

Communicating with Anecdotes*

Nika Haghtalab
UC Berkeley

Nicole Immorlica
Microsoft Research

Brendan Lucier
Microsoft Research

Markus Mobius
Microsoft Research

Divyarthi Mohan
Tel Aviv University

July 2, 2024

Abstract

We study a communication game between a sender and receiver. The sender chooses one of her signals about the state of the world (i.e., an *anecdote*) and communicates it to the receiver who takes an action affecting both players. The sender and receiver both care about the state of the world but are also influenced by personal preferences, so their ideal actions can differ. We characterize perfect Bayesian equilibria. The sender faces a *temptation to persuade*: she wants to select a biased anecdote to influence the receiver’s action. Anecdotes are still informative to the receiver (who will debias at equilibrium) but the attempt to persuade comes at a cost to precision. This gives rise to *informational homophily* where the receiver prefers to listen to like-minded senders because they provide higher-precision signals. Communication becomes *polarized* when the sender is an expert with access to many signals, with the sender choosing extreme outlier anecdotes at equilibrium (unless preferences are perfectly aligned). This polarization dissipates all the gains from communication with an increasingly well-informed sender when the anecdote distribution is heavy-tailed. Experts can therefore face a *curse of informedness*: receivers will prefer to listen to less-informed senders who cannot pick biased signals as easily.

JEL Classification: D82, D83

Keywords: anecdotes, communication, sender-receiver games, memory, language

*This work is partially supported by the National Science Foundation under grant CCF-2145898, the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 866132), the Office of Naval Research under grant N00014-24-1-2159, an Alfred P. Sloan fellowship, a Schmidt Sciences AI2050 fellowship, and a Google Research Scholar award. Some parts of this work were done when the authors visited the Simons Institute for the Theory of Computing and when Haghtalab and Mohan were employed at Microsoft Research, New England. We are grateful to Nageeb Ali, Lukas Bolte, Vincent Crawford, Ben Golub, Matt Jackson, Pooya Molavi, Harry Pei, Joel Sobel, Leeat Yariv and seminar participants at University of Chicago, Northwestern, Penn State, Stanford Institute for Theoretical Economics (SITE), Stony Brook Game Theory Festival, and Virtual Seminars in Economic Theory, among others, for helpful comments.

Half the truth is often a great lie.

– *Benjamin Franklin*

1 Introduction

As an economist, explaining the beginning of a recession to a friend who isn’t familiar with economic jargon can be difficult. While the economist usually analyzes trends using metrics like GDP growth and investment levels, her friend might have a different approach, such as noticing whether specific companies are growing or reducing their workforce. Both approaches can be valid representations of the same underlying state of the world, but the mismatched representations can create a communication challenge.

One effective approach involves both the sender and receiver reaching a consensus on a shared representation. Crafting such a shared “language” is a practice commonly adopted by experts within a given field, enabling them to communicate efficiently. A different method entails the sender transmitting raw data (or “anecdotes”) to the receiver, allowing them to interpret and map it to their own internal representation. Such *anecdotal communication* is common when the sender or receiver are non-experts. Newspapers, for example, assemble articles from factual snippets such as reporting the quarterly job creation figures¹, highlighting a recent high-profile layoff at a major tech company², or conducting a series of interviews with individuals who have recently resigned from their jobs³. Similarly, politicians often try to appeal to a broad spectrum of voters by incorporating stories into their speeches, such as initiatives that helped their districts in the past.

Anecdotal communication is versatile but less efficient than communication between experts because the sender cannot simply summarize diverse data points into a single number. The sender must select among the many anecdotes she could potentially report. For example, the New York Times must sift through a vast array of possible anecdotes from the work of its journalists and wire services to determine what to include in a 500-word article on its homepage. This need for selection creates a *temptation* for the sender to influence the receiver. For instance, a current officeholder may predominantly incorporate favorable anecdotes about the economy into their speeches instead of the most representative ones, motivated by their bid for re-election. Conversely, their opponent may utilize anecdotes to portray a less optimistic outlook, aiming to cast doubt on the incumbent’s abilities. This *persuasion temptation* makes anecdotal communication more noisy: a rational receiver can correct for the sender’s bias but not for the greater noise that is embedded when listening to less representative anecdotes.


