilibrium payoff can be strictly lower when the sender is more ignorant. Because ignorance always hurts the sender, an implication of this is that an increase in sender's ignorance can hurt *both* the sender and the receiver.

The reason that the sender's ignorance can hurt the receiver is that the sender designs the information structure to make the worst types in the set of priors just indifferent between taking the action and not. If a receiver with some prior strictly prefers to act upon seeing the signal recommending the action, so that the receiver's expectation given the signal is strictly above the threshold of action, the sender's payoff is not affected by the distance between the expectation and the threshold. Because of this, the impact of the change in the information structure due to the greater ignorance of the sender on the agents with priors that are not the worst in the set of priors that the sender considers possible can be ambiguous.

To better understand the intuition for the possible non-monotonicity of the receiver's payoff, recall the two comparative statics effects that increasing the sender's ignorance has: the support effect and the probability effect. The support effect is bad for the receiver because expanding support means that more innocent defendants are convicted. The probability effect is good for the receiver because a decreasing probability of approval on the support means that innocent defendants are convicted with a lower probability. My analysis shows that the support effect can dominate, causing greater ignorance of the sender to reduce the receiver's payoff.

One reason why the support effect can dominate is that the priors of the sender and the receiver may differ, and the sender makes the decision as to which states to add to the support based in part on her own prior F_s . Thus if, for example, the sender's prior puts a sufficiently high weight on state w , the sender will add this state to the support no matter how much she has to reduce the probability with which approval is recommended in other states. If adding w to the support reduces the probability of approval in other states a lot, then the probability effect dominates and the receiver's payoff increases, while if it reduces the probability of approval only a little bit, then the support effect dominates and the receiver's payoff decreases.

The final observation is that complete ignorance of the sender is the best possible circumstance for the receiver. This is because as the sender's ignorance a converges to 1, the sender becomes extremely cautious and recommends that the receiver take the high action only on the states where the sender and the receiver agree. Thus as the prosecutor becomes very ignorant about the judge's prior, she recommends that the judge convict the defendant if and only if the judge would convict this defendant under complete information, which is the first-best outcome for the judge.

These results may have implications for the assignment of judges to cases. For example, in order to limit the information that a prosecutor has about the judge when collecting evidence, one may want to delay revealing which judge will be assigned to the case. Prosecutors may know that the judge will be drawn from a given set but not exactly which judge from this set will be assigned. Another way to limit prosecutors' information may be by matching them with judges with whom they did not work in the past.

8. SUPPORT OF THE OPTIMAL SIGNAL

In this section, I show that recommending approval with a strictly positive probability in every state is a robust property of the signal chosen by an ignorant sender. I also consider the limits of the optimal information structure as the ignorance of the sender approaches its maximal and minimal values, respectively.

8.1 *Detail-free knowledge and full support*

I next introduce a special class of sets of the receiver's priors. Definition 1 says that the sender's knowledge is detail-free if the reference prior G puts zero mass on states below the receiver's threshold, where the sender and the receiver disagree about the optimal action.

DEFINITION 1. The sender's knowledge is said to be detail-free if $\mu_G([l, w^*)) = 0$.

The reason that a set of priors satisfying the condition in Definition 1 represents detail-free knowledge is that the reference prior G quantifies the limited knowledge that the sender has about the receiver's prior. If g varies on the states where the sender and the receiver disagree, then the sender knows more about the receiver's prior on some such states than on others: in particular, the sender knows more about the prior on the states where g is high. If, on the other hand, g is 0 on all states below the threshold, then a given probability mass can be moved freely below the threshold, and the sender is equally ignorant about the receiver's prior on all states in which there is a conflict of interest. Thus the sender only has a vague idea of how pessimistic the receiver can be: the receiver's prior can put no more than a certain mass on states below the threshold, but the sender lacks knowledge of the details of how the receiver's prior may vary on the states where there is disagreement.

Proposition 3 shows that the distribution over the receiver's actions induced by an optimal information structure is independent of the sender's prior if and only if the sender's knowledge is detail-free. That is, the distribution is independent of the sender's prior if and only if the receiver may consider all states on which he disagrees with the sender with regard to the optimal action impossible.

PROPOSITION 3. *An optimal information structure has full support and induces a distribution over the receiver's actions that is independent of the sender's prior F_s if and only if the sender's knowledge is detail-free.*

To understand the intuition, consider a simple set of priors where probability 0.1 has to be put on some state $w < w^*$ and a mass of no more than 0.5 can be put on states below w^* . Then there is a benefit to giving up and not recommending approval in state w : if the sender gives up on w , then a mass of no more than 0.4, rather than 0.5, can be put on the support below w^* of the signal realization recommending approval. The cost of giving up is that the sender would not collect the payoff from approval in state w . The sender will not want to give up on a state if her prior on that state is high enough and might want to give up otherwise. On the other hand, if the sender's knowledge is detail-free, so that a mass of no more than 0.5 can be put on states below w^* and there are no other restrictions, then there is no benefit to giving up on the states. This is because no matter which states the sender gives up on, a mass of 0.5 can still be put on states below w^* .

8.2 Maximal ignorance and full support

The next proposition describes how the support of the optimal signal behaves as the sender becomes very ignorant.

PROPOSITION 4. *As the ignorance index a goes to 1, $\Omega(t)$ converges to $[l, w^*)$. Moreover, if f_s is bounded away from zero, then $\Omega(t) = [l, w^*)$ for all $a < 1$ sufficiently close to 1.*

Proposition 4 shows that, as the ignorance of the sender approaches its maximal value of 1, in the limit, the support of the optimal signal converges to the set of all states (at the same time, the probability with which the high action is recommended below w^* converges to 0). Moreover, if the density of the sender's prior is bounded away from zero, then the signal has full support, not just in the limit, but for all sufficiently high levels of the sender's ignorance. This clarifies the sense in which recommending the high action with a strictly positive probability in all states is a robust property of the signal chosen by an ignorant sender.

The intuition for the result is that as the set of the receiver's priors converges to the whole set, the probability with which approval is recommended in the states in which the sender and the receiver disagree converges to zero. Intuitively, this is because only very convincing evidence can guard against very pessimistic priors. As the probability of approval recommendation decreases, a receiver who was willing to approve before now strictly prefers to approve. Then we can expand the support of the signal realization recommending approval a little bit, and this receiver is still willing to approve. This explains why, as the probability of approval recommendation decreases to zero, the support converges to the whole state space.

8.3 Minimal ignorance and continuity

Here I ask how the optimal information structure behaves as the ignorance of the sender converges to zero. I let π_0 denote the sender-optimal information structure when $a = 0$, so that the receiver has a commonly known prior. It can be shown that the optimal signal π_0 is partitionial, so in each state approval happens with probability 0 or 1. Moreover, the set of states where approval happens with probability 1 is $\Omega(t)$ and all states above w^* , where the threshold t is pinned down by the requirement that the receiver is indifferent after she sees the recommendation to approve. Proposition 5 establishes a continuity result as the ignorance of the sender converges to zero.

PROPOSITION 5. *As $a \rightarrow 0$, the signal that is optimal under an unknown prior converges to π_0 .*¹²

APPENDIX

A.1 Road map

The proofs are structured as follows. Lemmas 1, 2, and 3 prove that there exists a sender-optimal signal with two realizations such that after one realization, receivers with all possible priors take action 1. Lemmas 4 and 5 characterize the optimal signal in the class of signals with this property. The proof of Theorem 1 relies on these lemmas.

In greater detail, Lemma 1 shows that for all signals (satisfying some conditions any sender-optimal signal would satisfy), there exists a feasible prior such that under this

¹²Assumption 1 guarantees that the signal π_0 is essentially unique.

prior the probability action 1 is taken in state $w < w^*$ is below a certain bound and provides a formula for this bound (see the sketch of the proof of Lemma 3 to understand where the formula for the bound came from).

Lemma 2 shows that if the receiver's expectation after a signal realization is greater than w^* for all priors that are extreme points of a set of priors, then the receiver's expectation after the signal realization is greater than w^* for all priors in this set.

