

**Economics Department
University of Southampton
Southampton SO17 1BJ, UK**

**Discussion Papers in
Economics and Econometrics**

Babel: Persuasion with Anonymous Senders*

Maksymilian Kwiek, University of Southampton

No. 2502

This paper is available on our website

<http://www.southampton.ac.uk/socsci/economics/research/papers>

ISSN 0966-4246

Babel: Persuasion with Anonymous Senders^{*}

Maksymilian Kwiek[†]

1 Sep 2025

Abstract

In communication games with multiple senders, I refer to *babel* as a setting in which the receiver cannot attribute messages to specific senders. This captures features of election campaigns and other public debates, where some senders may benefit from sabotaging the shared message. A binary model with threshold utility functions is both rich and relatively tractable. While best-response dynamics is often cyclical, pure-strategy equilibria exist when senders are small and there is a common order of signals. Even when senders strongly disagree over outcomes, partial communication compromises can emerge. Conversely, when preferences are nearly aligned in outcomes, communication behavior can exhibit runaway divergence: either one sender contributes only noise, or the other one attempts to contribute full information.

^{*}Draft

[†]Department of Economics, University of Southampton. E-mail: m.kwiek@soton.ac.uk

(...) You just have to flood a country's public square with enough raw sewage. You just have to raise enough questions, spread enough dirt, plant enough conspiracy theorizing that citizens no longer know what to believe.

Barak Obama¹

The point of modern propaganda isn't only to misinform or push an agenda. It is to exhaust your critical thinking, to annihilate truth.

Garry Kasparov²

1 Introduction

Visitors to Tate Modern in London may come across a room containing an installation by Brazilian artist Cildo Meireles: a large circular tower made of approximately 800 second-hand analogue radios. Each radio is tuned to a different station and played at low volume, producing a mix of overlapping voices, music, and static. Occasionally, a distinct fragment of a song or broadcast becomes audible before fading back into the overall buzz. Titled *Babel 2001*, the piece evokes the difficulties of communication in an environment saturated with information sources and competing signals.³

I study a model of Bayesian persuasion with multiple senders and a single receiver. The key assumption is that the receiver cannot tell who said what, creating confusion about individual voices in public discourse—a situation I refer to as *babel*.

Anonymity is common on modern social media, but public discourse has long been shaped by back-channels, gossip, and anonymous tips. Today, intermediaries such as AI-driven news aggregators or anonymous bots that relay others' messages make the problem worse. The same role can be played by indoctrinated but otherwise ordinary individuals. When sources cannot be identified, any sender can steer the conversation, either increasing or reducing

¹“Disinformation Is a Threat to Our Democracy” at <https://barackobama.medium.com>

²Twitter, Dec 13, 2016

³Tate Modern audio description: <https://www.tate.org.uk/art/artworks/meireles-babel-t14041/audio-description-babel>

its informativeness. This contrasts with much of the literature, where senders can add information but not subtract from it.

Since there is an incentive to undermine the other sender’s messages, flooding public square with this kind of noise—as Barak Obama eloquently put it in the initial quote—is a common feature of equilibrium. One might attribute this behavior to information-processing costs or capacity limits which can be strategically exploited by senders who prefer the public to remain in the dark. However, this paper assumes neither is present. Instead, the opportunity to contaminate discourse stems from the senders’ anonymity.

Electoral example. As a concrete application, I adapt the prosecutor example from the seminal paper by Kamenica and Gentzkow (2011) to an electoral setting with two senders. Imagine a representative voter deciding whether to punish an incumbent politician suspected of corruption. The voter’s belief about the probability of corruption may shift over the course of the electoral campaign. If the belief remains low, the voter reelects the politician. If it rises moderately, the voter removes the politician from office but allows them to remain active in politics. If it becomes very high, the voter not only votes the politician out but also ends their prospects for a future political career.

In this setting, the voter plays the role of the *receiver*, and the electoral outcome depends on how her belief about corruption is shaped by the messages of two political groups, who act as *senders* and attempt to influence perceptions of corruption during the campaign.

The senders have misaligned objectives, and I consider two distinct scenarios. In the first, both senders oppose the politician, but with different levels of intensity. Sender *A* is pragmatic, focused only on defeating the politician in the upcoming election. Sender *B*, driven by a deeper antipathy, aims to raise the belief of corruption high enough to force the politician out of public life entirely.

The second scenario is different: while sender *B* continues to seek permanent removal, sender *A* is now a supporter, hoping to secure the politician’s reelection.

I analyze pure-strategy Nash equilibria in both scenarios, treating the campaign as a simultaneous-move game between two senders who strategically shape their informational policies. Despite the model’s binary structure and simple threshold utility functions, its behavior is rich, yielding clear and interpretable results.

Results. In the first scenario, both senders oppose the politician, but

they differ in how much information they wish to reveal. *Their preferred levels of disclosure follow a clear order.* This leads their behavior to diverge in equilibrium, setting off a kind of runaway process: one sender provides information, the other responds by adding noise, prompting the first to counter with even more informativeness, and so on. With no individual cost to this escalation, the process continues until one of the senders hits a boundary, either trying to reveal everything or flooding the discourse with noise.

The results are different in the second scenario, where the *senders' preferences over information are not mutually ordered.* One sender supports the politician while the other seeks their removal, yet both may still share an interest in informing the voter, even as each tries to shift beliefs in their own direction. In the most symmetric equilibria, neither resorts to flooding the public square with noise, but their efforts partially cancel each other out.

Implications for debate design. Formal models of debate design are rare in the literature, and even describing the tools available to a designer is not straightforward. For instance, requiring a moderator to filter arguments and block obstructionist ones might seem helpful in theory, but in practice it is prone to abuse. By contrast, measures that reduce anonymity can be incorporated into the design more naturally, since they do not appear to handicap either side in advance.

Consider two examples. First, society addresses divisive issues through various mechanisms such as referendums or citizen assemblies. A referendum is typically preceded by public debate in traditional and social media, much of it mediated and unattributed. A citizen assembly, by contrast, involves structured discussion, with expert input and full transparency. While the two mechanisms differ in many ways, the degree of anonymity is among key distinctions. Second, competition among social media platforms, such as Twitter and Bluesky, can be seen as a contest over debate design: some allow anonymous trolling in reply sections, while others promote transparency by providing tools such as robust blocking features.

We can evaluate the debate's performance by using the Blackwell informativeness as a normative benchmark and assess the role of senders' anonymity. It turns out that relying on conventional measures of "balance" in debates may lead counter-intuitive conclusions.

In the first scenario, a sender can sustain a high level of informativeness even when challenged by an obstructionist entrant, provided the former

has some room to counter the latter and the latter is not too prevalent.⁴ To an external observer, such as an analyst, media overseer, or parliamentary commission, this may look like an unbalanced public square: one side provides information while the other floods it with noise. By contrast, the public square in the second scenario might appear more vibrant and symmetric, since each sender, though pulling beliefs in their preferred direction, does provide some information. Yet even minimal participation by a second sender strictly reduces equilibrium informativeness. In the limit, as this scenario converges to a zero-sum game, no information may be revealed at all.

The marketplace of ideas is an alluring concept⁵, but the details of its design matter. For example, solutions aimed at preventing echo chambers, such as by exposing the audience to many senders, may inadvertently encourage anonymity, and paradoxically reduce the quality of debate.

Literature. Several papers modify the basic Bayesian persuasion framework of Kamenica and Gentzkow (2011) by adding multiple senders: Gentzkow and Kamenica (2017a), Gentzkow and Kamenica (2017b), Ravindran and Cui (2022). In those extensions, several senders simultaneously design their signals, but in all of them the signal aggregation technology implies that each sender can only add to informativeness of the signal of other senders and can add any more informative signal. The current paper departs from this assumption in both ways: first, because individual signals are not observable, senders can reduce the intended informativeness of signals of others; second, the information environment is not Blackwell-connected, meaning individual senders cannot unilaterally achieve any outcome that is more informative for a given behavior of other senders. Li and Norman (2018) provide examples in which departing from assumptions of Gentzkow and Kamenica (2017a, 2017b) can lead to less informativeness. Ravindran and Cui (2022) study an environment with multiple senders choosing information independently—

⁴Conversely, when a public square is joined by an informative sender, the quality of debate may not improve initially as the existing noisy senders counter the efforts of that entrant. The new sender will start to have a positive effect when she becomes sufficiently prevalent.

⁵The marketplace of ideas is the belief that open discourse allows truth to prevail through the competition of diverse viewpoints. This alluring metaphor is attributed to Justice Holmes dissenting in the *Abrams v. United States* (1919): “(...) the best test of truth is the power of the thought to get itself accepted in the competition of the market (...)”, <https://tile.loc.gov/storage-services/service/lh/usrep/usrep250/usrep250616/usrep250616.pdf>

like in the current paper—and they focus on zero-sum games. Their result is sharp: under some conditions, full information is provided in equilibrium. By contrast, in the anonymous sender environment of this paper, equilibria in zero-sum limit reveal bounded amount of information and in some cases none.

Mylovanov and Zapechelnyuk (2024) study obfuscation in debates. They consider a sequential Bayesian persuasion game in which the second sender can augment or obfuscate the information structure chosen by the first sender. They treat this as an exogenous feature of the environment and study its effect on equilibrium disclosure. In contrast, the current paper analyzes a simultaneous-move game designed to capture the long-run equilibrium performance in a public debate. More importantly, disclosure and obfuscation are not exogenously imposed but instead emerge endogenously from the model’s parameters, particularly from the preferences of the two senders.

Apart from the most relevant literature mentioned above, there are other papers which study multi-sender environments.⁶

2 Model

I begin by outlining the baseline Bayesian persuasion model with multiple senders, following the approach of Gentzkow and Kamenica (2017a, 2017b).

A *state of nature* is denoted $\omega \in \Omega$. The prior belief is common knowledge, $\mu \in \Delta(\Omega)$, with $\mu(\omega)$ denoting the probability of state ω .

