COMMUNICATION WITH PARTIALLY VERIFIABLE ENDOGENOUS INFORMATION*

[JOB MARKET PAPER]

[LINK TO LATEST DRAFT]

Matteo Escudé†

16 November 2019

Abstract

An expert can covertly acquire information about the state of the world before communicating with a decision maker in order to influence her action. The expert's information acquisition is unrestricted and costless but her ability to prove to the decision maker what she privately learnt is limited. I study how the verifiability of the expert's acquired information affects equilibrium information acquisition and transmission. Even when acquired information is only partially verifiable, I prove an unravelling result: all equilibria in which the expert influences the decision maker involve *full revelation* of the expert's private information. I then study optimal verifiability environments, giving necessary and sufficient conditions for optimality for each of the two agents. Expert-optimal environments are *credibly rich* in the sense that, even when facing a sceptical decision maker, the expert has access to a rich language to communicate her information. I show that this is akin to her having a large amount of commitment power. The optimum for the decision maker restricts the expert's ability to credibly communicate intermediate results, inducing the expert to acquire and disclose full information in equilibrium.

*I am grateful to Piero Gottardi for invaluable guidance and constant support. I would also like to thank Nina Bobkova and Giacomo Calzolari for their advice and support on this project. I thank Wouter Dessein, Alfredo Di Tillio, Alessandro Ferrari, Mathijs Janssen, Frédéric Koessler, Nenad Kos, David K. Levine, Andrea Mattozzi, Meg Meyer, Pietro Ortoleva, Marco Ottaviani, Antonio Penta, Ludvig Sinander and Alessandro Spiganti for helpful comments and enlightening conversations. I also thank audiences at Bocconi and EUI for useful feedback.

†Department of Economics, European University Institute, Italy. Contact: matteo.escude@eui.eu.

1 Introduction

An expert (sender) wants to influence the action of a decision maker (receiver). The action preferred by receiver depends on the unknown state of the world. At the outset, neither agent possesses any private information about the state. Sender, however, can *covertly* acquire information and then communicate with receiver in order to influence her decision. Sender's information acquisition is unrestricted and costless but her ability to prove to sender what she learnt is limited. How does sender's ability to prove what she learns affect information acquisition? What will ultimately be revealed to receiver?

There are many settings in which partial verifiability plays an important role in shaping information acquisition and transmission. Consider, for instance, a journalist who wants to write about the alleged misbehaviour of a politician. The journalist is biased: she would like her readers to think the politician is likely to be innocent. A reader knows the journalist's bias and would like to know the truth (e.g. he would like to support the politician if and only if he is sufficiently confident that he is innocent). The journalist starts out uninformed, just like her reader, but privately runs an investigation about the politician. After that, she writes her article. When writing, the journalist chooses a narrative tying together alleged findings and events. She may have access to credible evidence that she can present to her reader (e.g. legal documents or camera footage). She may however not be able to prove all that she learnt to the reader. For example, it might be prohibitively costly to prove a fact (e.g. because it would involve exposing an anonymous source, forever damaging the journalist's reputation), illegal (e.g. if a document is classified), or physically impossible (e.g. if the journalist is an eye-witness of an event that she was not able to record).

The reader understands this. Articles are not entirely cheap talk but also can't fully be taken at face value. He also understands that the journalist, anticipating the constraints and limitations she would face when writing, might acquire information selectively to avoid being in situations where she has to reveal information that harms her interests. In the language of this example, I address the following questions. How does the available evidence affect how the journalist conducts her inquiry? What will the reader ultimately learn? To what extend does the journalist want to be able to prove her acquired information? How would the reader want the journalist to be constrained?

Similar issues arise in other settings. A division manager collects information about a project's profitability in order to persuade headquarters to invest. The manager is biased in favour of the investment and her ability to prove her acquired information to headquarters is limited. A seller of an asset of uncertain quality can gather information in order to convince a prospective buyer to purchase. The seller is not necessarily able to prove to the buyer everything she learns.

I study this problem by building a game of communication between two players: sender (the journalist) and receiver (the reader). Sender covertly acquires information about a binary state of the world. She then communicates with receiver, who chooses an action. A tension arises because the two players' preferences are not aligned. I consider a transparent and stark conflict of interest between the two: it is common knowledge that sender (weakly) benefits from receiver holding a higher belief that the state is high.

Partial verifiability is modelled by a fixed map associating each outcome of sender's information acquisition with a set of statements she can make to receiver. I restrict attention to maps in which the information about the state obtained by acquiring information fully determines what statements can be made to receiver. I call this map the *verifiability structure*. Since the set of statements that sender can make varies with the acquired information, each statement constitutes partial proof of sender's private information. Proof is only partial because the same statement may be, potentially, available at different outcomes of sender's covert information acquisition.

To illustrate the role of the verifiability structure, consider two extreme examples. On one hand consider an environment where every outcome of the information acquisition can be proven and must be revealed. This setting can be interpreted as one in which information is publicly acquired, since the communication friction is entirely removed. An alternative interpretation is that information is covertly acquired by sender, but she is *fully committed* to revealing what she learns to receiver. In this setting, sender can only manipulate receiver's action by choosing what information to acquire, since there are no strategic considerations at the information transmission stage.[1] On the other hand consider an environment where any statement can be made, regardless of what sender learnt from acquiring information. This is a setting where all acquired information is 'soft', in the sense that it does not come in the form of evidence that is credible in the eyes of receiver.[2]

In between these extreme cases, following any outcome of the information acquisition, sender may have some evidence to present, but might not be able to fully prove what she learnt. This means that what she communicates to receiver will have both some 'face value' meaning, as specified by the verifiability structure, and some 'equilibrium' meaning, determined by sender's equilibrium information acquisition and communication strategies. The model I build is flexible enough to encompass all of these environments.

The first result (Proposition 1) describes the structure of 'persuasive' equilibria (these are the equilibria in which sender does not get the payoff she would get if there were no interaction with receiver) for a large class of verifiability structures. It shows that, in such equilibria, sender's acquired private information is always *fully revealed* to receiver: there can be no pooling between distinct outcomes of sender's information acquisition. Furthermore, receiver is necessarily *sceptical*, interpreting any on-path statement to mean that sender obtained the worst possible information that is consistent with that statement. Full revelation and scepticism are also key aspects of the classical 'unravelling' result of Milgrom (1981) and Grossman (1981) when sender's (exogenous) private information is fully verifiable.[3] When sender's (exogenous) private information is not fully verifiable, however, the classical result fails: pooling of sender types can occur in equilibrium.[4] I prove that, for persuasive equilibria, the unravelling result is restored if sender's private information is covertly acquired, rather than exogenously given, *even when* sender's information is only partially verifiable.

[1]This is analogous to the model of 'Bayesian persuasion' of Kamenica and Gentzkow (2011). Section 7 further discusses the connection.

[2]This is analogous to the model of 'cheap talk' of Crawford and Sobel (1982) but with sender's private information endogenously acquired rather than exogenously given.

[3]With 'fully verifiable' I mean that sender is able to prove exactly what her type is to receiver.

[4]See, for example, Okuno-Fujiwara, Postlewaite, and Suzumura (1990), Mathis (2008) and Hagenbach, Koessler, and Perez-Richet (2014).

The fact that, in persuasive equilibria, all acquired information is revealed does not mean that sender acquires and transmits *full* information. Indeed, Proposition 1 also shows that the verifiability structure plays a key role in determining *what* information is acquired in persuasive equilibria: sender never acquires more information that what she can prove. The key notion turns out to be *lowest consistency* of the outcomes of sender's information acquisition. An outcome is lowest consistent with some verifiable statement if it is the lowest outcome that allows sender to make that statement, as specified by the verifiability structure. This means that a lowest consistent outcome gives sender access to a piece of evidence that only just credibly separates it from all lower outcomes (but not necessarily from higher ones). In persuasive equilibria only outcomes of sender's information acquisition that are lowest consistent with some verifiable statement can be on-path.

Having determined that the verifiability structure has a crucial role in determining equilibrium acquisition and transmission I study the properties of verifiability structures that make them desirable, in terms of the equilibrium outcomes they induce, by each of the two players.

The second result (Proposition 2) characterizes the set of verifiability structures that allow sender to attain the highest equilibrium payoff among all structures. I first illustrate the equivalence between my model under full commitment (i.e. when the verifiability structure not only allows but *forces* sender to reveal all of her private information) and the model of 'Bayesian persuasion' in Kamenica and Gentzkow (2011). It is well-known that having full commitment power is optimal for sender. Proposition 2 extends this insight, giving necessary and sufficient conditions on the verifiability structure that allow sender to attain the same payoff as under full commitment in any equilibrium. The key condition turns out the be a 'credible richness' property of the verifiability structure: *every* possible outcome of sender's information acquisition must be lowest consistent with some verifiable statement. This means that even if receiver is sceptical, sender still has access to a language that is both rich and credible in order to communicate all outcomes of her inquiry, thereby doing as well as under full commitment.

I then turn to studying verifiability structures that are desirable to receiver, providing a characterization of receiver-optimal structures. The third result (Proposition 3) shows that provided communication can have *some* value to sender, a simple and 'coarse' verifiability structure allows receiver to obtain full information in any equilibrium of the game. It does so by harnessing receiver's scepticism in order to provide high-powered incentives for information acquisition. Specifically, the optimal verifiability structure allows sender to only prove that the state is high and makes all other claims unverifiable. This is optimal for receiver because, in equilibrium, he is maximally sceptical (i.e. he infers that the state is low) when not presented with the only piece of verifiable information that proves that the state is high. This in turn leads him to take the *worst* action for sender when no evidence is presented and the *best* action for sender when evidence is presented. Acquiring full information thus maximises the probability of presenting evidence and minimizes the probability of presenting no evidence.

Layout. Section 2 discusses the relation between this paper and the existing literature. Section 3 illustrates most of the results in the paper with an example in the context of the journalist-reader story. Section 4 describes the model. In Section 5 I present the solution concept and provide existence results. Section 6 contains the equilibrium structure result, showing that full revelation and scepticism necessarily occur in 'persuasive' equilibria. Section 7 is concerned with sender-optimal verifiability structures. Section 8 studies receiver-optimal verifiability structures. Section 9 concludes.

2 Related literature

I study a sender-receiver game of communication augmented with a stage of covert information acquisition. The literature on games of communication is broadly divided into two strands. One is the vast literature of games of communication with evidence.[5] In these games sender is endowed with some private information which she wishes to communicate to receiver in order to influence her action. To do so she has access to a set of messages that depend on her private information. This is in contrast with the literature of 'cheap talk', initiated by Crawford and Sobel (1982), where the reports sender can make to receiver are type-independent. Unlike in these literatures, in my paper sender's private information is not exogenously determined, but covertly and costlessly acquired at the beginning of the game. Additionally, in my environment sender's private information *can* shape what she can communicate with receiver, just like in an evidence game, but may also not affect it, like in a cheap talk game. Indeed, this paper studies how the 'verifiability' of information (that is, the map between sender's private information and the statements she can make) affects the incentives to acquire and transmit information. The class of verifiability structures I study is rich enough to encompass both the standard ones in evidence and cheap talk games.[6] In my game the state of the world is binary and sender's type is the privately observed outcome of her covert information acquisition. Like much of the evidence games literature I maintain a monotonicity assumption on payoffs whereby sender (weakly) benefits from inducing a high belief about the state in receiver.

It is well-known that if sender's (exogenous) private information is only partially verifiable, the classical unravelling result fails. In particular, the literature on partial verifiability with exogenous information (e.g. Mathis (2008), Lipman and Seppi (1995), Okuno-Fujiwara et al. (1990) and Hagenbach et al. (2014)) identifies conditions similar to lowest consistency of all sender types as necessary and sufficient for full revelation of sender's private information in equilibrium.[7] In contrast, I show that when information is endogenously acquired, *even if* information is only partially verifiable, full revelation obtains in all persuasive equilibria. Which signal realizations are lowest consistent in turn determines *what* information is acquired in equilibrium. I also show (Proposition 2) that lowest consistency of every outcome is necessary and sufficient for sender to obtain the largest ex ante equilibrium payoff among all possible verifiability structures.

[5]These games are also known as 'persuasion games' or 'disclosure games'. I will refer to them as 'evidence games', as is done in much of the recent literature, to avoid confusion with the 'Bayesian persuasion' literature. Seminal contributions in the evidence games literature include Grossman (1981), Milgrom (1981), Jovanovic (1982), and Milgrom and Roberts (1986). See Milgrom (2008) for a survey.

[6]Note that models in which evidence is stochastic, such as Dye (1985) and Jung and Kwon (1988), are *not* nested in the way I model verifiability.

[7]Rappoport (2017) provides a general result for constructing equilibria under partial verifiability. Green and Laffont (1986) initiated the study of the role of partial verifiability in mechanism design, establishing conditions under which the revelation principle holds. See Ben-Porath and Lipman (2012) for a study of the implementation problem with partially verifiable information.

The more recent 'Bayesian persuasion' literature, initiated by Kamenica and Gentzkow (2011), considers, like this paper, an environment in which information is endogenous and sender has no private information at the outset.[8] She can freely and unrestrictedly produce *public* information about the state. In an alternative interpretation of their model, she can produce *private* information about the state, but can *fully commit* to disclosing what she learnt to receiver. In my environment sender starts with no private information and can *privately* acquire information about the state. Importantly, she does not have full commitment power of disclosing her results. Both Kamenica and Gentzkow (2011) and Brocas and Carrillo (2007) observe the equivalence between their models of *public* information provision and models of *private* information acquisition in which sender can *fully prove* what she learnt with the same verifiability structure used in standard evidence games.[9] Proposition 2 characterizes sender-optimal evidence in my environment and thus extends these observations. It offers necessary and sufficient conditions on what sender must be able to prove in order for her to attain the full-commitment payoff in equilibrium, in the context of a monotone conflict of interest between sender and receiver.

Relatedly, my environment can also be interpreted as a 'Bayesian persuasion' problem with limited sender commitment. This question has received attention in the literature with approaches different from mine, examples include Fréchette, Lizzeri, and Perego (2018) and Lipnowski, Ravid, and Shishkin (2019).

A closely related paper is DeMarzo, Kremer, and Skrzypacz (2019).[10] The authors also consider a game in which sender covertly and privately acquires information about the state of the world by performing a test. Having observed the result, sender has the option of disclosing or withholding it. Importantly, in a similar way to Dye (1985)'s seminal work, every test may or may not produce a certifiable result. They provide a characterization of equilibrium and prove a result analogous to Dye (1985)'s in the case where information is endogenous: in equilibrium non-disclosure is met with maximal scepticism. A similarity with my paper is that non-observability of information acquisition plays a key role in determining the structure of equilibrium (Proposition 1). However, in their paper the result is driven by the possibility that sender may not have any evidence to present while I explore how different types of deterministic verifiability environments affect the equilibrium outcomes. The implications are very different. They show that the extension of the Dye set-up to endogenous information leads to a similar equilibrium structure. In contrast, in Proposition 1 I show that in persuasive equilibria full revelation of acquired information obtains even under partial verifiability, which does not occur when information is exogenous.

Pei (2015) studies a problem of cheap talk communication preceded by a stage of overt and costly information acquisition. The author also shows that sender communicates all her information in equilibrium. The fact that information acquisition is overt and costly, which is not the case in this paper, is key for the result. Furthermore the paper focuses on purely cheap talk communication, while I study various verifiability environments.

[8]Brocas and Carrillo (2007) study a discrete-time sequential sampling version of the 'Bayesian persuasion' problem and obtain similar results as Kamenica and Gentzkow (2011). The relation between the two approaches is discussed in Morris and Strack (2019).

[9]Gentzkow and Kamenica (2017a) and Gentzkow and Kamenica (2017b) build on this to analyse problems with multiple agents.

[10]See also Acharya, DeMarzo, and Kremer (2011) and Guttman, Kremer, and Skrzypacz (2014) for related models.

A recent literature studies the role of sequential sampling and selective disclosure, addressing related issues in different settings. Recent papers include Argenziano, Severinov, and Squintani (2016), Felgenhauer and Loerke (2017), Di Tillio, Ottaviani, and Sørensen (2019), Janssen (2018), Herresthal (2019).

3 An example

There is uncertainty about a politician's involvement in some scandal. Denote by $\omega = 1$ the state of the world in which the politician is innocent and by $\omega = 0$ the state in which he is guilty. An overtly partisan journalist wants to persuade a reader to vote for the politician. At the outset, both are uninformed and share the common prior belief that attaches probability $p_0 = 1/3$ to $\omega = 1$. The journalist privately gathers information and then writes an article to persuade the reader. After reading the article, the reader can choose among three actions: 'oppose', 'abstain' and 'support'. For concreteness suppose that he opposes if he believes the politician to be innocent with probability below $2/5$, he supports if he believes the politician to be innocent with probability of at least $4/5$, otherwise he abstains. The journalist is overtly partisan: it is common knowledge that she obtains a payoff of 0 if the reader opposes, a payoff of 1 if he abstains and a payoff of 3 if he supports. The actual guilt or innocence of the politician does not affect the journalist's payoff. Figure 1 depicts the journalist's payoff as a function of the reader's belief at the time of decision making. In the general model (described in Section 4) I allow the journalist's payoff to be any nondecreasing and upper semi-continuous function of the reader's belief.

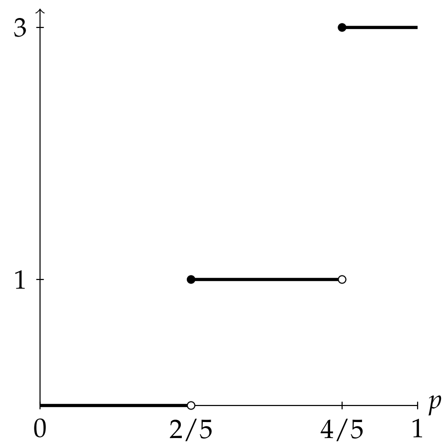

Figure 1: Journalist's value of the reader's beliefs.