Lemma 3 uses Lemmas 1 and 2 to show that there exists a sender-optimal signal with two realizations such that after one realization, receivers with all possible priors take action 1. Specifically, the lemma shows that the probability of action 1 being less than the bound in Lemma 1 is equivalent to the receiver taking action 1 under the binary signal with the above property for every prior that is an extreme point of the receiver's set of priors. This implies that under the binary signal, action 1 is induced in every state with probability at least as large as under the arbitrary signal. Then Lemma 2 implies that under the binary signal, action 1 is induced with probability at least as large as under the arbitrary signal for every possible receiver's prior, which concludes the proof.

Lemma 4 shows that under the optimal binary signal, the probability of the signal realization recommending approval is a hyperbola as a function of the state and provides the formula for the constant of proportionality of the hyperbola. The proof consists in showing that the receiver's expectation being equal to the threshold for extreme points of the set of priors is equivalent to the probability of action 1 being a hyperbola as a function.

Lemma 5 characterizes the support of the signal realization recommending approval under the optimal binary signal, showing that the support is $\Omega(t)$, the set of all states on which the importance index exceeds a threshold t . Lemma 5 also provides the formula for the threshold t . The proof proceeds by computing the marginal change in the sender's payoff from expanding the support and setting this marginal change to zero to find the unique optimum.

A.2 Proofs

In the proofs that follow, all statements hold a.e. with respect to the sender's prior F_s . To ease the exposition, I will write "for all" in most of the proofs instead.

Let $\mathcal{D}$ denote information structures with a finite number of realizations and let $\mathcal{D}_2$ denote information structures π with two realizations such that $E_{F\pi}[w|\sigma_1] \geq w^*$ for all $F \in C_{ag}$. Without loss of generality, the set of priors is $\mathcal{F}_{ag} := \{F : \mu_F = a\mu_{\hat{F}} + (1 - a)\mu_G, \mu_{\hat{F}}([l, w^*)) = 1\}$. Given an information structure π and a prior $F \in \mathcal{F}_{ag}$, define $S(F, \pi) = \{\sigma : E_{F\pi}[w|\sigma] \geq w^*\}$ and $S(\pi) = \{\sigma : \forall F \in \mathcal{F}_{ag}, E_{F\pi}[w|\sigma] \geq w^*\}$. $S(F, \pi)$ is the set of signal realizations after which the expectation of the receiver given the signal structure π and the prior F is above the threshold w^* , and $S(\pi)$ is the set of signal realizations after which the expectation of the receiver given the signal structure π is above the threshold w^* for all priors F in the set $\mathcal{F}_{ag}$. Let

$$p(w; F, \pi) = \begin{cases} \sum_{\sigma \in S(F, \pi)} \pi(\sigma|w) & \text{if } w < w^*, \\ 1 & \text{if } w \geq w^*, \end{cases}$$

$$p(w; \pi) = \begin{cases} \sum_{S(\pi)} \pi(\sigma|w) & \text{if } w < w^*, \\ 1 & \text{if } w \geq w^*. \end{cases}$$

$p(w; F, \pi)$ is the probability that the high action is taken in state ω given the prior F and the signal structure π . $p(w; \pi)$ is the probability that the high action is taken in state ω given the signal structure π and given that the receiver's prior is chosen from the set $\mathcal{F}_{ag}$ after the signal realization to minimize the probability of the high action in state ω .

I now provide a road map for the proof of Lemma 1. The lemma shows that, given an arbitrary information structure, there exists a feasible receiver's prior such that the probability of action 1 in each state is below a certain bound (depending on the state and the information structure).

I split the proof into several steps. Step 1 partitions the support of signal realizations resulting in action 1 for some prior into a finite number of sets. I then choose a finite set of states, one from each set in the partition. The goal of Step 1 is to ensure that we can work with a finite set of states in the subsequent proof. Note that Step 1 is only needed because we start with a continuum of states; if the state space was finite, we could just work with the whole state space.

Steps 2–4 define the receiver's prior. Two objects are important here: a vector of weights for signal realizations $\{\epsilon_\sigma\}$ and a matrix of weights for states and signal realizations $\nu = \{\nu_w^\sigma\}$. The weights ϵ are defined in Step 2. We can interpret ϵ_σ as the proportion of the receiver's expected loss from taking the wrong action after signal realization σ relative to the receiver's expected loss. The weights ν are defined in Step 3 via a fixed-point formula. We can interpret ν_w^σ as the proportion of the receiver's loss from taking the wrong action after signal realization σ in state w relative to the receiver's total loss after σ . Step 4 puts these two objects together to define the receiver's prior. Step 5 shows that the receiver's prior is well defined: because the weights ν are defined via a fixed-point formula, a fixed-point theorem ensures existence.

Step 6 establishes that with the prior that we have constructed, the probability action 1 is below a certain bound in every state in the finite set of states we have selected. The proof consists of rearranging the condition $E_{F\pi}[\omega|\sigma] \geq w^*$ that must hold for any signal realization resulting in action 1 to obtain a bound on $\pi(\sigma|w)$, the probability that action 1 is recommended in state w , and then summing these probabilities over all signal realizations. Step 7 concludes the proof by showing that the probability of action 1 is below the bound in all states. Note that, as with Step 1, we only need Step 7 because we have a continuum of states. If the state space was finite, Step 6 would have been enough to conclude the proof.

LEMMA 1 (Bound). *For all $\pi \in \mathcal{D}$ such that $|S(\pi)| \geq 1$ and for all $\sigma \in S(\pi)$, $\pi(\sigma|\omega) > 0$ for some $\omega \in [l, w^*)$, there is $F \in \mathcal{F}_{ag}$ such that for all $w \in [l, w^*)$,*

$$a(w^* - w)p(w; F, \pi) \leq (1 - a) \int_l^h (w' - w^*)p(w'; F, \pi)g(w')dw'.$$

PROOF. Fix $\pi \in \mathcal{D}$. Step 1 defines the set of states that are going to be the support of the receiver's prior below the threshold w^* . If the state space was finite, we could have used all states below w^* as the support. However, because the state space is an interval, a more complicated construction is required to obtain finite support.

STEP 1 (Partitioning the Set $\bigcup_{S(\pi)} \text{supp } \pi(\sigma|w) \cap [l, w^*)$). Given $A \subseteq S(\pi)$, define

$$X_A = \frac{1-a}{a} \left(\int_{[w^*, h]} (w - w^*) g(w) dw - \int_{[l, w^*]} g(w) \left(\sum_A \pi(\sigma|w) \right) (w^* - w) dw \right).$$

For each $\omega \in [l, w^*)$, define $C(\omega) = \{A \subseteq S(\pi) : \sum_A \pi(\sigma|\omega) > X_A/(w^* - \omega)\}$. $C(\omega)$ is the set of all subsets of $S(\pi)$ such that the probability of signals in each of those subsets in state ω exceeds $X_A/(w^* - \omega)$.

Define a finite number¹³ of sets $Y_1, \dots, Y_J$ such that $Y_i \subseteq \bigcup_{S(\pi)} \text{supp } \pi(\sigma|w) \cap [l, w^*)$ for all i as follows:

- (i) $\omega, \omega' \in Y_i$ implies that, for all $\sigma \in S(\pi)$, if $\pi(\sigma|\omega) > 0$, then $\pi(\sigma|\omega') > 0$;
- (ii) $\omega, \omega' \in Y_i$ implies that $C(\omega) = C(\omega')$;
- (iii) $\omega_i \in Y_i, \omega_j \in Y_j$ for $i \neq j$ implies that
 - (a) either there exists $\sigma \in S(\pi)$ such that either $\pi(\sigma|\omega_i) > 0, \pi(\sigma|\omega_j) = 0$ or $\pi(\sigma|\omega_j) > 0$ and $\pi(\sigma|\omega_i) = 0$
 - (b) or $C(\omega_i) \neq C(\omega_j)$.

Next, we choose a finite set of states as follows. From each $Y_i, i \in \{1, \dots, J\}$, choose exactly one ω_i . Let W denote the set of states $\{\omega_i\}_{i=1}^J$ chosen in this manner.