There is one receiver and a finite set N of senders. Senders know the state of nature and have access to the same set of *signal realizations*, S . Each sender $i \in N$ selects a *signal* (an informational policy, information structure, experiment), represented by $\pi_i : \Omega \rightarrow \Delta(S)$, where $\pi_i(s|\omega)$ is the probability that i announces *interim signal realization* $s \in S$, given state ω . An observer who can see realization s updates the prior belief to an *interim posterior* given by

$$x_i^s(\omega) = \frac{\mu(\omega) \pi_i(s|\omega)}{\sum_{\omega'} \mu(\omega') \pi_i(s|\omega')} \quad (1)$$

The receiver is initially uninformed about the state of nature which nevertheless determines her optimal action. She benefits from learning from the senders, but she learns a *public signal realization*, rather than interim

⁶E.g. Atakan et al. (2024), Gilligan and Krehbiel (1989)

signal realizations which remain hidden. I explain how the interim signal realizations determine the public one in a few paragraphs.

After observing the public signal realization r , the receiver updates the prior belief μ to a *public posterior* belief $y^r \in \Delta(\Omega)$ and on this basis chooses an action, thereby determining an outcome that matters to the senders.

The game’s timing is this: The senders first simultaneously commit to their strategies, π_i (later we will represent strategies by means of posteriors rather than these conditional distributions). The receiver observes these strategies. Then, the state of nature is realized and the interim signal realizations are drawn, which in turn determine the public signal realization. The receiver updates her prior belief and subsequently chooses an action.⁷

The receiver is treated as a black box—what matters is that the public posterior belief induced by the signal realization r determines senders’ payoffs. This can be captured by assuming that sender i ’s utility depends directly on the receiver’s belief, represented by the function $u_i : \Delta(\Omega) \rightarrow \mathbb{R}$. Since the receiver’s action is mechanistic, this entire interaction reduces to a simultaneous-move game between the senders.

Most of the analysis will be conducted under the binary assumption: two states of nature $\Omega = \{0, 1\}$, with scalar μ denoting the probability of state 1, two possible signal realizations called low and high, $S = \{L, H\}$, and two senders, labeled A and B , namely $N = \{A, B\}$.

Babel communication. The last part of the model yet to be explained determines how the interim signal realizations generate the public signal realization. Each time interim signal realizations are drawn, nature randomly selects one of them to reveal to the receiver *without disclosing which sender it came from*.

Let α_i denote the exogenous probability that sender i ’s interim signal realization is observed by the receiver, where $\sum_i \alpha_i = 1$. That is, if $s_i \in S$ is an interim signal realization for sender i , then the public signal realization is $r = s_i$ with probability α_i .

⁷The assumption that the sender can commit to a signaling strategy π_i is open to criticism, especially here, where anonymity rules out reputation-based justifications. Here, the choice of signal is understood to be a long-term and infrequent decision. A propagandist might hire journalists, columnists, and influencers, training them with carefully crafted materials. She may indoctrinate them so that they go on to distort the public discourse according to some specious script. In more modern context, senders might deploy bots, fine-tune aggregation algorithms, or use artificial intelligence to shape the information environment without continuous oversight.

This feature of the model is its key innovation, capturing the idea of babel—individual voices remaining indistinguishable except through indirect inference of the overall buzz. This assumption introduces an externality in communication strategies, as sender i 's signal directly interferes with or distorts the informational content of sender j 's signal. In doing so, it departs from the common assumption in other multi-sender models that information provided by sender i remains intact regardless of j 's actions.

To summarize, μ , x_i^s and y^s describe the same thing—the probability belief about the state of nature—but at different histories:

1. Common *prior* probability belief $\mu(\omega)$ is the belief of ω before observing any signal realization.
2. *Interim posterior* probability belief of ω conditional on sender i 's announcement s is given by (1). This hypothetical posterior is unavailable to the receiver because she cannot observe i 's interim signal realization. In addition, we can define $p_i^s = \sum_{\omega} \mu(\omega) \pi_i(s|\omega)$ to be the unconditional probability that sender i sends signal realization s .
3. *Public posterior* probability belief of ω after observing the *public signal* realization s is⁸

$$y^s(\omega) = \frac{\mu(\omega) \sum_i \alpha_i \pi_i(s|\omega)}{\sum_{\omega'} \mu(\omega') \sum_i \alpha_i \pi_i(s|\omega')}. \quad (2)$$

The probability of the receiver seeing this public signal realization is $p^s = \sum_i \alpha_i p_i^s$.

Finally, the Bayesian plausibility constraints state that the expected posterior for every state must be equal to the prior. They must hold for all interim as well as public posteriors:

$$\begin{aligned} \mu(\omega) &= \sum_s p_i^s x_i^s(\omega) \\ \mu(\omega) &= \sum_s p^s y^s(\omega) \end{aligned} \quad (3)$$

Common-order assumption. There are different ways to interpret communication, and the choice matters for the analysis. The fully flexible

⁸Public signal realizations were denoted r previously, but since they come from the same set as senders' signal realizations, we can index them as s .

approach treats signal realizations as messages without any natural meaning, allowing senders complete freedom in selecting the probability distributions $\pi_i(\cdot|\omega)$. Alternatively, signal realizations may have an inherent meaning resembling information disclosure: while senders have considerable flexibility in shaping the distribution of signal realizations, they cannot construct one that reverses its natural meaning.

To formalize this restriction for the binary case $S = \{L, H\}$, let signal realization H be allowed to only weakly increase the posterior:

Assumption 1. (*Common order*) $x_i^L \leq \mu \leq x_i^H$ for all i .

To emphasize, relabeling of signal realizations will not allow us to somehow bypass this assumption. It is not about inequalities going in this particular direction, but about them going in the same direction for all senders. This imposes a *common* order on signals.

Benchmark analysis will be conducted assuming a common order, but the assumption will be examined in some detail in section 6.

Equilibrium concept. The solution concept generalizes the standard pure-strategy Nash equilibrium.

Assumption 2. (*Small sender*) *The solution concept is local pure-strategy Nash equilibrium (LPSNE).*

The qualifier *local* indicates that only small deviations are considered. That is, a strategy profile may be an LPSNE even if there exists a profitable deviation, but this deviation must be of a size bounded away from zero.

At first glance, it may seem unnatural to rule out large deviations in a game involving two large senders, as in the setting described above. However, the preferred interpretation of the babel game is that it involves a public square populated by a large number of small senders, each exerting only a minor influence on the aggregate outcome.⁹

Technically, one can define a K -replica economy, in which N represents the number of *sender types* rather than individual senders, and each sender type consists of K identical individuals, each of which may be selected with equal probability. As K approaches infinity, the influence of any single sender on the aggregate outcome becomes infinitesimal. In this framework, a large deviation by one of the two original senders corresponds to a coordinated

⁹Incidentally, this perspective also supports the realism of the babel assumption, namely that individual voices are indistinguishable.




Figure 1: KG concavification of $u_i(\cdot)$

deviation by an entire sender type in the K -replica game, which is assumed away in a non-cooperative setting. By contrast, if a strategy profile fails to satisfy the conditions of a local Nash equilibrium in the original game, then even a single sender in the K -replica game for arbitrary large K could profitably deviate.

Rather than modeling the K -replica game explicitly, I adopt the simpler approach of analyzing local Nash equilibria in the base game. The implications of relaxing Assumption 2 will be addressed in section 5.

Single-sender benchmark. Here, I briefly recall the binary framework of Kamenica and Gentzkow (2011) with a single-sender, which is the polar case of the multi-sender environment I consider below, when $\alpha_i = 1$.

The step function with threshold at t on Figure 1 depicts utility of the single sender as a function of the probability belief of corruption induced in the voter. If the sender succeeds in rising this belief from the prior μ to a posterior equal to t or higher, the voter removes the politician, leading to utility one for the sender; otherwise, the voter reelects the politician and sender's utility remains zero.

With the single sender,¹⁰ the interim posterior is also the public posterior, and, with two states and signals, it can be denoted as $x_i = (x_i^L, x_i^H)$. For $x_i^H \neq x_i^L$, probability p_i^s is pinned down by Bayesian plausibility constraint (3) which simplifies to

$$\mu = p_i^L x_i^L + (1 - p_i^L) x_i^H. \quad (4)$$

If $x_i^H = \mu = x_i^L$, the probability p_i^s is free.

¹⁰In this section, index i is redundant, but it is not removed to keep notation consistent.

Figure 1 shows a possible configuration of posteriors on the horizontal axis. Each signal realization generates a utility level according to the utility curve. Graphically, the expected utility is the value reached by the dotted chord at μ , because this is what expected posterior must be, by (4).

Clearly, in this example, the sender can do better by lowering both x_i^L and x_i^H , moving to the dashed chord as indicated by arrows, and increasing expected utility.

3 Bayesian inference in a babel

This section begins the formal analysis of the multi-sender model. It introduces relationships derived purely from Bayes' rule, which apply regardless of the senders' preferences. Since senders' payoffs depend only on the public posterior, understanding how sender i can influence it is essential for analyzing incentives in the game between senders.

The following lemma states that public posterior defined in (2) as a function of senders' choices π_i can alternatively be stated as a function of (x_i^s, p_i^s) .

Lemma 1. *The public posterior belief after observing public signal realization s is given by:*

$$y^s(\omega) = \sum_i \lambda_i^s x_i^s(\omega) \quad (5)$$

where

$$\lambda_i^s = \frac{\alpha_i p_i^s}{\sum_j \alpha_j p_j^s}. \quad (6)$$

It is tempting to view the public posterior in (5) as a linear function of interim posteriors. However, the weights λ_i^s , which are the probability that a given signal realization s is sent by sender i ,¹¹ depend not only on the exogenous frequency of senders, α_i , but also on the endogenous probability of messages senders use, p_i^s . This latter component means that λ_i^s may depend on x_i^s through the Bayesian plausibility constraint (3).¹²

Strategies. Instead of π_i , I will take the strategy of sender i to be the pair of interim posteriors and signal realizations $(x_i^s, p_i^s)_{s \in S}$ satisfying Bayesian plausibility (3).