3.1 Exogenous information

I consider here the case in which the journalist's private information is exogenous i.e. she privately observes the realization of a signal correlated with the state whose distribution is common knowledge. This is a standard evidence game in the spirit of Milgrom (1981) and

Grossman (1981). This paper is about the case where information is endogenous but for the purposes of this example the exogenous case is a useful benchmark.

Information. The journalist privately observes the outcome of a *signal* $\tilde{S}$ about the state. $\tilde{S}$ is a random variable whose realization s (the journalist's type) takes values in $\{1/6, 1/2, 5/6\}$. Let $\pi_{\tilde{S}}(s|\omega)$ denote the probability that realization $s \in \{1/6, 1/2, 5/6\}$ occurs, conditional on the state being $\omega \in \{0, 1\}$. Write $\tau_{\tilde{S}}(s) \equiv \pi_{\tilde{S}}(s|1)p_0 + \pi_{\tilde{S}}(s|0)(1 - p_0)$ to denote the *marginal* probability mass function of $\tilde{S}$. Naturally, it must be that

$$\pi_{\tilde{S}}(1/6|\omega) + \pi_{\tilde{S}}(1/2|\omega) + \pi_{\tilde{S}}(5/6|\omega) = 1$$

for $\omega \in \{0, 1\}$. Fix $\pi_{\tilde{S}}(1/6|1) = 7/20$, $\pi_{\tilde{S}}(1/2|1) = 3/20$, $\pi_{\tilde{S}}(5/6|1) = 10/20$, $\pi_{\tilde{S}}(1/6|0) = 35/40$, $\pi_{\tilde{S}}(1/2|0) = 3/40$ and $\pi_{\tilde{S}}(5/6|0) = 2/40$.[11] The distribution of $\tilde{S}$ is common knowledge.

These numbers have been chosen so that signal $\tilde{S}$ is *unbiased* i.e. the label of the journalist's type *coincides* with her posterior belief that $\omega = 1$.[12] That is, if she observes $s = 5/6$ she attributes probability $5/6$ to $\omega = 1$: this is 'good' news about the politician's innocence. If she observes $s = 1/6$ she attributes probability $1/6$ to $\omega = 1$: this is 'bad' news about the politician's innocence. If she observes $s = 1/2$ she attributes probability $1/2$ to $\omega = 1$: this is 'mixed' news about the politician's innocence.

Communication. Having observed the signal realization, the journalist must write an article.[13] What she learnt shapes what she can say: communication is not purely cheap talk. Specifically assume there are 2 articles she can write: regardless of the news she obtains ($s \in \{1/6, 1/2, 5/6\}$) she can write a 'lukewarm' article m_L; if 'mixed' or 'good' news ($s \in \{1/2, 5/6\}$) arrive, however, she also has the option of writing a 'laudatory' article m_M. For example, this may be because when she gets 'mixed' or 'good' news she also obtains a piece hard irrefutable evidence (e.g. a legal document, or camera footage) which can be credibly conveyed to the readers in the article. When she gets 'bad' news she obtains no such piece of evidence, so she is stuck with writing the 'lukewarm' article. Observe that when the journalist gets 'good' news ($s = 5/6$) what she can write is the same as when she gets 'mixed' news ($s = 1/2$). For example, this may because whatever distinguishes 'good' from 'mixed' news may be impossible or prohibitively costly to credibly write about. For example the journalist might have acquired such information in classified documents, which she legally cannot reveal; or it might involve exposing anonymous sources, which would forever damage her reputation. Observe that, since what the journalist can say depends on what she learnt, articles constitute (partial) proof to the reader of her private information. For example, article m_M proves to the reader that the journalist observed $s \in \{1/2, 5/6\}$. Figure 2 depicts sender's three possible types and associated feasible articles.

[11] The associated marginal probability mass function is given by $\tau_{\tilde{S}}(1/6) = 7/10$, $\tau_{\tilde{S}}(1/2) = 1/10$, $\tau_{\tilde{S}}(5/6) = 2/10$.

[12] In general, a signal $\tilde{S}$ is unbiased iff $\mathbf{E}(\omega|s) = s$ for all $s \in \operatorname{supp} \tilde{S}$. $\tilde{S}$ being unbiased is a normalization equivalent to imposing $\tau_{\tilde{S}}(s)s = \pi_{\tilde{S}}(s|1)p_0$ for all $s \in \operatorname{supp} \tilde{S}$.

[13] Adding the option of not writing an article does not affect the result.

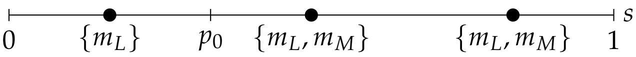

Figure 2: Journalist's types and feasible articles.

Equilibrium. Having fixed the reader's decision rule, an equilibrium is an article-writing strategy and a receiver belief function; the choice of article must be optimal given the belief function and the belief function must obey Bayes' rule when possible and always be consistent with the evidence.

In all equilibria, 'mixed' and 'good' news types ($s = 1/2$ and $s = 5/6$) *pool* and write the laudatory article m_M. The 'bad' news type ($s = 1/6$) separates and writes m_L. Types $s = 1/2$ and $s = 5/6$ induce a belief of $13/18 < 4/5$ in receiver[14], thus leading to the intermediate action 'abstain' (with probability $3/10 = \tau_{\tilde{S}}(1/2) + \tau_{\tilde{S}}(5/6)$). Type $s = 1/6$ induces belief $1/6 < 2/5$, leading to the bad action 'oppose' (with probability $7/10 = \tau_{\tilde{S}}(1/6)$). The journalist's equilibrium expected payoff is therefore $3/10$. Figure 3 depicts the equilibrium articles and induced action.

To see why, observe that in equilibrium article m_L must necessarily induce action 'oppose'. That's because journalist type $s = 1/6$ always writes it and there is no combination of types that would lead the reader to hold a belief at or above $2/5$.[15] It is then immediate that both $s = 1/2$ and $s = 5/6$ will send message m_M in any equilibrium and separate from $s = 1/6$ to induce action 'abstain'.

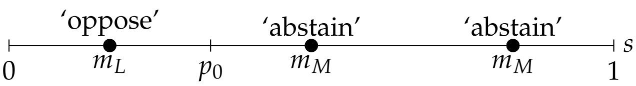

Figure 3: Equilibrium messages (below) and reader action (above).

3.2 Endogenous information

Information. Now a signal S is *covertly* chosen by the journalist among all possible unbiased signals about the state. Just like in the exogenous case, its realization is *privately* observed by the journalist. For example she can run a fully informative investigation S^F (with supp $S^F = \{0, 1\}$) or a fully uninformative one S^U (with supp $S^U = \{p_0\}$). Figure 4 depicts these two signals. She can also choose the signal $\tilde{S}$ from the exogenous information benchmark, for example.

(a) S^U (b) S^F

Figure 4: Fully uninformative and informative signals.

[14] $13/18 = (1/3) \times (1/2) + (2/3) \times (5/6)$ is receiver's posterior belief conditional on knowing that either $s = 1/2$ or $s = 5/6$ realized.

[15] In more detail. If both $s = 1/6$ and $s = 1/2$ send m_L the reader's posterior is $5/24 < 2/5$. If both $s = 1/6$ and $s = 5/6$ send m_L the reader's posterior is $17/63 < 2/5$. If all types send m_L the posterior is just the prior $1/3 < 2/5$.

Communication. Just as in the exogenous case, having observed the signal realization, the journalist writes an article. I extend the constraint from the exogenous case in the following way. The lukewarm article m_L can be written for any realization $s \in [0, 1]$; the laudatory article m_M can be chosen for any $s \in [1/2, 1]$. So, as in the exogenous benchmark, m_L proves nothing to the reader while article m_M proves to the reader that the journalist's interim type is at least 1/2. Figure 5 illustrates this. Again, articles constitute (partial) proof of what the journalist learnt. In the general model (described in Section 4) I study *any* possible map between what the journalist privately learns and what she can communicate with receiver.

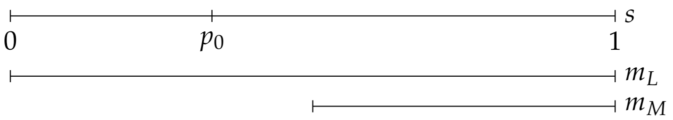

Figure 5: Available articles at each realization.

Equilibrium. The journalist's strategy involves acquiring information and writing an article that is feasible given what she learnt. The reader forms a belief about the state after reading an article and chooses an action. Equilibrium imposes that the journalist's information acquisition and article choice are optimal given how the reader forms beliefs and chooses actions. The reader's action given her belief is specified by the given decision rule, his belief must be formed using Bayes' rule when possible. When Bayes' rule cannot be used, the reader's belief must still be consistent with the evidence (i.e. she cannot hold a belief below 1/2 when she sees article m_M, even if m_M were not written with positive probability in equilibrium).[16]

In all equilibria, the journalist chooses signal S^* with $\operatorname{supp} S^* = \{0, 1/2\}$. That is, the equilibrium information acquisition strategy involves the extreme 'bad' news outcome (in which the journalist learns that the politician is guilty for sure) and a 'mixed' news outcome (in which the journalist attributes equal probability to each state). In equilibrium, when she learns 'mixed' news the journalist writes the laudatory article m_M while when she learns 'bad' news she writes the lukewarm article m_L. In equilibrium, after reading article m_L the reader updates to posterior $p = 0$ and chooses to 'oppose', when he reads m_M he updates to posterior $p = 1/2$ and 'abstains'. The journalist's equilibrium payoff is therefore 2/3.

The argument goes as follows. Observe first that following m_L the reader must 'oppose' in equilibrium. If he did not (i.e. he held any belief at or above 2/5) the journalist would be able to never have the reader choose 'oppose', since m_L is always available. But the reader must *sometimes* 'oppose' in equilibrium, since at the prior $p_0 = 1/3 < 2/5$ he would.[17]

Consider next the reader's response following the laudatory article, m_M. He cannot 'oppose', since his belief following m_M cannot be below $1/2 > 2/5$. Could he 'support' (i.e. hold a belief at or above 4/5)? If he did sender would choose an information acquisition strategy that maximises the probability of writing a laudatory article. It is straightforward to check that this is the signal S^* with $\operatorname{supp} S^* = \{0, 1/2\}$ as it allows the journalist to write the laudatory article

[16]See Section 5 for a detailed definition of equilibrium.

[17]Otherwise his expected posterior belief would differ from the prior; this cannot be if he uses Bayes' rule following articles that are are written with positive probability.

(inducing action 'support') with probability 2/3 and the lukewarm article (inducing 'oppose') with complementary probability 1/3. But this means that the lukewarm article is written when the journalist observed realization $s = 1/2$, so the reader holding a belief at or above 4/5 is not in line with Bayes' rule. It must therefore mean that following the laudatory article the reader chooses to 'abstain' in equilibrium. The journalist's unique best reply is, again, to choose S^*, as it maximises the probability of writing the laudatory article. It is immediate that the only equilibrium beliefs of receiver are necessarily $p = 0$ following m_L and $p = 1/2$ following m_M.

3.2.1 Discussion

Observe some salient characteristics of equilibrium. There is *full revelation* of the journalist's acquired information i.e. there is *no pooling* of journalist interim types. This holds *even if* there can be pooling when information is exogenous, because of partial verifiability, as illustrated in the benchmark. Put differently, interim type distributions that would lead to equilibrium pooling in the exogenous benchmark cannot arise in equilibrium when the distribution is covertly chosen by the journalist. The reason is intuitive. Consider for example the exogenous information benchmark signal $\tilde{S}$ where both 'mixed' and 'good' news arise with positive probability. Given the limitations of the available evidence, the journalist has no way of proving to the reader that news are 'good' rather than 'mixed', so these two types pool in equilibrium. In the case where information is endogenous, obtaining with some probability 'good' news carries an implicit cost: since it's a better outcome than 'mixed' news it is also less likely to arise, so the *total* probability of inducing action 'abstain' is lower than if she just sought 'mixed' news with positive probability. Another related feature of the equilibrium is that the reader is *sceptical* following each article: he (correctly) interprets each as the worst possible news that is consistent with it; this another feature of equilibrium which is not necessarily true when information is exogenous.

Proposition 1 generalizes these insights to any monotone conflict of interest between the two players (i.e. when the journalist weakly benefits from inducing a higher belief) and any possible map between what the journalist can prove and what she learnt.

Finally, observe that what the journalist can and cannot prove directly shapes equilibrium information acquisition and transmission. In particular, it is the *lowest* signal realizations that are consistent with *some* article ($s = 0$ with m_L and $s = 1/2$ with m_M) that occur with positive probability in equilibrium: the journalist never attempts to obtain 'better information' than what she is able to prove. Also this insight generalizes and naturally leads to the questions: given that all that is learnt by the journalist is necessarily revealed in equilibrium, what kind of evidence is desirable from the journalist's perspective? From the reader's? This is the focus of the next sections.

3.3 Optimal evidence when information is endogenous

3.3.1 Journalist

Does the journalist benefit from being in an environment with 'more evidence'? Or would she rather have a lot of 'leeway' in how she represents her acquired information? Consider adding a third, 'jubilant', possible article: the journalist can only write it if the outcome of her information acquisition is a realization at or above 5/6. This could be, for example, because some document proving the politician's very likely innocence is no longer classified and can therefore be reported in a (jubilant) article. Figure 6 depicts the new communication capabilities of the journalist.

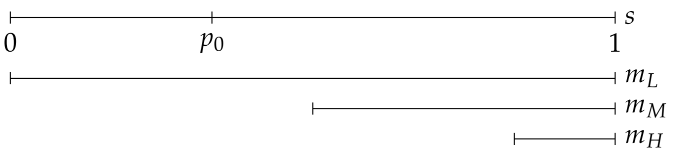

Figure 6: Available articles at each realization.

By an argument analogous to the one in the previous section, one can show that the only equilibrium signal is S^{**} with supp $S^{**} = \{0, 5/6\}$. Following realization $s = 0$ ('bad' news) the journalist writes the lukewarm article m_L and the reader, correctly updating to $p = 0$, chooses 'oppose'. Following realization $s = 5/6$ ('good' news) the journalist writes the jubilant article m_H and the reader, correctly updating to $p = 5/6$, chooses 'support'. Sender's equilibrium expected payoff is higher than when she only had access to m_L and m_M: it is now 6/5 while before it was 2/3.

The journalist can now 'prove more', in the sense that she has access to a (credibly) richer language to communicate her acquired information, and this makes her better off. What if the journalist's has access to as many credible messages at there are possible outcomes of the investigation? For example, what if at realization $s \in [0, 1]$ she can report any $m \in [0, s]$? In Proposition 2 I show that this kind of evidence is ex ante optimal for the journalist in the sense that in any equilibrium she attains the highest payoff across all equilibria for any possible verifiability structure. This is the 'full commitment' payoff i.e. the payoff she would obtain if she was forced to disclose the outcome s of the information acquisition; it coincides with the equilibrium payoff in the 'Bayesian persuasion' model of Kamenica and Gentzkow (2011). In this example this is achieved with a signal S^{FC} with supp $S^{FC} = \{0, 4/5\}$.

More generally, Proposition 2 provides necessary and sufficient conditions that the evidence needs to satisfy to be optimal for the journalist. The condition is precisely that every possible realization of the journalist's information acquisition is the lowest consistent (as specified by the available evidence) with *some* article.

3.3.2 Reader

Mandated disclosure is optimal for the journalist, but may lead her to acquire little information in equilibrium. What kind of evidence is desirable by the reader? Suppose there are again two available articles: the lukewarm one as usual, that can be written by any journalist type, and a 'conclusively jubilant' one, that can *only* be written if the journalist obtains outcome $s = 1$ i.e. if she learns that the politician is innocent with certainty. This could be, for example, because the only piece of evidence available in this environment is one that proves certain innocence; no other evidence that can be presented to the reader is available. Figure 7 depicts the new communication capabilities of the journalist.

Figure 7: Available articles at each realization.

By an argument analogous to the one in the previous sections, one can show that the only equilibrium signal is S^F with supp $S^F = \{0, 1\}$: the journalist acquires full information. Following realization $s = 0$ the journalist writes the lukewarm article m_L and the reader, correctly updating to $p = 0$, chooses 'oppose'. Following realization $s = 1$ the journalist writes the conclusively jubilant article m_H offering the definitive proof of the politician's innocence; the reader, correctly updating to $p = 1$, chooses 'support'. Notice that the reader chooses the action knowing the value of the state i.e. under full information. If the reader is an expected utility maximiser, for example, choosing the action under full information allows her to attain the highest ex ante expected payoff in the decision problem.[18] With this kind of evidence the journalist is stripped from the ability of credibly communicating any intermediate results. Following the lukewarm article the reader is necessarily sceptical, and 'opposes'. The only way the journalist can convince the reader to 'support' with positive probability is to obtain certain proof of innocence. The optimal way of doing so is to acquire full information, as this maximises the probability of obtaining result $s = 1$ and being able to write the conclusively jubilant article.

Proposition 3 generalizes this insight to any monotone conflict of interest, provided that the journalist is not already attaining the highest possible payoff at the prior (in this example, this holds since $p_0 = 1/3 < 4/5$). Allowing the journalist to only prove the 'best possible' news (i.e. $s = 1$) leads to, at every equilibrium outcome, *full* information acquisition and transmission. Furthermore, only allowing 'best possible' news to be provable is also necessary for the reader to obtain the full-information payoff in all equilibria. This occurs because, in equilibrium, any claim sender makes about having obtained an intermediate result is not credible. So, while she can freely choose any investigation, it is only the fully informative one that can arise in equilibrium.

