Screening Information

Falak Arrora*

University of Warwick

September 13, 2025

Click here for the latest version

Abstract

How does the presence of fake news affect incentives to acquire legitimate information? I study a model of costly information acquisition where either an honest or a fake sender communicates with a receiver through a platform. The honest sender sends a true but noisy signal, whereas the fake sender sends a false and uninformative signal. The platform can verify the signal's authenticity; however, it faces a tradeoff. Fake news, although harmful for the receiver, makes her more skeptical and increases the honest sender's incentives for acquiring information. The platform commits to a policy that indicates the screening probability and a disclosure rule. My central finding is that the screening policy that maximizes the receiver's welfare often requires tolerating fake news, *even when such screening is costless*. Moreover, not informing the receiver of the screening outcome is sometimes better than full transparency. These findings suggest that complete moderation and fact-checking of content may inadvertently leave the receiver worse off.

Keywords: Information acquisition, communication game, fake news, platforms
JEL-Classification: C72, D82, D83

*Department of Economics, University of Warwick, UK, falak.arrora@warwick.ac.uk, Personal Website

1 Introduction

The past 20 years have seen social media platforms emerge as a primary avenue for information transmission.[1] However, the concurrent rise of fake news on these platforms has plagued communication through this channel (Mont'Alverne et al., 2022). Fake news and social media, in turn, have been blamed for exacerbating polarization (Barrett, Hendrix and Sims, 2021), undermining democratic norms (Allcott and Gentzkow, 2017), popularizing incorrect medical advice (Do Nascimento et al., 2022), and occasionally sparking violence (Thomas and Sardarizadeh, 2024). Consequently, the users of these platforms report low trust in the content they see online.[2]

However, despite the loss of trust, can a discerning audience be better off? Consider a platform that could costlessly screen out all fake news and retain only true information. A naive view would be that screening every message for its authenticity is optimal from the user's perspective. In this paper, I demonstrate that such a view is incorrect because it fails to properly account for the sender's incentives to acquire information in the first place, especially when such acquisition is costly in itself. While fake news is indeed harmful, tolerating a degree of it can incentivize senders of true information to acquire more precise information. Doing so would credibly distinguish their content from fake news and overcome the receiver's increased skepticism. Once aware of this trade-off, we are naturally led to ask: Can the users of these platforms be better off with fake news? Moreover, what would be the optimal screening policy of such a platform?[3]

To answer these questions, I study a game of costly (verifiable) information acquisition and transmission (Che and Kartik, 2009) with two main innovations. First, I introduce a qualitative distinction in the type of information that directly affects the receiver's payoff. Information can be either *true* or *fake*. True information becomes more valuable as its precision increases; for instance, higher-quality evidence increases the receiver's expected payoff. In contrast, fake information is equivalent to an uninformative and fabricated piece of evidence, and if acted upon, may yield a strictly negative payoff to the receiver. Second, I introduce an intermediary, i.e., a platform, that can verify the information costlessly, while the receiver

[1] As per Newman, Fletcher, Robertson, Arguedas and Nielsen, 2024, social media was the main gateway for online news in 2024 for respondents across markets (29%)

[2] Six in 10 (59%) reported being concerned about what is real and what is fake on the internet, see footnote 1 for reference

[3] Multiple social media platforms have recently rolled back efforts to moderate content in support of free speech. Facebook, for example, stopped its fact-checking program for the same reason (Bobrowsky and Vipers, 2025). Here, I provide an alternate view on why less content moderation may be desirable.

herself cannot distinguish between the two types of information. The inability to do so could be due to several reasons. For example, a user of such a platform might lack the time or expertise to verify each message individually. The emergence of professional fact-checkers has also revealed fact-checking to be a non-trivial task (Graves and Cherubini, 2016).

I develop a simple model that features three players: a sender who may be *honest* or *fake*, a receiver who seeks information from the sender, and a platform that can verify the authenticity of any message. A payoff-relevant state of the world is initially unknown to all players, and the receiver would like to match her action to the state. The honest sender wants the receiver to take the correct action and can acquire a costly signal. He can either acquire a low-precision or a high-precision signal. The signal could be interpreted as a piece of evidence, such as a legal document, scientific experiment, or photographic evidence. The fake sender randomly provides fake evidence that *appears* to be of a certain precision level, but is entirely false. The receiver can observe the *apparent* precision of the message, but cannot verify if the evidence is authentic.

The platform commits to a screening policy at the game's outset. The policy consists of a probability with which it screens the authenticity of each message and an *opaque* disclosure rule. Under this disclosure rule, the receiver is informed if a message is screened *and* found to be fake. Otherwise, the receiver only observes the message and its apparent precision.

Theorem 1 illustrates the main result of the paper and characterizes the optimal screening probability that maximizes the receiver's payoff. I show that such a screening policy often requires tolerating fake news. When the cost of high-precision information is low, the main trade-off disappears, and complete screening is always optimal. As this cost rises, the honest sender switches to low-precision information despite having the same objective as the receiver. By allowing some fake news on the platform and thereby making the receiver skeptical, the honest sender's precision choice can be restored in a way that sometimes also improves the receiver's welfare. Specifically, when the gain in precision is sufficiently large or the probability of high-precision evidence being fake is low enough, partial screening of messages is optimal. I also provide closed-form characterizations.

In section 5, I extend the model to allow the *transparent* disclosure rule. Under this disclosure rule, the receiver is fully informed about the platform's screening choice and the outcome if the message is screened. Proposition 1 shows that each disclosure rule outperforms the other for different intervals of high precision information. For intermediate values of this cost, the *transparent* rule trumps the *opaque* rule, whereas as the cost goes on rising, the *opaque* regime is better. This is because under the transparent rule, the threat of not screening the message needs to increase as the cost increases. Specifically, it goes from full screening to no screening. However, this threat remains constant under the opaque rule. Eventually, for some cost cutoff, one becomes greater than the other. Moreover, randomizing between the two disclosure rules proves to be even more welfare-improving as characterized in proposition 2.

The existing literature on strategic information acquisition and transmission has demonstrated how different preferences (Argenziano, Severinov and Squintani, 2016) or different priors of senders (Che and Kartik, 2009) can provide greater incentives to acquire information, albeit the forces driving these results are different from mine. In these papers, senders are motivated by either off-path threats or, because the receiver discounts the advice in equilibrium, due to the possibility of some information being hidden.[4] By contrast, I introduce a 'fake' sender who provides no informational value and simply injects noise. His presence imposes a complex externality on the honest sender. On the one hand, the fake messages may drown out some true messages that are not very precise. On the other hand, the honest sender now has greater incentives to acquire more precise information. I further show *how* and *when* this externality can be exploited to make the receiver better off.

While the main application mostly speaks to social media and fake news, the model can be adapted to incorporate related settings. These would include any application where a receiver cannot directly verify the authenticity of the sender's information and must rely on an intermediary to inform her. For example, consider a CEO who wishes to invest in a start-up and is advised by an advisor. The advisor can decide how thoroughly to vet the information being shared by the target firm. The CEO is often better off if the adviser does not screen every start-up, even though occasionally investing in a fraudulent company is harmful for him. Or, as another application, consider the increasing cases of data fraud and fabrication in academia (Bik, 2022). Journals and editors can decide how frequently to vet or reproduce the information being supplied by the authors. By not checking every paper and tolerating some degree of bad science, information precision and transparency can sometimes be improved. As a last example, Ponzi or pyramid schemes are often floated along with genuine investment opportunities. Financial advisors and regulators can decide how frequently to verify the information being shared. While the exact protocol of communication in these examples may differ, nevertheless, the main forces identified in the paper are present.

The next section discusses the related literature. In section 3, I develop a model of strategic information acquisition and communication that formalizes the trade-off. Section 4

[4] These results are discussed in greater detail in section 2

characterizes the platform's optimal screening policy. In section 5, I explore some extensions to the model. Section 6 provides some discussion of the main modeling assumptions before concluding in section 7. The Appendix contains all proofs not shown in the main text.

2 Related Literature

This paper is related to different strands of the literature. First, it is related to the literature on communication games that incorporate costly lying by biased senders in a cheap-talk setting. Two related papers here include Kartik, Ottaviani and Squintani, 2007 and Kartik, 2009. In an unbounded type space, Kartik, Ottaviani and Squintani, 2007 characterizes a separating equilibrium that illustrates language inflation in equilibrium, but the (sophisticated) receivers infer the correct type correctly, while Kartik, 2009 shows incomplete separation in a bounded type space. In the current paper, however, fake news is transmitted by only one type of sender, while the other type always reports the truth. Moreover, the focus of the study is not so much on the amount of information communicated and inferred in equilibrium, but rather on the tradeoff between improving information precision and the harm from allowing fake news.

More closely related is the literature on information acquisition. Che and Kartik, 2009 study how having different priors can incentivize better information acquisition in a game of disclosure. In Che and Kartik, 2009, the expert chooses a higher effort to convince the receiver that her opinion (mean of her subjective prior) is wrong, and to overcome the subsequent discounting of advice in equilibrium by the receiver. In this paper, the reason the sender exerts higher effort is to overcome the skepticism induced by the receiver's inability to distinguish between the sender's type. In a similar vein, Argenziano, Severinov and Squintani, 2016 study whether a DM should delegate costly information acquisition to a possibly biased expert in a cheap-talk setting. They provide conditions under which delegating information acquisition to a biased expert can result in more precise action taken by the DM than if the DM were to acquire information herself, given the bias is sufficiently small. However, the forces here are different; in the overt case, this is driven by an off-path babbling equilibrium, while in the covert case, the inflexibility of the equilibrium language helps sustain over-investment in information acquisition. Another related paper incorporating information acquisition is Kartik, Lee and Suen, 2017 with multiple biased experts. They find that having too many experts reduces the incentive to acquire information, even if such experts are like-minded. In another departure from the models studied here, I also introduce an intermediary that can screen the receiver's type and characterize the optimal screening policy.