STEP 2 (Defining a Vector of Weights for Signal Realizations). In this step, I define a collection of weights $\{\epsilon_\sigma\}_{\sigma \in S(\pi)}$ for signal realizations with indices in $S(\pi)$ satisfying $\epsilon_\sigma \in [0, 1]$ for all $\sigma \in S(\pi)$ and $\sum_{S(\pi)} \epsilon_\sigma = 1$. For each $\sigma \in S(\pi)$, define a weight ϵ_σ as

$$\epsilon_\sigma = \frac{E_G[(w - w^*) \pi(\sigma|w)]}{E_G[(w - w^*) p(w; \pi)]}.$$

Note that $\epsilon_\sigma \geq 0$. The fact that $\sum_{S(\pi)} \pi(\sigma|w) = 1$ for all $w \in [w^*, h]$ and f_s has full support implies that $\sum_{S(\pi)} \epsilon_\sigma = 1$.

STEP 3 (Defining a Matrix of Weights for States and Signal Realizations). Let $T = \times_{\sigma \in S(\pi)} \Delta(W)$, where $\Delta(W)$ denotes the simplex over W and $\times$ denotes the Cartesian product. Let $\nu = \{\nu_w^\sigma\}_{\sigma \in S(\pi), w \in W}$. Define a mapping $H : T \rightarrow T$ by

$$H_{\sigma w}(\nu) = \frac{\pi(\sigma|w)(w^* - w) \sum_{k \in S(\pi)} \nu_w^k \epsilon_k}{\sum_{w' \in W} \pi(\sigma|w')(w^* - w') \sum_{k \in S(\pi)} \nu_{w'}^k \epsilon_k}.$$

¹³The number of sets is finite because $S(\pi)$ is finite.

Note that the denominator of $H_{\sigma w}(\nu)$ is strictly positive because for all $\sigma \in S(\pi)$, $\pi(\sigma|\omega) > 0$ for some $\omega \in [l, w^*)$.

For each $w \in W$ and $\sigma \in S(\pi)$, choose ν_w^σ such that $\nu_w^\sigma = H_{\sigma w}(\nu)$.

STEP 4 (Defining the Receiver's Prior). Consider prior μ_F that places $(1-a)\mu_G$ on states $[l, h]$ and, in addition, for all $\sigma \in S(\pi)$, places the mass of $\nu_w^\sigma \epsilon_\sigma a$ on each $\omega \in W$ such that $\nu_w^\sigma \in [0, 1]$ for all $\omega \in W$ and $\sum_{\omega \in W} \nu_w^\sigma = 1$.

Note that $\sum_{\sigma \in S(\pi)} \sum_{\omega \in W} \nu_w^\sigma \epsilon_\sigma a = \sum_{S(\pi)} \epsilon_\sigma a = a$, where the last equality follows from the fact that $\sum_{S(\pi)} \epsilon_\sigma = 1$. The mass of $a + (1-a) \int_{[l, w^*)} g(w) dw$ is placed on states in $[l, w^*)$. Observe that $\mu_F \in \mathcal{F}_{ag}$, so this prior is feasible.

STEP 5 (The Receiver's Prior is Well Defined). Next, we will show that the collection of weights $\{\nu_w^\sigma\}_{\sigma \in S(\pi), w \in W}$ is well defined. Observe that $\sum_{w \in W} \nu_w^\sigma = 1$. We will show that there exists a collection of weights satisfying the formula in the previous step such that $\nu_w^\sigma \in [0, 1]$ for all $\sigma \in S(\pi)$ and for all $w \in W$.

The weights $\{\nu_w^\sigma\}_{\sigma \in S(\pi), w \in W}$ are defined by the equation $\nu = H(\nu)$. Observe that $\Delta(W)$ is a compact and convex set. Because the Cartesian product of convex sets is convex and the Cartesian product of compact sets is compact, this implies that $T = \times_{S(\pi)} \Delta(W)$ is a compact and convex set. Thus $H(\cdot)$ is a continuous self-map on a compact and convex set T . Then the Brouwer fixed-point theorem implies that H has a fixed point.

STEP 6 (Bound on the Sum of Signal Probabilities $\sum_{\sigma \in S(F, \pi)} \pi(\sigma|\omega)$). Suppose that the receiver's prior is chosen as in Step 4 with weights defined by the fixed-point equation $\nu = H(\nu)$. By construction, since $E_{F\pi}[\omega|\sigma] \geq w^*$ for all $\sigma \in S(F, \pi)$ and $w \in W$, we have $\sum_{\sigma \in S(F, \pi)} \pi(\sigma|w) \leq c/(w^* - w)$ for $c = \frac{1-a}{a} \int_l^h (w' - w^*) p(w'; F, \pi) g(w') dw'$.

STEP 7 (Conclusion of the Proof). Suppose for the sake of contradiction that for some $w \in [l, w^*)$, we had $\sum_{\sigma \in S(F, \pi)} \pi(\sigma|w) > c/(w^* - w)$. By definition of $C(w)$, this is equivalent to $S(F, \pi) \in C(w)$. Let $Y(w)$ denote the element Y_i of the partition $Y_1, \dots, Y_J$ such that $w \in Y_i$. Observe that, because $\sum_{S(F, \pi)} \pi(\sigma|w) > c/(w^* - w)$, the fact that $\sum_{S(F, \pi)} \pi(\sigma|w_0) \leq c/(w^* - w_0)$ for all $w_0 \in W$ by Step 6 implies that $w \notin W$. Let w_1 denote the element of W such that $w_1 \in Y(w)$. Then, because $C(w) = C(w')$ for all $w, w' \in Y(w)$, it must be the case that $S(F, \pi) \in C(w_1)$. By definition of $C(w_1)$, this is equivalent to $\sum_{S(F, \pi)} \pi(\sigma|w_1) > c/(w^* - w_1)$. However, this contradicts the fact that $\sum_{S(F, \pi)} \pi(\sigma|w_0) \leq c/(w^* - w_0)$ for all $w_0 \in W$ by Step 6. Therefore, $\sum_{S(F, \pi)} \pi(\sigma|w) > c/(w^* - w)$ for all $w \in [l, w^*)$, as required. $\square$

LEMMA 2 (Extreme Points). Let $\mathcal{F}_{ag}^e$ be the extreme points of $\mathcal{F}_{ag}$ and fix $\pi \in \mathcal{D}_2$. If $E_{F\pi}[\omega|\sigma] \geq w^*$ for all $F \in \mathcal{F}_{ag}^e$, then $E_{F\pi}[\omega|\sigma] \geq w^*$ for all $F \in \mathcal{F}_{ag}$.

PROOF. Fix $F \in \mathcal{F}_{ag}$. Because $\mathcal{F}_{ag}^e$ is the set of the extreme points of $\mathcal{F}_{ag}$, we have $\mathcal{F}_{ag}^e = \{F_w : \mu_{F_w} = (1-a)\mu_G + a\delta_w, w \in [l, w^*)\}$. Then there exists a collection $\{\alpha_w\}_{w \in [l, w^*)}$ such that $\int_{w \in [l, w^*)} \alpha_w dw = 1$, $\mu_F = \int_{w \in [l, w^*)} \alpha_w \mu_{F_w} dw$, and $\alpha_w \in [0, 1]$, $\mu_{F_w} = (1-a)\mu_G + a\delta_w$

for all $w \in [l, w^*)$. Because $E_{F\pi}[\omega|\sigma] \geq w^*$ for all $F_w \in \mathcal{F}_{ag}^e$, for all $w \in [l, w^*)$, we have that $\frac{\int_l^h w\pi(\sigma|w) d\mu_{F_\omega}(w)}{\int_l^h \pi(\sigma|w) d\mu_{F_\omega}(w)} \geq w^*$. Then $\frac{\int_l^h w\pi(\sigma|w) d\mu_F(w)}{\int_l^h \pi(\sigma|w) d\mu_F(w)} \geq \min_{\omega \in [l, w^*)} \frac{\int_l^h w\pi(\sigma|w) d\mu_{F_\omega}(w)}{\int_l^h \pi(\sigma|w) d\mu_{F_\omega}(w)} \geq w^*$. $\square$

LEMMA 3 (Binary Optimal Signal). *For all $\hat{\pi} \in \mathcal{D}$ such that $|S(\hat{\pi})| \geq 1$ and for all $\sigma \in S(\hat{\pi})$, $\hat{\pi}(\sigma|\omega) > 0$ for some $\omega \in [l, w^*)$, there exists $\pi \in \mathcal{D}_2$ such that $U(\pi) \geq U(\hat{\pi})$.*

PROOF. If $\hat{\pi} \in \mathcal{D}$ is such that $|S(\hat{\pi})| = 1$, then we can obtain $\pi \in \mathcal{D}_2$ by merging all signal realizations inducing action 0 into one signal realization. Thus suppose that $|S(\hat{\pi})| > 1$. Then, by Lemma 1, there exists $F \in \mathcal{F}_{ag}$ such that for all $w' \in [l, w^*)$, $\hat{\pi}$ satisfies $a(w^* - w)p(w; F, \pi) \leq (1 - a) \int_l^h (w' - w^*)p(w'; F, \pi)g(w')dw'$.