4 Model

I study a game of incomplete information between two players, called sender (she) and receiver (he). The state of world ω can be either 0 or 1, drawn by nature at the beginning of the game such that the probability that $\omega = 1$ is $p_0 \in [0, 1]$. At the outset the two players share the common prior belief p_0. Sender first covertly acquires information about the state and then communicates with receiver; this part of the game is described in detail in the following subsections. Following communication, receiver updates his belief about the state, given what he believes sender privately learnt and what she communicated, and takes an action.

[18]The reader's decision rule in this example can be micro-founded with expected utility preferences. Let the payoff from 'oppose' be 2 if $\omega = 0$ and -3 if $\omega = 1$. Let the payoff from 'support' be -4 if $\omega = 0$ and 1 if $\omega = 1$. Let the payoff from 'abstain' be 0 regardless of the state.

4.1 Preferences

I assume that if receiver holds belief $p \in [0,1]$ about the state at the time of choosing the action sender earns a payoff of $v(p)$. I call the function $v : [0,1] \to \mathbb{R}$ the *value of induced posteriors*. I focus on a specific form of transparent and stark conflict of interest between sender and receiver: it is common knowledge that sender weakly benefits from receiver holding a higher belief about the state. Formally, I assume that v is nondecreasing.[19] I also assume that v is upper semi-continuous. For most of the paper I shall directly work with v rather than with one of its possible micro-foundations. Two natural micro-foundations are described in the examples below.

Example (Expected utility and single-crossing)**.** Let receiver's action set be $[0,1]$. Receiver has a Bernoulli utility function $u_R : \{0,1\} \times A \to \mathbb{R}$ while sender has a nondecreasing, state-independent Bernoulli utility function $u_S : A \to \mathbb{R}$; they both maximise expected utility. Assume that $u_R(\omega, \cdot)$ is continuous for each $\omega \in \{0,1\}$ and that u_R satisfies the strict single-crossing property i.e. $u_R(0, a') \geq u_R(0, a)$ implies $u_R(1, a') > u_R(1, a)$ for $a' > a$. By Milgrom and Shannon (1994)'s monotone selection theorem every selection of optimal receiver actions is nondecreasing in his belief p. Choose any selection of optimal receiver actions $\hat{a}$ that breaks ties in favour of sender. Then $v \equiv u_S \circ \hat{a}$ is nondecreasing and upper semi-continuous.

Example (Receiver is a 'market')**.** Let receiver's action set be $A = [0,1]$. Receiver maximises $u_R(a, p) = -(a-p)^2$ while sender maximises a nondecreasing and upper semi-continuous $u_S(a)$. Receiver's optimal decision rule is given by $\hat{a}(p) = p$ and therefore $v \equiv u_S \circ \hat{a}$ is nondecreasing and upper semi-continuous. An interpretation is that sender is the seller of a good of unknown quality ω, about which she can acquire information. The action is the payment that sender obtains from selling the good in some market. The market—for which receiver is a stand-in—is willing to pay the good its expected quality. Sender draws no value from holding the good, but she (weakly) values money.

In Section 8 I study receiver's welfare so there I will model receiver's preferences explicitly. Until then, I shall work with a generic nondecreasing and upper semi-continuous v, which captures receiver's behaviour given his belief and the implied payoff for sender.

4.2 Information acquisition

The game starts with sender covertly and costlessly acquiring information about the state. She does so by choosing an unbiased, finite-support *signal* S. S is a random variable correlated with the state of the world ω, defined by its conditional probability mass function $\pi_S(\cdot|\omega)$ over $[0,1]$. Signal S is unbiased if and only if

$$\mathbf{E}(\omega|s) = s \quad \text{for all} \quad s \in \operatorname{supp} S$$

i.e. s coincides with sender's private posterior belief about the posterior mean. Sender's information acquisition is unrestricted in the sense that she can choose S from set Σ, which is

[19] An analogous monotonicity assumption is also maintained in much of the evidence games literature. Notable examples that relax this assumption are Seidmann and Winter (1997), Giovannoni and Seidmann (2007) and Hagenbach et al. (2014).

the set of all unbiased, finite-support signals.[20] Let $\tau_S(\cdot)$ denote its marginal probability mass function. Restricting attention to unbiased signals rules out the possibility that two distinct signal realizations lead to the same posterior belief about the state for sender. It does not rule out anything else in terms of sender's ability to acquire information as signal realizations can simply be relabelled with the associated posterior belief.[21]

After choosing $S \in \Sigma$, nature draws a realization $s \in \operatorname{supp} S$, which sender privately observes. Note that at this point sender's private information is her hidden choice of signal S and its realization s.

4.3 Information transmission

After information acquisition comes information transmission: sender can communicate what she learnt about the state to receiver. I make a substantial assumption about how what sender learnt affects what she can say to receiver: it is only the information about the state that a signal provides that influences sender's ability to make certain statements. That is, only the signal realization (which coincides with sender's private posterior belief about state, since sender can only choose unbiased signals) determines what statements she can make, not the choice of signal itself. This means that, in the model, it is not possible for sender to acquire the same information about the state in ways that vary in their verifiability.

Formally, let $M : [0,1] \rightrightarrows 2^K$ be a nonempty-valued correspondence associating with each possible signal realization $s \in [0,1]$ a set of messages $M(s) \subseteq 2^K$ for some compact $K \subseteq \mathbb{R}$, $K \supseteq [0,1]$. Let $\mathcal{M} \equiv \cup_{s \in [0,1]} M(s)$ be the set of all possible messages in this environment for a given M. If sender covertly chooses signal $S \in \Sigma$ and observes realization $s \in S$ she must choose a message m from $M(s)$, which is observed by receiver. The interpretation is that message $m \in \mathcal{M}$ *proves* to receiver that sender's private information about the state lies in the set

$$M^{-1}(m) \equiv \{s \in [0,1] : m \in M(s)\}.$$

I refer to the correspondence M as the *verifiability structure* of this environment. Receiver, having observed m, updates his belief about the state of the world and chooses an action.

4.3.1 Examples of verifiability structures

The verifiability structure determines precisely what sender can prove to receiver. On one extreme consider the 'full commitment' verifiability structure, where $M(s) = \{s\}$ for each $s \in [0,1]$. For every signal realization there is a single message (which sender must send) that fully reveals the realization to receiver. This makes the communication stage completely mechanical, and is therefore akin to the signal realization being publicly observed. The only tool available to sender to affect receiver's decision is the choice of signal.[22] On the other extreme consider a 'cheap talk' verifiability structure, where $M(s) = [0,1]$ for each $s \in [0,1]$.

[20] That is, sender chooses a finite subset of $[0,1]$ (the support of S) and conditional probability mass functions $\pi_S(\cdot|0)$ and $\pi_S(\cdot|1)$ over the chosen set with the restriction that $[\pi_S(s|1)p_0 + \pi_S(s|0)(1-p_0)]s = \pi_S(s|1)p_0)$ for all s in the set.

[21] It will become clear in the next subsection how this relates to sender's ability to prove her acquired information.

[22] This verifiability structure is closely connected to the model of 'Bayesian persuasion' of Kamenica and Gentzkow (2011). The connection is spelled out in detail in Section 7.

For every signal realization there is the same set of messages available: each message does not prove anything to receiver.

The verifiability structure that describes 'classical' evidence à la Milgrom (1981) and Grossman (1981) is $M(s) = \{m \in 2^{[0,1]} : s \in m\}$ for each $s \in [0,1]$. At every realization sender can report a subset of $[0,1]$ provided that the true realization is in that set. The interpretation is that sender must speak the 'truth' but not necessarily 'the whole truth'.

Other examples include the 'leeway' verifiability structure with $M(s) = [s-\epsilon, s+\epsilon]$ for some $\epsilon > 0$: every sender can misreport the signal realization by exaggerating/understating it by at most ϵ and the 'exaggeration' ('understatement') one with $M(s) = [s,1]$ ($M(s) = [0,s]$): for any realization sender can report any realization higher (lower) that the one she obtained.

As an example of a finite verifiability structure consider the 'ordered intervals' one. Let $s_0 < s_1 < \cdots < s_n$ denote elements of $[0,1]$, with $s_0 = 0$ and $s_n = 1$. Let there be a total of n messages $\{m_1, m_2, \ldots, m_n\}$. Let the verifiability structure be so that $m_i \in M(s)$ iff $s \in [s_{i-1}, s_i]$ for all $i \in \{1, 2, \ldots, n\}$. This verifiability structure gives sender the ability to coarsely prove her acquired information.

5 Solution concept and existence

5.1 Definition of equilibrium

Fix a prior $p_0 \in [0,1]$, verifiability structure M and a nondecreasing, upper semi-continuous $v : [0,1] \to \mathbb{R}$. A pure strategy for sender is a feasible information acquisition policy $S \in \Sigma$ and a messaging rule $\mu : \Sigma \times [0,1] \to \mathcal{M}$ such that $\mu(S,s) \in M(s)$. A belief function for receiver is a map $p_R : \mathcal{M} \to [0,1]$ associating with each message a belief *about the state*, so that $p_R(m) \in [0,1]$ is the probability that receiver attributes to $\omega = 1$ after seeing message m. The solution concept is in the spirit of perfect Bayesian equilibrium in pure strategies; it is defined as follows. Consider a sender strategy (S^*, μ^*) and a receiver's belief function p_R^*. I say the triple (S^*, μ^*, p_R^*) is an equilibrium if it satisfies the following conditions. The belief function p_R^* must be consistent with the verifiability structure and formed using Bayes's rule whenever possible. Formally, if for some $m \in \mathcal{M}$ there exists a $s \in \operatorname{supp} S^*$ such that $\mu^*(S^*, s) = m$ (i.e. m is on-path) it must be that

$$p_R^*(m) = \frac{\sum_{\{s \in \operatorname{supp} S^* : \mu^*(S^*,s) = m\}} \tau_{S^*}(s) s}{\sum_{\{s \in \operatorname{supp} S^* : \mu^*(S^*,s) = m\}} \tau_{S^*}(s)}. \tag{1}$$

If message $m \in \mathcal{M}$ is such that there is no $s \in \operatorname{supp} S^*$ with $\mu^*(S^*, s) = m$ (i.e. m is off-path) the belief is only required to be consistent with the verifiability structure: it must be that

$$p_R^*(m) \in \operatorname{conv} M^{-1}(m) \tag{2}$$

where conv X denotes the convex hull of set X.[23] Given how receiver forms beliefs about the state and how she behaves at each belief sender's messaging strategy must be optimal so

$$v(p_R^*(\mu^*(S,s))) \geq v(p_R^*(m')) \tag{3}$$

for all $S \in \Sigma$, all $s \in \operatorname{supp} S$ and all $m' \in M(s)$. Finally, sender's information acquisition must be optimal given beliefs and strategies in the rest of the game,

$$\sum_{s \in \operatorname{supp} S^*} \tau_{S^*}(s) v(p_R^*(\mu^*(S^*,s))) \geq \sum_{s \in \operatorname{supp} S'} \tau_{S'}(s) v(p_R^*(\mu^*(S',s))) \tag{4}$$

for any $S' \in \Sigma$.[24]

5.2 Equilibrium existence

I provide an equilibrium existence result by providing three jointly sufficient conditions on M. The conditions are technical and satisfied, as discussed after the statement of Corollary 1, in most of the examples I consider.

Define the function $f_M : [0,1] \to [0,1]$, $f_M(s) = \sup_{m \in M(s)} \inf M^{-1}(m)$. The interpretation is that f_M is a map associating to each possible signal realization a belief for receiver. Specifically, it is the map obtained by choosing at each s the message that induces the highest belief in receiver, given that receiver is sceptical. I provide three conditions on the verifiability structure that are jointly sufficient for equilibrium existence. A discussion of each follows the statement of the result.

Assumption 1. *$M(s)$ is finite for all $s \in [0,1]$.*

Assumption 2. *$M^{-1}(m)$ is closed for all $m \in \mathcal{M}$.*

Assumption 3. *f_M is upper semi-continuous.*

Lemma 1 (Existence)**.** *If M satisfies Assumptions 1-3 an equilibrium in pure strategies exists.*

Proof. See Appendix B. □

The interpretation of Assumption 1 is immediate: at each signal realization sender has access to a finite set of messages she can send to receiver. Assumption 2 states that if some converging sequence of signal realizations is able to send message m then also the limit signal realization is able to send message m. Assumption 3 is also a continuity requirement. It is easily interpreted if Assumption 2 also holds: then it states that if receiver is always sceptical (i.e. $p_R(m) = \min M^{-1}(m)$ for all $m \in \mathcal{M}$) the map associating with each signal realization the highest possible inducible belief is upper semi-continuous.[25]

In the special case where the set of all messages in the environment is finite Assumption 1 holds a fortiori. It is also straightforward to show that Assumption 3 holds as well. The following corollary thus states the existence result when $\mathcal{M}$ is finite.

[23] That is, following any off-path message m receiver can hold any belief *over signal realizations* with support on $M^{-1}(m)$, his belief *about the state* thus necessarily lies in conv $M^{-1}(m)$.

[24] Observe that it is not necessary to consider double deviations for sender since μ^* is optimal given receiver's belief, which does not respond to sender deviating to a different signal since such a deviation is not observed.

[25] Appendix I provides an example in which no equilibrium exists when Assumption 3 fails.

Corollary 1 (Finite case)**.** *If M satisfies Assumption 2 and is such that $\mathcal{M}$ is finite an equilibrium in pure strategies exists.*

Proof. See Appendix B. $\square$

Remark. Since v is assumed to be upper semi-continuous, the results immediately give existence for the 'full commitment', 'cheap talk', 'leeway', 'exaggeration', 'understatement' and 'ordered intervals' verifiability structures from Section 4.3.1. For the 'classical' verifiability structure it is further necessary to assume that each $m \in 2^{[0,1]}$ is closed to apply the existence result.

In the rest of the paper Assumptions 1-3 are not maintained. Whenever they are needed they are explicitly invoked.

6 Equilibrium structure: full revelation and scepticism

I study the properties of pure-strategy equilibria of this game given a nondecreasing and upper semi-continuous v, verifiability structure M and prior $p_0 \in [0,1]$. Say an equilibrium is *unpersuasive* if sender's equilibrium payoff is equal to $v(p_0)$ with probability 1.[26] In an unpersuasive equilibrium sender obtains with certainty the same payoff she would obtain if the possibility of information acquisition and communication did not exist. Call any other equilibrium *persuasive*.[27]

I first provide some definitions that make the statement of the result more straightforward. For a message $m \in \mathcal{M}$ I say that signal realization $s \in [0,1]$ is *consistent with m* iff $s \in M^{-1}(m)$. I say that signal realization $s \in [0,1]$ is *lowest consistent with m* iff $\min M^{-1}(m)$ exists and $s = \min M^{-1}(m)$. For a given M I say that signal realization $s \in [0,1]$ is *lowest consistent* iff there exists an $m \in M(s)$ such that s is lowest consistent with m; no ambiguity regarding which verifiability structure M is under consideration should ever arise. Lowest consistency will turn out to be a key property in all of the results that follow.[28] The following result describes all persuasive equilibria.

Proposition 1 (Equilibrium structure)**.** *Let (S^*, μ^*, p_R^*) be a persuasive equilibrium. For all $s \in \operatorname{supp} S^*$:*

(I) *s is lowest consistent with $\mu^*(S^*, s)$ i.e. $\mu^*(S^*, s) \in \{m \in M(s) : s = \min M^{-1}(m)\} \neq \emptyset$;*

(II) *$p_R^*(\mu^*(S^*, s)) = \min M^{-1}(\mu^*(S^*, s)) = s$.*

The result provides necessary conditions that must be satisfied by sender's strategy and receiver's belief function in any persuasive equilibrium. Observe first that the verifiability structure plays an important role in shaping equilibrium information acquisition. Condition (I) says that only signal realizations that are lowest consistent with some message can be part

[26] That is, equilibrium (S^*, μ^*, p_R^*) is unpersuasive iff $v(p_R^*(\mu^*(S^*, s))) = v(p_0)$ for all $s \in \operatorname{supp} S^*$.

[27] In Section 6.1 I provide sufficient conditions for equilibria of each kind.

[28] Some of the literature on evidence games (e.g. Hart, Kremer, and Perry (2017), Ben-Porath, Dekel, and Lipman (2019) and Rappoport (2017)) works with the notion of disclosure order rather than with message sets. Each verifiability structure induces a disclosure order over $[0,1]$ but there is no natural analogue of the concept of lowest consistency when using the disclosure order. See Appendix H for a brief discussion.

sender's information acquisition strategy in any persuasive equilibrium. It further states that at any on-path signal realization s sender *must* be sending one of the messages with which s is lowest consistent. Condition (II) states the *full revelation* and *scepticism* result; it is an immediate consequence of condition (I). The first equality states that, in any persuasive equilibrium, receiver must necessarily interpret all on-path messages with scepticism: he associates with such messages the lowest signal realization that is consistent with each of them. The second equality illustrates that such belief *coincides* with the signal realization privately observed by sender (i.e. sender's private posterior belief about the state): in every persuasive equilibrium, at the time of decision making, receiver has the same information about the state as sender. Put differently, there can be no pooling of information acquisition outcomes in equilibrium. Naturally, the fact that everything learnt by sender is revealed to receiver in equilibrium does not mean that receiver learns the value of the state, since sender's equilibrium signal may not be the fully informative one.