Less closely related is the literature on Bayesian persuasion, incorporating lies and manipulation in the signal transmitted. Ederer and Min, 2022 study a Bayesian persuasion setting where the sender can lie about the state he privately knows, and with some probability, the lie will be detected. They find that as the probability of the lie being detected increases, the sender lies more frequently. Lipnowski, Ravid and Shishkin, 2022 presents a geometric approach that characterizes how the sender's informativeness and the receiver's payoff evolve as the message's credibility increases. Unlike these papers, the sender in my paper acquires costly information about the state, and the type of the sender doesn't directly impact the receiver's payoff, but rather does so only through her posterior beliefs.

Another related paper is by Glazer, Herrera and Perry, 2021, who study the receiver's learning about a product in the presence of fake reviews. They demonstrate the platform's inability to manipulate reviews (e.g., pooling messages) in a way that would aid in the receiver's learning. In contrast, I study the information acquisition incentives of the honest sender and allow the platform to check the sender's type directly. In the current paper, the platform's ability to screen improves the receiver's welfare and learning.

A vast literature on the economics of social media has emerged. Acemoglu, Ozdaglar and Siderius, 2024 provide a detailed model of social media platforms and demonstrate the emergence of filter bubbles. I further direct the reader to Aridor, Jiménez-Durán, Levy and Song, 2024 and the references therein for a comprehensive survey.

The paper is also related to the literature on information aggregation in models of social networks in the presence of fake news. Acemoglu, Ozdaglar and ParandehGheibi, 2010 study a model of information aggregation in the presence of misinformation. They characterize the extent of misinformation in the presence of *forceful* agents who may spread misinformation. Azzimonti and Fernandes, 2023 study how social media networks and fake news affect polarization in a society. Stiglitz and Kosenko, 2024 also provide a nice discussion on the economics of information in the presence of disinformation.

3 Model

In this section, I develop a model of a communication game with 3 players: a Receiver (she), a Platform (it), and a Sender (he). Messages take the form of verifiable information and are transmitted from the sender to the receiver through the platform. Payoffs depend on an unknown state of the world, $\theta \in \{0,1\}$ with a common prior $P(\theta = 1) = p$. The receiver would like to match her action $a \in \{0,1\}$ to the state. However, she cannot acquire information about the state on her own and must rely on the sender, whose type — *honest* or *fake* — is unknown to her.

Depending on the type of the sender, the information can be *true* or *fake.* A *fake* message would mean fabricated or incorrect information that is uninformative about the state. While the receiver can observe the message and its *apparent* precision, she cannot verify the message herself and must rely on the platform to screen it for her. The platform commits to a screening probability at the beginning that determines how frequently a message is screened.

Sender. Sender is one of two types, $t \in \{H, F\}$: *Honest* or *Fake* and this is privately known to him. The common prior that the type is honest is denoted by $p\,(t = H) = q \in (0, 1)$

The *honest* sender performs a costly investigation about the state. Specifically, he chooses a precision level $\kappa \in \{\alpha, \beta\}$, $(max\{p, 1-p\} < \alpha < \beta \leq 1)$ of the evidence he will acquire.[5,6] Precision here can be understood as the evidence supporting a message or the credibility of the message. Formally, it is the probability that the signal equals the true state $P\,(x = 1|\theta = 1) = P\,(x = 0|\theta = 0) = \kappa$. A higher κ leads to a higher shift in the posterior.

While the cost of low precision $\kappa = \alpha$ is normalized to 0, high precision evidence $\kappa = \beta$ carries a strictly positive cost $0 < c \leq \bar{c}$. Throughout the paper, the term 'cost' refers to the expense of acquiring high-precision information.

A signal $x \in \{0, 1\}$ with the chosen level of precision is then realized. Importantly, if a third party were to verify the evidence supporting the message, it would be validated.

The *fake* sender is assumed to be non-strategic and sends a fake signal (that is verifiable). If this information were to be verified by a third party, it would not be validated. Fake messages are assumed to be entirely uninformative about the state.

The fake sender sends a message $x \in \{1, 0\}$ that only *appears to be* of precision $\kappa \in \{\alpha, \beta\}$ but if were to be verified, would not change the receiver's prior. The fake sender randomly draws a message precision pair that does not depend on θ. Specifically, the fake sender randomizes with equal probability between the four possible messages (x, κ) where $x \in \{0, 1\}$ and $\kappa \in \{\alpha, \beta\}$.[7] For notational convenience, I denote $f_\alpha := f\,(1, \alpha) = f\,(0, \alpha)$ and

[5] Alternatively, we could assume that the sender chooses between low effort or high effort $e \in \{e_l, e_h\}$ where low effort leads to a low precision signal and high effort leads to a high precision signal.

[6] The restriction $max\{p, 1-p\} < \alpha$ ensures that low precision information is always informative and would otherwise induce an action.

[7] This setup is equivalent to assuming two types of fake sender, one that always sends a message 1 and the other that always sends a message 0. The prior probabilities are $q_1 + q_2 = 1 - q$ and the probability they

$f_\beta := f(1,\beta) = f(0,\beta)$.

Message space consists of the signal $x \in \{0,1\}$ and the claimed precision level $\kappa \in \{\alpha, \beta\}$. It is denoted by $\mathcal{M} = \{(x,\kappa)\}$.

Platform: The platform cannot observe the precision of a message, but can observe that a message has arrived.[8] At the beginning of the game, it publicly commits to a screening probability that indicates the likelihood with which a message will be screened and is denoted by $\gamma \in [0,1]$. Screening a message here means that the platform will verify the information. This is equivalent to verifying whether the sender is fake or honest. I assume screening to be completely *costless* for the platform.[9]

If a message is screened and found to be fake, the receiver is always informed. Otherwise, the receiver only observes the message and its apparent precision. I call this the *opaque* disclosure rule.[10] To summarize, the receiver observes the following:

$$(x,\kappa,t=F) \in \{0,1\} \times \{\alpha,\beta\}, \text{ if message is screened and the sender is Fake },$$
$$(x,\kappa) \in \{0,1\} \times \{\alpha,\beta\}, \text{ otherwise}$$

Payoffs: The receiver chooses an action $a \in \{0,1\}$ after observing the message.

If the information is true, taking the correct action gives her a payoff of 1, while the incorrect action gives her a payoff of -1. Without loss, I assume $p \in [0, 0.5]$ such that at the prior belief her preferred action is 0.

Crucially, acting according to a fake message, $x = 1$, yields a negative payoff, $-k < 0$. That is, whenever a message persuades her to take an action different from her prior action and that message turns out to be fake, the receiver gets a strictly negative utility. If $x = 0$ and the message is fake, she gets the same payoff as at her prior.[11]

draw a message of precision κ is given by f_κ

[8]In section 5, I discuss the optimal policy if the platform could also observe the precision level before verifying the message.

[9]Assuming screening to be costly would only make partial screening more likely than full screening

[10]In section 5, I discuss the *transparent* disclosure rule where the receiver is informed of both the ex-post screening choice, $\gamma = 1$ or $\gamma = 0$, and the outcome if the message is screened

[11]It is sufficient to assume that the *expected* payoff from acting on fake news is negative. The receiver may still get a positive payoff from taking the correct decision, but must get a negative payoff large enough in the incorrect state.

The receiver's payoffs are as follows:

$$U_r(\theta, a, t) = \begin{cases} -k, & \text{if } a = 1 \text{ and } t = F, \\ 1, & \text{if } a = \theta \text{ and otherwise,} \\ -1, & \text{if } a \neq \theta \text{ and otherwise} \end{cases}$$

The honest sender's payoffs are:

$$U_h(\theta, a, c(\kappa)) = \begin{cases} 1 - c(\kappa), & \text{if } \theta = a \\ \quad - c(\kappa), & \text{if } \theta \neq a \end{cases}$$

The honest sender would like the receiver to take the correct action; he gets a payoff of 1 if $\theta = a$, 0 if $\theta \neq a$. From this, his cost of acquiring the precision is subtracted. Hence, truthful reporting is strictly dominant, and without loss of generality, we can restrict attention to equilibria in which he always reports the truth. As such, the strategies are only specified to reflect the equilibrium choice of precision.

The platform's objective is to maximize the receiver's welfare. This is a natural assumption in a competitive market. Platforms compete with one another for users and strive to maximize their retention.

Timeline: The game proceeds as follows:

(i) Nature draws $\theta \in \{0, 1\}$, unobserved by all players, and the sender privately observes his type $t \in \{H, F\}$

(ii) The platform commits to a screening probability, γ, observed by all players

(iii) If $t = H$, the sender decides on a precision level: α or β, and then draws a signal $x \in \{0, 1\}$. If $t = F$, the sender draws a message-precision pair according to $f(x, \kappa)$

(iv) The platform determines whether or not to screen the transmitted message

(v) The receiver observes the message and the apparent precision. She also observes the sender type if the message was screened and found to be fake

(vi) The receiver then takes an action $a \in \{0, 1\}$

(vii) Payoffs are realized. The state and the sender's type become common knowledge

Strategies and equilibrium: The honest sender's strategy maps the platform's screening probability choice to a precision level. The receiver's strategy maps her observed tuple, $(x, \kappa, t = F)$ or (x, κ), to an action. The equilibrium solution concept is Perfect Bayesian Equilibrium (PBE, hereafter) in which (i) each player's strategy maximizes their expected payoff given their beliefs (ii) the receiver updates her posterior by Bayes rule wherever possible and (iii) there are no off-path beliefs specified since every message-precision pair arises in equilibrium with positive probability. I further restrict attention to pure strategies.[12]

3.1 Preliminaries

Simplifying Assumptions: For the rest of the paper, I assume the following to simplify the exposition. General results are stated in the appendix and do not change the main insight:

$$\beta = 1, p = 0.5$$

The problem also becomes trivial if the receiver always finds it optimal to follow the signal, even though there is some probability that the message may be fake. This would happen when the disutility from fake news is low enough that it does not interfere with the receiver's decision-making. The main trade-off then disappears because the honest sender no longer needs to distinguish their content from fake information to convince the receiver to take an action, and their effort choice would be independent of the screening probability choice. Then, full screening is always optimal.