In this paper, we build a simple model of anecdotal communication that explores the tension between the sender choosing representative anecdotes that are most informative and tail anecdotes that are meant to persuade the receiver to take a more preferred action.

¹E.g., WSJ article “November Employment Report Shows U.S. Economy Added 263,000 Jobs” (Chaney Cambon, 2022)

²E.g., NYTimes article “Meta Lays Off More Than 11,000 Employees” (Frenkel et al., 2022)

³E.g., NYTimes article “We Revisited People Who Left Their Jobs Last Year. Are They (Still) Happy?” Needelman (2023)

Figure 1: Illustration of translation invariant equilibrium



1.1 Summary of Results

We summarize the key findings using the most straightforward version of our model as shown in Figure 1. There are two agents - a sender S and a receiver R . The receiver takes an action that affects both agents but the state of the world θ is only observed by the sender. The sender would like the receiver to take action $a_S^* = \theta + M_S$ (shown in blue) while the receiver's preferred action is $a_R^* = \theta + M_R$ (shown in red). The offsets M_S and M_R indicate the agents' *personal preferences* and in the above example we have $M_S < M_R$ such that the sender would always like the receiver to take an action slightly more to the left. The sender has access to a set of signals which are drawn from a distribution that is centered around the state θ (shown as gray circles). We call these signals *anecdotes* because the sender is constrained in her communication with the receiver and can only send one such anecdote to the receiver.

Targeting Equilibrium. We will focus on *translation invariant* strategies where the receiver takes action $x + \sigma$ after observing the sender's anecdote x . Intuitively, the receiver acts as if he knows the sender's bias and corrects for it; for example, the receiver knows that a news publisher will tend to select more left-leaning anecdotes and correct for this bias by choosing an action further to the right ($\sigma > 0$). In equilibrium, the sender would therefore ideally want to send an anecdote exactly equal to $\theta + r$ where $r = M_S - \sigma$, as this would ensure that the receiver takes the sender's preferred action. However, the sender will generally not have this exact anecdote available. Our first main result is to show that she will select the anecdote closest to $\theta + r$ as the second-best alternative (see Theorem 1). Next, we observe that the anecdote closest to $\theta + r$ (shown as a green circle here) is more likely to lie to the right of the offset than to the left because the density of the

anecdote distribution is higher there. Therefore, the expected value of the sent anecdote is equal to $\theta + \beta(r)$ where $\beta(r) > r$ is the *bias* of the sender’s communication. In equilibrium, the receiver will choose his shift σ to undo this bias and then take her preferred action such that $\sigma = -\beta(r) + M_R$. Combining the sender’s best response condition with the receiver’s, we obtain:

$$\begin{aligned} r &= M_S - \underbrace{(M_R - \beta(r))}_{\sigma} \\ \beta(r) - r &= \underbrace{M_R - M_S}_{\Delta} \end{aligned} \tag{1}$$

This is the key equilibrium condition in our model and all our main results follow from it: for a given difference Δ in the personal preferences of receiver and sender there exists a *targeting communication scheme* with offset r satisfying the above equation such that the sender will choose the anecdote closest to $\theta + r$ (see Theorem 5).

Informational Homophily. When the sender and receiver have common interests, i.e., $M_R = M_S$, the above condition boils down to $\beta(r) = r$. For a single-peaked and symmetric distribution around the state θ , as shown in Figure 1, this implies that the sender will choose offset $r = 0$ and therefore select the anecdote closest to the θ . Intuitively, this anecdote is the most representative and conveys the most information to the receiver. When the sender is more left-leaning than the receiver ($\Delta > 0$) we instead obtain an equilibrium offset $r < 0$ and the sender will select biased anecdotes with a left-leaning bias. While the receiver is able to correct for this bias, he is worse off because these anecdotes have higher variance. In Section 4, we explore how this gives rise to *informational homophily* where both the sender and receiver prefer to talk to like-minded partners because lack of persuasion temptation makes the sender minimize the information loss from anecdotal communication.