Proof sketch of Proposition 1. See Appendix C for a full proof. I illustrate the logic of the proof for some candidate persuasive equilibrium (S^*, μ^*, p_R^*) where S^* is binary. Let $\operatorname{supp} S^* = \{s_L, s_H\}$ with $s_H > p_0 > s_L$; it is straightforward to check that $\tau_{S^*}(s_H) = (p_0 - s_L)/(s_H - s_L)$ and $\tau_{S^*}(s_L) = (s_H - p_0)/(s_H - s_L)$ since S^* is unbiased. Let m_H and m_L denote the on-path messages, respectively. $m_H \neq m_L$ since the equilibrium is persuasive. Observe that $v(p_R^*(m_H)) > v(p_R^*(m_L))$. This is because v is nondecreasing so $v(p_R^*(m_H)) \leq v(p_R^*(m_L))$ would imply $p_R^*(m_H) \leq p_R^*(m_L)$, contradicting that receiver uses Bayes' rule to update after on-path messages. Suppose, towards a contradiction, that s_H is not lowest consistent with m_H, so that there exists some $s < s_H$ such that $m_H \in M(s)$. There are three possibilities, each with a profitable deviation. (i) $s > p_0$. Then signal S' with $\operatorname{supp} S' = \{s_L, s\}$ is in Σ and is such that $\tau_{S'}(s) = (p_0 - s_L)/(s - s_L) > (p_0 - s_L)/(s_H - s_L) = \tau_{S^*}(s_H)$. Since $m_H \in M(s)$ and $v(p_R^*(m_H)) > v(p_R^*(m_L))$ sender's expected payoff from choosing S' and using the same on-path messages as in the candidate equilibrium is a profitable deviation. (ii) $s = p_0$. Then the uninformative signal (with $\operatorname{supp} S^U = \{p_0\}$) gives sender a payoff of $v(p_R^*(m_H))$ for sure, so this constitutes a profitable deviation. (iii) $s < p_0$. Then signal S' with $\operatorname{supp} S' = \{s, s_H\}$ is in Σ and gives sender a payoff of $v(p_R^*(m_H))$ for sure, so is a profitable deviation. So it must be that m_H is a message with which s_H is lowest consistent. Similarly, towards a contradiction, suppose that s_L is not lowest consistent with m_L, so that there exists some $s < s_L$ such that $m_L \in M(s)$. Then signal S' with $\operatorname{supp} S' = \{s, s_H\}$ is in Σ and is such that $\tau_{S'}(s) = (s_H - p_0)/(s_H - s) < (s_H - p_0)/(s_H - s_L) = \tau_{S^*}(s_L)$. Since $m_L \in M(s)$ and $v(p_R^*(m_H)) > v(p_R^*(m_L))$ sender's expected payoff from choosing S' and using the same on-path messages as in the candidate equilibrium is a profitable deviation. So it must be that m_L is a message with which s_L is lowest consistent. Point (II) is immediate since if receiver uses Bayes' rule on-path it must mean that $p_R^*(m_L) = s_L = \min M^{-1}(m_L)$ and $p_R^*(m_H) = s_H = \min M^{-1}(m_H)$. □

Observe the similarities and differences with evidence games à la Milgrom (1981) and Grossman (1981). Scepticism and full revelation are also key features of equilibria in those games. However, the classical 'unravelling' result of those papers does not obtain when sender's private information is only partially verifiable. My result shows that, even if sender's private information is only partially verifiable, full revelation must occur in persuasive equilibria of communication games with endogenous information. This is because only information acquisition strategies such that the *acquired* information *is* verifiable can be part of persuasive equilibria, and therefore full revelation ensues.

Observe that the reason why this obtains is different from the argument in standard evidence games with exogenous information. There the highest sender type (as measured by the payoff she would get if receiver knew it) obviously separates from the rest, as she can get the highest attainable payoff. Once it has separated, the second-highest type faces a similar situation, and so on. It is clear if verifiability is only partial (e.g. because all the evidence held by a higher type is also held by some lower one) some types will pool in equilibrium. I show that—given the verifiability structure—sender chooses, in persuasive equilibria, a distribution of signal realizations such that pooling cannot occur. The reason is that, in equilibrium, sender never chooses to learn more than what she is able to prove. Doing so would carry an implicit cost, due to more extreme realizations being less likely, and no gain, since there isn't any evidence to credibly convey what was learnt. While receiver does not observe sender's information acquisition, she correctly anticipates this. She is therefore sceptical when presented with a message and interprets it as the 'worst news' consistent with it.[29]

Remark. For the 'classical' ($M(s) = \{m \in 2^{[0,1]} : s \in m\}$), 'understatement' ($M(s) = [0, s]$) and 'leeway' ($M(s) = [s - \epsilon, s + \epsilon]$) verifiability structures *all* signal realizations are lowest consistent. Hence (I) in Proposition 1 only specifies *which* messages are sent in persuasive equilibria. These are, at each $s \in [0, 1]$: any m such that $\min m = s$ for the 'classical' case, $m = s$ for the 'understatement' case and $m = s + \epsilon$ for the 'leeway' case.

Remark. Consider the 'ordered intervals' verifiability structure defined in Section 4.3.1. (I) implies that in all persuasive equilibria the support of the equilibrium signal is a subset of $\{s_0, s_1, \ldots, s_{n-1}\}$. The associated messages are $m_1, m_1, m_2, \ldots, m_n$ respectively.

In the context of the journalist-reader example, messages are 'articles' that lay out a narrative and facts discovered by the journalist. Signal realizations correspond to different outcomes of the investigation, privately observed by the journalist. The verifiability structure partially anchors the meaning of what the journalist writes to what she privately learnt. In this context, condition (I) states that the journalist will never seek better news (in terms of the likelihood of the politician's innocence) than what she can prove. This is because information that does not come with hard evidence (e.g. classified documents, anonymous sources, eyewitness accounts) cannot be credibly transmitted to the reader because of the conflict of interest. The journalist thus only writes articles that are 'just good enough' (in terms of the evidence they present) to separate from lower investigation outcomes. More importantly, (II) states that—correctly, in equilibrium—any article is met with scepticism by the reader. This means that the reader

[29] A related logic drives the results in DeMarzo et al. (2019). There sender has the option of fully proving the outcome of her information acquisition or remaining silent. Just like in Dye (1985) there is always the possibility that sender has no evidence to present. In my problem sender's signal choice is unrestricted and the verifiability structure is flexible but deterministic. DeMarzo et al. (2019) extend Dye (1985)'s classical result showing that sender's expected payoff from nondisclosure is minimized in equilibrium. In contrast I show that full revelation obtains in a setting where it would not with exogenous information.

always assumes the worst possible investigation outcome that is consistent with the article written was obtained by the journalist. A direct consequence is that the reader learns all the information about the state discovered by the journalist.

6.1 Sufficient conditions for persuasive and unpersuasive equilibria

I now provide sufficient conditions on the primitives v, M and p_0 such that all equilibria are either unpersuasive or persuasive.

Lemma 2 (Unpersuasive equilibria)**.** *If* $\min M^{-1}(m) \leq p_0$ *for all* $m \in \mathcal{M}$ *then all equilibria are unpersuasive.*

Proof. See Appendix D. □

If no signal realization above the prior belief is lowest consistent then there cannot be a persuasive equilibrium. Intuitively, sender has no way of proving to sender that she obtained a realization above the prior. The conflict of interests between the two players therefore removes all possibility of persuasive communication. The remarks below illustrate implications for some verifiability structures.

Remark. The 'cheap talk' verifiability structure, $M(s) = [0, 1]$ for all $s \in [0, 1]$, does not admit persuasive equilibria. This is because $\min M^{-1}(m) = 0$ for all $m \in [0, 1]$.

Remark. The 'exaggeration' verifiability structure, $M(s) = [s, 1]$ for all $s \in [0, 1]$, also does not admit persuasive equilibria. This is because $\min M^{-1}(m) = 0$ for all $m \in [0, 1]$.

These remarks are intuitive. Since the conflict of interest between sender and receiver is stark, persuasive communication can occur only if the verifiability structure alleviates the incentive compatibility problem at the communication stage. Neither the 'cheap talk' nor the 'exaggeration' structures help since attractive deviations at the interim stage, where sender pretends to have obtained a higher realization, are not ruled out.

Now I turn to sufficient conditions for persuasive equilibria. Given a message $m \in \mathcal{M}$ define $\underline{v}_m \equiv v(\inf M^{-1}(m))$. This is a lower bound on sender's equilibrium payoff from sending message $m \in \mathcal{M}$. This is because equilibrium condition (2) implies that $p_R(m) \geq \inf M^{-1}(m)$ for all $m \in \mathcal{M}$. Since v is nondecreasing this implies that $v(p_R(m)) \geq v(\inf M^{-1}(m)) \equiv \underline{v}_m$ for all $m \in \mathcal{M}$.

Lemma 3 (Persuasive equilibria)**.** *If there exist* $m, n \in \mathcal{M}$ *and* $s_m \in M^{-1}(m)$, $s_n \in M^{-1}(n)$ *such that*

$$(s_m - p_0) \times (p_0 - s_n) > 0 \tag{5}$$

and

$$\frac{s_m - p_0}{s_m - s_n}\underline{v}_n + \frac{p_0 - s_n}{s_m - s_n}\underline{v}_m > v(p_0) \tag{6}$$

then all equilibria are persuasive.

Proof. See Appendix E. □

Condition (5) states that realizations s_m and s_n must lie on different sides of p_0, so that a signal supported on them is a feasible information acquisition strategy for sender. Condition (6) implies that for any belief function consistent the verifiability structure (equilibrium condition (2)) sender does better by choosing a signal supported on $\{s_m, s_n\}$ than she would if there was no possibility of communication.

Remark. Consider the case where v is strictly convex and $p_0 \in (0,1)$. Then the conditions are satisfied for the 'classical' evidence case ($M(s) = \{m \in 2^{[0,1]} : s \in m\}$); the 'understatement' case ($M(s) = [0,s]$); the 'leeway' case $M(s) = [s-\epsilon, s+\epsilon]$.

7 Sender-optimal verifiability structures

Proposition 1 illustrates that in persuasive equilibria sender reveals all acquired information to receiver. She still manages, however, to influence receiver's decision by obtaining information strategically. The extent to which she can do so, as illustrated by Proposition 1, depends on the available evidence, as only lowest consistent signal realizations can occur in persuasive equilibria.

In this section I study the extent to which sender benefits from the ability to prove her acquired information. The answer is not obvious for the following reason. Evidence ties up sender's hands at the interim stage (by allowing information acquisition outcomes to separate), thereby restricting her ability to manipulate the reporting of acquired information. This, however, adds credibility to her reports, and may therefore be beneficial ex ante.

I first spell out the analogy between the question of sender-optimal verifiability structures and the commitment problem of sender. I observe that a verifiability structure that completely ties sender's hands in the communication stage (by forcing her to reveal all acquired information) is optimal for sender. The meaning of sender's messages is fully pinned down by the disclosure requirement so she faces no credibility problem. The second subsection contains the main result, extending this insight by providing necessary and sufficient conditions that the verifiability structure must satisfy in order to be sender-optimal. It shows that sender's communication capabilities must be *credibly rich*, meaning that even when facing a sceptical receiver the available evidence allows her to finely and credibly covey the information she obtained.

7.1 Full commitment

Say a verifiability structure M exhibits *full commitment* if $M(s) = \{s\}$, for all $s \in [0,1]$: every signal realization is associated with a unique message. Observe that sender's message set (after having chosen the signal and observed its realization) is always a singleton, so the communication stage is mechanical. Observe also that each message can only sent by a single signal realization, so equilibrium condition (2) fully pins down receiver's belief following any message: after hearing message s receiver must hold belief $p_R(s) = s$ on- and off-path in any equilibrium. This means that if M exhibits full commitment it is as if receiver directly observed the signal realization.[30]

Under full commitment my model reduces to the 'Bayesian persuasion' problem of Kamenica and Gentzkow (2011). There sender overtly chooses a signal with a public outcome. The authors geometrically characterize sender's optimal value. They show that given a value of induced posteriors v sender's equilibrium payoff is $(\operatorname{cav} v)(p_0)$, where $\operatorname{cav} v$ denotes the smallest concave function that majorises v. A standard alternative interpretation of the 'Bayesian persuasion' model is that sender's information acquisition is private, as in my model, but that she is fully committed to revealing the outcome, as in my model under full commitment. The following observation states this connection.

Observation 1 (Full commitment)**.** *An unbiased public signal S^* is an equilibrium signal in the 'Bayesian persuasion' model if and only if S^* is an equilibrium signal in this model when the verifiability structure exhibits full commitment. Sender's equilibrium payoff is the* full commitment value $(\operatorname{cav} v)(p_0)$.

Proof. Omitted. □

7.2 Characterization of sender-optimal verifiability structures

I now characterize the set of sender-optimal verifiability structures. Recall that I am considering verifiability structures of the form $M : [0,1] \rightrightarrows 2^K$ where K is a fixed compact subset of $\mathbb{R}$, $K \supseteq [0,1]$. The set of verifiability structures under consideration is therefore the set of all such maps. Recall that, for a given M, signal realization s is lowest consistent iff there exists a $m \in M(s)$ such that $s = \min M^{-1}(m)$.

Proposition 2 (Sender-optimal structures)**.** *Fix a verifiability structure M satisfying Assumption 2. The following are equivalent:*

(i) *All signal realizations $s \in [0,1]$ are lowest consistent.*
(ii) *For any nondecreasing and upper semi-continuous v and any $p_0 \in [0,1]$ an equilibrium exists and in any equilibrium sender attains the full commitment value.*
(iii) *For any nondecreasing and upper semi-continuous v and any $p_0 \in [0,1]$ an equilibrium exists and in any equilibrium sender attains the highest ex ante equilibrium payoff among all equilibria for any verifiability structure.*

Proof. See Appendix F. □

The result characterizes sender-optimal verifiability structures in this environment. The key property is the lowest consistency of all signal realizations. This allows sender to attain the full commitment value for any nondecreasing and upper semi-continuous value of induced posteriors and any prior belief. Any failure of this property means that, for some admissible choice of primitives v and p_0, sender is not able to attain the commitment value.

Equivalence of (ii) and (iii) is not surprising. There exists no verifiability structure that allows sender to do better than full commitment and it is always possible to find a verifiability

[30] Observe that under full commitment it is irrelevant whether the choice of signal is overt or covert since the meaning of messages is entirely pinned down by the verifiability structure.

structure such that sender attains this value in equilibrium. To see why, consider any equilibrium for some fixed verifiability structure and let τ_R^* denote the equilibrium distribution of receiver's posterior beliefs about the state. It is straightforward to check that there exists a signal $S \in \Sigma$ such that $\tau_S = \tau_R^*$ so, under full commitment, the distribution of posteriors τ_R^* can be reproduced by simply choosing signal S. The full commitment value is attainable because v is upper semi-continuous. Intuitively, while sender weakly prefers a larger set of messages *at the interim stage* (after observing the signal realization), she weakly benefits *ex ante* from restricting the message sets. This is the case because the credibility of messages improves, thereby allowing sender to do, in expectation, better.

Equivalence between (i) and the other points is less obvious. In order to persuade receiver sender must have access to a *credibly rich* language. When all signal realizations are lowest consistent—even if receiver is maximally sceptical—sender has access to a nuanced language: every realization has a message that distinguishes it from the others. This allows sender to always credibly communicate the realization of the signal, and therefore to attain the same value she would get if she was committed to reveal it. For the converse, observe that condition (II) in Proposition 1 states that on-path scepticism is necessary in all persuasive equilibria. But then if some realization s is not lowest consistent it cannot be part of any persuasive equilibrium. It suffices then to find a suitable value of induced posteriors and prior such that the full commitment value is only attainable in a persuasive equilibrium in which sender must choose a signal with s in its support.

Remark. Observe that the 'understatement' verifiability structure ($M(s) = [0, s]$) is such that all signal realizations are lowest consistent, the proposition thus tells us it allows sender to always attain the full commitment value in equilibrium.[31] The 'exaggeration' verifiability structure ($M(s) = [s, 1]$), on the other hand, is such that only $s = 0$ is lowest consistent.[32]

Remark. Both the 'classical' verifiability structure ($M(s) = \{m \in 2^{[0,1]} : s \in m\}$) and the 'leeway' verifiability structure ($M(s) = [s - \epsilon, s + \epsilon]$) are such that every signal realization is lowest consistent.[33] The proposition thus tells us they allows sender to always attain the full commitment value in equilibrium.

Proposition 2 illustrates how sender's commitment power can be relaxed without affecting her ability to persuade receiver. What is necessary is that, even under maximal scepticism, each possible signal realization is still able to credibly separate. To gain some intuition consider the exaggeration/understatement example. In the 'understatement' case sender can attain the full commitment value. In the partial commitment interpretation of the model this means that sender is able to commit to *never over-reporting* the signal realization. Given the nondecreasing conflict of interest, sender will never want to under-report, so committing to never over-reporting is tantamount to committing to reporting truthfully. More generally what is necessary and sufficient is that each realization gives access to a message that can *never* be interpreted as an 'over-report' by a lower realization. Lowest consistency characterizes precisely this.

[31] In more detail. For $M(s) = [0, s]$, $M^{-1}(m) = [m, 1]$ for all $m \in [0, 1]$ so $\min M^{-1}(m) = m$ for all $m \in [0, 1]$: each $s \in [0, 1]$ is lowest consistent.

[32] In more detail. For $M(s) = [s, 1]$, $M^{-1}(m) = [0, m]$ for all $m \in [0, 1]$ so $\min M^{-1}(m) = 0$ for all $m \in [0, 1]$: only $s = 0$ is lowest consistent.

[33] In the 'classical' case there are many messages with which each s is lowest consistent; for example each s is lowest consistent with $m = s$. In the 'leeway' case each s is (only) lowest consistent with $m = s + \epsilon$.