To make the problem non-trivial, I make the following assumption for the rest of the analysis,

Assumption 1 (Non-triviality)**:**

$$k_\alpha = \frac{q\,(2\alpha - 1)}{(1-q)\,f_\alpha} \leq k$$

This assumption ensures that the consequences of acting on fake information are severe enough that the receiver finds it optimal to ignore the low-precision message $x = 1$ and take the prior preferred action 0 when no screening takes place.

If $k_\beta = \frac{q(2\beta-1)}{(1-q)f_\alpha} < k$, then the receiver also finds it optimal to ignore a high-precision

[12]This is without loss because the optimal policy would ensure that the receiver never mixes. Indifference only arises in a knife-edge case with parameter values of zero measure.

signal $x = 1$. This threshold could be greater or lower than k_α. I solve separately for both the cases when high precision information is not ignored $k_\alpha < k \leq k_\beta$, and when high precision information is also ignored, that is $\max\{k_\alpha, k_\beta\} < k$.

4 Optimal Screening

The optimal screening probability maximizes the receiver's welfare and is implemented by a benevolent platform. The next result presents the main finding of the paper and identifies the optimal γ for different cost intervals. This is illustrated in figure 1.

Theorem 1: *The optimal screening policy is,*

(1) If $\alpha \leq 0.75$ and $1-\alpha < c \leq 0.5$, then partial screening, $\gamma_o < 1$, is optimal and the sender chooses high-precision information. This screening probability is given by $\gamma_o = 1 - \frac{k_\alpha}{k}$.

(2) Otherwise, full screening is better. The sender chooses high precision when costs are low ($c \leq 1 - \alpha$) and low precision when costs are high ($c > 1 - \alpha$)

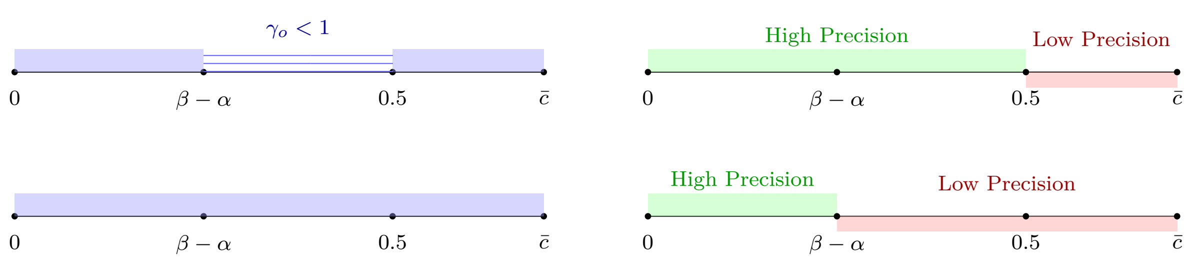

Figure 1: Optimal screening probability and precision when: (i) $\alpha \leq 0.75$ (ii) $\alpha > 0.75$

Proof. The proof is shown through a series of lemmas.

The first lemma analyzes the benchmark case of full screening and establishes the cost values where full screening is always optimal.

Lemma 1: *Full screening is optimal for $c \leq 1 - \alpha$ and the sender chooses high precision information*

Proof. Full screening is equivalent to a game of complete information with two players: a sender (honest or fake) and a receiver. Specifically, now the receiver can observe the type of the sender. When the sender is honest and since $0.5 < \alpha < \beta$, the receiver finds it optimal to take action 1 when $x = 1$ and action 0 when $x = 0$ at any precision level. When the sender is fake, she acts according to her prior information and takes action 0.

The honest sender's payoff from choosing low precision information is α, whereas his payoff from choosing high precision information is $1 - c$. So for $c \leq 1 - \alpha$, the cost is low enough for the honest sender to choose high precision information and hence full screening is optimal. This is illustrated in figure 2. ■

Figure 2: Sender's precision choice for different cost intervals when $\gamma = 1$

The next lemma turns to the receiver's decision and characterizes the minimum screening required by the receiver to benefit from high-precision information despite occasionally acting on fake information.

Lemma 2:

(1) The receiver is better off with high precision information for all $\max\{\gamma_r, 0\} \leq \gamma \leq 1$ *than full screening and low precision information*

(2) Where, $\gamma_r = 1 - \frac{8q(1-\alpha)}{(1-q)k} < 1$

Proof. The receiver's payoff with full screening and low precision information is $q(2\alpha - 1)$. Her payoff when the sender chooses high-precision information and the platform screens messages with some screening probability γ is given by $q(2\beta - 1) - 0.25(1 - \gamma)(1 - q)k$. Comparing the two payoffs provides the required expression. Moreover, for any finite $k < \infty$, the expression $\frac{8q(1-\alpha)}{(1-q)k}$ is strictly positive, ensuring a screening probability less than full screening can be welfare improving. ■

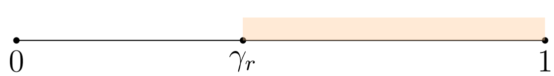

Figure 3: Values of γ for which the receiver is better off with high precision information

The following lemma considers the sender's precision choice and characterizes the maximum screening probability that would incentivize the sender to choose high precision information.

Lemma 3: *The sender's precision choice for different values of* γ *and* c *are as follows:*

(1) If $k_\alpha \leq k < k_\beta$ and $1-\alpha < c \leq 0.5$, the sender chooses high precision for all $\gamma \leq \gamma_o < 1$

(2) If $k_\alpha < k_\beta \leq k$ and $1-\alpha < c \leq 0.5$, the sender chooses high precision for all $\gamma'_o \leq \gamma_o < 1$. Otherwise, high precision cannot be incentivized for $1-\alpha < c \leq 1$

(3) Where, $\gamma'_o = 1 - \frac{k_\beta}{k} < 1, \gamma_o = 1 - \frac{k_\alpha}{k} < 1$

(4) For $c > 0.5$ the sender chooses low precision for any $\gamma \in [0, 1]$

Proof. I will first discuss the case where $k_\alpha \leq k < k_\beta$. Let the screening probability chosen by the platform be γ

If a fake message is screened, the receiver is informed and she takes action 0, acting according to her prior information. If instead the message is not screened or is found to be true, the receiver only observes the message and its apparent precision level. If the message is $x = 0$, she is better off taking the action 0 even if the message is fake. If instead the message is $x = 1$, her expected payoff from following the advice directly depends on γ. Specifically, she ignores the signal and chooses her prior preferred action 0 if,

$$q\,(\alpha - 0.5) - 0.25\,(1 - q)\,(1 - \gamma)k \leq q(0.5 - \alpha)$$

For all $\gamma \leq 1 - \frac{k_\alpha}{k}$, the receiver disregards a low precision signal of $x = 1$ and takes actions 0.

For these values of γ, the sender anticipates the receiver will choose $a = 0$. So his expected payoff from choosing low precision is his expected utility from this action, which is $1-p = 0.5$. His payoff from choosing high precision information is instead given by $1 - c$. Comparing the two provides the cutoff cost $c = 0.5$ below which the sender is better off choosing high-precision information. For $c > 0.5$, low precision is always preferred. Furthermore, since $\alpha > 0.5$, $1 - \alpha < 0.5$, so there exists a non-empty interval where high precision can be incentivized.

If instead $k_\beta \leq k$, then with no screening, i.e. $\gamma = 0$, the receiver ignores $x = 1$ even when the precision is high. To incentivize the sender to choose high-precision information, the screening probability will need to be such that the receiver finds it optimal to follow high-precision advice while ignoring the low-precision message, $x = 1$. For all $\gamma \geq 1 - \frac{k_\beta}{k}$, the receiver follows $x = 1$ with high precision and for all $\gamma < 1 - \frac{k_\alpha}{k}$, the receiver ignores low precision. Only when $k_\alpha \leq k_\beta$ is the former lower than the latter, allowing the platform to incentivize the sender to switch to high precision information. When $\alpha < \beta = 1$ and $f_\alpha = f_\beta$, $k_\alpha < k_\beta$.

By assumption since $k_\alpha \leq k < k_\beta$ or $\max\{k_\alpha, k_\beta\} \leq k$, γ_o and γ_o' are always strictly less than 1. ■

Figure 3 and figure 4 illustrate the following the previous two lemmas.

Figure 4: Values of γ for which the sender chooses high precision information when $k_\alpha < k \leq k_\beta$

The previous two lemmas provide the minimum screening probability that makes the sender better off, γ_r, and the maximum screening probability that would incentivize the sender to switch to high precision information, γ_o. Putting these pieces together, partial screening is welfare-improving and feasible if $\gamma_o \geq \gamma_r$, which holds when $\alpha \leq 0.75$ and $1 - \alpha < c \leq 0.5$, where the sender can be incentivized. For all such low-precision values, the platform would optimally allow some fake news when $1 - \alpha < c \leq 0.5$. Outside of these regions, full screening is optimal. This establishes the theorem. The strategies for the platform, sender, and receiver constitute a Perfect Bayesian Equilibrium, as each player best responds given their beliefs, and the receiver's beliefs are consistent with Bayes' rule on the equilibrium path.

■

For the sake of completeness, I also briefly discuss the second benchmark case with no screening, $\gamma = 0$. This is equivalent to analyzing an incomplete information game between a sender of an unknown type and a receiver. When $k_\alpha \leq k < k_\beta$, i.e., a low precision message $x = 1$ is not persuasive, but a high precision message is. The receiver takes action 0 on observing any low-precision message. In this case, the sender gets a payoff of 0.5 from choosing low precision information and a payoff of $1 - c$ from choosing high precision. Hence, for $0.5 < c$, the sender prefers low precision and high precision otherwise, as shown in figure 5. When $\alpha > 0.75$, the optimal policy requires full screening. Improving the receiver's welfare comes at the cost of incentivizing low-precision information for $1 - \alpha < c \leq 0.5$, where the sender chose high precision under no screening.