Polarization. Unlike models in the cheap talk and rational persuasion literature, the sender in our model is constrained in her message space to select an existing anecdote (Crawford and Sobel, 1982; Kamenica and Gentzkow, 2011). This enables us to define the degree of polarized communication by the offset r of the sender’s targeting communication scheme: the further to the left (or right) this offset lies, the more polarized is the sender’s communication. As explained in the example above, polarization increases with the amount of misalignment between the sender and receiver (as measured by $|\Delta|$). However, in Section 5 we show that polarization also increases with the number of anecdotes available to the sender: the more informed the sender is, the more polarized her communication becomes. Intuitively, the information loss decreases with the number of anecdotes at any given offset r (formally, $\beta(r) - r$ decreases). This makes it less costly for the sender to choose a more extreme anecdote and the equilibrium targeting scheme will become more polarized. Interestingly, while any misalignment between the sender and receiver will generate some degree of polarization this phenomenon is more pronounced for more informed experts (i.e., senders with access to more anecdotes).

Curse of Informedness. Section 5.2 explores whether a receiver prefers to listen to a less informed but more like-minded sender or a better-informed expert. While communication with the expert is more polarized (which lowers utility) the expert is also better informed (which lowers the information loss). We show that for heavy-tailed distributions (with tail densities declining at a less than exponential rate) receivers prefer to seek out non-experts. Intuitively, experts in such an environment have access to too many outlier anecdotes such that the variance in the sent anecdote increases with the informedness of the expert. We say that such experts suffer from the *curse of informedness* - they have an incentive to collect fewer anecdotes in order to be listened to more by less aligned receivers.

Commitment. Section 6 analyzes a version of our model where the sender can commit to a targeting communication scheme. Commitment eliminates persuasion temptation and therefore gives rise to most informative communication schemes that maximize welfare. For example, in the example above the sender will send the most representative anecdote which is the anecdote closest to the posterior mean.

1.2 Related Work

Our model is related to the economic literature on strategic communication which includes both cheap talk games and verifiable disclosure games (also called persuasion games). In the canonical paper in the literature on cheap talk games, Crawford and Sobel (1982) consider a setting similar to ours in which the sender and receiver have state-dependent but misaligned preferences. The major difference in their model is that the sender is unrestricted in what signal she can send and hence can pool states arbitrarily. The authors characterize the set of perfect Bayesian equilibria and show the most informative of these equilibria pools only nearby states. The resulting precision of the receiver's posterior belief, interpreted as a coarseness of the message space or vagueness of the chosen language in equilibrium, depends on the degree of misalignment between the sender and receiver's preferences. We observe a similar phenomenon in our setting: more aligned preferences induce more information transmission in equilibrium. Our results however point to a different driver of this phenomenon. Namely, the friction that prevents communication from devolving to the least informative signal is the exogenous limitation of communicating anecdotes rather than the endogenous equilibrium choice of language coarseness.

Our exogenous limitation of communicating anecdotes is reminiscent of the literature on verifiable disclosure games or persuasion games, introduced by Grossman and Hart (1980), Grossman (1981), and Milgrom (1981) (see Milgrom (2008) for a survey of this literature). These papers consider the setting of a seller who can choose whether to disclose information about a product to a buyer and wishes to maximize the buyer's posterior belief about the value or quality of the product. Similar to our model of communication via anecdotes, the seller in these papers cannot arbitrarily distort information about the product. The classic results show that in every perfect Bayesian equilibrium, the seller fully discloses her information or, if she's limited in how much information

she can disclose, reveals the most favorable information. In our model, this sort of unraveling to extreme signals is tempered by the fact that the sender’s preferences are state dependent and so she does not simply wish to maximize the receiver’s belief. Other mechanisms that limit unraveling include the presence of naive receivers (Kartik et al. (2007)), exogenous costs to misrepresentation (Kartik (2009)), uncertainty in the informedness of the sender (Dye (1985); Jung and Kwon (1988); Dziuda (2011)), multi-dimensionality of the state (Martini (2018)), or the introduction of commitment power (discussed below), among others.