The result strengthens the connection between evidence games à la Milgrom (1981) and Grossman (1981) and the 'Bayesian persuasion' game of Kamenica and Gentzkow (2011). A connection was already pointed out in Brocas and Carrillo (2007) and Kamenica and Gentzkow (2011). In the language of my paper, these authors consider the 'classical' verifiability structure ($M(s) = \{m \in 2^{[0,1]} : s \in m\}$). They show this is *sufficient* for sender to obtain the full commitment value for any conflict of interest between sender and receiver. I prove that, for a monotone conflict of interest, the weaker notion of lowest consistency of all signal realizations implies that sender is able to always obtain the full commitment value.[34] I also show that lowest consistency is necessary, therefore offering a tight characterization of what verifiability structures allow sender to do as well as under full commitment.

In the context of the journalist-reader example the result can be interpreted as follows. The journalist would like to be have access to a rich set of verifiable statements she can present in her articles. Specifically, she would always (i.e. for every outcome of her investigation) like to be able to prove that what she learnt is not worse than what she actually did learn. The presence of such proof disciplines the reader's scepticism and therefore allows the journalist to credibly convey any acquired information. This, is turn, allows her to manipulate the reader's decision by strategically choosing what information to acquire and then, credibly, disclosing it.

8 Receiver-optimal verifiability structures

Proposition 1 illustrates that, in persuasive equilibria, all covertly acquired information is necessarily revealed to receiver. It also shows that what sender can prove to receiver plays a key role in determining what she chooses to learn and, in equilibrium, reveal. In this section I address the question of what verifiability structures are desirable, in terms of the equilibrium outcomes they induce, by receiver. Section 7 illustrated that a 'credibly rich' verifiability structure is desirable by sender. In particular, mandated disclosure of all acquired information is 'credibly rich', and thus sender-optimal.

Mandated disclosure (and, more generally, lowest consistency of all outcomes of sender' information acquisition) may, however, lead sender to acquire little information in equilibrium. Consider the following modified version of the example from Section 3 as an illustration.

Example (Three actions)**.** Modify the example from Section 3 so that the payoff to sender from action 'abstain' is 2 instead of 3. Additionally, increase the prior to $p_0 = 3/4$. Microfound the reader's decision rule with expected utility preferences as follows. Let the payoff from 'oppose' be 2 if $\omega = 0$ and -3 if $\omega = 1$. Let the payoff from 'support' be -4 if $\omega = 0$ and 1 if $\omega = 1$. Let the payoff from 'abstain' be 0 regardless of the state. Figure 8 illustrates sender's and receiver's expected payoff as a function of receiver's belief at the time of choosing the action. Suppose the verifiability structure is such that all signal realizations are lowest consistent (e.g. disclosure of each signal realization is mandated). Proposition 2 tells us sender attains the full commitment value of $(\operatorname{cav} v)(p_0) = 23/8$, just as in the 'Bayesian persuasion'

[34] More generally, when v is not required to be nondecreasing but only upper semi-continuous the necessary and sufficient condition on M is that each $s \in [0, 1]$ be both lowest and highest consistent. Any verifiability structure such that $s \in M(s)$ for all $s \in [0, 1]$ (e.g. full commitment and 'classical' evidence) clearly satisfy this, so these structures are sufficient.

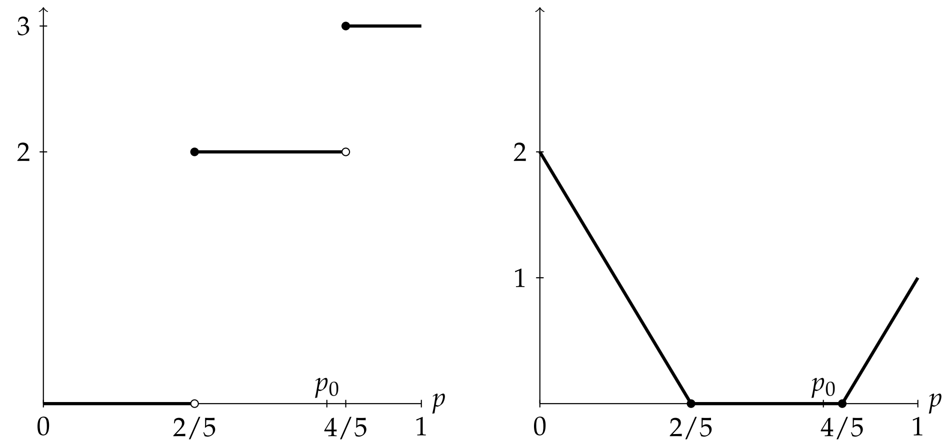

Figure 8: Sender's and receiver's value of posteriors.

problem of Kamenica and Gentzkow (2011). It is straightforward to check that here the only equilibrium information acquisition strategy is choosing signal S^* with supp $S^* = \{2/5, 4/5\}$. All equilibria are persuasive (the conditions of Lemma 3 are satisfied) so receiver's equilibrium belief distribution is also τ_{S^*} by Proposition 1, point (II). It immediately follows that receiver's equilibrium payoff is always 0, since when his belief is $2/5$ he chooses action 'abstain' and obtains a payoff of 0, when his belief is $4/5$ he 'opposes' and also obtains a payoff of 0. This is the *lowest* possible payoff he can obtain when behaving optimally given his information.

As illustrated by the example above, mandated disclosure may well have no value to receiver because sender acquires too little information in equilibrium. So, which verifiability structures are desirable from his perspective? Can receiver remove credibility from sender's claims to improve the incentives to acquire and disclose information? These questions are of practical interest in all settings where it may be infeasible or undesirable to restrict sender's information acquisition activity but it is possible to regulate what outcomes of information acquisition are certifiable and which are not. The main result in this section offers a simple and intuitive verifiability structure that is necessary and sufficient for *fully informative* information acquisition and transmission in any equilibrium of the game.

As in the rest of the paper, I consider verifiability structures of the form $M : [0,1] \rightrightarrows 2^K$ where K is a fixed compact subset of $\mathbb{R}$, $K \supseteq [0,1]$. The set of verifiability structures under consideration is therefore the set of all such maps.

8.1 Preliminaries: valuable communication and information

Since the object of interest of this section is receiver's welfare I am now explicit about his preferences. Let him choose an action from a compact set A. Let $w_R(p, a)$ denote receiver's payoff when he chooses action $a \in A$ while holding belief $p \in [0,1]$ about the state. Let $w_R(p, \cdot)$ be continuous for all $p \in [0,1]$. I say that receiver *values information* if $w_R(\cdot, a)$ is convex for all $a \in A$. This is clearly satisfied if receiver maximises expected utility as in that case $w_R(\cdot, a)$ is affine.[35] Given a posterior belief $p \in [0,1]$ of receiver define

$$v_R(p) \equiv \max_{a \in A} w_R(p,a), \tag{7}$$

which is his payoff from behaving optimally at belief p. Observe that, if receiver values information, $v_R : [0,1] \to \mathbb{R}$ is convex as it is the pointwise maximum of a family of convex functions. Let $S^F \in \Sigma$ denote the full-information unbiased signal, that is: supp $S^F = \{0,1\}$ so that $\pi_{S^F}(1|1) = 1$ and $\pi_{S^F}(0|0) = 1$. Define

$$v_R^F \equiv \sum_{s \in \operatorname{supp} S^F} \tau_{S^F}(s) v_R(s) = p_0 v_R(1) + (1-p_0) v_R(0) \tag{8}$$

to be the full-information ex ante expected payoff for receiver.

Results are more easily stated if an additional assumption on primitives is made to rule out the case where sender cannot possibly benefit from communication. Formally, I assume that at the prior receiver is not already taking sender's favourite action, that is

$$v(p_0) < \sup_{s \in [0,1]} v(s). \tag{9}$$

If (9) holds say that *communication is potentially valuable* for sender.[36] If communication is *not* potentially valuable for sender it is immediate that for any verifiability structure M there always exists an equilibrium with no information acquisition and transmission.[37]

8.2 Characterization of receiver-optimal verifiability structures

The following result shows that if communication *is* potentially valuable for sender then we can find a M such that full information acquisition and transmission arises in all equilibrium outcomes. Furthermore the *same* M can be used to achieve this for any nondecreasing and upper semi-continuous v and any $p_0 \in [0,1]$.

Proposition 3 (Receiver-optimal structures). *Fix a verifiability structure M satisfying Assumption 2. The following are equivalent:*

(a) Signal realization s is lowest consistent iff $s \in \{0,1\}$.

(b) For any nondecreasing and upper semi-continuous v and any $p_0 \in [0,1]$ such that communication is potentially valuable for sender an equilibrium exists and in any equilibrium receiver obtains the full-information expected payoff.

(c) For any nondecreasing and upper semi-continuous v and any $p_0 \in [0,1]$ such that communication is potentially valuable for sender an equilibrium exists and—if receiver values information—in any equilibrium receiver attains the highest ex ante equilibrium payoff among all equilibria for any verifiability structure.

[35] If receiver maximises expected utility with Bernoulli utility function $u_R : \{0,1\} \times A \to \mathbb{R}$, $(\omega, a) \mapsto u_R(\omega, a)$ we have that $w_R(p,a) = p u_R(1,a) + (1-p) u_R(0,a)$, which is obviously affine.

[36] Since v is nondecreasing, this is equivalent to assuming that $v(p_0) < v(1)$.

[37] For example, fix an $m \in M(p_0)$ and let $p_R(m) = p_0$. It's straightforward to check that the uninformative signal S^U and a messaging strategy with $\mu(S^U, p_0) = m$ on-path are part of an equilibrium.

The result characterizes receiver-optimal verifiability structures in this environment. There must be at least one message that is available to $s = 1$ only, this proves to receiver that $\omega = 1$. Messages that are not available to $s = 1$ only *must* be also available to $s = 0$. If this holds, provided communication is potentially valuable for sender, all equilibrium outcomes involve full information acquisition and transmission. Any failure of this property means that, for some admissible choice of primitives v and p_0, receiver is not able to obtain full information in equilibrium.

Such coarse verification is optimal for receiver because it fully aligns sender's ability to provide evidence with the direction of the conflict of interest. In equilibrium, receiver is maximally sceptical (i.e. he believes that $\omega = 0$) when not presented with evidence and thus takes the *worst* action for sender. If he is provided with evidence, he believes that $\omega = 1$ and takes the *best* action for sender. Coarse verification, because of equilibrium scepticism, provides high-powered incentives for information acquisition.

Proof sketch of Proposition 3. See Appendix G for a full proof. That (b) implies (c) is immediate: receiver cannot do better than choosing the action under full information, provided that he values information. The proof that (c) implies (b) is constructive. If $p_0 = 0$ the result is trivially true, so let $p_0 > 0$. A simple verifiability structure with two messages is such that in the only equilibrium receiver attains the full-information payoff. This verifiability structure involves a message which proves nothing, m_0, which is available to sender at all signal realizations, and a message that proves that $\omega = 1$, m_1. Proposition 1 implies that if there is a persuasive equilibrium, it must involve full information acquisition and revelation. It is immediate that such an equilibrium exists. A similar but slightly more elaborate argument shows that (a) implies (b). A crucial additional step is to show that there cannot be an unpersuasive equilibrium. Suppose there were. Then observe that whichever messages are being used in equilibrium (to attain unpersuasive payoff of $v(p_0)$) are also available at $M(0)$, since only 0 and 1 are lowest consistent. Then a deviation to the full information signal would attain the payoff $p_0 v(1) + (1 - p_0) v(p_0)$, which is strictly higher than $v(p_0)$ since communication is potentially valuable to sender ($v(1) > v(p_0)$). So no unpersuasive equilibrium can exist.

To show that (b) implies (a) I prove the contrapositive; I use Proposition 1 and counterexamples. If $s = 1$ is not lowest consistent, Proposition 1 implies that it cannot be a signal realization in any persuasive equilibrium. It therefore suffices to show that no unpersuasive equilibria exist and that a persuasive equilibrium does. It immediately follows that there exists an equilibrium in which full information is not acquired and transmitted. If some other $s \notin \{0, 1\}$ is lowest consistent one can construct a suitable value of induced posteriors and prior such that receiver obtains less than full information in some equilibrium. $\square$

Example (Three actions, continued)**.** Now fix verifiability structure with $M(s) = \{m_0\}$ for all $s \in [0, 1)$ and $M(s) = \{m_0, m_1\}$ for $s = 1$. Communication is clearly potentially valuable for sender since $v(p_0) = 2 < 3 = v(1)$. Equilibria exist (the conditions of Corollary 1 are satisfied) and all equilibria are persuasive (the conditions of Lemma 3 are satisfied[38]). Hence, by Proposition 1, the only candidate equilibrium signal is the fully informative S^F, since only

[38]Only $p_R(m_1) = 1$ is consistent with equilibrium condition (2) so the fully informative signal S^F supported on $\{0, 1\}$ must lead to a payoff for sender of at least $p_0 \times 3 + (1 - p_0) \times 0 = 9/4 > 2 = v(p_0)$.

$s = 0$ and $s = 1$ are lowest consistent. Constructing the equilibrium is straightforward. Sender chooses m_1 whenever possible and $p_R(m_1) = 1$, $p_R(m_0) = 0$. S^F is a best reply signal as it maximises the probability of sending m_1, and receiver's belief is formed using Bayes' rule on path. Receiver's equilibrium payoff is $5/4$, which is the highest he could attain as he cannot do better than having full information at the time of choosing the action.

Observe that preventing all intermediate outcomes from being certifiable can be interpreted as a form of commitment to scepticism for receiver. In this interpretation, receiver commits to not believing any intermediate claim about the state and only accepts proof that the state is high with certainty. This is optimal because it leaves full information acquisition and transmission as the best sender can do to (sometimes) induce a high belief.

9 Conclusion

Many decisions are made relying on the advice of experts with superior information but conflicting motives. I set out to understand how—in the presence of a conflict of interest—an expert's capability to prove her private information to a decision maker shapes her incentives to acquire and transmit such information. I do this by studying a sender-receiver game augmented with two key features: (i) sender's private information is covertly and costlessly acquired, rather than exogenously given and (ii) sender's acquired private information is only partially provable to receiver.

I show that, whenever sender manages to influence receiver, she never chooses to learn more than what she is able to prove. A stark consequence is that, in equilibrium, she *fully reveals* all of her acquired information. This is in contrast to the benchmark where sender's private information is exogenous and partially verifiable, in which case full revelation does not necessarily occur. I then study what kind of evidence is desirable by each of the two players. I characterize in what sense having 'rich evidence' is optimal for sender by highlighting the connection with the problem of commitment to disclose acquired information. I then study what kind of verifiability is desirable from receiver's perspective. Receiver-optimal evidence is 'binary' and aligned with the conflict of interest in order to provide strong incentives for information acquisition.

Appendix A Preliminaries

A.1 Unbiased signals and belief distributions

Consider an unbiased finite-support signal $S \in \Sigma$ with conditional probability mass function $\pi_S(\cdot|\omega)$ and marginal probability mass function $\tau_S(\cdot)$. Since S is unbiased, it is immediate that $\mathbf{E}(S) = p_0$. In detail,

$$\mathbf{E}(S) = \sum_{s \in \text{supp } S} \tau_S(s)s = \sum_{s \in \text{supp } S} \tau_S(s) \frac{\pi_S(s|1)p_0}{\tau_S(s)} = \sum_{s \in \text{supp } S} \pi_S(s|1)p_0 = p_0$$

where the second inequality follows from the fact that S is unbiased, the third is manipulation, the last holds since conditional probabilities sum to 1.

Now consider any finite-support mean-p_0 distribution τ over $[0,1]$. Let $T(p_0)$ denote the set of all such distributions. Let S_τ be a signal with $\text{supp } S_\tau = \{s \in [0,1] : \tau(s) > 0\}$ with condition distribution given by

$$\pi_{S_\tau}(s|1) = \tau(s)\frac{s}{p_0} \quad \text{and} \quad \pi_{S_\tau}(s|0) = \tau(s)\frac{1-s}{1-p_0}$$

for all $s \in \text{supp } S_\tau$. It is straightforward to observe that this signal is unbiased since

$$\mathbf{E}(\omega|s) = \frac{\pi_{S_\tau}(s|1)p_0}{\pi_{S_\tau}(s|1)p_0 + \pi_{S_\tau}(s|0)(1-p_0)} = \frac{\tau(s)s}{\tau(s)s + \tau(s)(1-s)} = s$$

where the first equality follows from Bayes' rule, the second from the definition of $\pi_{S_\tau}(s|\omega)$ and the last from manipulation. Putting these observations together we obtain the following well-known result.

Lemma 4. *Every unbiased finite-support signal $S \in \Sigma$ has a finite-support mean-p_0 marginal distribution $\tau_S \in T(p_0)$ on $[0,1]$. Every finite-support mean-p_0 distribution on $[0,1]$ $\tau \in T(p_0)$ is the marginal distribution of some unbiased finite-support signal $S_\tau \in \Sigma$.*

A.2 Canonical deviations

In this section I introduce a deviation for sender that is useful in most of the proofs that follow. Let $S \in \Sigma$ be an unbiased finite-support signal and let τ_S denote its marginal probability mass function. Suppose that $\text{supp } S \geq 2$ so that there exist a $s > p_0$ in the support of S (since S is unbiased).

I call a *canonical deviation* for sender a deviation from signal S to a signal S' that removes all probability mass from s and 'shifts it' to some other signal realization $s' > p_0$ (or, similarly shifts it from some $r < p_0$ so some $r' < p_0$), leaving all other odds ratios between results unchanged. I also (abusing terminology) refer to it as a canonical deviation from s to s', in the sense that it deviates from some original signal S that puts mass on s to another signal S' that shifts probability mass away from s and on to s'. Such a deviations are particularly convenient to work with as it is easy to characterize when they are profitable.