We define $\pi \in \mathcal{D}_2$ as follows. We set $\pi(\sigma_1|w) = \sum_{S(F, \hat{\pi})} \hat{\pi}(\sigma|w)$ for all $w \in [l, w^*)$, $\pi(\sigma_1|w) = 1$ for all $w \in [w^*, h]$, and $\pi(\sigma_0|w) = 1 - \pi(\sigma_1|w)$ for all $w \in [l, h]$.

Observe that $a(w^* - w)p(w; \pi) \leq (1 - a) \int_l^h (w' - w^*)p(w'; \pi)g(w')dw'$ for all $w' \in [l, w^*)$. This implies that $E_{F_w\pi}[\omega|\sigma_1] \geq w^*$ for all priors F_w of the form $\mu_{F_w} = (1 - a)\mu_G + a\delta_w$ for some $w \in [l, w^*)$. Priors F_w are the extreme points of $\mathcal{F}_{ag}$. Because $E_{F_w\pi}[\omega|\sigma_1] \geq w^*$ for all such priors, Lemma 2 implies that $E_{F\pi}[\omega|\sigma_1] \geq w^*$ for all $F \in \mathcal{F}_{ag}$. This implies that $U(\pi) = E_{F_s}[\pi(\sigma_1|w)]$. Then $U(\pi) \geq E_{F_s}[\sum_{S(F, \hat{\pi})} \hat{\pi}(\sigma|w)]$.

Observe that $U(\hat{\pi}) \leq \min_{\tilde{F} \in \mathcal{F}_{ag}} E_{F_s}[\sum_{S(\tilde{F}, \hat{\pi})} \hat{\pi}(\sigma|w)]$. In particular, $U(\hat{\pi}) \leq E_{F_s}[\sum_{S(F, \hat{\pi})} \hat{\pi}(\sigma|w)]$.

Then the above implies that $U(\hat{\pi}) \leq U(\pi)$, as required. $\square$

Note that, given an optimal $\pi \in \mathcal{D}_2$, there is exactly one signal realization such that receivers with all feasible priors take action 1 after it. Henceforth, I will use σ_1 to denote this signal realization.

LEMMA 4 (Hyperbola). *If π is optimal in $\mathcal{D}_2$, then for all $w < w^*$ such that $p(w; \pi) > 0$, $\pi(\sigma_1|w) = \min\{c/(w^* - w), 1\}$, where*

$$c = \frac{1-a}{a} \int_l^h (w - w^*)p(w; \pi)g(w)dw.$$

PROOF. Let $C := \frac{1-a}{a} \int_l^h (w - w^*)p(w; \pi)g(w)dw$. Given $A, B \subseteq [l, w^*)$, define

$$Y(A, B) = \frac{(1-a) \int_{A \cup [w^*, h]} (w - w^*)p(w; \pi)g(w)dw}{a + (1-a)\mu_G(B)}.$$

Note that Lemmas 1 and 3 imply that if π is optimal in $\mathcal{D}_2$, then for all $w < w^*$ such that $p(w; \pi) > 0$, $\pi(\sigma_1|w) \leq \frac{c}{w^* - w}$ for $c = C$.

CLAIM 4.1. *Let (A, B) be the partition of $\text{supp } \pi(\sigma_1|w) \cap [l, w^*)$ such that $\pi(\sigma_1|w) = \min\{\frac{c}{w^* - w}, 1\}$ for $w \in A$, $\pi(\sigma_1|w) < \min\{\frac{c}{w^* - w}, 1\}$ for $w \in B$ for $c = C$. Then $c = Y(B \cup (A \cap [w^* - c, w^*]), A \cap [l, w^* - c])$.*

PROOF. The formula for c follows from substituting $\pi(\sigma_1|w) = \frac{c}{w^* - w}$ for $w \in A$ into $\frac{1-a}{a} \int_l^h (w - w^*) p(w; \pi) g(w) dw$ and rearranging the expressions. We next show that there exists c satisfying the formula in the lemma. Consider a mapping $c \mapsto L(c)$ given by $L(c) = Y(B \cup (A \cap [w^* - c, w^*]), A \cap [l, w^* - c))$. $L(0) \leq 0$ implies that $E_{F\pi}[w|\sigma_1] < w^*$ for all $F \in \mathcal{F}_{ag}$, which contradicts the hypothesis that π is optimal. Therefore, we must have $L(0) > 0$. Moreover, $L(\bar{c}) < 0$ for $\bar{c}$ sufficiently large. Then, because $L(0) \geq 0$, $L(\bar{c}) < 0$, and $c \mapsto L(c)$ is continuous, there exists $c \in [0, \bar{c}]$ such that $L(c) = 0$. $\square$

Suppose π is optimal in $\mathcal{D}_2$. Since π is optimal in $\mathcal{D}_2$, $w < w^*$ such that $p(w; \pi) > 0$, we must have $\pi(\sigma_1|w') \leq C/(w^* - w)$. Suppose for the sake of contradiction that there exists π that is optimal in $\mathcal{D}_2$ such that $\pi(\sigma_1|w') < C/(w^* - w)$ for some $B \subseteq [l, w^*)$ of a strictly positive measure.

Define $A = (\text{supp } \pi(\sigma_1|w) \cap [l, w^*)) \setminus B$. Fix $x \in B$ such that for all intervals I satisfying $x \in \text{int}(I)$ (where int denotes the interior), we have $\mu(I \cap B) > 0$. Because, by the hypothesis, $\mu_{F_s}(B) > 0$, such x exists. Given $\epsilon > 0$ sufficiently small, fix an interval $I_{x,\epsilon} = [x - \frac{\epsilon}{2}, x + \frac{\epsilon}{2}]$.

We define $\pi_1 \in \mathcal{D}_2$ as follows. For all $w \in I_{x,\epsilon} \cap B$, we let $\pi_1(\sigma_1|w) = \pi(\sigma_1|w) + \eta$ for some $\eta > 0$. We let $\pi_1(\sigma_1|w) = \pi(\sigma_1|w)$ for all $w \in B \setminus I_{x,\epsilon}$. We choose η small enough such that $\pi_1(\sigma_1|w) < 1$ and $E_{F_w\pi_1}[\omega|\sigma_1] > w^*$ for all $w \in B$. This is feasible because $\pi(\sigma_1|w) < 1$ for all $w \in B$ and because, if $\pi(\sigma_1|w) < C/(w^* - w)$, then $E_{F_w\pi}[\omega|\sigma_1] > w^*$, which implies that $E_{F_w\pi_1}[\omega|\sigma_1] > w^*$ for all $w \in B$.

Let us write $\pi(\sigma_1|w) = \min\{\frac{c}{w^* - w}, 1\}$ for all $w \in A$ and $\pi_1(\sigma_1|w) = \min\{\frac{c_1}{w^* - w}, 1\}$ for all $w \in A$. To complete the construction of π_1 , we require that $\pi_1(\sigma_1|w) = C/(w^* - w)$ for all $w \in A$. To ensure that $\pi_1(\sigma_1|w) = C/(w^* - w)$ for all $w \in A$, we choose c_1 satisfying the formula in Claim 4.1.

Let $\mathbb{1}^c(w) = 1$ if $w > w^* - c$ and $\mathbb{1}^c(w) = 0$ otherwise.