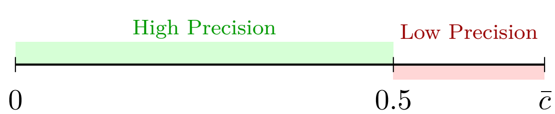

Figure 5: Sender's precision choice for different cost intervals when $\gamma = 0$ and $k_\alpha \leq k < k_\beta$

The second case corresponds to the situation where $\max\{k_\beta, k_\alpha\} \leq k$ and especially highlights the value of partial screening. Now, even high precision is not enough to convince the receiver, and she always chooses action 0 no matter the precision level. In this case, the sender can do no better than always choose low precision, as depicted in figure 6. While full screening would restore high precision information for $c \leq 1-\alpha$, it wouldn't do so for $c > 1-\alpha$. When $\alpha \leq 0.75$, partially screening messages incentivizes high precision information, which wouldn't be the sender's choice either under full screening or no screening.

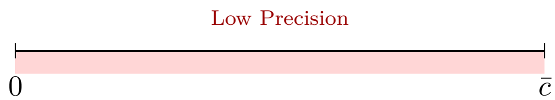

Figure 6: Sender's precision choice for different cost intervals when $\gamma = 0$ and $\max\{k_\alpha, k_\beta\} \leq k$

5 Extensions

5.1 Transparent Disclosure Rule

The main analysis assumed the opaque disclosure rule, where the receiver is only informed if a message is screened *and* found to be fake. Another possibility is the *transparent disclosure rule*, where the receiver is informed of the ex-post screening decision of the platform $\gamma = 1$ or $\gamma = 0$, and its outcome, $t = \{H, F\}$ if the message is screened. That is, the receiver observes the following:

$$\begin{aligned} (x, \kappa, t) &\in \{0,1\} \times \{\alpha, \beta\} \times \{H, F\} \text{ if } \gamma = 1, \\ (x, \kappa) &\in \{0,1\} \times \{\alpha, \beta\} \text{ if } \gamma = 0 \end{aligned}$$

When the disclosure rule is transparent, the sender will choose high-precision information only if the *threat* of his message not being screened is just high enough compared to the cost of high-precision information. The next result characterizes the maximum of such screening probability under the transparent rule that would incentivize the sender to choose high-precision information.

Lemma 4: *Let the disclosure rule be transparent,*

(1) If $k_\alpha \leq k < k_\beta$ *and* $1 - \alpha < c \leq 0.5$*, the sender chooses high precision for all* $\gamma \leq \gamma_{tr}$

(2) Where, $\gamma_{tr} = \frac{0.5-c}{\alpha-0.5} < 1$

(3) If $\max\{k_\alpha, k_\beta\} \leq k$*, high precision cannot be incentivized for any* γ

Proof. When $k_\alpha \leq k < k_\beta$, if the sender chooses low precision information, his payoff would be $\gamma(\alpha) + (1-\gamma)(0.5)$. This is because with probability γ his message will be screened and he gets a payoff of α whereas with probability $1-\gamma$, the message is not screened and the receiver takes action 0, giving the sender an expected payoff of $p(\theta = 0) = 0.5$. If he instead chooses high precision information, then his payoff is $1-c$. Comparing the two payoffs provides the value for γ_{tr}.

When $\max\{k_\beta, k_\alpha\} \leq k$, then the receiver always chooses action 0 for any precision level. Then the sender's payoff from low precision information is $\gamma(\beta) + (1-\gamma)(0.5)$ and from low precision information is $\gamma(\alpha) + (1-\gamma)(0.5)$. Comparing the two payoffs, high precision is better than low precision only for $\frac{c}{1-\alpha} \leq \gamma$. But for all $c > 1-\alpha$, this is always greater than 1.

Furthermore, $\gamma_{tr} = \frac{0.5-c}{\alpha-0.5} < 1 \implies 0.5 - c < \alpha - 0.5 \implies 1-\alpha < c$ which holds. ■

The opaque disclosure rule provides stronger incentives for the sender to choose high precision than the transparent regime. This is because under the opaque regime, the receiver is *always* skeptical of low-precision information, whereas under the transparent regime, the skepticism only arises when a message is not screened. The ex-ante choice of the screening probability directly impacts the receiver's decision to ignore or follow any information under the opaque regime. However, under the transparent regime, it is only the ex-post decision to not screen, $\gamma = 0$, that determines the receiver's choice to act. Nevertheless, for a certain cost interval, the transparent rule proves to be more effective than the opaque rule.

The maximum screening probability under the transparent rule depends on the cost of high-precision information. As the cost goes from $1-\alpha$ to 0.5, this screening probability also falls from full screening ($c = 1-\alpha$) to no screening ($c = 0.5$). That is, as the cost of information becomes higher, the threat of *not screening* the message also needs to increase. The maximum screening probability under the opaque regime, however, is independent of the cost but strictly greater than 0 (for all $k_\alpha < k$). Eventually, for cost above some cutoff $c^* \leq c$, γ_{tr} becomes lower than γ_o, while being greater for all costs below this cutoff.

The next result characterizes when the transparent regime might be better than the opaque regime and is shown in figure 7. When $\max\{k_\alpha, k_\beta\} \leq k$, then only the opaque disclosure rule can be optimal.

Proposition 1: *Let* $\tilde{c} = 1 - \alpha + \frac{8q(\alpha-0.5)(1-\alpha)}{k(1-q)}$ *and* $c^* = 1 - \alpha + \frac{4q(\alpha-0.5)(2\alpha-1)}{k(1-q)}$. *The optimal screening policy that maximizes the receiver's welfare is as follows,* *(1) If $k_\alpha \leq k < k_\beta$ and $1-\alpha < c \leq \min\{c^*, \tilde{c}\}$, then partial screening and transparent disclosure rule is optimal*

(2) If $k_\alpha \leq k < k_\beta$ and $\tilde{c} \leq c^* < c < 0.5$, then partial screening and opaque disclosure rule is optimal. If $k_\beta \leq k$ and $\alpha \leq 0.75$, then partial screening and opaque rule is optimal

(3) Otherwise, full screening is optimal

Proof. The transparent rule is better than the opaque rule if $\gamma_{tr} \geq \gamma_o$ and $\gamma_{tr} \geq \gamma_r$. The first condition translates to $c \leq 1-\alpha+\frac{4q(\alpha-0.5)(2\alpha-1)}{k(1-q)}$ whereas the second condition translates to $c \leq 1-\alpha+\frac{8q(\alpha-0.5)(1-\alpha)}{k(1-q)}$. Putting the two requirements together gives $c \leq \min\{\tilde{c}, c^*\}$. For all such c, the transparent disclosure rule improves the receiver's welfare and provides higher screening than the opaque rule.

If $\tilde{c} > c^*$, then $\alpha > 0.75$ and the opaque rule is never optimal. Otherwise, the opaque rule is optimal and better than the transparent disclosure rule for all $c^* \leq c \leq 0.5$, as in the main theorem — the same proof as before follows. ■

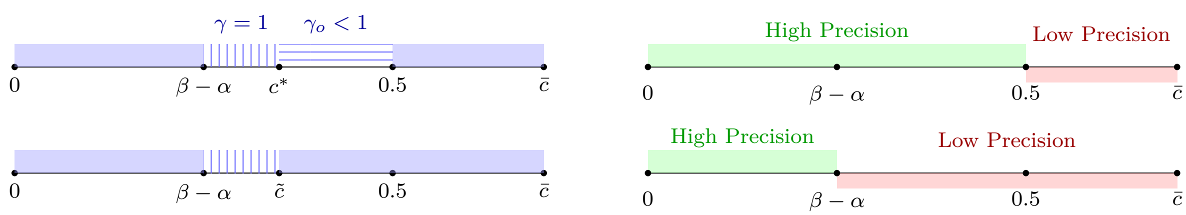

Figure 7: Optimal policy and precision when: (i) $\alpha \leq 0.75$ (ii) $\alpha > 0.75$

An immediate consequence of the above result is that whenever $k_\alpha \leq k < k_\beta$ partial screening and transparent disclosure rule is necessarily optimal for some cost interval, for any low-precision value. This can be seen by examining the cost cutoffs c^* and $\tilde{c}$, which are both strictly above $1-\alpha$. From theorem 1, while $\alpha \leq 0.75$ is a sufficient condition for partial screening, it is not necessary.

Corollary 1: *If $k \leq k_\alpha < k_\beta$, there always exists a non-empty cost interval where partial screening is optimal and the disclosure rule is transparent, for any value of α.*

Proof. Let $c^* \leq \tilde{c}$, then the cost interval where partial screening and transparent rule is optimal is given by $1-\alpha < c < 1-\alpha+\frac{4q(\alpha-0.5)(2\alpha-1)}{k(1-q)}$. The second term of c^* is strictly positive for any finite $k < \infty$, $\alpha > 0.5$, and $q \in (0,1)$, all of which hold. If $\tilde{c} < c^*$, then this interval is given by, $1-\alpha < c < 1-\alpha+\frac{8q(\alpha-0.5)(1-\alpha)}{k(1-q)}$. Since $0.5 < \alpha < 1$, the second term of $\tilde{c}$ is also always positive. ■

5.1.1 Randomizing between opaque and transparent

The platform could also choose to randomize between the two disclosure rules. Specifically, at the beginning of the game, the platform commits to a screening probability $\gamma \in [0,1]$ and a disclosure probability $d \in [0,1]$. The disclosure probability indicates the frequency with which the disclosure rule will be transparent, and with the complementary probability, the disclosure rule will be opaque. The next result characterizes the optimal policy when randomization is feasible.