In detail, construct signal S' as follows. Let $\tau_{S'}(s) = 0$ and define $k \equiv (s' - p_0) + (s - s')\tau(s)$. Let

$$\tau_{S'}(s') = \frac{s - p_0}{k}\tau_S(s) + \frac{s' - p_0}{k}\tau_S(s')$$

and

$$\tau_{S'}(t) = \frac{s' - p_0}{k}\tau_S(t)$$

for $t \in \operatorname{supp}\tau$, $t \neq p, p'$. It is straightforward to check that S' is a finite-support unbiased signal i.e. that for all $s \in \operatorname{supp} S'$, $\mathbf{E}(\omega|s) = s$. This in turn implies that $S' \in \Sigma$ so that signal S' is feasible.

Suppose that receiver forms beliefs about the state according to the function $p_R : \mathcal{M} \to [0,1]$. I now provide sufficient conditions for a canonical deviation to be profitable. Start with a signal $S \in \Sigma$ with $\operatorname{supp} S \geq 2$. Then there exists a $s > p_0$ and a $r < p_0$ in the support of S (since S is unbiased). I consider a canonical deviation from s to some $s' > p_0$ (or, alternatively, a canonical deviation from r to some $r' > 0$). Let $w(q) \equiv \sup_{m \in M(q)} v(p_R(m))$ for $q \in \{r, r', s, s'\}$ and let $\alpha, \beta \in \mathbb{R}$ denote the solution to simultaneous equations $w(r) = \alpha + \beta r$ and $w(s) = \alpha + \beta s$.

Lemma 5. *Such a canonical deviation is profitable if $w(s') > \alpha + \beta s'$ (alternatively, $w(r') > \alpha + \beta r'$).*

Proof. Immediate from Kamenica and Gentzkow (2011, Corollary 2). □

Appendix B Proofs of the existence results

B.1 Preliminaries

I first establish some preliminary results that hold under Assumptions 1-3. Fix a receiver belief about the state $p_R : \mathcal{M} \to [0,1]$. Let $\eta : [0,1] \to \mathcal{M}$ such that $\eta(s) \in M(s)$ for all $s \in [0,1]$ be a map associating with each realization a message. Write $(v \circ p_R \circ \eta)(s) = v(p_R(\eta(s)))$ to denote the payoff to sender when s realizes, she messages according to η and receiver uses p_R to update his belief.

I provide a result establishing existence of a best reply signal. Let

$$s_-(p_0) \equiv \sup\{s \in [0, p_0] : (\operatorname{cav} v \circ p_R \circ \eta)(s) = (v \circ p_R \circ \eta)(s)\}$$

and

$$s_+(p_0) \equiv \inf\{s \in [p_0, 1] : (\operatorname{cav} v \circ p_R \circ \eta)(s) = (v \circ p_R \circ \eta)(s)\}$$

where $\operatorname{cav} g$ denotes the smallest concave function that majorises g.

Lemma 6 (Sender's best reply signal: existence condition)**.** *Let receiver's belief function be p_R and $\mu(S, s) = \eta(s)$ for all $S \in \Sigma$ and all $s \in \operatorname{supp} S$. If $v \circ p_R \circ \eta$ is upper semi-continuous a best reply signal exists; $S^* \in \Sigma$ is a best reply signal if* $\operatorname{supp} S^* = \{s_-(p_0), s_+(p_0)\}$.

Proof. The result is immediate from Kamenica and Gentzkow (2011). □

Say that receiver's belief function p_R is *always sceptical* iff $p_R(m) = \min M^{-1}(m)$ for all $m \in \mathcal{M}$. This belief function is well-defined since $M^{-1}(m)$ is closed for all $m \in \mathcal{M}$ under Assumption 2.

Lemma 7 (Sender's best reply: existence under scepticism)**.** *Let receiver form beliefs according to the always sceptical belief function p_R and*

$$\mu(S,s) = \eta^{p_R}(s) \in \arg\max_{m \in M(s)} \min M^{-1}(m)$$

for all $S \in \Sigma$ and all $s \in \operatorname{supp} S$. A best reply signal exists; $S^{p_R} \in \Sigma$ is a best reply signal if $\operatorname{supp} S^{p_R} = \{s_-(p_0), s_+(p_0)\}$.

Proof. Observe that $p_R(\eta^{p_R}(s)) = \max_{m \in M(s)} \min M^{-1}(m)$. So $p_R \circ \eta^{p_R}$ is upper semi-continuous under Assumption 3. Since v is nondecreasing and upper semi-continuous, so is $v \circ p_R \circ \eta^{p_R}$. The result follows from Lemma 6. □

I introduce some notation. Define $v^{p_R}(s) = \sup_{m \in M(s)} v(p_R(m))$. Since $M(s)$ is finite for all $s \in [0,1]$, $\sup_{m \in M(s)} v(p_R(m)) = \max_{m \in M(s)} v(p_R(m))$. Define $\eta^{p_R} : [0,1] \to \mathcal{M}$ to be a selection from $s \mapsto \arg\max_{m \in M(s)} v(p_R(m))$, which is nonempty since $M(s)$ is finite. By definition of v^{p_R} and η^{p_R}, it holds that $v^{p_R}(s) = v(p_R(\eta^{p_R}(s)))$ for all $s \in [0,1]$.

B.2 Proof of Lemma 1

The proof is constructive and divided into two cases.

Case 1: For all $m \in \mathcal{M}$, $\min M^{-1}(m) \leq p_0$.

I first construct a candidate equilibrium (S^*, μ^*, p_R^*). Fix a message $m \in M(p_0)$ and let $p_R^*(m) = p_0$. For all $m' \neq m$ let $p_R^*(m') = \min M^{-1}(m')$. Let $\operatorname{supp} S^* = \{p_0\}$, $\mu^*(S^*, p_0) = m$ and μ^* off-path be any optimal message given p_R^*, which exists since $M(s)$ is finite. I now argue that this profile is indeed an equilibrium. Observe that since for all $m' \in \mathcal{M}$, $\min M^{-1}(m') \leq p_0$ we have that $p_R^*(m') = \min M^{-1}(m') \leq p_0$ for all $m' \in \mathcal{M}$. Since v is nondecreasing, $v(p_R^*(m')) \leq v(p_R^*(m))$ for all $m' \in \mathcal{M}$. It follows that $v^{p_R^*}(p_0) \geq v^{p_R^*}(s)$ for all $s \in [0,1]$. It is therefore immediate from equilibrium condition (4) that S^* is a best reply signal. A fortiori, it also holds that for all messages $m' \in M(p_0)$, $v(p_R^*(m')) \leq v(p_R^*(m))$, so $\mu^*(S^*, p_0) = m$ is a best reply message at (S^*, p_0). m is the only message on-path and $p_R^*(m) = p_0$ is correct. Off-path messages are optimal by construction and off-path beliefs are consistent with the verifiability structure, so the profile is indeed an equilibrium.

Case 2: There exists an $m \in \mathcal{M}$ such that $\min M^{-1}(m) > p_0$.

Start the equilibrium construction by imposing always sceptical receiver beliefs: let $p_R^*(m) = \min M^{-1}(m)$. Let $\eta^{p_R^*} : [0,1] \to \mathcal{M}$ be a selection from $s \mapsto \arg\max_{m \in M(s)} p_R^*(m)$. It is nonempty since $M(s)$ is finite, it selects an optimal message at s since v is nondecreasing. Let $\mu^*(S,s) = \eta^{p_R^*}(s)$ for all $S \in \Sigma$ and all $s \in \operatorname{supp} S$. By Lemma 7, a best reply signal exists and $S^* \in \Sigma$ with support on $\{s_-(p_0), s_+(p_0)\}$ is a best reply signal. From now until the end of the proof, lighten notation by writing s_- (s_+) instead of $s_-(p_0)$ ($s_+(p_0)$).

Lemma 8. $v^{p_R^*}(s_+) \geq v^{p_R^*}(s_-)$.

Proof. Suppose towards a contradiction that $v^{p_R^*}(s_+) < v^{s_R^*}(s_-)$. Then, optimality of signal S^* implies that

$$v^{p_R^*}(s_-) > v^{p_R^*}(s) \quad \text{for all} \quad s > s_-. \tag{10}$$

To see why, suppose there is some $s' > s_-$ violating (10). Consider canonical deviations as discussed in Appendix A.2. If $s' < p_0$ a canonical deviation from s_- to s' is profitable. If $s' > p_0$ a canonical deviation from s_+ to s' is profitable. If $s' = p_0$ then a deviation to S^{**} with $\operatorname{supp} S^{**} = \{p_0\}$ is profitable. These deviations contradict that S^* is a best reply signal, so if $v^{p_R^*}(s_+) < v^{s_R^*}(s_-)$ holds then also (10) holds.

Rewrite (10) using the definitions of $v^{p_R^*}$ and $\eta^{p_R^*}$ to obtain

$$v(p_R^*(\eta^{p_R^*}(s))) < v(p_R^*(\eta^{p_R^*}(s_-))) \quad \text{for all} \quad s > s_-.$$

By optimality of $\eta^{p_R^*}$,

$$v(p_R^*(m)) \leq v(p_R^*(\eta^{p_R^*}(s))) \quad \text{for all} \quad m \in M(s) \quad \text{and} \quad s \in [0,1]$$

so

$$v(p_R^*(m)) < v(p_R^*(\eta^{p_R^*}(s_-))) \quad \text{for all} \quad m \in M(s) \quad \text{and} \quad s > s_-.$$

Since v is nondecreasing this in turn implies that

$$p_R^*(m) < p_R^*(\eta^{p_R^*}(s_-)) \quad \text{for all} \quad m \in M(s) \quad \text{and} \quad s > s_-.$$

By definition of p_R^*, this is equivalent to

$$\min M^{-1}(m) < \min M^{-1}(\eta^{p_R^*}(s_-)) \quad \text{for all} \quad m \in M(s) \quad \text{and} \quad s > s_-.$$

Obviously $s_- \in M^{-1}(\eta^{p_R^*}(s_-))$ so $\min M^{-1}(\eta^{p_R^*}(s_-)) \leq s_- \leq p_0$, where the last inequality follows from the definition of s_-. Substituting into the last display we obtain

$$\min M^{-1}(m) < p_0 \quad \text{for all} \quad m \in M(s) \quad \text{and} \quad s > p_0.$$

This contradicts that there exists an $m \in \mathcal{M}$ such that $\min M^{-1}(m) > p_0$, completing the proof. □

Consider now two sub-cases:

Sub-case 1: $v^{p_R^*}(s_+) > v^{p_R^*}(s_-)$. Let sender use on-path messages $\mu^*(S^*, s_-) = \eta^{p_R^*}(s_-)$ and $\mu^*(S^*, s_+) = \eta^{p_R^*}(s_+)$, which is a best reply by definition of $\eta^{p_R^*}$. Let her use any optimal message off-path. I claim that $p_R^*(\eta^{p_R^*}(s_-)) = s_-$ and $p_R^*(\eta^{p_R^*}(s_+)) = s_+$ i.e. that $\min M^{-1}(\eta^{p_R^*}(s_-)) = s_-$ and that $\min M^{-1}(\eta^{s_R^*}(p_+)) = s_+$. The argument is identical to the argument in the proof of Proposition 1 and is therefore omitted. Observe that the argument used there applies since we have already established that S^* is a best reply and this is all is needed to apply the argument. It follows that receiver's on-path beliefs are formed using Bayes' rule, so the candidate profile is an equilibrium.

Sub-case 2: $v^{p_R^*}(s_+) = v^{p_R^*}(s_-)$. Two sub-sub-cases:

(1) $s_- = s_+ = p_0$. By definition of p_R^*, $p_R^*(\eta^{p_R^*}(p_0)) = \min M^{-1}(\eta^{p_R^*}(p_0))$. Suppose first that $\min M^{-1}(\eta^{p_R^*}(p_0)) = p_0$. Then $p_R^*(\eta^{p_R^*}(p_0)) = p_0$ so the only on-path message induces a correct belief in receiver. On-path messaging according to $\mu^*(S^*, p_0) = \eta^{p_R^*}(p_0)$ is a best

reply so we have constructed an equilibrium.
Suppose next that $\min M^{-1}(\eta^{p_R^*}(p_0)) < p_0$. We shall construct another signal $S^{**} \in \Sigma$ that is also a best reply given p_R^* and a messaging strategy that is a best reply and leads to correct on-path receiver beliefs. To lighten notation define $s_L \equiv \min M^{-1}(\eta^{p_R^*}(p_0))$ and $m_L \equiv \eta^{p_R^*}(p_0)$. Recall that in this case of the proof, there exists an m_H such that $\min M^{-1}(m_H) > p_0$. Define $s_H \equiv \min M^{-1}(m_H)$.

Lemma 9. $v^{p_R^*}(s_L) \geq v^{p_R^*}(p_0)$.

Proof. $v^{p_R^*}(p_0) = v(p_R^*(m_L))$ by definition of m_L. $m_L \in M(s_L)$ by definition of s_L, so $v^{p_R^*}(s_L) \geq v(p_R^*(m_L)) = v^{p_R^*}(p_0)$ follows. □

Lemma 10. $v^{p_R^*}(s_H) = v^{p_R^*}(p_0)$.

Proof. $s_H > p_0$ by definition so $p_R^*(m_H) > p_R^*(m_L)$. Since v is nondecreasing, it follows that $v(p_R^*(m_H)) \geq v(p_R^*(m_L))$. By definition of $\eta^{p_R^*}$ we also have

$$v^{p_R^*}(s_H) = v(p_R^*(\eta^{p_R^*}(s_H))) \geq v(p_R^*(m_H)).$$

It follows that $v^{p_R^*}(s_H) \geq v(p_R^*(m_L)) = v^{p_R^*}(p_0)$.
I now want to show that $v^{p_R^*}(p_0) \geq v^{p_R^*}(s_H)$. Suppose otherwise. Then signal $\tilde{S} \in \Sigma$ with support $\{s_L, s_H\}$ is a profitable deviation as, by Lemma 9 it gives a an expected payoff strictly higher than $v^{p_R^*}(p_0)$. This constradicts that S^* is a best reply signal.
Having shown that $v^{p_R^*}(p_0) \leq v^{p_R^*}(s_H)$ and $v^{p_R^*}(p_0) \geq v^{p_R^*}(s_H)$, the proof is complete. □

Lemma 11. $v^{p_R^*}(s_L) \leq v^{p_R^*}(p_0)$.

Proof. I want to show $v^{p_R^*}(s_L) \leq v^{p_R^*}(p_0)$. Suppose otherwise. By Lemma 10 a signal $\tilde{S} \in \Sigma$ with support $\{s_L, s_H\}$ would be a profitable deviation, contradicting that S^* is a best reply signal. □

Combining the three previous lemmas we have that

$$v^{p_R^*}(s_L) = v^{p_R^*}(s_L) = v^{p_R^*}(p_0).$$

It follows that signal S^{**} with support on s_L and s_H is also a best reply signal. Sending messages m_L and m_H respectively is an best reply on-path messaging strategy. Let sender send any optimal message off-path. On-path beliefs are now correct by construction so the profile is an equilibrium.

(2) $s_- < s_+$. Define $v^* \equiv v^{p_R^*}(s_+) = v^{p_R^*}(s_-)$ and observe that

$$v^* \geq v^{p_R^*}(s) \tag{11}$$

for all $s \in [0, 1]$ as otherwise S^* could not be a best reply signal. Also, observe that

$$v^* \leq v(s) \tag{12}$$

for all $s \geq s_-$. This is because for any $m \in M(s_-)$, $p_R^*(m) = \min M^{-1}(m) \leq s_-$ so $v(p_R^*(m)) \leq v(s_-)$ for any $m \in M(s_-)$ and $v(s_-) \leq v(s)$ for all $s \geq s_-$ since v is nondecreasing. It follows that $v(p_R^*(m)) \leq v(s)$ for all $s \geq s_-$ and all $m \in M(s_-)$. Since $\eta^{p_R^*}(s_-) \in M(s_-)$ it immediately follows that $v^* = v(p_R^*(\eta^{p_R^*}(s_-))) \leq v(s)$ for all $s \geq s_-$. Recall that in this case of the proof, there exists an m_H such that $\min M^{-1}(m_H) > p_0$. Define $s_H \equiv \min M^{-1}(m_H)$. Clearly

$$v^{p_R^*}(s_H) = v(s_H) \tag{13}$$

since $p_R^*(m_H) \geq p_R^*(m)$ for all $m \in M(s_H)$ by definition of p_R^* and v is nondecreasing. It follows that sending message m_H at s_H is a best reply. Combining (13) with (11) and (12) and since $s_H > p_0 > p_-$ we obtain

$$v^* \geq v^{p_R^*}(s_H) = v(s_H) \geq v^*$$

and therefore

$$v^{p_R^*}(s_H) = v(s_H) = v^*. \tag{14}$$

Consider now signal realization $s_L \equiv \min M^{-1}(\eta^{p_R^*}(s_-))$ and let $m_L \equiv \eta^{p_R^*}(s_-)$. Observe that

$$v^* \leq v^{p_R^*}(s_L) \tag{15}$$

since $\eta^{p_R^*}(p_-) \in M(s_L)$ and $v(p_R^*(\eta^{p_R^*}(p_-))) = v^*$. Combining (11) with (15) we obtain

$$v^{p_R^*}(s_L) = v^*.$$

Observe that $s_L < p_0$ since $s_- < p_0$ and $s_L \leq s_-$. We may choose signal $S^{**} \in \Sigma$ with support on $\{s_L, s_H\}$, which is a best reply since it attains the same payoff as S^*. On-path messaging is given by $\mu^*(S^{**}, s_L) = m_L$ and $\mu^*(S^{**}, s_H) = m_H$. Off-path let sender choose any optimal message. On-path beliefs are correct by construction, so we have found an equilibrium.