CLAIM 4.2. $\lim_{\epsilon \rightarrow 0} \frac{U(\pi_1) - U(\pi)}{\eta\epsilon} = I(x) - \frac{E[I(w)\mathbb{1}_{w \in A}(1 - \mathbb{1}^c(w))]}{\frac{a}{1-a} + E_G[\mathbb{1}_{w \in A}(1 - \mathbb{1}^c(w))]}.$

PROOF. The proof follows from the formula for c_1 in Claim 4.1, approximating the integrals, the fact that g and f_s are C^1 , and the fact that $U(\pi_1) - U(\pi) = \eta \int_{I_{x,\epsilon} \cap B} f_s(w) dw - E_{F_s}[\frac{c-c_1}{w^* - w} \mathbb{1}_{w \in A}(1 - \mathbb{1}^c(w))] - E_{F_s}[\mathbb{1}_{w \in A}(\mathbb{1}^{c_1}(w) - \mathbb{1}^c(w))(\frac{c_1}{w^* - w} - 1)]$. $\square$

In order for the sender to not have a strictly improving deviation, we need that for all $x \in B$ such that for all intervals I satisfying $x \in \text{int}(I)$ we have $\mu(I \cap B) > 0$, either $\lim_{\epsilon \rightarrow 0} \frac{U(\pi_1) - U(\pi)}{\eta\epsilon} = 0$ or $\lim_{\epsilon \rightarrow 0} \frac{U(\pi_1) - U(\pi)}{\eta\epsilon} > 0$ and $\pi(\sigma_1|x) = 1$ (note that if $\pi(\sigma_1|x) = 0$, then $x \notin B$). Thus, by Claim 4.2, we need $I(x) - c_0 = 0$ for some constant c_0 for almost all $x \in B$ satisfying $\pi(\sigma_1|x) < 1$. However, Assumption 1 implies that this fails. Thus if $\pi(\sigma_1|w) \leq C/(w^* - w)$ holds strictly on a set of a strictly positive measure, then the sender has a strictly improving deviation, which contradicts the optimality of π . Therefore, $\pi(\sigma_1|w) \leq C/(w^* - w)$ must be satisfied with equality on $\text{supp } \pi(\sigma_1|w)$ below w^* . $\square$

Let $\mathbb{1}^c(w) = 1$ if $w > w^* - c$ and $\mathbb{1}^c(w) = 0$ otherwise. Let $\mathbb{1}_t(w) = 1$ if $w \in \Omega(t)$ and $\mathbb{1}_t(w) = 0$ otherwise. Let $\mathbb{1}(w) = 1$ if $w \in [w^*, h]$ and $\mathbb{1}(w) = 0$ otherwise. Define also

$$L(c, t) = \frac{(1-a)E_G[(\mathbb{1}_t(w)\mathbb{1}^c(w) + \mathbb{1}(w))(w - w^*)]}{a + (1-a)E_G[\mathbb{1}_t(w)(1 - \mathbb{1}^c(w))]}.$$

Let $c(t)$ denote the unique solution to $c = L(c, t)$. Define

$$y(t) = t(a/(1-a) + E_G[\mathbb{1}_t(w)(1 - \mathbb{1}^c(w))]) - E_{F_s}[\mathbb{1}_t(w)(1 - \mathbb{1}^c(w))(1/(w^* - w))].$$

LEMMA 5 (Threshold). *If π is optimal in $\mathcal{D}_2$, then there exists a unique threshold $t \geq \underline{t}$ such that $\text{supp } \pi(\sigma_1|w) \cap [l, w^*) = \Omega(t)$ and either $y(t) = 0$ or $t = \underline{t}$ and $y(\underline{t}) > 0$.*

PROOF. We first show that a threshold with required properties exists. Define $\Omega^\pi = \{w \in [l, w^*) : \pi(\sigma'|w) > 0 \text{ for some } \sigma' \in S(\pi)\}$.

Given $x \in [l, w^*)$ and $\epsilon > 0$ satisfying $x - \frac{\epsilon}{2} \geq l$, $x + \frac{\epsilon}{2} < w^*$, define $I_{x,\epsilon} = [x - \frac{\epsilon}{2}, x + \frac{\epsilon}{2}]$. Consider $\pi_1 \in \mathcal{D}_2$ obtained by adding an interval $I_{x,\epsilon}$ to the support of σ_1 below the threshold w^* . That is, $\pi_1 \in \mathcal{D}_2$ satisfies $\Omega^{\pi_1} = \Omega^\pi \cup I_{x,\epsilon}$ and $\mu(I_{x,\epsilon} \cap \Omega^\pi) < \epsilon$. Then, by Lemma 4, there exist constants c and c_1 such that $\pi(\sigma_1|w) = \min\{\frac{c}{w^* - w}, 1\}$ for all $w \in \Omega^\pi$ and $\pi_1(\sigma_1|w) = \min\{\frac{c_1}{w^* - w}, 1\}$ for all $w \in \Omega^{\pi_1} = \Omega^\pi \cup I_{x,\epsilon}$.

CLAIM 5.1. *If $y(I(x)) > 0$, then $U(\pi_1) > U(\pi)$. If $U(\pi_1) \geq U(\pi)$, then $y(I(x)) \geq 0$.*

The proof follows from approximating the integrals.

CLAIM 5.2. *If $w_1 \notin \Omega^\pi$, then $w_2 \notin \Omega^\pi$ for all $w_2 \in [l, w^*)$ such that $I(w_1) > I(w_2)$.*

PROOF. If $w_1 \in [l, w^*)$ is such that $w_1 \notin \Omega^\pi$ and π is optimal, then adding $I_{w_1,\epsilon}$ to the support does not strictly benefit the sender. That is, then $U(\pi_1) \leq U(\pi)$, where the support of π_1 is obtained by adding $I_{w_1,\epsilon}$ to Ω^π . By Claim 5.1, this implies that $y(I(w_1)) > 0$. Then for any $w_2 \in [l, w^*)$ such that $I(w_1) > I(w_2)$, we have $y(I(w_2)) < 0$. By Claim 5.1, this implies that $U(\pi_2) < U(\pi)$, where the support of π_2 is obtained by adding $I_{w_2,\epsilon}$ to Ω^π . $\square$

Claim 5.2 implies that for $w_1, w_2 \in [l, w^*)$, if $w_2 \in \Omega^\pi$ and $I(w_1) > I(w_2)$, then $w_1 \in \Omega^\pi$. Therefore, $\text{supp } \pi(\sigma_1|w) \cap [l, w^*) = \Omega(t)$, as required.

We next show that the threshold is unique. Note that the proof of Lemma 4 and the fact that $\Omega^\pi = \Omega(t)$ imply that c satisfying $\pi(\sigma_1|w) = \min\{\frac{c}{w^* - w}, 1\}$ is given by $c = c(t)$.

CLAIM 5.3. *$t \mapsto c(t)$ is a continuous and strictly increasing function for $t \geq \underline{t}$.*

PROOF. Define a function $(c, t) \mapsto x(c, t)$ as $x(c, t) = E_G[(\mathbb{1}(w) + \mathbb{1}_t(w)\mathbb{1}^c(w))(w - w^*)] - c(a/(1-a) + E_G[\mathbb{1}_t(w)(1 - \mathbb{1}^c(w))])$. We have $x(c + \eta, t) - x(c, t) = O(\eta^2) - \eta(a/(1-a) + E_G[\mathbb{1}_t(w)(1 - \mathbb{1}_{c+\eta}(w))])$ because g is C^1 . Then $\lim_{\eta \rightarrow 0} \frac{x(c+\eta, t) - x(c, t)}{\eta} < 0$.

Assumption 1 and the fact that f_s and g are continuous imply that $\Omega(t + \eta) \cap [l, w^* - c) \subset \Omega(t) \cap [l, w^* - c)$ for all $t \geq \underline{t}$ and for all $\eta > 0$ sufficiently small. Then $x(c, t + \eta) - x(c, t) = E_G[\mathbb{1}^c(w)(\mathbb{1}_t(w) - \mathbb{1}_{t+\eta}(w))(w^* - w)] + cE_G[(\mathbb{1}_t(w) - \mathbb{1}_{t+\eta}(w))(1 - \mathbb{1}^c(w))] > 0$.