Proposition 2: *Let $\tilde{c} = 1-\alpha+\frac{8q(\alpha-0.5)(1-\alpha)}{k(1-q)}$. The optimal policy that maximizes the receiver's welfare is as follows:*

(1) If $\alpha \leq (1-0.5c)$ and $1-\alpha < c \leq 0.5$, platform chooses transparent disclosure rule with probability $d = \frac{0.5-c}{\alpha-0.5}$ and fully screens, i.e., $\gamma_{tr} = 1$, and with the complementary probability chooses opaque disclosure with screening probability $\gamma_0 = 1-\frac{k_\alpha}{k}$

(2) If $\alpha \geq (1-0.5c)$ and $1-\alpha < c \leq \tilde{c}$, the optimal screening probability is $\gamma_{tr} = \frac{0.5-c}{\alpha-0.5}$ and the disclosure rule is transparent, $d = 1$

(3) Otherwise, full screening is optimal

Proof. The platform now commits to a probability d of choosing a transparent rule and a probability $1-d$ of selecting the opaque disclosure rule. When the disclosure rule is transparent, the platform screens with probability γ_{tr}; when the rule is opaque, it screens with probability γ_o.

The sender's payoff from choosing low-precision information is $d\left(\gamma_{tr}\alpha+(1-\gamma_{tr})(0.5)\right)+(1-d)\,0.5$ while payoff from high precision information would be $1-c$. For all $\gamma_{tr} \leq \frac{0.5-c}{d(\alpha-0.5)}$ the sender chooses high precision information. The maximum screening probability under the transparent regime is dependent on d, $\frac{0.5-c}{d(\alpha-0.5)}$. Since $0 \leq \gamma_{tr} \leq 1$, this implies $d \in [\frac{0.5-c}{\alpha-0.5}, 1]$. At the lowest value of d, $c = 1-\alpha$, and the highest value of $d = 1$, $c = 0.5$, so high precision can be incentivized only for $1-\alpha < c \leq 0.5$.

The screening probability under the opaque regime is given by $\gamma_o = 1-\frac{k_\alpha}{k}$, since on observing only $\{x, \kappa\}$ the receiver knows the disclosure rule is opaque i.e. ex-post $d = 0$ and will disregard $x = 1$ if $q\,(\alpha-0.5)-0.25\,(1-q)\,(1-\gamma)k \leq q(0.5-\alpha)$ as before.

The platform chooses d to minimize the loss to the receiver

$$\min_{d\in[\frac{0.5-c}{\alpha-0.5},1]} d\left(1-\gamma_{tr}\left(d\right)\right)\left(1-q\right)k0.25+\left(1-d\right)\left(1-\gamma_{o}\right)\left(1-q\right)k0.25$$

No interior solution exists and because the equation increases in d, the minimum is at $d=\frac{0.5-c}{\alpha-0.5}$ which gives $\gamma_{tr}=1$.

The receiver's minimum screening probability is now given by $\gamma_{r}=1-\frac{8q(1-\alpha)}{(1-q)(1-d)k}<\gamma_{tr}=1$. Comparing the screening probability under the opaque rule, $\gamma_{o}\geq\gamma_{r}$ gives a quadratic equation. Specifically, for all α such that $4\alpha^{2}-6\alpha+2\alpha c+2-c\leq 0$, $\gamma_{o}\geq\gamma_{r}$. The quadratic roots $\alpha\in[\frac{6-2c-\sqrt{(2c-6)^{2}-16(2-c)}}{8},\frac{6-2c+\sqrt{(2c-6)^{2}-16(2-c)}}{8}]$ can be simplified to $[0.5,1-0.5c]$. Since $0.5<\alpha$, we only require $\alpha\leq 1-0.5c$ for partially screening to be feasible.

If $\alpha>1-0.5c$, the transparent disclosure rule can still be implemented for some cost interval. Now $d=1$. The remaining result follows directly from lemma 4 and proposition 1.

■

Notably, when the conditions for the first point are met, randomizing between the two policies is significantly better than choosing either transparent or opaque. The loss from screening now is only incurred $\left(1-d\right)\left(1-\gamma_{o}\right)$ times while still incentivizing the sender to acquire high precision information.

5.2 Continuous Signals

I also extend the model to allow for continuous signals. The main insight goes through with this more general version. Specifically, now the sender can send any message $x\in[0,1]$. The honest sender chooses a precision level $\kappa\in\{\alpha,\beta\}$. Conditional on the state $\theta\in\{0,1\}$ and the choice κ a message is realized according to conditional density $g_{\theta,\kappa}\left(x\right)$, with corresponding CDF, $G_{\theta,\kappa}\left(x\right)$. The fake sender also draws a fake message that appears to be of a certain precision with probability $f\left(x,\kappa\right)$ that is uninformative of the underlying state. The screening probability under the opaque disclosure may now also decrease as the cost increases. However, there still exists a cutoff c^{*} such that below it, the transparent disclosure rule is optimal, while above it, the opaque rule is better. The detailed analysis for this extension is provided in the Appendix.

5.3 Conditioning on the Precision

I assume that the platform cannot observe the precision of the message. If instead the platform could correctly infer the exact precision just as well as the receiver, then the optimal policy can be conditional on the precision level. For $1-\alpha < c < 0.5$, it would choose to fully screen only *high precision* messages while never screening *low precision* messages. The sender then always prefers choosing high-precision information for this cost interval. For costs below $1-\alpha$ only high precision messages would be screened, while for costs above 0.5 only low precision messages would be screened.

6 Discussion

Disutility from Fake News $(k < 0)$**:** This assumption captures the fact that the underlying basis of any decision-making can also impact a decision-maker's payoff. Making decisions based on factually untrue evidence may impose psychological costs or reputational loss. Agents may feel greater regret, guilt, or loss of credibility with others. It may even permanently increase mistrust and impact all future decision-making. k captures all such negative consequences that are more likely to follow when the information, based on which a decision was made, is fake.

Fake Sender: I also assume that the precision (α, β) draw of the fake sender is exogenously specified. The ability to duplicate evidence depends on several factors, including the level of technology, institutional safeguards, and the regulatory environment. As such, there is some inherent randomness to how successfully a duplication of any evidence can be pulled off. Stronger evidence should ideally be more difficult to falsify or fabricate. For example, how well an AI technology can produce an image or a video would vary with the exact prompt it is given and with the underlying model quality. Or how easy it might be to falsify documents would depend on the complexity of the original document, such as the security features or formatting standards, and how frequently they change.

7 Conclusion

In this paper, I develop a simple and tractable model of costly information acquisition in the presence of fake news. While social media platforms have played a huge role in ensuring greater information equality, not all information is equal. Some of it is just pure noise. This distinction between the underlying authenticity of information has so far been understudied.

In this paper, I show that the presence of such information can have consequences for the sender's information acquisition strategy and the receiver's learning. Greater moderation of online content may incentivize content creators to be more complacent about the information they share. The optimal moderation policy will often require some fake news to be allowed on the platform to incentivize the sender to distinguish their content. The increase in welfare gains due to the improved information precision is often enough to offset the loss from acting on fake news. This is even when the harm from fake news becomes very high. These insights extend to a more general model with a continuous message space and binary precision.

The analysis attempts to highlight the importance of the underlying quality of information and raises further questions. Even otherwise, *true* information can often be misinterpreted or selectively disclosed to support a narrative. A binary classification of true and fake information may then not be sufficient to capture such subtleties. I also assume the fake sender to be a non-strategic player. Endogenizing the fake sender's technology and imposing costs for fabricating evidence can provide new insights. Another potential extension is introducing multiple sellers and allowing for the aggregation of information. Honest senders may have an incentive to free ride if they are convinced that other information providers are also legitimate. However, uncertainty about the other sender's type may help undermine some of the free-riding effect.

References

Acemoglu, D., Ozdaglar, A., & ParandehGheibi, A. (2010). Spread of (mis) information in social networks. *Games and Economic Behavior*, *70*(2), 194–227.

Acemoglu, D., Ozdaglar, A., & Siderius, J. (2024). A model of online misinformation. *Review of Economic Studies*, *91*(6), 3117–3150.

Allcott, H., & Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of economic perspectives*, *31*(2), 211–236.

Argenziano, R., Severinov, S., & Squintani, F. (2016). Strategic information acquisition and transmission. *American Economic Journal: Microeconomics*, *8*(3), 119–155.

Aridor, G., Jiménez-Durán, R., Levy, R., & Song, L. (2024). The economics of social media. *Journal of Economic Literature*, *62*(4), 1422–1474.

Azzimonti, M., & Fernandes, M. (2023). Social media networks, fake news, and polarization. *European journal of political economy*, *76*, 102256.

Barrett, P., Hendrix, J., & Sims, G. (2021). *How tech platforms fuel u.s. political polarization and what government can do about it*. Brookings.

Bik, E. (2022). Science has a nasty photoshopping problem. *The New York Times*. Retrieved October 29, 2022, from https://www.nytimes.com/interactive/2022/10/29/opinion/science-fraud-image-manipulation-photoshop.html

Bobrowsky, M., & Vipers, G. (2025). Meta ends fact-checking on facebook, instagram in free-speech pitch. *WSJ*. Retrieved January 7, 2025, from https://www.wsj.com/tech/meta-ends-fact-checking-on-facebook-instagram-in-free-speech-pitch-8e46ad52

Che, Y.-K., & Kartik, N. (2009). Opinions as incentives. *Journal of Political Economy*, *117*(5), 815–860.

Do Nascimento, I. J. B., Pizarro, A. B., Almeida, J. M., Azzopardi-Muscat, N., Gonçalves, M. A., Björklund, M., & Novillo-Ortiz, D. (2022). Infodemics and health misinformation: a systematic review of reviews. *Bulletin of the World Health Organization*, *100*(9), 544.