¹¹Note the difference with p_i^s , which is the probability that a given sender i sends s .

¹²Ravindran and Cui (2022) study a similar persuasion model, with the key difference that the independent senders are non-anonymous. Suppose the receiver observes a pair of signal realizations $r = (s_A, s_B)$. In the current notation, the public posterior in their

In a single sender case, we can drop the second component because it is either determined by Bayesian plausibility, or if it is not, it is irrelevant. For instance, in the binary example, there are many values of p_i^L consistent with posteriors being equal to the prior, but this is immaterial, as the expected utility remains $u_i(\mu)$ in this case.

This is no longer true if there is more than one sender. Specifically, even though i 's interim posteriors may be equal to the prior, sender i may choose many different probabilities of signal realizations and in this way affects the public posterior y^s .¹³

I will say that i attempts to provide full information if $x_i^s = 1$ for some s . Sender i will provide no information if $x_i^s = \mu$ for some $p_i^s \in [0, 1]$.

Public effect of individual choice. When sender i changes one of her interim posteriors for a given signal realization and state, the Bayesian plausibility condition (3) must still hold. This requires a corresponding adjustment to another component of her strategy. The next two lemmas examine how the public posterior changes under different assumptions about which element is adjusted. The analysis focuses on the two-signal case, where $s, z \in \{L, H\}$, $s \neq z$.

In the first result, probabilities of signal realizations are adjusted and remaining interim posterior are kept constant:

Lemma 2. *Let the set of signals be $\{L, H\}$. Fix the strategies of all senders except for sender i . Suppose sender i increases her interim posterior $x_i^s \neq \mu$, and adjusts (p_i^s, p_i^z) according to Bayesian plausibility constraint (4), so that $x_i^z \neq \mu$ remains constant. Then public posteriors react as follows:*

1. *The same-signal effect: y^s increases (i.e., $\frac{dy^s}{dx_i^s} > 0$) if and only if*

$$(x_i^z - \mu) \left(\alpha_i (x_i^z - \mu) + \sum_{j \neq i} \alpha_j p_j^s (x_i^z - x_j^s) \right) > 0 \quad (7)$$

model can be derived to be

$$y^r(\omega) = \frac{\frac{1}{\mu(\omega)} x_A^{sA}(\omega) x_B^{sB}(\omega)}{\sum_{\omega'} \frac{1}{\mu(\omega')} x_A^{sA}(\omega') x_B^{sB}(\omega')}.$$

¹³Only if $x_i^H \neq x_i^L$ when p_i^s is determined by the Bayesian plausibility constraint, I may skip mentioning it as part of the strategy.

2. The cross-signal effect: y^z increases (i.e., $\frac{dy^z}{dx_i^s} > 0$) if and only if

$$(x_i^z - \mu)(x_{-i}^z - x_i^z) > 0, \quad (8)$$

where x_{-i}^z is the public posterior after signal z , computed from the interim posteriors of all senders as if sender i was absent:

$$x_{-i}^z = \frac{\sum_{j \neq i} \alpha_j p_j^z x_j^z}{\sum_{j \neq i} \alpha_j p_j^z}.$$

Under the common-order Assumption 1, the same-signal effect is unambiguously positive.¹⁴

The cross-signal effect is less straightforward, as it depends on the interim posterior of the other sender. In essence, the effect is *positive* if and only if the opposite-signal interim posterior of the other sender is *more extreme*, i.e. further from the prior. For example, with two senders, the condition (8) simplifies to $(x_i^z - \mu)(x_j^z - x_i^z) > 0$. Increasing sender i 's low-signal interim posterior (i.e. $s = L, z = H$), thus moving it closer to the prior, makes her high signal realization less likely, by Bayesian plausibility constraint. As a result, sender j 's high-signal interim posterior gains more weight in shaping public belief. If it is higher than sender i 's—in other words more extreme—it pulls the public posterior upward, and the effect is positive. If it is lower then it pulls the public posterior downward, and the effect is negative.

In the previous lemma, we increased x_i^s while holding x_i^z constant, which required adjusting p_i^s to satisfy Bayesian plausibility. We now consider a different approach: keeping p_i^s fixed while varying x_i^z instead.

Lemma 3. *Let the set of signals be $\{L, H\}$. Fix the strategies of all senders except for sender i . Suppose sender i increases her interim posterior $x_i^s \neq \mu$, and adjusts $x_i^z \neq \mu$ to maintain (p_i^s, p_i^z) constant, according to Bayesian plausibility constraint (4). Then public posteriors react as follows:*

1. The same-signal effect: $\frac{dy^s}{dx_i^s} = \lambda_i^s > 0$

2. The cross-signal effect: $\frac{dy^z}{dx_i^s} = -\lambda_i^z \frac{p_i^s}{p_i^z} < 0$.

¹⁴If $s = L$ and $z = H$ then $x_i^z - \mu > 0$ and $x_i^z - x_j^z > 0$, confirming (7) is positive. If $s = H$ and $z = L$ then these terms are negative but the whole expression is still positive.

This lemma is simple to prove, given that weights λ_i^s are fixed by assumption. For example, an increase in $x_i^L < \mu$ while keeping p_i^s constant translates into a decrease of informativeness in the interim posterior sense; by the Lemma, informativeness decreases also in the public posterior sense.

Finally, let us record a simple result that if sender A is sufficiently likely, she has a non-common-order strategy that “undoes” any attempt by sender B to inform the receiver.

Lemma 4. *Let $\alpha_i \geq 1/2$. For any pure strategy of sender j such that $x_j^s \leq \mu \leq x_j^z$, sender i can respond with a non-common-order strategy $x_i^z \leq \mu \leq x_i^s$ such that the public posterior is the same as the prior, $y^s = \mu = y^z$.*

4 Local equilibria in electoral scenarios

This section leverages the structure of threshold utility functions in the binary case to characterize equilibria under assumptions 1 and 2. While the formal statements may appear tedious due to many cases, I summarize their key insights after presenting them.

Recall the electoral environment sketched up in the Introduction. The voter reelects the politician if her corruption belief is less than t_A , and otherwise votes the politician out of power. Moreover, if the belief reaches an even greater threshold t_B , the politician permanently crashes out of politics.

In the first scenario, both senders oppose the politician, with sender A aiming only to ensure the politician’s defeat in the election, while sender B prioritizing ending the politician’s future career. Their preferences can be represented by an increasing *threshold* utility function of the public probability belief that the politician is corrupt, defined as follows:

$$u_i(y^s) = \begin{cases} 0 & \text{if } y^s < t_i \\ 1 & \text{if } y^s \geq t_i \end{cases}$$

where t_i represents the corresponding receiver’s critical threshold for the level of belief. All our examples assume $t_A < t_B$.

Scenario 1. Senders’ utility functions have a threshold form with t_i , and they are both increasing.

Figure 2 illustrates the utilities of the two senders as functions of the probability belief of corruption. The left panel shows the case when the prior




Figure 2: Some scenario 1 cases.

probability belief falls between the thresholds, $t_A < \mu < t_B$, while the panel on the right depicts $\mu < t_A < t_B$. In either case, the public posterior that sender B considers ideal is too informative for sender A . In the left case, sender A wants to reveal nothing, while on the right, her ideal policy is somewhat informative but it is still less informative than the ideal policy for B .

Figure 2 also illustrates possible interim posteriors, though at this stage it is unclear whether these strategy profiles arise in equilibrium; the public posteriors are also shown. The expected payoffs these strategies generate are also depicted as middle dots on the dashed line segments.

The equilibrium analysis for different parameter configurations is next. Define $t^* = \frac{\mu}{\alpha_B \mu + \alpha_A}$ and define $D_1(t_A) = \frac{t_A \alpha_B \mu}{t_A - \alpha_A \mu}$ for $t_A > \mu \alpha_A$. Both are shown in Figure 3, with t_A on the horizontal axis and $t_B > t_A$ on the vertical one. Scenario 1 equilibria for all parameter constellations are described as follows.

Proposition 1. *Assume Scenario 1 with $t_A < t_B$. Under Assumptions 1 and 2, equilibrium exists:*

1. If $t_B \leq t^*$, then the following is a LPSNE: $x_A = (0, \mu)$ and $x_B = \left(0, \frac{t_B \alpha_B \mu}{\mu - t_B \alpha_A}\right)$, with public posterior $y = (0, t_B)$.
2. If $t^* < t_B$ and $t_A \leq t^*$, then the following is a LPSNE: $x_A = (0, \mu)$

and $x_B = (0, 1)$, with public posterior $y = \left(0, \frac{\mu}{\alpha_A + \mu\alpha_B}\right)$.

3.  If $\mu < t_B < D_1(t_A)$, then the following is a LPSNE: $x_A = (\mu, 1)$ and $x_B = (0, t_B)$, with public posterior $y = \left(\frac{\alpha_A t_B \mu}{t_B - \alpha_B \mu}, t_B\right)$.
4.  If $t^* < t_A$, then the following is a LPSNE: $x_A = \left(0, \frac{t_A \alpha_A}{1 - t_A \alpha_B}\right)$ and $x_B = (0, 1)$, with public posterior $y = (0, t_A)$.
5.  If $t_B \leq \mu$, then the following is a LPSNE: $x_A = (\mu, 1)$ and $x_B = (\mu, 1)$, with public posterior $y = (\mu, 1)$.

Equilibria described in Proposition 1 are depicted in the left panel of Figure 3. Although the analysis does not aim to find all possible equilibria, the categories identified in the proposition cover together the entire parameter space. In some regions they overlap, indicating the possibility of multiple equilibria.

Despite the multitude of cases, the interpretation is straightforward. Recall that senders' preferences can be ordered. Sender A prefers to provide less information than sender B , although this misalignment may be tiny. In all equilibria, sender A reduces the provision of information while sender B does the opposite. This is a runaway race that eventually hits a bound: either sender A provides no information, flooding the public square with noise, or sender B attempts to provide full information.