Much of the literature on strategic communication studies the impact of commitment power. Kamenica and Gentzkow (2011) characterize the optimal signaling scheme of a sender with commitment power in cheap talk games. Their characterization implies that a sender with quadratic loss (as in our model) who is not restricted to communicating anecdotes would communicate her belief about the state of the world thereby completely avoiding the uninformative babbling equilibrium of Crawford and Sobel (1982). In Section 6, we show similarly that, subject to the restriction of sending an anecdote, the sender wishes to communicate the most informative anecdote. This result suggests that in our model, commitment power can move the equilibrium from a partially informative one to the most informative one. Other recent work demonstrates investigates commitment power in verifiable disclosure games, studying the impact of partial commitment (Lipnowski et al. (2022); Min (2021); Nguyen and Tan (2021); Lin and Liu (2022)) or characterizing conditions under which commitment has limited impact on achievable outcomes in equilibrium (Glazer and Rubinstein (2008); Hart et al. (2017); Sher (2011); Zhang (2022)). Importantly, in all of these models the sender’s incentives are purely to persuade, with utility that depends on the receiver’s action but not the state of the world. In our model the sender has an incentive to inform as well as to persuade. Thus in light of the aforementioned literature, our results can be interpreted as showing that a desire to persuade can lead to a significant loss of communication fidelity (and hence welfare) in the absence of commitment. Our perfect Bayesian equilibrium characterization in Section 4 shows that this loss can happen and be significant with even with a small desire to persuade (i.e., even with a small preference misalignment).

2 A Model of Anecdotal Communication

In this section we formally introduce our model and characterize the sender’s and receiver’s best response strategies which will allow us to characterize the equilibria of the game in Section 4.

2.1 Model Setup

We consider a communication game played by two players, a sender (“she”) and a receiver (“he”). The sender has information about a payoff-relevant state of the world $\theta \in \mathbb{R}$ drawn from a common prior. The receiver, in turn, chooses a payoff-relevant action $a \in \mathbb{R}$ (for example, how much to invest in the current market).

Preferences. Players' preferences over their actions depend on the state of the world. However, their preferences can differ.

We model this by introducing personal preferences $M_R \in \mathbb{R}$ and $M_S \in \mathbb{R}$ for the receiver and sender respectively, which are shifts of the ideal action relative to the state of the world. More formally, the receiver's utility is

$$u_R(a, \theta) = -(a - (\theta + M_R))^2,$$

and the sender's utility is

$$u_S(a, \theta) = -(a - (\theta + M_S))^2.$$


We assume personal preferences are publicly known and write $\Delta = M_R - M_S$ for the known difference in personal preferences. Intuitively, Δ captures the preference misalignment between sender and receiver.

Sender's knowledge. The sender has access to noisy signals about the state of the world that she can potentially share with the receiver. Given a distribution F over the reals, we model these shareable signals as a set of n samples x_1, \dots, x_n where each $x_i = \theta + \epsilon_i$ for $\epsilon_i \sim F$ drawn independently. We will write $\vec{x} = (x_1, \dots, x_n)$ for the profile of samples which we will refer to as *anecdotes* from now on. We think of these anecdotes as immutable facts about the world which the sender can decide to share, but which she cannot otherwise manipulate. For example, the receiver might not know about the survey or research paper until the sender chooses to reveal it but he can subsequently look up the survey or paper and fact-check it. While \vec{x} is known only to the sender, we assume the anecdote distribution F as well as the number of anecdotes, n , is common knowledge.⁴

The sender might have additional information that cannot be easily shared or fact-checked at low cost. For example, the sender's knowledge about the state of the world might be informed by her own detailed research and modeling efforts. We model such side information as an additional signal $y \in \mathbb{R} \cup \{\emptyset\}$. Given a distribution G over $\mathbb{R} \cup \{\emptyset\}$ and $\gamma \sim G$, the sender has either access to no additional information (if $\gamma = \emptyset$) or to a signal $y = \theta + \gamma$ (if $\gamma \in \mathbb{R}$). Most of our intermediate results hold for general distributions G . However, we pay special attention to two cases: the *foresight* setting where $\gamma = 0$ with probability 1, and the sender has full information about the state of the world; second, the setting with *no foresight*, where $\gamma = \emptyset$ and the sender has no information beyond the set of anecdotes. As with \vec{x} , we assume that y is private knowledge of the sender but that the distribution G is common knowledge. The sender uses her anecdotes \vec{x} and side information y to form a posterior belief over the state of the world θ . We will denote by $\theta_S(\vec{x}, y)$ the posterior mean of θ given (\vec{x}, y) .

⁴This shuts down a common pathway for partial information transmission: in our model, there is no uncertainty about *how much* information the sender has.