Ederer, F., & Min, W. (2022). *Bayesian persuasion with lie detection* (tech. rep.). National Bureau of Economic Research.

Glazer, J., Herrera, H., & Perry, M. (2021). Fake reviews. *The Economic Journal*, *131*(636), 1772–1787.

Graves, L., & Cherubini, F. (2016). The rise of fact-checking sites in europe.

Kartik, N. (2009). Strategic communication with lying costs. *The Review of Economic Studies*, *76*(4), 1359–1395.

Kartik, N., Lee, F. X., & Suen, W. (2017). Investment in concealable information by biased experts. *The RAND Journal of Economics*, *48*(1), 24–43.

Kartik, N., Ottaviani, M., & Squintani, F. (2007). Credulity, lies, and costly talk. *Journal of Economic theory*, *134*(1), 93–116.

Lipnowski, E., Ravid, D., & Shishkin, D. (2022). Persuasion via weak institutions. *Journal of Political Economy*, *130*(10), 2705–2730.

Mont'Alverne, C., Badrinathan, S., Arguedas, A. R., Nielsen, R. K., Toff, B., & Fletcher, R. (2022, September 22). *The trust gap: how and why news on digital platforms is viewed more sceptically versus news in general*. Retrieved September 22, 2022, from https://reutersinstitute.politics.ox.ac.uk/trust-gap-how-and-why-news-digital-platforms-viewed-more-sceptically-versus-news-general

Newman, N., Fletcher, R., Robertson, C. T., Arguedas, A. R., & Nielsen, R. K. (2024). *Digital news report 2024*. Reuters Institute.

Stiglitz, J. E., & Kosenko, A. (2024). *The economics of information in a world of disinformation: a survey part 2: direct communication* (tech. rep.). National Bureau of Economic Research.

Thomas, E., & Sardarizadeh, S. (2024). How a deleted linkedin post was weaponised and seen by millions before the southport riot. *BBC*. Retrieved October 25, 2024, from https://www.bbc.co.uk/news/articles/c99v90813j5o

Appendix

A General Result

In this section, I state the general results for any $p \in [0, 0.5]$ such that the prior preferred action is $a = 0$. The minimum k required such that the receiver ignores low-precision information when $\gamma = 0$ is given by

Assumption 2 (Non-triviality)**:**

$$k_\alpha = \frac{2q\left(\alpha - (1-p)\right) - (1-q)\left(1-2p\right)}{(1-q)\,f_\alpha} \leq k$$

Theorem 2: *The optimal screening policy is,*

(1) If $\alpha \leq \alpha^*$ and $\beta - \alpha < c \leq \beta - (1-p)$, then partial screening, $\gamma_o < 1$, is optimal and the sender chooses high-precision information. This screening probability is given by $\gamma_o = 1 - \frac{k_\alpha}{k}$.

(2) where $\alpha^* = \frac{2qk\beta + 2q(1-p)k + (1-2p)(1-q+6q(1-p))}{2q(3(1-2p)+2k)}$ and $1 - p < \alpha^*$

(3) Otherwise, full screening is better. The sender chooses high precision when costs are low ($c \leq \beta - \alpha$) and low precision when costs are high ($c > \beta - \alpha$)

Proof. The proof is shown through a series of lemmas.

The first lemma analyzes the benchmark case of full screening and establishes the cost values where full screening is always optimal.

Lemma 5: *Full screening is optimal for $c \leq \beta - \alpha$ and the sender chooses high precision information*

Proof. When the sender is honest and since $1 - p < \alpha < \beta$, the receiver finds it optimal to take action 1 when $x = 1$ and action 0 when $x = 0$ at any precision level. When the sender is fake, she acts according to her prior information and takes action 0. The honest sender's payoff from choosing low precision information is α, whereas his payoff from choosing high precision information is $\beta - c$. For $c \leq \beta - \alpha$, the sender chooses high precision information. ■

The next lemma characterizes the minimum screening required by the receiver to benefit from high-precision information despite occasionally acting on fake information.

Lemma 6:

(1) The receiver is better off with high precision information for all $\max\{\gamma_r, 0\} \leq \gamma \leq 1$ *than full screening and low precision information*

(2) Where, $\gamma_r = 1 - \frac{2q(\beta-\alpha)}{(1-q)(3(1-2p)+k)} < 1$

Proof. The receiver's payoff with full screening and low precision information is $q(2\alpha - 1) + (1-q)(1-2p)$. Her payoff when the sender chooses high-precision information and the platform screens messages with some screening probability γ is given by $q(2\beta - 1) + (1-q)((\gamma)(1-2p) + (1-\gamma)(1-2p)0.25 - (1-\gamma)k0.25)$. Comparing the two payoffs provides the required expression. Moreover, for any finite $k < \infty$, $\gamma_r < 1$ ∎

The following lemma considers the sender's precision choice and characterizes the maximum screening probability that would incentivize the sender to choose high-precision information.

Lemma 7: *The sender's precision choice for different values of* γ *and* c *are as follows:*

(1) If $k_\alpha \leq k < k_\beta$ *and* $\beta - \alpha < c \leq \beta - (1-p)$*, the sender chooses high precision for all* $\gamma \leq \gamma_o < 1$

(2) If $k_\alpha < k_\beta \leq k$ *and* $\beta - \alpha < c \leq \beta - (1-p)$*, the sender chooses high precision for all* $\gamma'_o \leq \gamma_o < 1$*. Otherwise, high precision cannot be incentivized for* $\beta - \alpha < c \leq 1$

(3) Where, $\gamma'_o = 1 - \frac{k_\beta}{k} < 1, \gamma_o = 1 - \frac{k_\alpha}{k} < 1$

(4) For $c > \beta - (1-p)$ *the sender chooses low precision for any* $\gamma \in [0,1]$

Proof. The proof is the same as in the main text and is hence omitted here. ∎

Comparing γ_o and γ_r provides the required formula. Let $t = 1 - 2p$, then

$$\implies 1 - \frac{2q(\alpha - (1-p)) - (1-q)(1-2p)}{(0.25)(1-q)k} \geq 1 - \frac{2q(\beta-\alpha)}{(1-q)((1-2p) + k0.25 - (1-2p)0.25)}$$

$$2q(\alpha - (1-p)) - (1-qt) \leq \frac{2q(\beta-\alpha)k}{3(1-2p)+k}$$

$$\implies \alpha \leq \frac{2q\beta k + (1-q)t + 2q(1-p)(3t+k)}{2q(3t+2k)}$$

To see that $\alpha^* > 1-p$ note

$$1-p \leq \frac{2q\beta k + (1-q)\,t + 2q\,(1-p)\,(3t+k)}{2q\,(3t+2k)}$$

$$4k\,(1-p)\,q + 6tq\,(1-p) \leq 2q\beta k + (1-q)\,t + 6tq\,(1-p) + 2q\,(1-p)\,k \implies 2qk\,(1-p)\,q \leq 2q\beta k + (1-q)\,t$$

which always holds since $\beta > \alpha > 1-p$

This establishes the theorem. The strategies for the platform, sender, and receiver constitute a Perfect Bayesian Equilibrium, as each player best responds given their beliefs, and the receiver's beliefs are consistent with Bayes' rule on the equilibrium path.

■

A.1 Mixed Strategies

The focus on pure strategies in the main model is without loss. For instance, the platform could instead commit to a distribution over $[0,1]$. Since payoffs for the sender and the receiver would depend only on the expected screening, a commitment to a distribution matters only through its expectation; therefore, I restrict attention to pure γ. The platform optimally either fully screens ($\gamma = 1$) or selects a partial screening. If the platform fully screens, neither the sender nor the receiver is indifferent except for specific parameter values of measure zero. When the optimal policy requires partial screening, the screening probability is chosen so that a message $(1, \alpha)$ is ignored. In contrast, $(1, \beta)$ is followed. Taking action 0 on observing a message 0 of either precision is always better for the receiver. As such, the receiver is never indifferent except for parameter values of measure zero. Given the platform's screening choice and the $c > 0$ cost of high-precision information, the sender is also never indifferent between the two precision levels, except for knife-edge cases of zero measure. Indifference for the receiver is broken in favor of action 0 and indifference for the sender is broken in favor of high-precision information.

B Continuous Signal

In this section, I amend the model to accommodate continuous signals $x \in [0,1]$. The core insight remains robust to the more general environment. I first state the optimal policy when the disclosure rule is transparent and then move to the opaque disclosure rule. For simplicity, I state the results for $p = 0.5$ where the prior preferred action is 0.

The honest sender first chooses a precision level $\kappa \in \{\alpha, \beta\}$. Conditional on the state $\theta \in \{0, 1\}$ and the choice κ a message is realized according to conditional density $g_{\theta,\kappa}(x)$, with corresponding CDF, $G_{\theta,\kappa}(x)$. I further impose some assumptions on the honest sender's signal structure:

Assumption 3 (Signal Structure)**:**

A3.1 $G_{1,\beta}(x)$ *FOSD* $G_{1,\alpha}(x)$, $\forall x \in [0,1]$, *with strict inequality for some* $x \in [1/2, 1]$

A3.2 $G_{0,\alpha}(x)$ *FOSD* $G_{0,\beta}(x)$, $\forall x \in [0,1]$ *with strict inequality for some* $x \in [1/2, 1]$

A3.3 *MLRP:* $\frac{g_{1,\kappa}(x)}{g_{0,\kappa}(x)}$ *is increasing in* x *for a given* κ

A3.4 *Symmetry:* $g_{1,\kappa}(x) = g_{0,\kappa}(1-x) \iff G_{1,\kappa}(x) = 1 - G_{0,\kappa}(1-x)$

The two FOSD assumptions imply that choosing high precision, β, makes observing higher signals more likely than low precision when the state is 1 and similarly for state 0. The third assumption ensures that higher signals are progressively more indicative of state 1 and that the receiver's strategy is a simple cutoff rule. The last assumption ensures symmetry of the two conditional distributions around $x = \frac{1}{2}$. However, Assumptions A3.3 and A3.4 are for notational convenience and can be relaxed.