Specifically, in cases 1 and 2 of the proposition, sender A withholds information entirely by sending only message H . Since a low signal realization cannot lead to electoral defeat, sending only H maximizes the chance of the outcome favored by sender A . Case 1 is depicted on the left panel of Figure 2. Here, sender B manages to provide enough information to induce a favorable action after H , achieving her optimal expected payoff, whereas in case 2, she attempts to do so by attempting fully revealing information, but is unable to sufficiently shift the outcome in her favor.

Case 3 equilibrium is different in nature. Sender A still avoids providing information, but does so in a completely different way: by sending only message L . This makes public realization L sufficiently noisy to push the corresponding posterior far enough from zero, so that $t_A \leq y^L$. As a result, the politician loses the election with probability one, which is sender A 's preferred outcome.




Figure 3: Types of LPSNE. Left: Scenario 1. Right: Scenario 2.

In case 4, even though sender B provides as much information as possible, sender A can completely offset this attempt by introducing carefully calibrated lies and obtains her first-best outcome. This case is shown on the right of Figure 2.

Case 5, provided here for completeness, is the easiest one, as the senders have common no-communication interest. Thus, there is an equilibrium in which they fully coordinate in withholding information.

Now, let us move to the second scenario. Recall that now, sender B continues to be a strong opponent of the politician, but sender A is a supporter.

Scenario 2. Sender B has an increasing threshold utility function with threshold t_B . Sender A has a *decreasing* threshold utility function with threshold t_A .

These utilities are shown in Figure 4. If $t_A < \mu < t_B$ and if no information is provided, the outcome of the election is a mere loss in the vote but the politician survives to live another day. This outcome does not satisfy either group: sender A does not get their preferred politician, and sender B does not succeed in banishing them. Despite the dramatic conflict of interest between the senders, they agree that information should be provided, as this may sometimes crush the politician's career if they are corrupt, or elect them if they are not, the average of which is better than no information.




Figure 4: Case 1 of scenario 2.

To analyze this scenario, define

$$D_2(t_A) = \frac{(1 - t_A)\mu - \alpha_A(1 - \mu)t_A}{(1 - t_A) - \alpha_A(1 - \mu)}.$$

Proposition 2. *Assume Scenario 2 with $t_A < t_B$. Under Assumptions 1 and 2, equilibrium exists:*

1. $\square$ If $D_2(t_A) \leq t_B \leq D_1(t_A)$, then

$$x_A = \left(\frac{t_A(t_B - \mu)}{t_A(t_B - \mu) + (1 - \mu)\alpha_A(t_B - t_A)}, 1 \right)$$

$$x_B = \left(0, \frac{(t_B - t_A)\alpha_B\mu}{(1 - \mu\alpha_A)(t_B - t_A) - (1 - t_A)(t_B - \mu)} \right)$$

is a LPSNE. The public posterior is $y = (t_A, t_B)$.

2. $\square$ If $D_1(t_A) \leq t_B$, then $x_A = (\mu, 1)$ and $x_B = (0, t_B)$ is a LPSNE. The public posterior is $y = \left(\frac{\alpha_A t_B \mu}{t_B - \alpha_B \mu}, t_B \right)$.

3. $\square$ If $\mu < t_A < t_B$ and $\frac{\mu}{\alpha_A + \alpha_B \mu} < t_B$ then $x_A = (0, \mu)$, $x_B = (0, 1)$ is a LPSNE. The public posterior is $y = \left(0, \frac{\mu}{\alpha_A + \mu\alpha_B} \right)$.¹⁵

¹⁵Symmetrically, cases for $t_A < t_B < D_2(t_A)$ are captured as follows. If $t_B \leq$

Despite the fact that there is no outcome the senders agree to support, there is one they are united in opposing, and this gives some scope for cooperation in providing information to the receiver, as shown in Case 1 of the Proposition and illustrated in Figure 4. Both senders provide information, although they gamble that this information will realize different outcomes: sender A prefers the low signal realization and election win, and sender B prefers the high one leading to a full crash of the politician. Consequently, sender A prefers to fully reveal if the politician is honest and lie slightly if she is corrupt to maximize the expected probability of winning the election. Sender B prefers the opposite: to always reveal the corrupt politician and sometimes lie if she is honest, to maximize the chances that the politician is permanently canceled. The resulting public message is noisy and it never achieves senders' respective most preferred outcomes. However, it achieves a compromise in which the worst outcome (a mere loss of the election) never occurs.

Case 2 is similar in nature except that it describes the situation in which thresholds are sufficiently asymmetric to trigger one of the senders to withhold information. Specifically, even though sender A would like the low signal realization to occur more often, she already does everything from her point of view to flood the public square with it.

Case 3 is somewhat different than the previous two. Since the outcome involves $y^H < t_A$, sender A obtains her first-best payoff of 1 with certainty, so the only chance for this not to be an equilibrium is if sender B can profitably deviate locally to shift the public posterior to above t_B . With such high threshold in this case, she is not able to.

5 Relaxing the small-sender assumption

The small-sender assumption 2 helps to sustain equilibria by allowing only local deviations. In this section, I consider a standard pure-strategy Nash equilibrium (PSNE) between two senders, A and B , where large deviations are also permitted. This can alternatively be interpreted as a setting with a large population of small independent senders of two types, A and B , where

$D_2(t_A)$ then $x_A = (t_A, 1)$ and $x_B = (0, \mu)$ is a LPSNE with the public posterior $y = \left(t_A, \frac{\alpha_A(\mu - t_A) + (1 - t_A)\alpha_B\mu}{\alpha_A(\mu - t_A) + (1 - t_A)\alpha_B}\right)$. If $t_A < t_B < \mu$ and $t_A < \frac{\alpha_B\mu}{\alpha_A(1 - \mu) + \alpha_B}$, then $x_A = (0, 1)$ and $x_B = (\mu, 1)$ is a LPSNE.




Figure 5: Types of PSNE. Left: Scenario 1. Right: Scenario 2.

each type can act as a coalition, capable of coordinating its strategy.

Figure 5 is a counterpart of Figure 3. Clearly, all PSNE are also LPSNE, but the reverse is not true. We see that some regions shrink in comparison to Figure 3 as newly available non-local deviations knock out certain categories of equilibria for given parameter values. All new constraints are defined in the Appendix in section 9.8, where the counterparts of propositions 1 and 2 are stated and proven.

Case 4 in the right panel is worth highlighting because it shows that a standard PSNE may not exist for some parameter values even though Assumption 1 holds. Here, sender A is conflicted by two different methods of designing information provision. On the one hand, she wants to flood the public square with the high signal realization to guarantee low public posterior, very much à la case 3. But then, with a low threshold t_B in case 4, sender B has incentives to provide information to move y^H sufficiently high to cancel the politician. In this case, sender A wants to carefully engineer communication to increase both public posteriors y^L and y^H , in the spirit of case 1. However, if this is part of an equilibrium, sender B must decrease both y^L and y^H in response, again like in case 1. Now, in contrast to case 1, sender A realizes she has an incentive to switch to no information provision as in case 3, because t_B is so close to t_A . We are back at the beginning of the best response cycle and thus no pure strategy Nash equilibrium exists. This switchback requires a large deviation, so it could not play any role in

the LPSNE of Proposition 2.

6 Relaxing the common-order assumption

This section provides some more general intuitions to clarify the idea that when non-common-order strategies can be used, the game with sufficiently conflicting interests has features of a Matching Pennies game, where best response cycles in pure strategies arise and no pure-strategy equilibrium exists. The logic goes as follows: Sender B , who wants to promote communication, constructs informative signals in response to a given signal strategy of sender A , but then sender A who wants to reduce informativeness, associates those signal realizations with precisely opposite outcomes, thus introducing noise. This cannot be an equilibrium either, because then sender B wants to deviate to match the sender A 's signals, restoring their informativeness.

Armed with Lemma 4 as a stepping stone for the case of the sufficiently frequent obstructionist, we obtain the following result, where cou_i is a concavification of function u_i .

Proposition 3. *Suppose $\alpha_A \geq \alpha_B$, $\text{cou}_A(\cdot)$ is strictly concave at μ , and $u_B(\cdot)$ is strictly convex at μ . Then no PSNE exists.*

The above result requires that sender A , the obstructionist, is sufficiently likely, but it rests on local curvature at μ .¹⁶

The next result does not require sender A to occur frequently, but relies on a stronger assumptions of global concavity and convexity.

Proposition 4. *Suppose $u_A(\cdot)$ is strictly concave everywhere; suppose $u_B(\cdot)$ is strictly convex everywhere. Then no PSNE exists.*

The common-order Assumption 1 may help to restore pure-strategy equilibria.

Proposition 5. *Suppose $u_A(\cdot)$ is strictly concave everywhere; suppose $u_B(\cdot)$ is strictly convex everywhere. Under Assumption 1,*

1. *a fully informative strategy, $x_B = (0, 1)$ is strictly dominant for sender B ,*

¹⁶Function $u_B(\cdot)$ was required to be strictly convex only around μ which is a local condition. Function $\text{cou}_A(\cdot)$ is strictly concave also only locally, but constructing this concavification requires the entire function $u_A(\cdot)$.

2. any strategy in which $x_A^L < \mu < x_A^H$ is strictly dominated by some strategy in which $x_A = (\mu, \mu)$ for sender A ,
3. a PSNE exists.

Proposition 5 asserts the existence of pure strategy equilibria in a very extreme case. However, the common-order assumption does not guarantee the existence of pure strategy equilibrium in general as we saw in the previous section.

These observations provide an interpretation of what “annihilating truth”—as Garry Kasparov stated in the initial quote—might mean. A sophisticated viewer who recognizes individual senders and observes their signal realizations might be alarmed by how much poorer public information is in comparison. From this perspective, information is destroyed. Of course, the Blackwell theorem holds because from its point of view, the public sees deteriorated signal $\sum_i \alpha_i \pi_i(s|\omega)$ rather than its higher quality components, $\pi_A(s|\omega), \pi_B(s|\omega)$.