The fact that $\lim_{\eta \rightarrow 0} \frac{x(c+\eta, t) - x(c, t)}{\eta} < 0$ and $x(c, t + \eta) - x(c, t) > 0$ for all $\eta > 0$ by the arguments above implies that $t \mapsto c(t)$ is a strictly increasing function, as required.

$(c, t) \mapsto x(c, t)$ is continuous because of Assumption 1 and the fact that f_s and g are continuous. The continuity of $t \mapsto c(t)$ then follows from the fact that $(c, t) \mapsto x(c, t)$ is continuous. $\square$

Because $t \mapsto c(t)$ is a strictly increasing function for $t \geq \underline{t}$ by Claim 5.3, we have $c(t) < c(t + \eta)$ for all $\eta > 0$ sufficiently small.

Define $m_\eta(w) = \mathbb{1}_t(w)(\mathbb{1}^{c(t+\eta)}(w) - \mathbb{1}^{c(t)}(w)) + (\mathbb{1}_t(w) - \mathbb{1}_{t+\eta}(w))(1 - \mathbb{1}^{c(t)}(w)) + (\mathbb{1}_t(w) - \mathbb{1}_{t+\eta}(w))(\mathbb{1}^{c(t+\eta)}(w) - \mathbb{1}^{c(t)}(w))$. Observe that $\mathbb{1}_{t+\eta}(w)(1 - \mathbb{1}^{c(t+\eta)}(w)) = \mathbb{1}_t(w)(1 - \mathbb{1}^{c(t)}(w))(1 - m_\eta(w))$. Then $y(t + \eta) - y(t) = \eta(a/(1 - a) + E_G[\mathbb{1}_t(w)(1 - \mathbb{1}^{c(t+\eta)}(w))]) - tE_G[m_\eta(w)] + E[m_\eta(w)I(w)]$, where E denotes the expectation with respect to the Lebesgue measure.

We claim that $E[m_\eta(w)I(w)] - tE_G[m_\eta(w)] \geq 0$ for all $\eta > 0$. The definition of $\Omega(t)$ implies that for all $w \in \Omega(t)$, we have $I(w) \geq t$. Because $m_\eta(w) \leq \mathbb{1}_t(w)$, this implies that for all w such that $m_\eta(w) > 0$, we have $I(w) \geq t$. Then $\frac{E[m_\eta(w)I(w)]}{E_G[m_\eta(w)]} \geq \inf_{w: m_\eta(w) > 0} I(w) \geq t$ implies the claim.

Because $a/(1 - a) + E_G[\mathbb{1}_{t+\eta}(w)(1 - \mathbb{1}^{c(t+\eta)}(w))] > 0$, this implies that $y(t + \eta) - y(t) > 0$ for all $\eta > 0$. Consider a point t^* such that $y(t^*) = 0$. The fact that $y(t + \eta) - y(t) > 0$ for all $t \geq \underline{t}$ and for all $\eta > 0$ implies that the function $t \mapsto y(t)$ can intersect the zero function at $t \geq \underline{t}$ at most one point. $\square$

PROOF OF THEOREM 1. Observe that for any sender-optimal information structure π we must have $|S(\pi)| \geq 1$ and, for all $\sigma \in S(\pi)$, $\pi(\sigma|\omega) > 0$ for some $\omega \in [l, w^*)$. Then Lemma 3 shows that there exists an optimal signal with two realizations such that after one realization, receivers with all possible priors take action 1. Lemma 4 shows that if π is an optimal signal with this property, then for all $w < w^*$ such that $p(w; \pi) > 0$, $\pi(\sigma_1|w) = \min\{c/(w^* - w), 1\}$ for some constant c . Lemma 5 and the proof of Lemma 4 imply that $c = c(t)$. Lemma 5 shows that there exists a unique threshold $t \geq \underline{t}$ such that $\text{supp } \pi(\sigma_1|w) \cap [l, w^*) = \Omega(t)$. The fact that the distribution over the receiver's actions induced by an optimal signal is unique up to the sets of measure zero under the sender's prior is then immediate. $\square$

PROOF OF PROPOSITION 1. Let us write $y_a(t) = t(\frac{a}{1-a} + E_G[\mathbb{1}_t(w)(1 - \mathbb{1}^{c(t)}(w))]) - E[\mathbb{1}_t(w)(1 - \mathbb{1}^{c(t)}(w))I(w)]$.

CLAIM 1.1. $y_{a'}(t) - y_a(t) < 0$ for all $t \geq \underline{t}$.

PROOF. $y_{a'}(t) - y_a(t) = t(\frac{a'}{1-a'} - \frac{a}{1-a}) < 0$ because $t > 0$ and $a' < a$. $\square$

CLAIM 1.2. $t_{a'} - t_a > 0$.

PROOF. Because $\mu(\Omega(t_a)) < \mu([l, w^*))$, we have $t_a > \underline{t}$. Then, because, by the proof of Lemma 5, we have $y(t + \eta) > y(t)$ for all $t \geq \underline{t}$ and for all η , to show that $t_{a'} - t_a > 0$, it is enough to show that $y_{a'}(t) - y_a(t) < 0$ for all $t \geq \underline{t}$. Claim 4.1 implies that this is satisfied. $\square$

The fact that $t_{a'} - t_a > 0$ by Claim 4.2 implies that $\Omega(t_{a'}) \subset \Omega(t_a)$. Finally, the fact that $c(t_{a'}) > c(t_a)$ follows from the fact that $t \mapsto c(t)$ is strictly increasing by Claim 5.3 in Lemma 5. $\square$

PROOF OF PROPOSITION 2. We will prove the following result.

Suppose that the receiver's prior is $F \in C_{a'g} \subset C_{ag}$ for $0 < a' < a < 1$. Then

- (i) If $V(a') = 0$, then $V(a) > 0$.
- (ii) Suppose that the optimal signal recommends approval in all states with a strictly positive probability under a and a' . Then $V(a) \geq V(a')$. Moreover, if the sender's knowledge is detail-free, then $V(a) = V(a')$, and if the reference prior g has full support, then $V(a) > V(a')$.
- (iii) There exist g, F_s such that $0 < V(a) < V(a')$.

Let $d \in \{0, 1\}$ denote the receiver's actions. Let us write the receiver's utility function as $u(d, w) = \mathbb{1}_{d=1}(w - w^*)$. Let π_a denote the optimal signal with two realizations when the ignorance index is a . Then the expected payoff of a receiver with prior F is $V(a) = E_F[\pi_a(\sigma_1|w)(w - w^*)]$.

Observe that $V(a) = E_F[\pi_a(\sigma_1|w)](E_F[\pi_a(w|\sigma_1)] - w^*)$. This implies that $V(a') = 0$ if and only if $E_F[\pi_{a'}(w|\sigma_1)] - w^* = 0$. Then, because $E_F[\pi_a(\sigma_1|w)] > 0$, to prove the first part of the proposition, it is enough to show that if $E_F[\pi_{a'}(w|\sigma_1)] = w^*$, then $E_F[\pi_a(w|\sigma_1)] > w^*$.

Observe that μ_F is a convex combination of priors of the form $\mu_{F_w} = a'\delta_w + (1 - a')\mu_G$. Because $E_{F_w}[\pi_{a'}(w|\sigma_1)] = w^*$ for some $w \in [l, h]$ and $a > a'$, we have $E_{F_w}[\pi_a(w|\sigma_1)] > w^*$ for all $w \in [l, h]$. Then $E_F[\pi_a(w|\sigma_1)] > w^*$, as required.