The definition of precision can be directly extended from the discrete case. Due to MLRP and symmetry, for all signals $x \in [0.5, 1]$ the receiver prefers to take action 1 and 0 otherwise. Then precision here can be defined as $\int_{0.5}^{1} P(x|\theta = 1) = \int_{0}^{0.5} P(x|\theta = 0) = \kappa$.

The fake sender, on the other hand, draws a message (x, κ) according to a joint density $f(x, \kappa)$ on $x \in [0, 1] \times \{\alpha, \beta\}$ with full support. The cumulative probability of observing a message x up to t with precision κ is $F_\kappa(t) = \int_0^t f(u, \kappa)\,du$ and and $\sum_{\kappa \in \{\alpha,\beta\}} \int_0^1 f(x, \kappa)\,dx = 1$.

As before, some simplifying assumptions are made for cleaner exposition. Specifically, $f(x, \kappa) = f_\kappa = 0.5$ for all $x \in [0, 1]$. The payoffs also remain the same as in the binary message space setup, the receiver suffers a disutility $-k$ if she takes an action different from her prior action *and* if the information turns out to be fake. If she takes the prior action 0, her payoff is 0 if the information is fake. If the information is true, then the correct action gives her 1 while the incorrect action gives her -1. The sender's payoff is also the same as in the binary message case; he gets $1 - c$ if the action is correct and $-c$ if the action is incorrect.

On observing any message $x \in [0, 0.5]$ of either precision level, the receiver would prefer to take action 0; her preferred action on observing a signal $x \in (0.5, 1]$ would, however, depend on the value of k.

To make the problem interesting, I again assume k is large enough so that at least some low-precision messages are always ignored when there is no screening, $\gamma = 0$.

Assumption 4 (Non-triviality)**:** *There exists* $x \in [1/2, 1]$ *such that,*

$$k_\alpha = \frac{q\left(g_{1,\alpha}(x) - g_{0,\alpha}(x)\right)}{(1-q)\, f_\alpha} \leq k$$

The first lemma analyzes the sender's and receiver's choices under full screening to identify the parameters at which the trade-off disappears.

Lemma 8: *For all* $c \leq c_o = G_{0,\beta}(0.5) - G_{0,\alpha}(0.5)$ *and* $c > \bar{c} = G_{0,\beta}(0.5) - 0.5$*, full screening is optimal*

Proof. Under full screening, $\gamma = 1$, the receiver takes action 1 when she observes a signal above 0.5 and action 0 when she observes a signal below 0.5 if the sender is honest. If the sender is fake, she takes her prior preferred action 0. The sender chooses high precision for all $c \leq G_{1,\alpha}(0.5) - G_{1,\beta}(0.5)$ and low precision otherwise. Hence, for $c \leq c_o = G_{1,\alpha}(0.5) - G_{1,\beta}(0.5)$, the trade-off disappears and full screening is optimal.

On the other extreme, the best case scenario would be if every low precision message $x \in [0.5, 1]$ is ignored while every high precision message is followed. The sender's payoff from low precision payoff then is 0.5, while high precision information payoff is $G_{0,\beta}(0.5) - c$. For all $c \geq G_{0,\beta}(0.5) - 0.5$ high precision cannot be incentivized. ■

B.1 Transparent Disclosure Rule

The next proposition characterizes the minimum screening probability that makes the receiver better off with high precision and some fake news compared to full screening.

When the disclosure rule is transparent and the receiver is informed of the screening decision, then when a message has been screened, i.e., $\gamma = 1$, all messages are admitted by the receiver with either precision. However, when the message is not screened, the receiver only admits messages that are less likely to be sent by the fake sender. For either precision, these messages are denoted by $[\underline{x}_\kappa, \bar{x}_\kappa]$. For notational convenience, I again assume $\bar{x}_\kappa = 1$, however, this can be relaxed to any generic $x > \underline{x}_\kappa$.

Lemma 9: *Let* $\Delta_\kappa(x) = G_{0,\kappa}(x) - G_{1,\kappa}(x)$*. Let the information regime be transparent,* *(1) The receiver is better off with high precision information for all $\gamma_{r,tr} \leq \gamma$ than with full screening and low precision information*

(2) Furthermore, $\gamma_{r,tr} = 1 - \frac{2qc_o}{q\left(\Delta_\beta(0.5) - \Delta_\beta\left(\underline{x}_\beta\right)\right) + (1-q)k\left(0.5\left(1-\underline{x}_\beta\right)\right)} < 1$

Proof. The receiver's payoff from full screening and low precision information is given by

$$q\left(G_{0,\alpha}(0.5) - G_{1,\alpha}(0.5)\right) = q\left(\Delta_\alpha(0.5)\right) \tag{1}$$

The second inequality is due to symmetry (assumption 3). The receiver's payoff from high precision information and partial screening is given by

$$\begin{aligned} & q\gamma\left(G_{0,\beta}(0.5) - G_{1,\beta}(0.5)\right) + q(1-\gamma)\left(G_{0,\beta}\left(\underline{x}_\beta\right) - G_{1,\beta}\left(\underline{x}_\beta\right)\right) - (1-q)(1-\gamma)k\left(0.5\left(1-\underline{x}_\beta\right)\right) \\ = & q\gamma\left(\Delta_\beta(0.5)\right) + q(1-\gamma)\left(\Delta_\beta\left(\underline{x}_\beta\right)\right) - (1-q)(1-\gamma)2k\left(0.5\left(1-\underline{x}_\beta\right)\right) \end{aligned} \tag{2}$$

Comparing equation (1)and equation (2) provides the required expression.

■

For the sender, the maximum screening probability that incentivizes him to switch to high precision would also depend on the cutoff x_β. When a message is screened, any message of either precision is obtained, whereas when a message is not screened, only some high-precision messages are screened. The threat needs to be high enough for the sender to switch. If all high-precision messages are ignored by the receiver, then the transparent disclosure rule cannot incentivize high-precision information.

Lemma 10: *Let the information regime be transparent,*

(1) When $\underline{x}_\beta < 1$ and $c_o < c \leq \bar{c}$ the sender chooses high precision for all $\gamma \leq \gamma_{tr}$

(2) Furthermore, $\gamma_{tr} = \frac{(0.5)\left(\Delta_\beta\left(\underline{x}_\beta\right) - \Delta_\alpha\left(\underline{x}_\alpha\right)\right) - c}{(0.5)\left(\Delta_\beta\left(\underline{x}_\beta\right) - \Delta_\alpha\left(\underline{x}_\alpha\right)\right) - c_o} < 1$ *for all $c > c_o$*

(3) Otherwise, high precision information cannot be incentivized

Proof. The proof follows the same logic as in the discrete case. By choosing low precision, the sender's payoff is given by

$$\gamma\left(G_{0,\alpha}(0.5)\right) + (1-\gamma)\left(0.5\left(G_{0,\alpha}\left(\underline{x}_\alpha\right)\right) + (0.5)\left(1 - G_{1,\alpha}\left(\underline{x}_\alpha\right)\right)\right)$$

The payoff from choosing high precision information instead is

$$\gamma\left(G_{0,\beta}\left(0.5\right)\right)+\left(1-\gamma\right)\left(0.5\left(G_{0,\beta}\left(\underline{x}_{\beta}\right)\right)+\left(0.5\right)\left(1-G_{1,\beta}\left(\underline{x}_{\beta}\right)\right)\right)-c$$

Comparing the two expressions provides the required expression. ∎

We can now state the optimal screening probability by comparing the two expressions. Let

Proposition 3: *Let* $c^{*}=c_{0}+\frac{2qc_{0}\left(0.5\left(\Delta_{\beta}\left(\underline{x}_{\beta}\right)-\Delta_{\alpha}\left(\underline{x}_{\alpha}\right)\right)\right)-c_{0}\right)}{q\left(\Delta_{\beta}(0.5)-\Delta_{\beta}\left(\underline{x}_{\beta}\right)\right)+(1-q)k\left(0.5\left(1-\underline{x}_{\beta}\right)\right)}$. *The optimal screening policy is as follows:*

(1) For all $c_{o}<c\leq c^{*}$, *partial screening,* γ_{tr}, *is optimal and the sender chooses high precision information*

(2) Otherwise, full screening is optimal, and the sender chooses low-precision information

Proof. This follows directly by comparing γ_{tr} and $\gamma_{r,tr}$, and hence is omitted here. ∎

B.2 Opaque Disclosure Rule

I first discuss the maximum screening probability that would incentivize the sender to choose high-precision information and then the receiver's minimum screening threshold.

Any choice of screening probability γ will determine the corresponding signal cutoffs for a given precision level. For a fixed (x,κ) the receiver's net payoff from choosing $a=1$ is

$$U_{r}\left(x,\kappa;\gamma\right)=q\left(g_{1,\kappa}\left(x\right)-g_{0,\kappa}\left(x\right)\right)-\left(1-q\right)\left(1-\gamma_{o}\left(c\right)\right)kf_{k}\leq 0$$

$U_{r}\left(x,\kappa;\gamma\right)$ is continuous and strictly increasing in γ and $U_{r}\left(x,\kappa;1\right)\geq 0$ for every $x\in(1/2,1]$ by assumption A3.3. If $U_{r}\left(x,\kappa;\gamma\right)\geq 0$ for some (x,κ), then a signal cutoff is present at $\gamma=0$. Otherwise, if $U_{r}\left(1,\kappa;0\right)<0$ then by the intermediate value theorem there exists a unique $\gamma\in(0,1)$ at which $U_{r}\left(1,\kappa;\gamma\right)=0$. Define $\tilde{\gamma}$ to be the smallest screening probability at which for at least one precision level, $U_{r}\left(1,\kappa;\gamma\right)=0$:

$$\tilde{\gamma}=\inf\{\gamma\in(0,1):U_{r}\left(1,\kappa;\gamma\right)=0,\kappa\in\{\alpha,\beta\}\}$$

For each $\gamma\in[\tilde{\gamma},1]$, let $x_{\kappa}\left(\gamma\right)$ denote the signal cutoff, such that for all $\left[x_{\kappa}\left(\gamma\right),1\right]$ the receiver takes action 1 and 0 otherwise. Due to assumption A3.3, as $\gamma\rightarrow 1$, both cutoffs decrease continuously and approach 0.5.