7 Normative performance of debates

Criteria. As a starting point for normative analysis, I focus on the welfare criterion of Blackwell *informativeness*, applied to equilibrium public signals when these can be ordered. The main comparison is between a two-sender babel and a non-anonymous environment represented by a single-sender case (or the average of two single-sender cases). Formally, this means comparing parameter configurations where $\alpha_i \in (0, 1)$ to cases where $\alpha_i = 1$ for some i .¹⁷

I also consider *balance*, understood as the degree of symmetry in the roles of the senders. This is an ad hoc but more conventional way to assess the quality of public debate. I do not offer a precise definition of this concept, but

¹⁷The current version of the babel model, in which one sender is selected at random, is not directly comparable to the non-anonymous environment studied in Gentzkow and Kamienica (2017b), Ravindran and Cui (2022) and related work, where the receiver observes both senders’ signal realizations. However, a babel variant in which both realizations are revealed but their sources remain anonymous would be comparable to their setup. In such a variant, the public signal realization is $r \in \{LL, LH, HH\}$ where $s_A = L$ and $s_B = H$ is indistinguishable from $s_A = H$ and $s_B = L$, and both are recorded simply as LH ; the resulting public posterior y^r lies in a two-dimensional simplex. I do not pursue this analysis here.

several possible specifications come to mind which I explain as the analysis unfolds.

Scenario 1. Consider Scenario 1 with some conflict of interests $\mu < t_B$, and examine two polar cases: when α_B is near 100% and when it is close to 0%.

Suppose the public square is initially dominated by the more informative senders B , and some less informative senders begin enter, so α_B drops slightly from 100%. In this situation, only case 1 of proposition 1 arises in equilibrium. The resulting public posterior remains $y = (0, t_B)$, the same when $\alpha_B = 1$, even as some contamination occurs. The reason is that sender A has enough capacity to neutralize the noise introduced by relatively sporadic obstructionists. This capacity is constrained by how sporadic they are, as defined by α_B , but also by her own honesty, as measured by the distance of t_B from one.¹⁸

Now, let's consider the opposite polar case. Suppose that the marketplace of ideas is initially populated only by senders A , and more informative senders B begin to enter, so that α_B increases from 0%. When $t_A > \mu$, there is a unique equilibrium given by case 4 of the proposition 1. By a similar argument to the one in the previous paragraph, sender A 's effect is fully neutralized even as α_B is positive, since sender B has enough capacity to cancel out A 's influence. When $t_A < \mu$, both cases 2 and 3 from proposition 1 characterize equilibria. In both, the entry of sender B may immediately improve the informativeness of the public square.

We can also view Scenario 1 through the lens of balance of the public square. The receiver is best off when the public square is unbalanced in the sense that it is fully dominated by the more informative sender, that is, when $\alpha_B = 100\%$. Any attempt to increase balance of such a marketplace of ideas will undermine these benefits.

On the other hand, invoking balance as a justification to increase the presence of more informative senders starting from a low-information environment such as $\alpha_B = 0\%$, may improve the informativeness of the debate. This would, however, require a designer to "pick the winner." This may be possible if the designer (for example a regulator, national broadcaster, or newspaper editor) can be trusted to identify anonymous but informative

¹⁸To see the latter point, note that sender B 's ability to counter the noise from sender A relies on her own reluctance to fully inform the receiver. If she were totally honest, so that $t_B = 1$, she would have no capacity to offset sender A 's noise and any such noise would immediately result in case 2 or 4 of the proposition 1.

sources. Still, the process remains vulnerable to manipulation. A more robust approach is to remove anonymity, where possible, and allow the receiver to pick the winner themselves without relying on gatekeepers.

Scenario 2. Here, the informativeness of babel is *strictly worse* than in the single-sender case, regardless of which sender is met at the public square. The outcome is also worse if a non-anonymous but randomly selected sender (say, with probability α_i) interacts with the receiver in the public square.

Comparative statics with respect to α_B is as expected: when α_B is close to 100%, the outcome approaches that of the single-sender case with sender B , and similarly for the opposite case.

It is also interesting to consider the case where t_A is close to t_B , as this situation represents a zero-sum game in the limit. The starkest result appears in case 1 of proposition 2, where t_A and t_B both converge to μ , satisfying the condition for this case. The resulting equilibrium public posterior is $y = (t_A, t_B)$, which becomes completely uninformative in the limit. This stands in sharp contrast with zero-sum games discussed in Gentzkow and Kamenica (2017a) and Ravindran and Cui (2022), where the equilibrium must be fully informative. Even if the situation is not so knife-edge, the equilibrium remains at best partially informative. For instance, when t_A converges to a fixed $t_B > \mu$, equilibrium is described by case 2 or 3 of the proposition and is not close to being fully informative.

To someone concerned with the balance of the debate, Scenario 2 might appear to represent a well-functioning marketplace of ideas. Consider case 1 in proposition 2 as the leading case. Although both senders try to steer the public posterior in their preferred direction, they are symmetric in their roles. Unlike in Scenario 1, each sender shares some information, and neither acts as a clear saboteur. However, the outcome is poor: each sender provides less information than they would if acting alone, and the resulting public posterior is less informative as well.

8 Conclusions

This study takes a standard multi-sender model of Bayesian persuasion and modifies it by making the senders anonymous. Consequently, the receiver cannot attribute the message to the author. This approach enables a nuanced study of endogenous information obfuscation and, in particular, the role of anonymity in debate design.

Unfortunately, the model is highly-stylized. The most problematic assumption is that the set of signal realizations is binary. In the information design analysis, the choice of signal realizations ought to be endogenous and the set from which they can be selected ought to be large. Other assumptions, such as binary state of nature or two agents seem less of an issue.

9 Appendix: proofs and additional results

9.1 Proof of Lemma 1

Proof. Manipulating interim posteriors (1) and using $p_i^s = \sum_{\omega} \mu(\omega) \pi_i(s|\omega)$, we obtain

$$\pi_i(s|\omega) = \frac{1}{\mu(\omega)} p_i^s x_i^s(\omega).$$

Substituting these into public posterior (2) and canceling $\mu(\omega)$ leads to

$$y^s(\omega) = \frac{\sum_i \alpha_i p_i^s x_i^s(\omega)}{\sum_{\omega'} \sum_i \alpha_i p_i^s x_i^s(\omega')}.$$

Using $\sum_{\omega'} x_i^s(\omega') = 1$ in the denominator, we obtain

$$y^s(\omega) = \frac{\sum_i \alpha_i p_i^s x_i^s(\omega)}{\sum_i \alpha_i p_i^s}.$$

which, after substituting in λ_i^s defined in (6), gives the desired result. $\square$

9.2 Proof of Lemma 2

Proof. Step 1. Using the formula for public posterior from Lemma 1 and isolating player i , we obtain

$$y^s = \frac{\alpha_i p_i^s x_i^s + D_i^s}{\alpha_i p_i^s + C_i^s} \tag{9}$$

where $C_i^s = \sum_{j \neq i} \alpha_j p_j^s$ and $D_i^s = \sum_{j \neq i} \alpha_j p_j^s x_j^s$ are independent of i 's strategy.

Step 2. Using Bayesian plausibility (4) and taking the same-signal derivative of p_i^s implies $\frac{dp_i^s}{dx_i^s} = \frac{x_i^z - \mu}{(x_i^z - x_i^s)^2}$. Similarly, the cross-signal derivative is $\frac{dp_i^z}{dx_i^s} = -\frac{x_i^z - \mu}{(x_i^s - x_i^z)^2}$.

Step 3. Calculate the same-signal derivative of (9), where step 2 and the definitions of C_i^s and D_i^s were used to simplify the numerator:

$$\frac{dy^s}{dx_i^s} = \frac{\alpha_i (x_i^z - \mu)}{(x_i^z - x_i^s)^2 (\alpha_i p_i^s + C_i^s)^2} \left(\alpha_i (x_i^z - \mu) + \sum_{j \neq i} \alpha_j p_j^s (x_i^z - x_j^s) \right).$$

This proves point 1 of the Lemma.

Step 4. Similarly, starting with y^z written in the form of (9), calculate the cross-signal derivative. Use step 2 and the definitions of C_i^z and D_i^z to simplify the derivative to

$$\frac{dy^z}{dx_i^s} = \frac{\alpha_i (x_i^z - \mu)}{(x_i^z - x_i^s)^2 (\alpha_i p_i^z + C_i^z)^2} \left(\sum_{j \neq i} \alpha_j p_j^z \right) \left(\frac{\sum_{j \neq i} \alpha_j p_j^z x_j^z}{\sum_{j \neq i} \alpha_j p_j^z} - x_i^z \right).$$

This proves point 2 of the Lemma. $\square$

9.3 Proof of Lemma 3

Proof. Strategy of sender j do not change, (x_j, p_j) . The probabilities of signals do not change for player i either, by assumption. Hence, λ_i^s defined in (6) remains constant for all senders and signal realizations. Increasing x_i^s leads to the increase of $y^s = \lambda_i^s x_i^s + \sum_{j \neq i} \lambda_j^s x_j^s$ and (through x_i^z) the decrease of $y^z = \lambda_i^z \left(\mu - \frac{p_i^s}{p_i^z} x_i^s \right) + \sum_{j \neq i} \lambda_j^z x_j^z$. $\square$

9.4 Proof of Lemma 4

Proof. With two senders $i \neq j$ and two signals $s \neq z$, equation (6) can be expressed as

$$\lambda_i^s = \frac{\alpha_i (x_i^z - \mu) (x_j^z - x_j^s)}{\alpha_i (x_i^z - \mu) (x_j^z - x_j^s) + \alpha_j (x_j^z - \mu) (x_i^z - x_i^s)} \quad (10)$$

and the public posterior (5) reduces to

$$y^s = \lambda_i^s x_i^s + \lambda_j^s x_j^s. \quad (11)$$

The case $\alpha_i = \alpha_j$ is particularly easy. Let sender i 's non-common-order strategy be explicitly defined in terms of j 's strategy: $(x_i^s, x_i^z) = (x_j^z, x_j^s)$.




Figure 6: How sender i 's strategy affects the public posterior and payoffs, for a fixed strategy of j .