We next consider the case in which the support of $\pi_a(\sigma_1|w)$ and $\pi_{a'}(\sigma_1|w)$ below w^* is $[l, w^*)$. Note that $V(a) = E_F[(w - w^*)\pi_a(\sigma_1|w)]$. Then $V(a) - V(a') = E_F[(w^* - w)(\pi_{a'}(\sigma_1|w) - \pi_a(\sigma_1|w))]$. Observe that $\pi_a(\sigma_1|w) = \min\{\frac{c_a}{w^* - w}, 1\}$ for some constant c_a if $w \in [l, w^*)$ and $\pi_a(\sigma_1|w) = 1$ if $w \in [w^*, h]$. Moreover, $c_a < c_{a'}$. Then we have $\pi_a(\sigma_1|w) < \pi_{a'}(\sigma_1|w)$ for $w \in [l, w^* - c_a]$ and $\pi_a(\sigma_1|w) \leq \pi_{a'}(\sigma_1|w)$ for $w \in [w^* - c_a, w^*)$. Thus $V(a) - V(a') \geq 0$. Moreover, $V(a) - V(a') > 0$ if $\mu_F([l, w^* - c_a]) > 0$ and $V(a) - V(a') = 0$ if $\mu_F([l, w^* - c_a]) = 0$. Therefore, if $\mu_F([l, w^*)) = 0$, then $V(a) = V(a')$, while if g has full support, then $V(a) > V(a')$, as required.

We now provide an example of the parameters under which $V(a) < V(a')$. For simplicity, suppose that $\mu_G = (1 - \kappa)\mu_{G_0} + \kappa\delta_h$ for some $\kappa \in (0, 1)$ and μ_{G_0} that is uniform on $[w_0, w_1]$ for some $l < w_0 < w_1 < w^*$.

CLAIM 2.1. There exist parameters such that $\Omega(t_a) = [l, w^*)$ and $[w_0, w_1] \cap \Omega(t_{a'}) = \emptyset$.

PROOF. First observe that, because $\mu_G([l, w_0]) = 0$, we have $[l, w_0] \subseteq \Omega(t_a)$, $\Omega(t_{a'})$. Fix $a' \in (0, 1)$. Consider F_s such that $f_s(w) = \eta$ for $w \in [w_0, w_1]$ and $f_s(w) = \eta_0$ for $w \in [l, h] \setminus [w_0, w_1]$. Then for η sufficiently small, we have $[w_0, w_1] \cap \Omega(t_{a'}) = \emptyset$. This is because $\lim_{\eta \rightarrow 0} I(w) = 0$ for $w \in [w_0, w_1]$ but $\lim_{\eta \rightarrow 0} t^* > 0$ for t^* satisfying $y(t^*) = 0$ because $\mu(\Omega(t) \cap [l, w_0]) > 0$ and $f_s(w) \geq \eta_0$ for all $w \in [l, w_0]$.

Observe next that $l < w^* - c(t_a)$ for $a > a'$ sufficiently large. This is because $\lim_{a \rightarrow 1} c(t_a) = 0$ and $l < w^*$.

The above and that fact that $[l, w_0] \subseteq \Omega(t)$ imply that $\mu(\Omega(t_a) \cap [l, w^* - c(t_a)]) > 0$ for $a > a'$ sufficiently large. Because $\lim_{a \rightarrow 1} \frac{a}{1-a} = \infty$, the fact that $\mu(\Omega(t) \cap [l, w^* - c(t_a)]) > 0$ for $a > a'$ sufficiently large implies that $\lim_{a \rightarrow 1} t^* = 0$ for t^* satisfying $y(t^*) = 0$. Because $I(w) = \frac{\eta}{g(w)(w^* - w)} > 0$ for $w \in [w_0, w_1]$, this implies that as $a \rightarrow 1$, $\Omega(t_a) \rightarrow [l, w^*)$. $\square$

Consider $\mu_F = (1 - a)(1 - \kappa)\mu_{G_0} + (1 - a)\kappa\delta_h + a\delta_h$. Because $V(a) = E_F[(w^* - w)\pi_a(\sigma_1|w)]$, we have $V(a') = \int_{[w^*, h]} (w - w^*) dF(w)$ since $[w_0, w_1] \cap \Omega(t_{a'}) = \emptyset$. Note that $\int_{[l, w^*)} (w^* - w)\pi_a(\sigma_1|w) dF(w) > 0$ because $\Omega(t_a) = [l, w^*)$. This implies that $V(a') > V(a)$, as required. $\square$

PROOF OF PROPOSITION 3. Suppose that $\mu_G([l, w^*)) = 0$. Then $\Omega(t) = [l, w^*)$ for all $t < \infty$. Thus Theorem 1 implies that an optimal signal induces the probability of action 1 in state w given by $s(w) = 1$ for $w \in [w^*, h]$ and $s(w) = \min\{\frac{c}{w^* - w}, 1\}$ for $w \in [l, w^*)$ for some constant c . Observe that s is independent of F_s , as required.

Next suppose that an optimal signal with two realizations induces a distribution over the receiver's actions that is independent of F_s . Suppose for the sake of contradiction that $\mu_G([l, w^*)) > 0$. Let $A \subseteq [l, w^*)$ denote the set satisfying $\mu_G(A) = \mu_G([l, w^*))$ (note that A is unique up to sets of measure zero).

For simplicity, I will provide a proof allowing for arbitrary priors F_s . I first show that $\Omega(t) = [l, w^*)$. Suppose for the sake of contradiction that $\Omega(t) \subset [l, w^*)$. Fix $w \in [l, w^*) \setminus \Omega(t)$. Choose $\mu_{F_s} = \delta_w$. Then any optimal signal must recommend action 1 with a positive probability in state w , which is a contradiction.

Next, I show that if $\mu_G([l, w^*)) > 0$, then $\Omega(t) \subset [l, w^*)$. Suppose for the sake of contradiction that $\Omega(t) = [l, w^*)$. Fix $w \in [l, w^*) \setminus A$ and choose $\mu_{F_s} = \delta_w$. Note that the sender's payoff is strictly higher when $\Omega(t) \cap A = \emptyset$ than when $\Omega(t) = [l, w^*)$, which is a contradiction. $\square$

PROOF OF PROPOSITION 4. Let us write the fixed-point equation pinning down the values of the threshold $t > \underline{t}$ as $y(t^*) = 0$. Observe that $\lim_{a \rightarrow 1} t^* = 0$.

The fact that g is C^1 and $[l, h]$ is a compact set implies that g attains a maximum M on $[l, h]$, which implies that $g(w) \leq M$ for all $w \in [l, w^*)$. It follows that if there exists $m > 0$ such that $f_s(w) \geq m$ for all $w \in [l, w^*)$, then $\underline{t} = \min_{w \in [l, w^*)} I(w) > 0$.

Thus, because $\lim_{a \rightarrow 1} t^* = 0$, there exists $\bar{a} \in (0, 1)$ such that for all $a \in (\bar{a}, 1)$, we have $t^* < \underline{t}$. Therefore, for all $a \in (\bar{a}, 1)$, the threshold is $t = \underline{t}$, so that $\Omega(t) = [l, w^*)$, as required. $\square$

PROOF OF PROPOSITION 5. We first characterize the optimal signal when $a = 0$. We will show that it is given by $\pi(\sigma_1|w) = 1$ for $w \in \Omega(t) \cup [w^*, h]$, $\pi(\sigma_1|w) = 0$ for $w \in [l, w^*) \setminus \Omega(t)$, where $t \geq 0$ is such that $E_{G\pi}[\omega|\sigma_1] = w^*$.

Because $a = 0$, the revelation principle applies. Thus, because the receiver has two actions, there is an optimal signal with two realizations, σ_1 and σ_0 . Then the sender's problem is $\sup_{\{\pi(\sigma_1|w) \in [0, 1]\}_{w \in [l, h]}} \int_l^h f_s(w)\pi(\sigma_1|w) dw$ subject to $E_{G\pi}[\omega|\sigma_1] \geq w^*$. The

constraint is equivalent to $\int_l^h (w^* - w)g(w)\pi(\sigma_1|w)dw \leq 0$. Observe that the constraint must bind.