B.3 Proof of Corollary 1

Since $\mathcal{M}$ is finite, so is $M(s)$ for all $s \in [0,1]$ a fortiori. All left to establish is that f_M is upper semi-continuous. I establish a slightly stronger result. Fix any map $p_R : \mathcal{M} \to [0,1]$ and define $g_M : [0,1] \to [0,1]$, $g_M(s) = \max_{m \in M(s)} p_R(m)$. It is immediate that if $p_R(m) = \min M^{-1}(m)$ then $g_M = f_M$. I shall prove that g_M is upper semi-continuous.

Define $g_M^m : [0,1] \to [0,1] \cup \{-1\}$ by

$$g_M^m(s) = \begin{cases} p_R(m) & \text{if} \quad m \in M(s), \\ -1 & \text{otherwise.} \end{cases}$$

Observe that g_M^m is an upper semi-continuous function for each $m \in \mathcal{M}$ since $M^{-1}(m)$ is closed for each $m \in \mathcal{M}$. Notice also that

$$g_M(s) = \max_{m \in \mathcal{M}} g_M^m(s)$$

for each $s \in [0,1]$, where the maximum exists since $\mathcal{M}$ is finite. The upper envelope of a finite family of upper semi-continuous functions is itself upper semi-continuous (see, for example, Bourbaki (1995, Chapter IV, §6, Proposition 2)) so g_M is upper semi-continuous, as required.

Appendix C Proof of Proposition 1

C.1 Preliminaries

I prove first three useful lemmas.

Lemma 12. *Let τ_R^* denote the distribution of receiver's posterior beliefs is some pure-strategy equilibrium. There exists a signal $S \in \Sigma$ with $\tau_S = \tau_R^*$.*

Proof. Fix a pure strategy equilibrium (S^*, μ^*, p_R^*). Consider a $p \in [0,1]$ such that $\tau_R^*(p) > 0$ and let $M_p^* \equiv \{m \in \mathcal{M} : p_R^*(m) = p\}$ denote the set of messages after which receiver holds belief p about the state. For any $m \in \mathcal{M}$ let $M_*^{-1}(m) \equiv \{s \in [0,1] : \mu^*(S^*, s) = m\}$ denote the set of signal realizations at which sender chooses message m. Therefore the probability that sender induces belief p in receiver is given by

$$\tau_R^*(p) = \sum_{m \in M_p^*} \sum_{s \in M_*^{-1}(m)} \tau_{S^*}(s).$$

Observe that since τ_{S^*} has finite support and sender is using a pure strategy also τ_R^* has finite support. We can therefore write

$$\sum_{p:\tau_R^*(p)>0} p\tau_R^*(p) = \sum_{p:\tau_R^*(p)>0} p \sum_{m \in M_p^*} \sum_{s \in M_*^{-1}(m)} \tau_{S^*}(s). \tag{16}$$

Note that if some message m is sent with positive probability in equilibrium it must be that

$$p = \frac{\sum_{m \in M_p^*} \sum_{r \in M_*^{-1}(m)} r\tau_{S^*}(r)}{\sum_{m \in M_p^*} \sum_{r \in M_*^{-1}(m)} \tau_{S^*}(r)}$$

since receiver's beliefs are updated using Bayes' rule whenever possible. We can substitute this expression into (16) to write

$$\sum_{p:\tau_R^*(p)>0} p\tau_R^*(p) = \sum_{p:\tau_R^*(p)>0} \frac{\sum_{m \in M_p^*} \sum_{r \in M_*^{-1}(m)} r\tau_{S^*}(r)}{\sum_{m \in M_p^*} \sum_{r \in M_*^{-1}(m)} \tau_{S^*}(r)} \sum_{m \in M_p^*} \sum_{s \in M_*^{-1}(m)} \tau_{S^*}(s).$$

Simplifying the right-hand side we obtain

$$\sum_{p:\tau_R^*(p)>0} p\tau_R^*(p) = \sum_{p:\tau_R^*(p)>0} \sum_{m \in M_p^*} \sum_{r \in M_*^{-1}(m)} r\tau^*(r) = \sum_{s:\tau_{S^*}(s)>0} s\tau_{S^*}(s) = p_0,$$

where the second equality follows from the definitions of M_p^* and M_*^{-1} and the last follows from the fact that $S^* \in \Sigma$. This shows that τ_R^* is a finite support mean-p_0 distribution over $[0,1]$. It follows from Lemma 4 that there exists some $S \in \Sigma$ with $\tau_S = \tau_R^*$. $\square$

Lemma 13. *If (S^*, μ^*, p_R^*) is a persuasive equilibrium then there exist $s, r \in \operatorname{supp} S^*$ such that $v(p_R^*(\mu^*(S^*, s))) \neq v(p_R^*(\mu^*(S^*, r)))$.*

Proof. We have to show that if sender's equilibrium payoff is constant, then it is equal to $v(p_0)$. Suppose otherwise that equilibrium (S^*, μ^*, p_R^*) is such that $v(p_R^*(\mu^*(S^*, s))) = v^* >(<)\ v(p_0)$ for all $s \in \operatorname{supp} S^*$. Since v is nondecreasing, min (max) $\{p \in [0,1] : \tau_R^*(p) > 0\} >(<)\ p_0$. But then it must be that

$$\sum_{p:\tau_R^*(p)>0} p\tau_R^*(p) >(<)\ p_0.$$

Since for all $S \in \Sigma$, $\mathbf{E}(S) = p_0$ there can be no $S \in \Sigma$ such that $\tau_S = \tau_R^*$, contradicting Lemma 12. $\square$

Let $v^*(s) \equiv v(p_R^*(\mu^*(S^*, s)))$ denote sender's interim equilibrium payoff if signal realization s occurs after choosing equilibrium signal S^*. Let $v^* \equiv \sum_{s \in \operatorname{supp} S^*} \tau_{S^*}(s) v^*(s)$ denote the sender's equilibrium expected payoff.

Lemma 14. *If (S^*, μ^*, p_R^*) is a persuasive equilibrium then for all $s \in \operatorname{supp} S^*$, $s < p_0 \Rightarrow v^*(s) < v^*$, $s = p_0 \Rightarrow v^*(s) = v^*$ and $s > p_0 \Rightarrow v^*(s) > v^*$.*

Proof. I prove first that there exist $\alpha, \beta \in \mathbb{R}$ such that, for all $s \in \operatorname{supp} S^*$, $v^*(s) = \alpha + \beta s$. If $\#\operatorname{supp} S^* \leq 2$ the result is trivially true so consider the case where $\#\operatorname{supp} S^* > 2$. Since $\#\operatorname{supp} S^* > 2$ and $S^* \in \Sigma$ there is a $s' < p_0$ and a $s'' > p_0$ in $\operatorname{supp} S^*$. Let α and β be defined by the solution of simultaneous equations $v^*(s') = \alpha + \beta s'$ and $v^*(s'') = \alpha + \beta s''$. Suppose some $r \in \operatorname{supp} S^*$ is such that $v^*(r) > \alpha + \beta r$. If $r < p_0$ a canonical deviation, defined in Appendix A.2, from s' to r is profitable by Lemma 5, if $r > p_0$ a canonical deviation from s'' to r is profitable by Lemma 5, if $r = p_0$ a signal S^{**} with $\tau_{S^{**}} = \tau_{S^*}$ except for $\tau_{S^{**}}(s') < \tau_{S^*}(s')$ and $\tau_{S^{**}}(s'') < \tau_{S^*}(s'')$ and $\tau_{S^{**}}(p_0) > \tau_{S^*}(p_0)$ can be obviously chosen to be feasible and is a profitable deviation. The case in which there is a $r \in \operatorname{supp} S^*$ such that $v^*(r) < \alpha + \beta r$ can be addressed with analogous arguments. This completes the proof that there exist $\alpha, \beta \in \mathbb{R}$ such that, for all $s \in \operatorname{supp} S^*$, $v^*(s) = \alpha + \beta s$.

Observe immediately that if $p_0 \in \operatorname{supp} S^*$, $v^*(p_0) = v^*$. To show the strict inequalities I prove that β is necessarily positive. $\beta = 0$ would imply that sender's payoff is constant in equilibrium, which contradicts Lemma 13. So suppose $\beta < 0$. Then for any $s, r \in \operatorname{supp} S^*$, $v^*(s) \neq v^*(r)$ so also $\mu^*(S^*, s) \neq \mu^*(S^*, r)$ and therefore $p_R^*(\mu^*(S^*, s)) = s$ and $p_R^*(\mu^*(S^*, r)) = r$ since receiver must use Bayes' rule on-path (equilibrium condition (1)) and each on-path message is sent at a single signal realization. This in turn implies that for $s, r \in \operatorname{supp} S^*$ and $s > r$, since $v^*(s) < v^*(r)$, also $v(s) < v(r)$ must hold, which contradicts that v is nondecreasing.

Having established that $\beta > 0$ the result immediately follows. $\square$

C.2 Proof of the proposition

Consider a $s \in \operatorname{supp} S^*$ and suppose otherwise that $\mu^*(S^*, s)$ is such that there exists a $s_{\mu^*(S^*,s)} < s$ with $\mu^*(S^*, s) \in M(s_{\mu^*(S^*,s)})$. There are three possible cases depending on where s lies. Note that we shall use extensively canonical deviations, defined in detail in Appendix A.2.

Case 1: $s > p_0$. There are three sub-cases depending on where $s_{\mu^*(S^*,s)}$ is.

Sub-case 1: $s_{\mu^*(S^*,s)} > p_0$. Consider a canonical deviation from s to $s_{\mu^*(S^*,s)}$ and call S^{**} the deviating test. By definition of $s_{\mu^*(S^*,s)}$ we have that $\mu^*(S^*, s) \in M(s_{\mu^*(S^*,s)})$ so, by optimality of μ^*, it must be that

$$v^*(s_{\mu^*(S^*,s)}) \geq v(p_R^*(\mu^*(S^*, s))) = v^*(s). \tag{17}$$

The sender's payoff from deviating is

$$\begin{aligned}
&\sum_{t \in \operatorname{supp} S^* \setminus \{s, s_{\mu^*(S^*,s)}\}} \tau_{S^{**}}(t) v^*(t) + \tau_{S^{**}}(s_{\mu^*(S^*,s)}) v^*(s_{\mu^*(S^*,s)}) \\
&= \frac{s_{\mu^*(S^*,s)} - p_0}{k} v^* - \frac{s_{\mu^*(S^*,s)} - p_0}{k} \tau_{S^*}(s) v^*(s) + \frac{s - p_0}{k} \tau_{S^*}(s) v^*(s_{\mu^*(S^*,s)}) \\
&\geq \frac{s_{\mu^*(S^*,s)} - p_0}{k} v^* + \tau_{S^*}(s) v^*(s) \left(\frac{s - p_0}{k} - \frac{s_{\mu^*(S^*,s)} - p_0}{k} \right) \\
&= \frac{s_{\mu^*(S^*,s)} - p_0}{k} v^* + \tau_{S^*}(s) v^*(s) \left(\frac{s - s_{\mu^*(S^*,s)}}{k} \right) \\
&= \frac{(s_{\mu^*(S^*,s)} - p_0) v^* + (s - s_{\mu^*(S^*,s)}) \tau_{S^*}(s) v^*(s)}{(s_{\mu^*(S^*,s)} - p_0) + (s - s_{\mu^*(S^*,s)}) \tau_{S^*}(s)} \\
&> v^*
\end{aligned}$$

Where the weak inequality follows from (17) and the strict inequality follows from Lemma 14 since $p > p_0$. This contradicts that S^* is an equilibrium signal.

Sub-case 2: $s_{\mu^*(S^*,s)} = p_0$. Then signal S^{**} with $\operatorname{supp} S^{**} = \{p_0\}$ is feasible and attains payoff of at least $v^*(s) > v^*$ by Lemma 14. It is therefore a profitable deviation.

Sub-case 3: $s_{\mu^*(S^*,s)} < p_0$. A signal $S^{**} \in \Sigma$ such that $\operatorname{supp} S^{**} = \{s_{\mu^*(S^*,s)}, s\}$ obviously exists and attains a payoff of at least $v^*(s) > v^*$ by Lemma 14. It is therefore a profitable deviation.

Case 2: $s < p_0$. By Lemma 14, $v^*(s) < v^*$. This case is analogous to the first sub-case of Case 1.

Case 3: $s = p_0$. By Lemma 14, $v^*(p_0) = v^*$. By Lemma 13 sender's equilibrium payoff is not constant she must induce in receiver, in equilibrium, a belief either strictly below or strictly above p_0. By Lemma 12 it follows that she must induce in receiver both a belief above and below p_0. Since on-path receiver must update using Bayes' rule there must therefore also be a realization $s'' > p_0$ in $\operatorname{supp} S^*$. By Lemma 14, $v^*(s'') > v^*$. A signal $S^{**} \in \Sigma$ with $\operatorname{supp} S^{**} = \{s'', s_{\mu^*(S^*,s)}\}$ obviously exists and has expected payoff strictly higher than v^* as can be seen with a construction analogous to the first sub-case of Case 1.

Having considered all cases, it follows that it cannot be that $\mu^*(S^*, s)$ is such that there exists a $s_{\mu^*(S^*,s)} < s$ with $\mu^*(S^*, s) \in M(s_{\mu^*(S^*,s)})$. If some $s \in [0, 1]$ is such that for all $m \in M(s)$ there exist a $r < s$ with $m \in M(r)$ there are no feasible messaging strategies at s that can be part of an equilibrium, so s cannot be in the support of any persuasive equilibrium signal strategy. So if $s \in \operatorname{supp} S^*$ it must mean that $\mu^*(S^*, s) \in \{m \in M(s) : s = \min M^{-1}(m)\} \neq \emptyset$, as required. This shows (I).

To show (II) consider an on-path message m. Since m is on-path, there must be a set of realizations $S_m \subseteq \operatorname{supp} S^*$ such that $\mu^*(S^*, s) = m$ for $s \in S_m$. By (I) we have that $S_m = \{\min M^{-1}(m)\}$ so $s_m = \min M^{-1}(m)$ is the only possible realization that can send message m in any persuasive equilibrium, concluding the proof.

Appendix D Proof of Lemma 2

There exists no $m \in \mathcal{M}$ such that $\min M^{-1}(m) > p_0$ so there cannot be a persuasive equilibrium with a signal that puts probability mass on realizations above p_0 by Proposition 1. If the equilibrium signal puts all probability mass on p_0 the equilibrium is obviously unpersuasive, completing the argument.

Appendix E Proof of Lemma 3

I show that there cannot be an equilibrium with equilibrium payoff equal to $v(p_0)$ so, a fortiori, there cannot be an unpersuasive equilibrium. This is because there exists a signal S_{mn} that guarantees a payoff strictly greater than $v(p_0)$. Let $\operatorname{supp} S_{mn} = \{s_m, s_n\}$, $S_{mn} \in \Sigma$ because of (5). By construction, for any receiver belief function p_R that satisfies equilibrium condition (2) sender's payoff from sending message m (n) is not below $\underline{v}_m$ $(\underline{v}_n)$. Sender's expected payoff from signal S_{mn} is therefore not below

$$\frac{s_m - p_0}{s_m - s_n}\underline{v}_n + \frac{p_0 - s_n}{s_m - s_n}\underline{v}_m$$

since $m \in M(s_m)$ and $n \in M(s_n)$. The claim follows.

Appendix F Proof of Proposition 2

(ii) ⇔ (iii). We want to show that the full commitment value, $(\operatorname{cav} v)(p_0)$, is the highest ex ante equilibrium payoff among all verifiability structures. Observation 1 shows that the full verifiability structure allows sender to attain the full commitment value. All left to show is that in any equilibrium with any verifiability structure sender's ex ante expected payoff is equal to or less than $(\operatorname{cav} v)(p_0)$. Let S^* denote a signal that attains the full commitment value in the 'Bayesian persuasion' model. That is

$$\sum_{s \in \operatorname{supp} S^*} \tau_{S^*}(s)v(s) \geq \sum_{s \in \operatorname{supp} S} \tau_S(s)v(s) \tag{18}$$

for all $S \in \Sigma$.

Consider some verifiability structure and any equilibrium (S', μ', p'_R) and let τ'_R denote receiver's belief distribution in equilibrium. Sender's ex ante equilibrium payoff is then

$$v' \equiv \sum_{p:\tau'_R(p)>0} \tau'_R(p)v(p)$$

By Lemma 12 there exists some $\tilde{S} \in \Sigma$ such that $\tau_{\tilde{S}} = \tau'_R$. But then

$$(\operatorname{cav} v)(p_0) = \sum_{s \in \operatorname{supp} S^*} \tau_{S^*}(s)v(s) \geq \sum_{s \in \operatorname{supp} \tilde{S}} \tau_{\tilde{S}}(s)v(s) = v'$$

by inequality (18), completing the argument.