The public posterior (5) reduces to $y^s = \mu$, because terms $x_j^z - x_j^s = x_i^s - x_i^z$ and $\alpha_i = \alpha_j$ cancel, and the numerator and the denominator further simplify.

The general case $\alpha_i \geq \alpha_j$ is shown indirectly. The public posterior is equal to the prior, $y^s = \mu$, if the following condition holds

$$\frac{\alpha_i (x_i^z - \mu) (\mu - x_i^s)}{(x_i^z - x_i^s)} = \frac{\alpha_j (\mu - x_j^z) (x_j^s - \mu)}{(x_j^s - x_j^z)}$$

For any α_j and μ , the term on the right is a number in the interval $[0, \alpha_j \mu (1 - \mu)]$ and any such number is generated by some strategy of sender j . The term on the left is a continuous function $[\mu, 1] \times [0, \mu] \rightarrow [0, \alpha_i \mu (1 - \mu)]$ of sender i strategy, x_i . Since $\alpha_i \geq \alpha_j$, for any right-hand side, there is a pair (x_i^s, x_i^z) that satisfies the equation, thus generating the desired posterior μ . $\square$

9.5 Expected utility in threshold examples

I begin by examining the general effect of changing a strategy component x_i^s on the sender's expected utility, holding fixed the opponent's strategy and the other components of sender i 's own strategy, x_i^z . Lemma 2 showed that while the direction of the same-signal effect is clear, the cross-signal effect can go either way, depending on other factors.

To see the issue, consider sender A in Figure 4, assuming strict inequality $y^L < t_A$. Increasing x_A^L raises y^L , which on its own increases expected utility. However, it may simultaneously reduce y^H , which lowers expected utility. On the surface, the total effect is ambiguous. I show that the structure imposed by the threshold utility function resolves this ambiguity.

Here, the expected utility is Eu_i . Written as a function $U_i(x_i^s, y^z)$, a correct way to decipher it is this: Suppose signal realization z is the one

for which i gets the higher payoff and s is the one for which she gets the lower payoff, i.e. $u_i(y^z) = 1$ and $u_i(y^s) = 0$. Then $U_i(\cdot, y^z)$ is the expected utility viewed as a function of x_i^s for a fixed y^z (i.e. when x_i^s changes, x_i^z must be adjusted to keep y^z constant; then both the initial change of x_i^s and the adjusted x_i^z affect the expected utility via y^s only, see Figure 6). By contrast, $U_i(x_i^s, \cdot)$ is the same expected utility but viewed as a function of y^z for a fixed x_i^s (i.e. when y^z changes, x_i^z is adjusted but then so is y^s , while x_i^s is kept constant; both the initial change of y^z and adjusted y^s affect the expected utility, both triggered by x_i^z only).

Lemma 5. *Suppose sender i has a threshold utility function. Suppose signal realization z is the one for which i gets the higher payoff, and s is the one for which she gets the lower payoff: $u_i(y^z) = 1$ and $u_i(y^s) = 0$. Fix the strategy of sender j . Then i 's expected utility is between 0 and 1 and*

1. *Expected utility as a function of (y^s, y^z) is*

$$Eu_i = \frac{y^s - \mu}{y^s - y^z}$$

2. *Expected utility as a function of (x_i^s, y^z) is*

$$U_i(x_i^s, y^z) = \frac{\alpha_i(x_i^s - \mu)(x_j^s - x_j^z) + \alpha_j(x_j^s - \mu)(x_i^s - x_j^z)}{(x_j^s - x_j^z)(x_i^s - y^z)}$$

Proof. Using (10) and (11), public posteriors become

$$y^L = \frac{\alpha_A(x_A^H - \mu)(x_B^H - x_B^L)x_A^L + (x_A^H - x_A^L)\alpha_B(x_B^H - \mu)x_B^L}{\alpha_A(x_A^H - \mu)(x_B^H - x_B^L) + (x_A^H - x_A^L)\alpha_B(x_B^H - \mu)} \quad (12)$$

$$y^H = \frac{\alpha_A(\mu - x_A^L)(x_B^H - x_B^L)x_A^H + (x_A^H - x_A^L)\alpha_B(\mu - x_B^L)x_B^H}{\alpha_A(\mu - x_A^L)(x_B^H - x_B^L) + (x_A^H - x_A^L)\alpha_B(\mu - x_B^L)} \quad (13)$$

Step 1. Equation (12) expresses y^L as a function of x_A^L . Invert this to express x_A^L as a function of y^L .

$$x_A^L = \frac{\psi_1}{\psi_1 + \psi_2} y^L + \frac{\psi_2}{\psi_1 + \psi_2} x_A^H$$

where $\psi_1 = \alpha_A (x_A^H - \mu) (x_B^H - x_B^L)$ and $\psi_2 = \alpha_B (x_B^H - \mu) (y^L - x_B^L)$. The two following differences can be derived immediately:

$$\begin{aligned}\mu - x_A^L &= \frac{\psi_1 (\mu - y^L) + \psi_2 (\mu - x_A^H)}{\psi_1 + \psi_2} \\ x_A^H - x_A^L &= \frac{\psi_1 (x_A^H - y^L)}{\psi_1 + \psi_2}\end{aligned}$$

Step 2. Equation (13) expresses y^H as a function of x_A^L . Using step 1, we eliminate x_A^L to rewrite y^H as a function of y^L . After simplifying and canceling terms, we obtain:

$$y^H = \frac{\xi_1 x_A^H + \xi_2 x_B^H}{\xi_1 + \xi_2}$$

where $\xi_1 = \alpha_A (\psi_1 (\mu - y^L) + \psi_2 (\mu - x_A^H)) (x_B^H - x_B^L)$ and $\xi_2 = \psi_1 (x_A^H - y^L) \alpha_B (\mu - x_B^L)$. We can also compute the following two differences directly:

$$\begin{aligned}y^H - \mu &= \frac{\xi_1 (x_A^H - \mu) + \xi_2 (x_B^H - \mu)}{\xi_1 + \xi_2} \\ y^H - y^L &= \frac{\xi_1 (x_A^H - y^L) + \xi_2 (x_B^H - y^L)}{\xi_1 + \xi_2}\end{aligned}$$

Step 3. Recall that the expected utility for the decreasing threshold utility function is given by $Eu_A = \frac{y^H - \mu}{y^H - y^L}$. Using step 2, we substitute the differences and cancel the common term $\xi_1 + \xi_2$:

$$Eu_A = \frac{\xi_1 (x_A^H - \mu) + \xi_2 (x_B^H - \mu)}{\xi_1 (x_A^H - y^L) + \xi_2 (x_B^H - y^L)} = \frac{E_n}{E_d}$$

Observe that the numerator E_n is a linear function of y^L , since it enters indirectly through ψ_2 , ξ_1 , and ξ_2 . In contrast, the denominator E_d is a quadratic function of y^L . The goal now is to factor out the term $\mu - y^L$ from both the numerator and the denominator, cancel it, and conclude that the expected utility depends on y^L only through the denominator.

Focus on the numerator E_n

$$\begin{aligned}E_n &= \alpha_A (\psi_1 (\mu - y^L) + \psi_2 (\mu - x_A^H)) (x_B^H - x_B^L) (x_A^H - \mu) \\ &\quad + \psi_1 (x_A^H - y^L) \alpha_B (\mu - x_B^L) (x_B^H - \mu)\end{aligned}$$

Substituting ψ_2 into the expression and some cleaning up allows the term $\mu - y^L$ to be factored out:

$$E_n = (\mu - y^L) \alpha_A (x_A^H - \mu) (x_B^H - x_B^L) C$$

where

$$C = \alpha_A (x_A^H - \mu) (x_B^H - x_B^L) + \alpha_B (x_B^H - \mu) (x_A^H - x_B^L)$$

Now focus on the denominator E_d :

$$\begin{aligned} E_d = & \alpha_A (\psi_1 (\mu - y^L) + \psi_2 (\mu - x_A^H)) (x_B^H - x_B^L) (x_A^H - y^L) \\ & + \psi_1 (x_A^H - y^L) \alpha_B (\mu - x_B^L) (x_B^H - y^L) \end{aligned}$$

As before, substituting for ψ_1 and ψ_2 and much cleaning out shows that the term $\mu - y^L$ can be factored out:

$$E_d = (\mu - y^L) \alpha_A (x_A^H - y^L) (x_B^H - x_B^L)^2 (x_A^H - \mu) (\alpha_A + \alpha_B)$$

Putting the expressions for E_n and E_d together, the factor $\mu - y^L$ along with several other common terms, cancels out. The expected utility simplifies to:

$$Eu_A = \frac{\alpha_A (x_A^H - \mu) (x_B^H - x_B^L) + \alpha_B (x_B^H - \mu) (x_A^H - x_B^L)}{(x_A^H - y^L) (x_B^H - x_B^L) (\alpha_A + \alpha_B)}$$

□

9.6 Proof of Proposition 1

Point 1 of the Proposition

Statement: If $t_B \leq \frac{\mu}{\alpha_B \mu + \alpha_A}$, then the following is a LPSNE: $x_A = (0, \mu)$ and $x_B = \left(0, \frac{t_B \alpha_B \mu}{\mu - t_B \alpha_A}\right)$. Proof:

Firstly, observe that these strategies are feasible; in particular, $x_B^H \leq 1$ is guaranteed by the required condition of the statement. Secondly, by applying equations (10) and (11), observe that these interim posteriors generate the public posterior $y = (0, t_B)$.

Sender B 's problem: Sender B obtains the first best, so she has no incentives to deviate.

Sender A 's problem: Suppose $x_B = \left(0, \frac{t_B \alpha_B \mu}{\mu - t_B \alpha_A}\right)$ and consider sender A 's expected utility when she uses $x_A = (0, \mu)$. Sender A obtains 1 when the

signal realization is high and 0 when it is low, so the expected utility is $Eu_A = \frac{\mu - y^L}{y^H - y^L}$. By Lemma 2, increasing x_A^L from zero increases both y^L and y^H , which reduces expected utility. Likewise, increasing x_A^H from μ does the same.