Figure 2: Timing of communication game



Communication. The sender communicates exactly one anecdote in \vec{x} to the receiver. While the anecdote is communicated honestly, the sender can cherry-pick from the set of anecdotes she has access to. Note, that the sender’s side information y cannot be communicated – only an anecdote can be shared. A strategy for the sender in our game is a communication scheme $\pi: \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}$ that maps every realization of n anecdotes \vec{x} and side information y to a choice of one of the n anecdotes. In particular, for all \vec{x} and y we have $\pi(\vec{x}, y) = x_i$ for some $i \in [n]$.

Equilibrium. Since the receiver does not observe the choice of nature, a strategy for the receiver is an action rule $\alpha: \mathbb{R} \rightarrow \mathbb{R}$ that maps the sender’s chosen anecdote to a choice of action. The timing of our game is shown in Figure 2. In round 0, nature chooses state θ , anecdotes x_1, \dots, x_n , and signal y . In round 1, the sender selects anecdote $\tilde{x} = \pi(\vec{x}, y)$; this choice is observed by the receiver. In round 2, the receiver selects action $\alpha(\tilde{x})$. Payoffs are then realized as described above. Given a communication scheme π , we will write $D_{\pi, x}$ for the posterior distribution of θ given that $\pi(\vec{x}, y) = x$. We are interested in the perfect Bayesian equilibrium of the game, that is, strategies for the sender and receiver that maximize payoffs under their (consistent) beliefs.⁵

Definition 1. A pair of strategies (π^*, α^*) , together with a belief function $B: \mathbb{R} \rightarrow \Delta(\mathbb{R})$ for the receiver mapping every observation to a distribution over the state of the world, form a perfect Bayesian equilibrium if:

1. For each observed anecdote x , action $\alpha^*(x)$ maximizes expected receiver utility given distribution $B(x)$ over θ , i.e.,

$$\alpha^*(x) \in \arg \max_a \left\{ \mathbb{E}_{\theta \sim B(x)} [u_R(a, \theta)] \right\}.$$

2. B is the rational belief with respect to π^* . That is, for each x , $B(x) = D_{\pi^*, x}$ is the posterior distribution of θ given that $\pi^*(\vec{x}, y) = x$.

⁵The perfect Bayesian equilibria we consider will not involve zero-probability events and will therefore also be sequential equilibria.

3. For each \vec{x} and y , $\pi^*(\vec{x}, y)$ maximizes sender utility given α^* , i.e.,

$$\pi^*(\vec{x}, y) \in \arg \max_{x_i \in \vec{x}} \left\{ \mathbb{E}_{\theta} [u_S(\alpha^*(x_i), \theta) \mid (\vec{x}, y)] \right\}.$$

Given communication scheme π , we denote by α_{π} the action rule that satisfies requirements (1) and (2) of Definition 1 and call it the *best response* to π .

Diffuse Prior. Expected payoffs in our game are driven by the posterior beliefs of the receiver and sender regarding the state of the world θ . These posterior beliefs are formally defined with respect to a prior distribution over θ . As we are interested in the setting where agents are initially uninformed about the state of the world, we will define posterior beliefs with respect to a *diffuse prior* over θ . That is, the common prior over θ reveals no information about it.

Assumption 1. *The prior over θ is diffuse, i.e., it is $N(0, \infty)$.*

We note that the exact form of the prior over θ is not important so long as it is a diffuse prior that reveals no information about θ , i.e., its density is almost uniform everywhere.

Of course, the diffuse prior is not a valid probability distribution and is not formally defined. We emphasize that it is being used to define posterior beliefs given the information available to the sender and receiver.⁶ As has been noted by Ambrus and Kolb (2021), one must take care when defining ex-ante payoffs in any game involving a diffuse prior. They describe sufficient conditions for ex ante payoffs to be well-defined and consistent with any appropriate limiting sequence of proper prior distributions. In Appendix A we verify that ex ante payoffs are likewise well-defined in our model and discuss further implications of Assumption 1, including the precise form of the posterior beliefs.

2.2 Characterizing Best Response Strategies

Our analysis focuses on a class of *translation invariant* strategies for both the sender and the receiver. Intuitively, a translation invariant communication scheme encodes a sender’s *communication posture* such as “always send the left-most anecdote” or “always send the right-most anecdote” that does not depend on the particular realization of the state of the world. This is an appealing property in the context of our diffuse prior assumption where specific numbers have no special meaning.