(i) ⇒ (ii). Prove this constructively. Fix an arbitrary nondecreasing v and prior $p_0 \in [0,1]$. Let receiver's belief function be maximally sceptical in the sense that in the candidate equilibrium $p^*_R(m) = \min M^{-1}(m)$ for any $m \in \mathcal{M}$. Define as usual $\eta^{p^*_R} : [0,1] \to \mathcal{M}$ to be a selection of optimal messages at every signal realization, given that receiver forms beliefs according to p^*_R. Let $m_s \in M(s)$ denote a message with which realization s is lowest consistent (i.e. $\min M^{-1}(m_s) = s$). Observe that, for all $m \in M(s)$ and all $s \in [0,1]$, $v(p^*_R(m_s)) \geq v(p^*_R(m))$ since $p^*_R(m_s) \geq p^*_R(m)$ by definition of p^*_R and v is nondecreasing. It follows that selection $\eta^{p^*_R}(s) = m_s$ for all $s \in [0,1]$ is optimal. Let $\mu^*(S,s) = \eta^{p^*_R}(s)$ for all $S \in \Sigma$ and all $s \in \operatorname{supp} S$. For all $s \in [0,1]$, let $v^{p^*_R}(s) \equiv v(p^*_R(\eta^{p^*_R}(s)))$ be sender's continuation value in the candidate equilibrium when she observes signal realization s. It is immediate that $v^{p^*_R} = v$ since $v(p^*_R(\eta^{p^*_R}(s))) = v(p^*_R(m_s)) = v(s)$ where the first equality follows from definition of $\eta^{p^*_R}$ and the second from the definition of p^*_R. Let S^* denote a signal that attains the full commitment value in the 'Bayesian persuasion' model. That is

$$\sum_{s \in \operatorname{supp} S^*} \tau_{S^*}(s)v(s) = (\operatorname{cav} v)(p_0).$$

It is immediate that S^* also allows sender to attain the full commitment value in the setting under consideration since we showed that $v^{p^*_R} = v$. Since S^* is optimal in the 'Bayesian persuasion' model it is also a best reply signal in the present setting, so (S^*, μ^*, p^*_R) is an equilibrium in which sender attains the full commitment value.

I claim that sender cannot do worse in any equilibrium. To see why consider any receiver belief function p_R that satisfies (2). Then, by definition of p^*_R, $p_R(m) \geq p^*_R(m)$ for all $m \in \mathcal{M}$. Since v is nondecreasing, $v(p_R(m)) \geq v(p^*_R(m))$ for all $m \in \mathcal{M}$. It follows that $\overline{v}(s) \equiv \sup_{m \in M(s)} v(p_R(m)) \geq \sup_{m \in M(s)} v(p^*_R(m)) = v^{p^*_R}(s) = v(s)$. If p_R is part of an equilibrium, it implies that sender's equilibrium payoff must weakly exceed $(\operatorname{cav} \overline{v})(p_0)$, as otherwise she would have a profitable deviation. But $(\operatorname{cav} \overline{v})(p_0) \geq (\operatorname{cav} v)(p_0)$ since $\overline{v} \geq v$, completing the proof sender cannot do worse in any equilibrium. Sender also cannot do better, as was shown in the proof of (ii) ⇔ (iii).

(ii) ⇒ (i). Prove the contrapositive. Suppose that some $s \in [0,1]$ is not lowest consistent so that for all $m \in M(s)$ there is some $s' < s$ with $m \in M(s')$. Consider $v(s) = \mathbf{1}_{[s,1]}$ and some prior $p_0 < q$. Observe that the only equilibrium with the full verifiability constraint involves signal $S^* \in \Sigma$ with $\operatorname{supp} S^* = \{0, s\}$. For sender to attain the full commitment value in equilibrium receiver's belief distribution must be τ_{S^*}. Observe that if receiver's equilibrium

belief distribution is τ_{S^*}, the equilibrium is persuasive. By Proposition 1, point (II) sender's equilibrium signal must be S^*. But this contradicts Proposition 1 point (I) since s is not lowest consistent.

Appendix G Proof of Proposition 3

If $p_0 = 0$ the proposition is trivially true so let $p_0 > 0$.

(a) ⇒ (b). If an equilibrium exists we have that it is either unpersuasive (leading to a payoff $v(p_0)$ for sender) or persuasive. By Proposition 1, if it is persuasive it must be supported on $\{0,1\}$, as these are the only lowest consistent realizations. Let m_1 denote a message with which realization $s = 1$ is lowest consistent. Obviously in any candidate equilibrium $p_R(m_1) = 1$.

I claim first that there cannot be any unpersuasive equilibria. Suppose otherwise and denote one by (S^*, μ^*, p_R^*). As argued in the previous paragraph, $p_R^*(m_1) = 1$. So $v(p_R^*(m_1)) = v(1) > v(p_0)$ since communication is potentially valuable for sender. It follows that $s = 1$ cannot be in the support of S^* since the equilibrium is unpersuasive. Consider now some $s \in \operatorname{supp} S^*$ and the associated equilibrium message $\mu^*(S^*, s)$. Since $s \neq 1$, $\min M^{-1}(\mu^*(S^*, s)) = 0$ as only 0 and 1 are lowest consistent. Consider the following deviation to signal $S^{**} \in \Sigma$ with $\operatorname{supp} S^{**} = \{0,1\}$ and $\mu^{**}(S^{**}, 0) = \mu^*(S^*, s)$ and $\mu^{**}(S^{**}, 1) = m_1$. Sender's expected payoff from this deviation is $p_0 v(1) + (1 - p_0) v(p_0) > v(p_0)$ since $p_0 > 0$, so it is profitable. This contradicts that (S^*, μ^*, p_R^*) is an equilibrium.

All equilibria must therefore be persuasive. By Proposition 1 in any persuasive equilibrium the signal must be the fully informative one $S^F \in \Sigma$ since only $s = 0$ and $s = 1$ are lowest consistent. By point (II) of Proposition 1 it follows that receiver learns the state in any persuasive equilibrium, and therefore obtains the full-information expected payoff.

The only thing left to show is that persuasive equilibria exist. Consider profile (S^*, μ^*, p_R^*) where $p_R^*(m) = \min M^{-1}(m)$ for all $m \in \mathcal{M}$ and $S^* = S^F$. Let m_1 (m_0) denote a message with which $s = 1$ ($s = 0$) is lowest consistent. Let $\mu^*(S^*, 0) = m_0$ and $\mu^*(S^*, 1) = m_1$ and choose any optimal message given p_R^* off-path. $S^* = S^F$ is clearly a best reply signal as it maximises the probability of realization $s = 1$. The messaging strategy is a best reply by construction and the belief function is consistent with the verifiability structure and obeys Bayes' rule on-path. The candidate profile is therefore an equilibrium, completing the argument.

(a) ⇐ (b). I prove the contrapositive. 0 must be lowest consistent so to negate (a) it must that either: (i) signal realization $s = 1$ is not lowest consistent or (ii) some $s \in (0,1)$ is also lowest consistent, or both.

Consider (i). Suppose there is an equilibrium in which receiver obtains the full-information expected payoff for some nondecreasing and upper semi-continuous v and some $p_0 \in [0,1]$. Then in such an equilibrium receiver learns the state so, since $v(1) > v(p_0) \geq v(0)$ by assumption, such an equilibrium is persuasive. But then by Proposition 1 sender's equilibrium signal must be the full information one S^F. But $s = 1$ is not lowest consistent, contradicting condition (I) of Proposition 1.

Consider (ii). Let $v(r) = \mathbf{1}_{[s,1]}$ and $p_0 \in (0, s)$. Let m_s be a message with which s is lowest consistent. Construct an equilibrium as follows. Choose the sceptical belief function $p_R^*(m) = \min M^{-1}(m)$ for all $m \in \mathcal{M}$. Let sender choose signal $S^* \in \Sigma$ with $\operatorname{supp} S^* = \{0, s\}$. Choose

$\mu^*(S^*, s) = m_s$ and otherwise let μ^* select any optimal message at each realization. Signal S^* maximises the probability of getting the payoff of $v(p_R^*(m_s)) = 1$, which is the highest payoff available to sender, and is therefore a best reply. The messaging strategy is a best reply by construction and the belief function is consistent with the verifiability structure and obeys Bayes' rule on-path. The candidate profile is therefore an equilibrium, completing the argument.

(b) ⇒ (c). Immediate since receiver cannot do better than knowing the state with certainty if he values information since v_R is convex.

(b) ⇐ (c). If he values information, the highest ex ante payoff for receiver corresponds to what he obtains when he chooses the action knowing the value of the state. So, if we show that the highest ex ante payoff among all verifiability structures coincides with the full-information expected payoff v_R^F, we are done. Consider verifiability structure with $M(s) = \{m_0\}$ for all $s \in [0, 1)$ and $M(s) = \{m_0, m_1\}$ for $s = 1$. Let $p_R^*(m_0) = 0$ and $p_R^*(m_1) = 1$ so sender faces continuation payoffs $v^{p_R^*}(p) = v(0)$ for $p \in [0, 1)$ and $v^{p_R^*}(1) = v(1)$ at the information acquisition stage. Since communication is potentially valuable for sender and v is nondecreasing $v(1) > v(p_0) \geq v(0)$ so signal $S^* = S^F$ is best reply for sender. On-path receiver beliefs following the two possible messages are correct, so this is an equilibrium which gives receiver the full-information ex ante expected payoff v_R^F.

Appendix H Verifiability structure and disclosure order

Define the M-induced disclosure order $s' \succeq_M s$ iff $M(s') \supseteq M(s)$. $s' \succeq_M s$ means that s' can mimic s since it can send any message available to s. It is straightforward to check that $([0, 1], \succeq_M)$ is a partially ordered set. Observe that if $s \in [0, 1]$ is lowest consistent (given M) then $s' < s \Rightarrow s' \not\succeq_M s$ i.e. no lower realization can mimic s. Lowest consistency is saying more than that, however. Let $M(s) = \{m_1\}$ for $s \in [0, 1/3)$, $M(s) = \{m_1, m_2\}$ for $s = 1/3$, $M(s) = \{m_2\}$ for $s \in (1/3, 2/3)$, $M(s) = \{m_2, m_3\}$ for $s = 1/3$ $M(s) = \{m_3\}$ for $s \in (2/3, 1)$, $M(1) = \{m_1, m_2, m_3\}$. Observe that M^{-1} is closed. Clearly $s \not\succeq_M 1$ for all $s < 1$. It is also immediate that $s = 1$ is *not* lowest consistent.

Appendix I Existence failure

Let $v(s) = \mathbf{1}_{[1/2,1]}$. Let $M(s) = \{m_s, \hat{m}\}$ for all $s \in [0, 1/2) \cup (1/2, 1]$ and $M(1/2) = \{\hat{m}\}$ with $m_s \neq m_{s'}$ if $s \neq s'$. The verifiability structure is constructed so that all realizations except $1/2$ can identify themselves unequivocally. Additionally, all realizations have a common message $\hat{m}$, which is also the only message available to realization $1/2$. Fix the prior $p_0 = 1/4$. Observe that conditions Assumptions 1 and 2 are satisfied while Assumption 3 is not since $f_M(s) = s$ if $s \in [0, 1/2) \cup (1/2, 1]$ and $f_M(1/2) = 0$, which is not upper semi-continuous.

Observe first that in any candidate equilibrium it must be that $p_R^*(m_s) = s$ by equilibrium condition (2). This implies that sender can secure a payoff arbitrarily close to the equilibrium payoff with the full verifiability structure[39] by choosing a signal with probability mass on

[39] This is the equilibrium payoff in the model with full commitment power i.e. with $M(s) = \{m_s\}$ for all $s \in [0, 1]$

realizations 0 and on $1/2 + 1/n$ for some large n. It follows that any candidate equilibrium must yield at least the full commitment expected payoff of $1/2$. The only receiver distribution of beliefs that allows sender to attain this payoff involves her having belief 0 with probability $1/2$ and belief $1/2$ with probability $1/2$. A necessary condition for this is that following message $\hat{m}$ receiver believes the state to be 1 with probability $1/2$ i.e. $p_R^*(\hat{m}) = 1/2$. This would a persuasive equilibrium so by Proposition 1 it would also be the distribution of the equilibrium signal chosen by sender. But $\hat{m} \in M(0)$ and $v(p_R^*(\hat{m})) = 1 > v(p_R^*(m_0)) = 0$ so sending message $\hat{m}$ is a profitable deviation. It follows that the only candidate equilibrium is in fact not an equilibrium, completing the argument for equilibria in pure strategies. It is immediate that also no equilibrium in mixed strategies exists.

with $m_s \neq m_{s'}$ if $s \neq s'$. Sender's equilibrium payoff is equal to $1/2$ and is attained by choosing a signal supported on 0 and on $1/2$, which are revealed truthfully by construction using messages m_0 and $m_{1/2}$, respectively. See Section 7.1 for details on this benchmark.

References

Acharya, Viral V., Peter DeMarzo, and Ilan Kremer. 2011. "Endogenous information flows and the clustering of announcements." *American Economic Review* 101 (7): 2955–2979.

Argenziano, Rossella, Sergei Severinov, and Francesco Squintani. 2016. "Strategic information acquisition and transmission." *American Economic Journal: Microeconomics* 8 (3): 119–55.

Ben-Porath, Elchanan, Eddie Dekel, and Barton L. Lipman. 2019. "Mechanisms with evidence: Commitment and robustness." *Econometrica* 87 (2): 529–566.

Ben-Porath, Elchanan and Barton L. Lipman. 2012. "Implementation with partial provability." *Journal of Economic Theory* 147 (5): 1689–1724.

Bourbaki, Nicolas. 1995. *General topology: Chapters 1-4*. Springer-Verlag.

Brocas, Isabelle and Juan D. Carrillo. 2007. "Influence through ignorance." *The RAND Journal of Economics* 38 (4): 931–947.

Crawford, Vincent P. and Joel Sobel. 1982. "Strategic information transmission." *Econometrica* 50 (6): 1431–1451.

DeMarzo, Peter M., Ilan Kremer, and Andrzej Skrzypacz. 2019. "Test design and minimum standards." *American Economic Review* 109 (6): 2173–2207.

Di Tillio, Alfredo, Marco Ottaviani, and Peter Norman Sørensen. 2019. "Strategic sample selection." Working Paper.

Dye, Ronald A. 1985. "Disclosure of nonproprietary information." *Journal of Accounting Research* 23 (1): 123–145.

Felgenhauer, Mike and Petra Loerke. 2017. "Bayesian persuasion with private experimentation." *International Economic Review* 58 (3): 829–856.

Fréchette, Guillaume R., Alessandro Lizzeri, and Jacopo Perego. 2018. "Rules and commitment in communication: An experimental analysis." Working Paper.

Gentzkow, Matthew and Emir Kamenica. 2017a. "Bayesian persuasion with multiple senders and rich signal spaces." *Games and Economic Behavior* 104: 411–429.

Gentzkow, Matthew and Emir Kamenica. 2017b. "Disclosure of endogenous information." *Economic Theory Bulletin* 5 (1): 47–56.

Giovannoni, Francesco and Daniel J. Seidmann. 2007. "Secrecy, two-sided bias and the value of evidence." *Games and Economic Behavior* 59 (2): 296–315.

Green, Jerry R. and Jean-Jacques Laffont. 1986. "Partially verifiable information and mechanism design." *Review of Economic Studies* 53 (3): 447–456.

Grossman, Sanford J. 1981. "The informational role of warranties and private disclosure about product quality." *The Journal of Law and Economics* 24 (3): 461–483.

Guttman, Ilan, Ilan Kremer, and Andrzej Skrzypacz. 2014. "Not only what but also when: A theory of dynamic voluntary disclosure." *American Economic Review* 104 (8): 2400–2420.

Hagenbach, Jeanne, Frédéric Koessler, and Eduardo Perez-Richet. 2014. "Certifiable pre-play communication: Full disclosure." *Econometrica* 82 (3): 1093–1131.

Hart, Sergiu, Ilan Kremer, and Motty Perry. 2017. "Evidence games: Truth and commitment." *American Economic Review* 107 (3): 690–713.

Herresthal, Claudia. 2019. "Hidden testing and selective disclosure of evidence." Working Paper.

Janssen, Mathijs. 2018. "The whole truth? Generating and suppressing hard evidence." Working Paper.

Jovanovic, Boyan. 1982. "Truthful disclosure of information." *The Bell Journal of Economics* 13 (1): 36–44.

Jung, Woon-Oh and Young K. Kwon. 1988. "Disclosure when the market is unsure of information endowment of managers." *Journal of Accounting Research* 26 (1): 146–153.

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101 (6): 2590–2615.

Lipman, Barton L. and Duane J. Seppi. 1995. "Robust inference in communication games with partial provability." *Journal of Economic Theory* 66 (2): 370–405.

Lipnowski, Elliot, Doron Ravid, and Denis Shishkin. 2019. "Persuasion via weak institutions." Working Paper.

Mathis, Jérôme. 2008. "Full revelation of information in Sender–Receiver games of persuasion." *Journal of Economic Theory* 143 (1): 571–584.

Milgrom, Paul. 1981. "Good news and bad news: Representation theorems and applications." *The Bell Journal of Economics* 12 (2): 380–391.

Milgrom, Paul. 2008. "What the seller won't tell you: Persuasion and disclosure in markets." *Journal of Economic Perspectives* 22 (2): 115–131.

Milgrom, Paul and John Roberts. 1986. "Relying on the information of interested parties." *The RAND Journal of Economics* 17 (1): 18–32.

Milgrom, Paul and Chris Shannon. 1994. "Monotone comparative statics." *Econometrica* 62 (1): 157–180.

Morris, Stephen and Philipp Strack. 2019. "The Wald problem and the equivalence of sequential sampling and ex-ante information costs." Working Paper.

Okuno-Fujiwara, Masahiro, Andrew Postlewaite, and Kotaro Suzumura. 1990. "Strategic information revelation." *Review of Economic Studies* 57 (1): 25–47.

Pei, Harry. 2015. "Communication with endogenous information acquisition." *Journal of Economic Theory* 160: 132–149.

Rappoport, Daniel. 2017. "Evidence and skepticism in verifiable disclosure games." Working Paper.

Seidmann, Daniel J. and Eyal Winter. 1997. "Strategic information transmission with verifiable messages." *Econometrica* 65 (1): 163–169.