For any $\gamma \in [\tilde{\gamma}, 1]$, the sender's marginal benefit from choosing high precision rather than low precision information is

$$MB(\gamma) = 0.5\left(G_{0,\beta}(0.5) - G_{0,\alpha}(0.5) + G_{1,\alpha}(x_\alpha(\gamma)) - G_{1,\beta}(x_\beta(\gamma))\right)$$

Let

$$c^\dagger = \max_{\tilde{\gamma} \leq \gamma \leq 1} MB(\gamma)$$

Since $x_\alpha(\gamma)$ and $x_\beta(\gamma)$ decrease continuously as γ increases and due to the continuity of the CDFs, $MB(\gamma)$ is continuous in $\gamma \in [0, 1]$. Due to the extreme value theorem, a maximum $c^\dagger$ is attained, and the set $\{\gamma : MB(\gamma) \geq c\}$ is nonempty and closed whenever $c \leq c^\dagger$.

Then for every $c \leq c^\dagger$, there is at least one γ such that the sender strictly prefers acquiring high precision information rather than low precision. Since $MB(\gamma)$ is continuous, and as $\gamma \to 1$, it decreases to c_o, i.e., it takes every value between $c^\dagger$ and c_o. So for any given $c_o < c \leq c^\dagger$, the corresponding maximum γ occurs exactly where $MB(\gamma) = c$. For $c^\dagger < c$, high precision cannot be induced for any screening probability. The above discussion is formalized as follows,

Definition 1: *When $c_o < c^\dagger$, the following definitions hold*

(1) Let $\gamma_o(c)$ be the maximum screening probability that incentivizes the sender to choose high precision information. Specifically,

$$\gamma_o(c) = \max\{\gamma \in [0, \tilde{\gamma}] : MB(\gamma) = c\}$$

i.e. $\gamma_o(c)$ is the rightmost feasible γ

(2) $\gamma_o(c)$ satisfies the fixed point relation

$$\gamma_o(c) = 1 - \left(\frac{q}{1-q}\right)\left(\frac{(g_{1,\beta}(x_\alpha^*) - g_{0,\beta}(x_\alpha^*))}{kf(x_\alpha^*, \alpha)}\right)$$

where $x_\alpha^* = x_\alpha(\gamma_o(c))$

The next lemma states the behaviour of $\gamma_o(c)$.

Lemma 11: $\gamma_o(c)$ *is strictly decreasing and left-continuous* $\forall\, c \in (c_o, c^{\dagger}]$

Proof. If for atleast some values of $\gamma \in [\tilde{\gamma}, 1]$, $MB(\gamma)$ lies above c_o then as γ approaches 1, $MB(\gamma)$ also approaches c_o. Specifically, this ensures that for any $c^{\dagger} \geq c_1 > c_2 \geq c_o$, $\gamma_o(c_2) > \gamma_o(c_1)$. By definition, $\gamma_o(c)$ is the rightmost point with $MB = c_1$. Since MB is continuous, $\lim_{\delta \downarrow 0} MB(\gamma_o(c_1) + \delta) = c_1$. Therefore, for any $c_2 < c_1$, there exists $\varepsilon > 0$ such that $MB(\gamma_o(c_1) + \varepsilon) > c_2$. Hence $\gamma_o(c_2) \geq \gamma_o(c_1) + \varepsilon > \gamma_o(c_1)$. So, $\gamma_o(c_2) > \gamma_o(c_1)$.

I next prove left continuity and provide a counterexample to show where right continuity might fail.

Fix $c \in (c_o, c^{\dagger}$ and let c_n be any sequence such that $c_n < c$ and $\lim_{n\to\infty} c_n = c$. Let $\gamma_o(c) = \gamma_c$ and $\gamma_o(c_n) = \gamma_n$. Since $c_n < c$ and because $\gamma_o(c)$ is decreasing in c, $\gamma_n \geq \gamma_o(c)$ for all n. Since γ_n is bounded above by 1 and below by γ_c, it is a bounded decreasing sequence that converges to some $\gamma^* \geq \gamma_c$. By definition, $MB(\gamma_n) \leq c_n$. Applying limits, $\lim_{n\to\infty} MB(\gamma_n) \geq \lim_{n\to\infty} c_n \implies MB(\gamma^*) \geq c$. So γ^* is feasible at c but by maximality of γ_c, it must be that $\gamma^* \leq \gamma_c$. Combining the two inequalities gives $\gamma^* = \gamma_c$. As this holds for any sequence $c_n \to c$ where $c_n < c$, $\gamma_o(c_n) \to c$ i.e. $\gamma_o(c)$ is left continuous at c.

To show right continuity may fail, consider the following counterexample. Assume $MB(\gamma)$ curve strictly decreases till γ_a, then has a plateau between $[\gamma_a, \gamma_b]$ and then strictly decreases again till $\gamma = 1$ and $MB(1) = c_o$. Let the value of $MB(\gamma) = c^*$, for $\gamma \in [\gamma_a, \gamma_b]$. Also, let c_n be any sequence, $c_n > c^*$ and $\lim_{n\to\infty} c_n = c^*$. By definition $\gamma_o(c^*) = \gamma_b > \gamma_a$. Since the MB curve is strictly decreasing between 0 and γ_a by assumption for each c_n there is a corresponding γ_n such that $MB\gamma(c_n) = c_n$. As $c_n \to c^*$, the corresponding $\gamma_n \to \gamma_a$, since MB is continuous. Therefor $\lim_{n\to\infty} \gamma_n = \gamma_a \neq \gamma_o(c^*) = \gamma_b$. This disproves right continuity. ∎

Lemma 12: *Let the information regime be opaque,*

(1) If $c_o < G_{1,\alpha}(x_\alpha(\gamma)) - G_{1,\beta}(x_\beta(\gamma))$ *for some* $\gamma \in [0,1]$, *then* $c_o < c^{\dagger}$. *For all such,* $c_o < c \leq c^{\dagger}$, *the sender chooses high precision information when the screening probability is* $\gamma_o(c)$ *and chooses low precision information otherwise*

(2) If $c_o > G_{1,\alpha}(x_\alpha(\gamma)) - G_{1,\beta}(x_\beta(\gamma))$ *for all* $\gamma \in [0,1]$, *high precision information cannot be incentivized*

Proof. For any $\gamma_o(c)$, the sender's payoff from low precision information is $0.5(G_{0,\alpha}(0.5)) + 0.5(1 - G_{1,\alpha}(x^*_\alpha))$ while the payoff from high precision information is $0.5(G_{0,\beta}(0.5)) + 0.5\left(1 - G_{1,\beta}\left(x^*_\beta\right)\right) - c$. If $G_{1,\alpha}(x_\alpha(\gamma)) - G_{1,\beta}(x_\beta(\gamma)) < c_o$ for all $\gamma \in [0,1]$, then $\max MB(\gamma) =$

$c^{\dagger} < c_o$, and high precision cannot be incentivized. Otherwise, for some $\gamma \in [\tilde{\gamma}, 1]$, $c^{\dagger} > c_0$ and high precision information can be implemented by the corresponding $\gamma_o(c)$. ■

The next lemma characterizes the receiver's minimum screening probability that would make her better off than full screening

Lemma 13:

(1) The receiver is better off with high precision information for all $\max\{\gamma_r, 0\} \leq \gamma \leq 1$ *than full screening and low precision information*

(2) Where, $\gamma_r = 1 - \frac{q\left(1-G_{1,\beta}\left(x_{\beta}^*\right)+G_{0,\beta}(0.5)-2G_{0,\alpha}(0.5)\right)}{(1-q)k\left(1-x_{\beta}^*\right)} < 1$

Proof. The receiver's payoff under full screening and low precision information is given by $q\left(2G_{0,\alpha}(0.5)-1\right)$ and her payoff from high precision information that is screened with probability $\gamma_o(c)$ is given by $q\left(G_{0,\beta}\left(x_{\kappa}^*\right)+G_{0,\beta}(0.5)-2G_{0,\alpha}(0.5)\right)-\left(1-\gamma_o(c)\right)(1-q)\,k\left(F_{\beta}(1)-F_{\beta}\left(x_{\kappa}^*\right)\right)$. Comparing the two expressions provides the given expression. ■

Proposition 4: *If* $c_o < c \leq c^{\dagger}$ *and* $\frac{0.5\left(G_{0,\beta}(0.5)+1-G_{1,\beta}\left(x_{\beta}^*\right)\right)-2G_{0,\alpha}(0.5)}{1-x_{\beta}^*} \geq \left(g_{1,\beta}\left(x_{\alpha}^*\right)-g_{0,\beta}\left(x_{\alpha}^*\right)\right)$ *then partial screening is optimal and the screening probability is given by* $\gamma_o(c)$*. Otherwise, full screening is optimal.*

Proof. For any given $c_o < c$, we calculate the corresponding $\gamma_o(c)$, x_{α}^* and x_{β}^*. Comparing $\gamma_o(c)$ and γ_r, the latter is greater than the former if $\frac{0.5\left(G_{0,\beta}(0.5)+1-G_{1,\beta}\left(x_{\beta}^*\right)\right)-2G_{0,\alpha}(0.5)}{1-x_{\beta}^*} \geq \left(g_{1,\beta}\left(x_{\alpha}^*\right)-g_{0,\beta}\left(x_{\alpha}^*\right)\right)$, which establishes the condition. ■