Similarly, a translation invariant action rule describes the receiver’s action as a fixed offset from the received anecdote. Such an action rule describes a receiver who simply believes the anecdote he hears is representative of the state, albeit potentially with a shift. We will show in this section that

⁶One could equivalently view the realization of θ as a (non-stochastic) choice of nature, and interpret Assumption 1 as a behavioral assumption that the agents form posterior beliefs consistent with having no prior knowledge of θ .

translation invariance is internally consistent in the sense that the best response of the sender to a translation invariant action rule is a translation invariant communication scheme and vice versa.

We formally define translation invariant action rules as follows.

Definition 2. *An action rule α is a translation if $\alpha(x + \delta) = \alpha(x) + \delta$ for all x and all $\delta \in \mathbb{R}$.*

Note that if action rule α is a translation, then there is a value $\sigma \in \mathbb{R}$ such that $\alpha(x) = x + \sigma$ for all x . We refer to σ as the *shift* of α , written $\sigma(\alpha)$. These receivers act as if they know the typical bias of a sender, e.g., a receiver who thinks the New York Times, being a slightly left-of-center paper, sends anecdotes shifted slightly left.

We next formally define translation invariant communication schemes. Given a profile of anecdotes \vec{x} and a constant $\delta \in \mathbb{R}$ we will write $\vec{x} + \delta$ for the shifted profile of anecdotes $(x_1 + \delta, x_2 + \delta, \dots, x_n + \delta)$.

Definition 3. *A communication scheme π is translation invariant if $\pi(\vec{x} + \delta, y + \delta) = \pi(\vec{x}, y) + \delta$ for all \vec{x}, y and all $\delta \in \mathbb{R}$.⁷*

Not all communication schemes are translation invariant. For example, the communication scheme that sends the anecdote closest to zero is not translation invariant, nor is the one that sends the minimum anecdote if that anecdote is an even number and the maximum otherwise. However, many natural communication schemes are translation invariant. The top panel of Figure 3 shows the simplest example: the *minimum scheme* $\pi_{\min}(\vec{x}, y)$ selects the minimum anecdote (analogously, the *maximum scheme* $\pi_{\max}(\vec{x}, y)$ selects the maximum anecdote).

A particularly important class of translation invariant communication schemes are those where the sender selects a signal that is closest to a shift r from her posterior mean. We call these *targeting schemes*.

Definition 4. *The targeting scheme with offset $r \in \mathbb{R}$ is a communication scheme that always returns the anecdote from \vec{x} that is closest to $\theta_S(\vec{x}, y) + r$.*

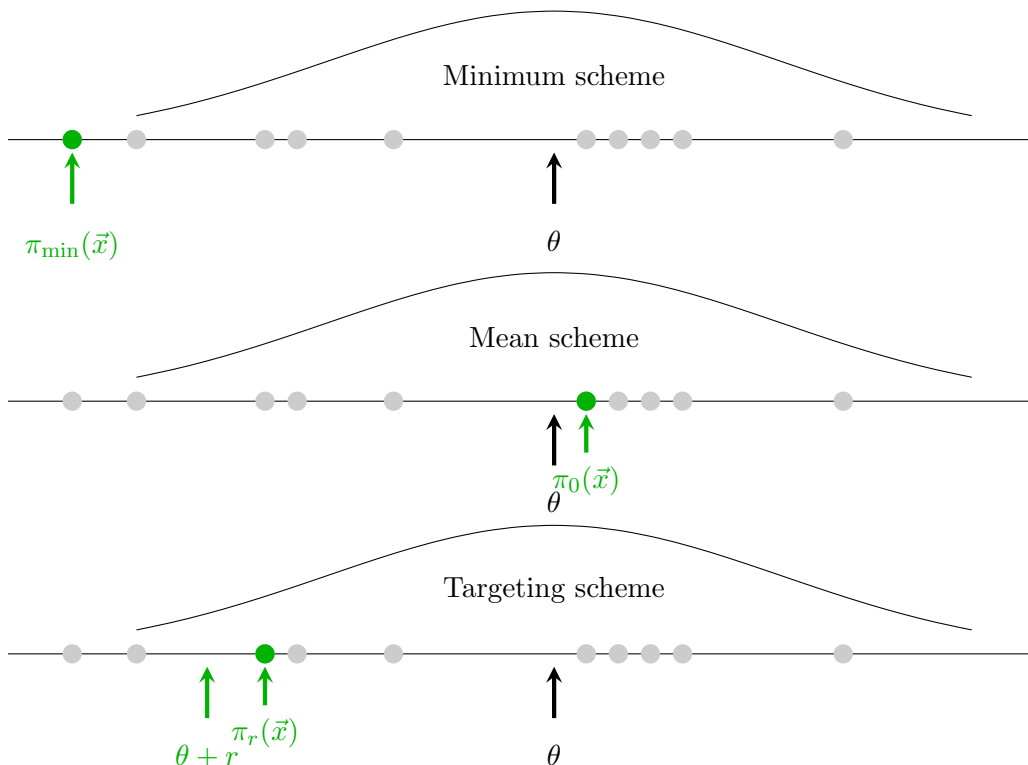
Note that since θ is drawn from a diffuse prior, we have $\theta_S(\vec{x} + \delta, y + \delta) = \theta_S(\vec{x}, y) + \delta$ (formalized in Appendix A). Hence a targeting scheme is translation invariant. The middle panel of Figure 3 shows the *mean targeting scheme* for $r = 0$ where the sender selects the signal closest to the posterior mean. The bottom panel illustrates a targeting scheme with $r < 0$. Note that as $r \rightarrow -\infty$ the targeting scheme approaches the minimum scheme.