We can write the Lagrangian as $\int_l^h f_s(w)\pi(\sigma_1|w)dw - \lambda \int_l^h (w^* - w)g(w)\pi(\sigma_1|w)dw$. Equivalently, $\int_l^h (f_s(w) - \lambda(w^* - w)g(w))\pi(\sigma_1|w)dw$. Then the solution is $\pi(\sigma_1|w) = 1$ if $f_s(w) \geq \lambda(w^* - w)$ and $\pi(\sigma_1|w) = 0$ if $f_s(w) < \lambda(w^* - w)$, where $\lambda > 0$ is such that $E_{G\pi}[\omega|\sigma_1] = w^*$. This implies that $\pi(\sigma_1|w) = 1$ for all $w \in [w^*, h]$. Moreover, for $w \in [l, w^*)$, we have $\pi(\sigma_1|w) = 1$ if $I(w) \geq \lambda$ and $\pi(\sigma_1|w) = 0$ if $I(w) \leq \lambda$, where λ is such that $E_{G\pi}[\omega|\sigma_1] = w^*$.

Next, we prove the convergence result in the proposition. Let π_a denote the optimal signal with two realizations when the set of priors is C_{ag} . Without loss of generality, suppose that $\lim_{a \rightarrow 0} \int_l^h |\pi_a(\sigma_1|w) - \pi(\sigma_1|w)|dF_s(w)$ and $\tilde{\pi} = \lim_{a \rightarrow 0} \pi_a(\sigma_1|w)$ exist. Suppose for the sake of contradiction that $\lim_{a \rightarrow 0} \int_l^h |\pi_a(\sigma_1|w) - \pi(\sigma_1|w)|dF_s(w) \neq 0$.

Note that π_a is $U(\pi_a) = \int_l^h \pi_a(\sigma_1|w)dF_s(w)$. Suppose first that $\lim_{a \rightarrow 0} U(\pi_a) = U(\pi)$. Note that we must have $E_{G\pi_a}[w|\sigma_1] \geq w^*$ for all $a > 0$. This implies that $\lim_{a \rightarrow 0} E_{G\pi_a}[w|\sigma_1] = E_{G\tilde{\pi}}[w|\sigma_1] \geq w^*$, so under $\tilde{\pi}$, the receiver with prior G takes action 1 after σ_1 . Then $\tilde{\pi}$ is a signal with two realizations such that $\int_l^h |\pi_a(\sigma_1|w) - \tilde{\pi}(\sigma_1|w)|dF_s(w) \neq 0$ and $U(\tilde{\pi}) = U(\pi)$ given that the receiver's prior is G . This contradicts the fact that π is unique up to sets of measure zero under F_s .

Next, suppose that $\lim_{a \rightarrow 0} U(\pi_a) \neq U(\pi)$. Because $U(\pi_a) < U(\pi)$ for all $a > 0$, this implies that $\lim_{a \rightarrow 0} U(\pi_a) < U(\pi)$. We will show that there exists a sequence of signals $\{\hat{\pi}_a\}$ with two realizations such that $E_{F\hat{\pi}_a}[w|\sigma_1] \geq w^*$ for all $F \in C_{ag}$ and $\lim_{a \rightarrow 0} U(\hat{\pi}_a) = U(\pi)$. Because $\lim_{a \rightarrow 0} U(\pi_a) < U(\pi)$, this would establish that $\{\pi_a\}_{a \in (0,1)}$ was not a collection of optimal signals, which is a contradiction.

Let $\hat{\pi}_a(\sigma_1|w) = \pi(\sigma_1|w) - \epsilon_a$ if $\pi(\sigma_1|w) = 1$ and $w < w^*$, $\hat{\pi}_a(\sigma_1|w) = \pi(\sigma_1|w) = 1$ if $w \geq w^*$ and $\hat{\pi}_a(\sigma_1|w) = 0$ if $\pi(\sigma_1|w) = 0$. Choose the minimal $\epsilon_a > 0$ such that $E_{F\hat{\pi}_a}[w|\sigma_1] \geq w^*$ for all $F \in C_{ag}$. Note that this is feasible because if $\epsilon_a = 1$, then $E_{F\hat{\pi}_a}[w|\sigma_1] > w^*$ for all $F \in C_{ag}$. Observe that $\lim_{a \rightarrow 0} \epsilon_a = 0$. This implies that $\lim_{a \rightarrow 0} U(\hat{\pi}_a) = U(\pi)$, as required. $\square$

REFERENCES

- Alonso, Ricardo and Odilon Câmara (2016a), “Bayesian persuasion with heterogeneous priors.” *Journal of Economic Theory*, 165, 672–706. [1076, 1078, 1079]
- Alonso, Ricardo and Odilon Câmara (2016b), “Persuading voters.” *American Economic Review*, 106, 3590–3605. [1078]
- Beauchêne, Dorian, Jian Li, and Ming Li (2019), “Ambiguous persuasion.” *Journal of Economic Theory*, 179, 312–365. [1080]
- Boleslavsky, Raphael and Kyungmin Kim (2018), “Bayesian persuasion and moral hazard.” Manuscript. [1079]
- Brocas, Isabelle and Juan D. Carillo (2007), “Influence through ignorance.” *RAND Journal of Economics*, 38, 931–947. [1078]

- Carrasco, Vinicius, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira (2017), "Optimal selling mechanisms under moment conditions." Manuscript. [1080]
- Carroll, Gabriel (2015), "Robustness and linear contracts." *American Economic Review*, 105, 536–563. [1080]
- Carroll, Gabriel (2017), "Robustness and separation in multidimensional screening." *Econometrica*, 85, 453–488. [1080]
- Chan, Jimmy, Seher Gupta, Fei Li, and Yun Wang (2019), "Pivotal persuasion." *Journal of Economic Theory*, 180, 178–202. [1078]
- Chassang, Sylvain (2013), "Calibrated incentive contracts." *Econometrica*, 81, 1935–1971. [1080]
- Depew, Briggs, Ozkan Eren, and Naci Mocan (2017), "Judges, juveniles, and in-group bias." *Journal of Law and Economics*, 60, 209–239. [1075]
- Dworczak, Piotr and Alessandro Pavan (2020), "Preparing for the worst but hoping for the best: Robust (Bayesian) persuasion." Manuscript. [1080]
- Ellsberg, Daniel (1961), "Risk, ambiguity, and the Savage axioms." *Quarterly Journal of Economics*, 75, 643–669. [1076]
- Ely, Jeffrey C. (2017), "Beeps." *American Economic Review*, 107, 31–53. [1079]
- Guo, Yingni and Eran Shmaya (2019), "The interval structure of optimal disclosure." *Econometrica*, 87, 653–675. [1079]
- Hedlund, Jonas (2017), "Bayesian persuasion by a privately informed sender." *Journal of Economic Theory*, 167, 229–268. [1079]
- Hu, Ju and Xi Weng (2018), "Robust persuasion of a privately informed receiver." Manuscript. [1079]
- Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101, 2590–2615. [1075, 1076, 1078]
- Kolotilin, Anton (2015), "Experimental design to persuade." *Games and Economic Behavior*, 90, 215–226. [1076]
- Kolotilin, Anton (2017), "Optimal information disclosure: A linear programming approach." *Theoretical Economics*, 13, 607–635. [1079]
- Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnjuk, and Ming Li (2017), "Persuasion of a privately informed receiver." *Econometrica*, 85, 1949–1964. [1079]
- Laclau, Marie and Ludovic Renou (2017), "Public persuasion." Manuscript. [1079]
- Ostrovsky, Michael and Michael Schwarz (2010), "Information disclosure and unraveling in matching markets." *American Economic Journal: Microeconomics*, 2, 34–63. [1078]
- Perez-Richet, Eduardo (2014), "Interim Bayesian persuasion: First steps." *American Economic Review: Papers & Proceedings*, 104, 469–474. [1079]

Perez-Richet, Eduardo and Vasiliki Skreta (2018), “Test design under falsification.” Manuscript. [1079]

Pollard, David (2002), *A User’s Guide to Measure Theoretic Probability*. Cambridge University Press, New York, NY. [1081]

Rayo, Luis and Ilya Segal (2010), “Optimal information disclosure.” *Journal of Political Economy*, 118, 949–987. [1078]

Rodina, David (2017), “Information design and career concerns.” Manuscript. [1079]

Rosar, Frank (2017), “Test design under voluntary participation.” *Games and Economic Behavior*, 104, 632–655. [1079]

Co-editor Marina Halac handled this manuscript.

Manuscript received 2 February, 2021; final version accepted 21 August, 2021; available online 30 August, 2021.