Point 2 of the Proposition

Statement: If $\frac{\mu}{\alpha_B \mu + \alpha_A} < t_B$ and $t_A \leq \frac{\mu}{\alpha_A + \mu \alpha_B}$, then the following is a LPSNE: $x_A = (0, \mu)$ and $x_B = (0, 1)$. Proof:

This strategy profile generates public posterior $y = \left(0, \frac{\mu}{\alpha_A + \mu \alpha_B}\right)$.

Sender B 's problem: The first restriction implies that $y^H < t_B$. Not only sender B 's payoff is zero with certainty, she cannot change this locally. Since sender B 's utility is considered strictly convex in this case, she must attempt fully informative strategy $x_B = (0, 1)$.

Sender A 's problem: The second restriction implies $t_A \leq y^H$, and hence sender A obtains 1 when the signal realization is high and 0 when it is low; the expected utility is $Eu_A = \frac{\mu - y^L}{y^H - y^L}$. By Lemma 2, increasing x_A^L from zero increases both y^L and y^H , which reduces expected utility. Likewise, increasing x_A^H from μ does the same.

Point 3 of the Proposition

Statement: If $\mu < t_B \leq D_1(t_A)$, then the following is a LPSNE: $x_A = (\mu, 1)$ and $x_B = (0, t_B)$. Proof:

This strategy profile implies that $y = \left(\frac{\alpha_A t_B \mu}{t_B - \alpha_B \mu}, t_B\right)$.

Sender B 's problem: When this strategy profile is played, sender B obtains payoff one after high signal and zero after low signal, hence her expected payoff is $Eu_B = \frac{\mu - y^L}{y^H - y^L}$. By Lemma 2, increasing x_B^L from zero increases y^L and does not affect y^H ; this lowers the expected utility. Increasing x_B^H from t_B increases both y^L and y^H which has a similar effect.

Sender A 's problem: The restriction in the statement implies that $t_A \leq y^L$, and therefore sender A obtains utility 1 with certainty which is her first best.

Point 4 of the Proposition

Statement: If $\frac{\mu}{\alpha_A + \mu\alpha_B} \leq t_A$, then the following is a LPSNE: $x_A = \left(0, \frac{t_A\alpha_A}{1-t_A\alpha_B}\right)$ and $x_B = (0, 1)$. Proof:

Firstly, observe that the restriction in the statement ensures that equilibrium $x_A^H \geq \mu$. Secondly, this strategy profile generates $y = (0, t_A)$.

Sender A 's problem: Public posterior is sender A 's first best when $\mu < t_A$, so she has no incentives to deviate.

Sender B 's problem: Sender B obtains payoff zero regardless of the public signal realization and therefore her utility is considered convex. Her strategy is already maximally informative, so no deviation is beneficial.

Point 5 of the Proposition

Statement: If $t_B \leq \mu$, then the following is a LPSNE: $x_A = (\mu, 1)$ and $x_B = (\mu, 1)$. This is a common no-communication interest case; no communication is a trivial equilibrium.

9.7 Proof of Proposition 2

Point 1 of the Proposition

Statement: If $D_2(t_A) \leq t_B \leq D_1(t_A)$, then $x_A = \left(\frac{t_A(t_B - \mu)}{t_A(t_B - \mu) + (1 - \mu)\alpha_A(t_B - t_A)}, 1\right)$ and $x_B = \left(0, \frac{(t_B - t_A)\alpha_B\mu}{(1 - \mu\alpha_A)(t_B - t_A) - (1 - t_A)(t_B - \mu)}\right)$ is a pure strategy equilibrium. Proof:

This strategy profile generates public posterior $y = (t_A, t_B)$.

Sender A 's problem: Suppose B uses strategy x_B . Sender A can achieve payoff 1 only after low signal realization and only when $y^L \leq t_A$ holds. By Lemma 5, sender A 's expected payoff in this case is $U_A(x_A^H, y^L)$. Observe that this function is increasing in y^L . It is also increasing in x_A^H if $t_B < D_1(t_A)$, which holds here. Thus, to characterize the best response of A , both those variables must be the highest possible in their domains, i.e. $y^L = t_A$ and $x_A^H = 1$. Component x_A^L is recovered from these.

Sender B 's problem: The argument is symmetric to the previous paragraph: Suppose A uses x_A . Sender B can achieve payoff 1 only after high signal realization and only when $y^H \geq t_B$ holds. By Lemma 5, sender B 's expected utility is $U_B(x_B^L, y^H)$. This function is decreasing in y^H . It is also decreasing in x_B^L if $D_2(t_A) < t_B$, which is the case. Both should be as low

as possible within their domains, hence $y^H = t_B$ and $x_B^L = 0$. The remaining variable x_B^H can be recovered from these.

Point 2 of the Proposition

Statement: If $D_1(t_A) \leq t_B$, then $x_A = (\mu, 1)$ and $x_B = (0, t_B)$ is a LPSNE.
Proof:

This strategy profile generates $y = \left(\frac{\alpha_A t_B \mu}{t_B - \alpha_B \mu}, t_B \right)$.

Sender A 's problem: Let B use strategy $x_B = (0, t_B)$. Changing x_A should not improve the expected payoff locally when $y^L \leq t_A < y^H$.

Let's identify the domain of y^L . By checking the formula for y^L for the given strategy x_B and for any x_A^H , we see that the resulting public posterior is $y^L \in \left[0, \frac{\alpha_A t_B \mu}{t_B - \alpha_B \mu} \right]$ as x_A^L varies in the feasible interval $[0, \mu]$. The condition $D_1(t_A) \leq t_B$ in the statement guarantees that the upper bound is not greater than t_A , and therefore utility from the low signal is guaranteed to be 1. Sender A wants to maximize the overall expected utility $U_A(x_A^H, y^L)$ in the domain. Since this function is increasing in its second argument, the solution must be the upper bound, $y^L = \frac{\alpha_A t_B \mu}{t_B - \alpha_B \mu}$. In this case, the second argument does not affect the payoff, so $x_A^H = 1$ is a best response. The remaining argument x_A^L can be recovered from this to be μ .

Sender B 's problem: Let sender A use $x_A = (\mu, 1)$. Sender B maximizes $U_B(x_B^L, y^H)$ over the domain which, for the given strategy $x_A = (\mu, 1)$, is $x_B^L \in [0, \mu]$ and $y^H \in [t_B, 1]$. This function is decreasing in both arguments, and hence the solution is $(x_B^L, y^H) = (0, t_B)$. The last variable, x_B^H , is recovered from this.

Point 3 of the Proposition

Statement: if $\mu < t_A < t_B$ and $\frac{\mu}{\alpha_A + \alpha_B \mu} < t_B$, then $x_A = (0, \mu)$, $x_B = (0, 1)$ is a LPSNE. Proof:

The public posterior is $y = \left(0, \frac{\mu}{\alpha_A + \mu \alpha_B} \right)$.

Sender A gets her best payoff of 1 with certainty, so she has no incentive to deviate. Because of the restriction in the statement, $y^H < t_B$, sender B gets her worst payoff of 0 with certainty and cannot change it.

9.8 Section 5 results

This section adds additional constraints to the one listed in Propositions 1 and 2, converting them from LPSNE to conventional PSNE.

Define

$$D_3(t_A) = \frac{t_A \mu}{t_A(1 + \alpha_A) - \alpha_A \mu}, \text{ for } t_A > \frac{\mu \alpha_A}{1 + \alpha_A}$$

Proposition 6. *Assume conditions for Proposition 1. Under the following additional conditions, LPSNE described in Proposition 1 become PSNE:*

1.  $D_3(t_A) < t_B$.
2.  $\frac{\alpha_A \mu}{1 - \alpha_B \mu} \leq t_A$.
3.  no additional conditions.
4.  no additional conditions.
5.  no additional conditions.

Proof. Case 1 of the proposition: When the strategy profile of this case is played, sender A can try to increase y^L to reach t_A , which would be beneficial as it would guarantee payoff of 1 with certainty (this is a different deviation than the local one already considered in the proof of Proposition 1). Using $x_B^L = 0$, (10) and (11), we have

$$y^L = \frac{\alpha_A (x_A^H - \mu) x_B^H x_A^L}{\alpha_A (x_A^H - \mu) x_B^H + (x_A^H - x_A^L) \alpha_B (x_B^H - \mu)}$$

which increases with x_A^L for any x_B^H . Hence, the best chance to reach threshold t_A is to set $x_A^L = \mu$, which implies $y^L = \frac{\alpha_A x_B^H \mu}{x_B^H - \alpha_B \mu}$. Plugging in sender B 's strategy x_B^H implies that this public posterior is

$$y^L = \frac{\alpha_A t_B \mu}{t_B - \mu + t_B \alpha_A}.$$

However, the additional condition of the statement assures that the right hand side is lower than t_A , so sender A cannot increase y^L to reach threshold t_A .

Case 2 of the proposition: When the strategy profile of this case is played, sender A can try to increase y^L up to t_A , which would be again beneficial

as it would guarantee payoff of 1 with certainty. To maximize y^L , sender A should increase x_A^L to μ , by Lemma 2. This leads to $y^L = \frac{\alpha_A \mu}{1 - \alpha_B \mu}$. However, the extra condition in the proposition means that y^L cannot reach t_A .

Other cases: No other non-local deviations are profitable. $\square$

Moving to the counterpart of Scenario 2, define

$$D_4(t_A) = \frac{t_A \alpha_B \mu}{\mu - t_A \alpha_A}, \text{ for } \mu < t_A < \frac{\mu}{\alpha_A}.$$

Proposition 7. *Assume conditions for Proposition 2. Under the following additional conditions, LPSNE described in Proposition 2 become PSNE:*

1.  no additional conditions,
2.  $D_4(t_A) < t_B$.
3.  no additional conditions.