It will be helpful to define the *bias* of a translation invariant communication scheme. Given a translation invariant communication scheme π , we will say the bias of π , $\beta(\pi)$, is equal to $\mathbb{E}_{\theta, \vec{x}, y}[\pi(\vec{x}, y) - \theta]$. If the anecdote distribution F is symmetric around 0 then the minimum scheme has a left bias ($\beta(\pi_{\min}) < 0$) and a targeting scheme π_r has a left bias for $r < 0$ and a right bias for $r > 0$ (and will be unbiased for $r = 0$).

We can now state the main result: the sender's best response to a translation invariant action rule is not just translation invariant but also a targeting scheme.

⁷We assume that $y + \delta = \emptyset$ for $y = \emptyset$.

Figure 3: Examples of translation invariant communication schemes



Theorem 1. *If action rule α is a translation, then the best response of the sender is translation invariant. More specifically, it is the targeting scheme with offset $M_S - \sigma(\alpha)$.*

Note that Theorem 1 characterizes the sender's best response among all (not necessarily translation invariant) communication schemes. All the proofs in this section are relegated to Appendix B. The intuition for the result is as follows: assume that the sender had exactly an anecdote at distance r from the posterior mean. The recipient will then take action $a = \theta_S(\vec{x}, y) + \sigma(\alpha)$. If $r = M_S - \sigma(\alpha)$ we get $a = \theta_S(\vec{x}, y) + M_S$ which is exactly the loss-minimizing action from the sender's perspective. In general, the sender will not have exactly an anecdote at the offset available and so she chooses the closest anecdote as a second-best – for this reason, the sender's best response is a targeting scheme.

Also, note that the sender's best response is generally unique *except* when there is foresight and the anecdote distribution F is bounded. For example, if \underline{x} is the lower bound of the domain of F and $\sigma(\alpha) \leq \underline{x}$ then the sender's optimal targeting scheme is the minimum scheme and any $r < \underline{x}$ will be optimal (including the one prescribed by Theorem 1).

We close out our analysis of translation invariant strategies by showing that the best response to a translation invariant communication scheme is a translation invariant action rule.

Theorem 2. *For any translation invariant communication scheme π , the best response of the*

receiver to π is a translation with shift $M_R - \beta(\pi)$.

The intuition is as follows. The receiver knows that the sent anecdote is on average a distance $\beta(\pi)$ away from the sender’s posterior mean. Hence, he will try to undo this bias by subtracting it from the sent anecdote which provides her with an unbiased estimate of the state of the world. He will then minimize her loss by taking an equal to this estimate plus her personal preference M_R .

We have now established that sender and receiver translation invariant strategies are self-consistent in the sense that they are best responses to each other. We have not yet shown that equilibria with translation invariant strategies exist. We will show their existence in Section 4.