Consequently, if $\mu < t_A < t_B$, $t_B < t^*$, and $t_B \leq D_4(t_A)$ then no PSNE exists.¹⁹

Proof. Case 2: When the strategy profile of this case is played, sender A can try to increase y^H up to t_A , which would be beneficial as it would guarantee payoff of 1 with certainty. The best chance to profit from this deviation is to flood the public square with the low signal, $x_A = (0, \mu)$. This strategy leads to $y^H = \frac{t_B \mu}{\alpha_A t_B + \alpha_B \mu}$, which reaches t_A as long as $t_B \geq D_4(t_A)$ for $t_A < \frac{\mu}{\alpha_A}$. If the inequality is reversed, this deviation is not possible.

Case 1 and 2: No other non-local deviations are profitable.

The rest of this proof deals with the last statement (new case 4): If $\mu < t_A < t_B$, $t_B < t^*$, and $t_B \leq D_4(t_A)$ then no PSNE exists. If $\mu < t_A < t_B$ then in any outcome, $y_L \leq t_A$. There are two possible kinds of PSNE: one in which $y^L < y^H < t_B$ (Sender B obtains payoff 0 with certainty), or $y^L < t_B \leq y^H$ (sender B obtains expected payoff between $(0, 1)$). Consider them in turn:

¹⁹Symmetrically, cases for $t_A < t_B < D_2(t_A)$ are captured as follows. Let $D_5(t_A) = \frac{\alpha_A(1-\mu)t_A + (1-t_A)\alpha_B\mu}{\alpha_A(1-\mu) + (1-t_A)\alpha_B}$. If $D_5(t_A) < t_B \leq D_3(t_A)$ then $x_A = (t_A, 1)$ and $x_B = (0, \mu)$ is a PSNE with the public posterior $y = \left(t_A, \frac{\alpha_A(\mu-t_A) + (1-t_A)\alpha_B\mu}{\alpha_A(\mu-t_A) + (1-t_A)\alpha_B}\right)$. If $t_A < t_B < \mu$ and $t_A < \frac{\alpha_B\mu}{\alpha_A(1-\mu) + \alpha_B}$, then $x_A = (0, 1)$ and $x_B = (\mu, 1)$ is a PSNE achieving $y^H = y^L = \mu$. If $t_A < t_B < \mu$, $\frac{\alpha_B\mu}{\alpha_A(1-\mu) + \alpha_B} < t_A$, and $t_B < D_5(t_A)$ then no PSNE exists.

Step 1: Suppose there is an equilibrium in which $y^L < y^H < t_B$. Sender B can deviate to attempt a fully informative strategy, $x_B = (0, 1)$, achieving the public posterior

$$y^H = \frac{\alpha_A (\mu - x_A^L) x_A^H + (x_A^H - x_A^L) \alpha_B \mu}{\alpha_A (\mu - x_A^L) + (x_A^H - x_A^L) \alpha_B \mu}$$

which is bounded below by $\frac{\mu}{\alpha_A + \alpha_B \mu}$ (achieved when $x_A^H = \mu$), which itself is greater than t_B . Hence this deviation guarantees $y^H > t_B$ and leads to a positive expected payoff for sender B . Hence, not an equilibrium.

Step 2. Suppose there is an equilibrium in which $y^L < t_B \leq y^H$. In this equilibrium, sender A considers two types of deviations: (i) can y^H be dropped to t_A or less, to ensure payoff of 1 with certainty, (ii) increase both y^L and y^H to increase Eu_A . Firstly, note that sender A must flood the public square with low signal realization, $x_A^L = \mu$. Suppose not: examining incentive (ii), we observe that $U_A(x_A^H, y^L)$ increases in y^L (in the domain $y^L \in [0, \mu]$, for given x_A^H), which can be increased by increasing x_A^L alone. Since $x_A^L = \mu$ and $x_A^H \geq \mu$, the public posterior must be

$$\begin{cases} y^L = \frac{\alpha_A (x_B^H - x_B^L) \mu + \alpha_B (x_B^H - \mu) x_B^L}{\alpha_A (x_B^H - x_B^L) + \alpha_B (x_B^H - \mu)} \\ y^H = x_B^H \end{cases}$$

Secondly, note that in this equilibrium, sender B 's strategy must involve $x_B^L = 0$ (lowering x_B^L lowers y^L and doesn't affect y^H which has an unambiguously positive effect expected utility). Thirdly, this equilibrium must have $x_B^H = t_B$ ($U_B(x_B^L, y^H)$ is decreasing in y^H , so in equilibrium $y^H \geq t_B$ should be as low as possible and so should x_B^H). Finally, we must check if sender A does not benefit from following deviation (i) in response to sender B playing $x_B = (0, t_B)$ in equilibrium. If this is what sender B plays, the high public posterior is

$$y^H = \frac{\alpha_A (\mu - x_A^L) t_B x_A^H + (x_A^H - x_A^L) \alpha_B \mu t_B}{\alpha_A (\mu - x_A^L) t_B + (x_A^H - x_A^L) \alpha_B \mu}$$

Sender A can lower it by setting $x_A^H = \mu$ and reach $y^H = \frac{t_B \mu}{\alpha_A t_B + \alpha_B \mu}$. This is not an equilibrium if $\frac{t_B \mu}{\alpha_A t_B + \alpha_B \mu} \leq t_A$, equivalent $t_B \leq D_4(t_A)$, which holds. $\square$

9.9 Proof of Proposition 3

Proof. Suppose there is a pure strategy equilibrium where $y^L < \mu < y^H$. Sender A expected utility from this outcome is strictly lower than $u_A(\mu)$ because this point is on the strictly concave part of the concave function $cou_A(\cdot)$. By Lemma 4, sender A can respond with a non-common-order strategy that generates utility $u_A(\mu)$. This cannot be a Nash equilibrium. Thus, a pure strategy equilibrium must lead to outcomes $y^L = \mu$ or $y^H = \mu$. However, for any $\alpha_B > 0$, sender B can deviate and make posteriors close to the prior but different than them, $y^L \neq \mu \neq y^H$. This deviation benefits sender B since her utility is strictly convex around point μ . Therefore, a pure strategy equilibrium does not exist. $\square$

9.10 Proof of Proposition 4

Proof. In any pure strategy equilibrium, there must be some information transmitted, i.e. $y^L \neq \mu \neq y^H$. For if we had $y^L = \mu$ or $y^H = \mu$, sender B could deviate and make posteriors close to the prior but different than them. Regardless of $\alpha_B > 0$, this deviation benefits sender B since her utility is strictly convex around point μ .

Therefore, suppose that in a pure strategy equilibrium, the outcome is $y^L < \mu < y^H$. The sender who wants to avoid communication, sender A , must use a non-common-order strategy, $x_A^L > \mu > x_A^H$ in which either $x_A^H = 0$ or $x_A^L = 1$. For if not, she could move y^L up and y^H down by moving simultaneously x_A^H or x_A^L (keeping p_H^A constant, by Lemma 3); this would be beneficial to her as her utility is strictly concave. Suppose in equilibrium $0 < x_A^H < \mu$, $x_A^L = 1$. In this case, a direct application of Lemma 2 implies that lowering x_A^H and keeping everything else constant will lower y^H and increase y^L . This would be strictly beneficial for A . Hence, in this equilibrium, $x_A = (1, 0)$. However, then sender B would prefer to match sender A 's strategy and select $x_B = (1, 0)$ resulting in public posterior $y = (1, 0)$. This contradicts the initial assumption that $y^L < y^H$. Other cases are similar. $\square$

9.11 Proof of Proposition 5

Proof. Sender B 's strategy $x_B = (0, 1)$ is dominant. Fix any strategy of A satisfying the common-order assumption $x_A^L \leq \mu \leq x_A^H$. Consider sender B 's strategy $0 < x_B^L \leq \mu$ and $\mu \leq x_B^H < 1$ and consider a more informative

strategy of B , so that p_B^s does not change as in Lemma 3. The resulting public posterior (y^L, y^H) is more informative. Since the utility function is convex, the expected utility is strictly greater. Hence, a strategy with either $x_B^L = 0$ or $x_B^H = 1$ strictly dominates the original one. To see that both equalities must hold, suppose w.l.o.g. that $x_B^L = 0$ and $x_B^H < 1$. By Lemma 2, increasing x_B^H increases y^H by the same-signal effect and decreases y^L by the cross-signal effect. Again, convexity of B 's utility implies this change is beneficial. Hence $x_B = (0, 1)$ is strictly dominant.

For point 2, a strategy of A such that $x_A^L < \mu < x_A^H$ is strictly dominated by $x_A = (\mu, \mu)$ and the same probability of signals p_A^s as in the original strategy. Fix any strategy of B . Posteriors x_A^L can be increased and x_A^H decreased towards μ , to keep p_A^s constant as in Lemma 3. This leads to a higher y^L and a lower y^H . Given that A 's utility is strictly concave, the payoff strictly increases.

For the final point, since in any equilibrium we have $x_A = (\mu, \mu)$ and $x_B = (0, 1)$, the only remaining variable to construct an equilibrium is p_A^L . Sender A has a well-defined maximization problem of a continuous expected concave utility function on a unit interval and hence a solution exists. $\square$

References

(tbc)

Atakan, Alp, Mehmet Ekmekci, and Ludovic Renou (2024), “Cross-verification and persuasive cheap talk”, *Journal of Economic Theory*, 222

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), “Competition in Persuasion”, *Review of Economic Studies*, 84, 300–322.

Gilligan, Thomas W, and Keith Krehbiel (1989), “Asymmetric information and legislative rules with a heterogeneous committee.” *American Journal of Political Science*, 33: 459–490.

Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion”, *American Economic Review*. Vol. 101, No. 6, 2590–2615.

Li, Fei and Peter Norman (2018), “On Bayesian persuasion with multiple senders”, *Economic Letters*, 170, 66–70.

Mylovanov, Tymofiy and Andriy Zapechelnjuk (2024), “Constructive Vs

Toxic Argumentation in Debates”, *American Economic Journal: Microeconomics*, vol. 16, no. 1, 262–92.

Ravindran, Dilip and Zhihan Cui (2022), “Competing Persuaders in Zero-Sum Games”, available at <https://arxiv.org/abs/2008.08517>