2.3 Broadcasting

Senders often communicate anecdotes to several receivers simultaneously. Experts or politicians, for instance, can appear on TV or are cited in newspapers, reaching many viewers or readers. Our model can easily be extended to allow for *broadcasting* to multiple receivers whose personal preference M_R is drawn from a distribution G over the real numbers. Each type M_R is taking an action $a(\tilde{x}|M_R)$ after observing anecdote \tilde{x} which is “broadcast” to all receiver types. The sender cares about the *mean* action of the receivers:

$$\bar{a} = \int a(\tilde{x}|M_R)dG \tag{2}$$

We denote the mean personal preferences of the receivers with \overline{M}_R :

$$\overline{M}_R = \int M_R dG \tag{3}$$

Our notion of translation invariant communication schemes readily generalizes to this setting.

Proposition 1. *In any translation invariant communication scheme of the broadcast model, the best response of the receiver type M_R is a translation with shift $\sigma(M_R) = M_R - \beta(\pi)$. The sender’s best response is the targeting scheme with offset $M_S - \bar{\sigma}$ where $\bar{\sigma} = \int \sigma(M_R)dG$ is the mean shift across receiver types.*

Therefore, the sender in the broadcast model behaves as if she is communicating with a single “representative” receiver, whose personal preference matches the average personal preferences of the entire audience.

2.4 Discussion of Modeling Assumptions

Crawford-Sobel (1982). Our model follows Crawford and Sobel (1982) except that we restrict communication to anecdotes. This restriction is most appropriate in situations where the sender or receiver are non-experts – for example, newspapers communicating news to readers or politicians talking to voters. Experts often share a *common language* that allows them to communicate beliefs efficiently. For example, experts on infectious diseases use the reproduction factor R_0 to

describe how quickly a disease such as measles can spread between people. An expert sender can therefore summarize her knowledge about measles, for example, by communicating this factor to the receiver. A parent, however, might find it easier to understand infectiousness by being told that a neighboring school had to cancel classes because too many teachers called in sick. A journalist might therefore relate the expert’s message by listing examples of school districts who experienced measles outbreaks and interviews with local school principals and teachers.

Common Knowledge. The assumption that sender’s preference is publicly known is justifiable in settings where the sender is a known entity, say a politician or newspaper. In such settings, the sender is often communicating with a known distribution of receiver types – the general public for instance. Our results would follow largely unchanged if the receiver’s preference is drawn from a known distribution.

No “Fake News”. We assume that anecdotes cannot be manipulated or falsified but can only be selected in a possibly biased manner. The non-falsification assumption is reasonable in public discourse: politicians and newspapers by-and-large report facts, or else risk being caught by fact-checkers. However, they have editorial control over the selection of those facts and can influence the listener this way. Unlike lying about and making up facts, biased selection of facts is typically accepted in public discourse. The non-manipulation assumption is stronger: certainly two senders can sometimes frame two facts differently. A full analysis of framing exceeds the scope of our paper but we note that our framework can accommodate some natural forms of framing: for example, if the sender can add a “spin” to the sent anecdote \tilde{x} by shifting its value by a fixed constant s then the receiver observes $\tilde{x} + s$ and can undo this spin as long as s is common knowledge (which is natural if the sender’s personal preference is also common knowledge).

Commitment. In Section 6 we will study a variant of the game where the sender can *commit* to a communication scheme. A sender may be able to commit in settings where her reputation precedes her. In this case an equilibrium does not need to meet requirement (3) of Definition 1, i.e., $\pi^*(\vec{x}, y)$ does not need to maximize expected sender utility given $\alpha^*(x)$. Rather, for every π , we fix a best response α_π of the receiver. We then require that

3'. π^* is such that

$$\pi^* \in \arg \max_{\pi} \left\{ \mathbb{E}_{\theta, \vec{x}, y} [u_S(\alpha_\pi(\pi(\vec{x}, y)), \theta)] \right\}. \quad (4)$$

Model of Memory. We can reinterpret our framework as a model of *memory* where the current self (sender) decides which data point to commit to her long-term memory that is available to her future self (receiver). This interpretation is particularly intuitive for the case of commitment or where the current and future self share the same personal preferences ($\Delta = 0$).

3 Most Informative Communication and Persuasion Temptation

Before characterizing the equilibria of the game it is helpful to highlight a fundamental tension in the sender’s objective: on the one hand she wants to send an informative anecdote so that the

receiver takes an action that reflects the state of the world. On the other hand, if the anecdote is too precise then the receiver will take an action close to his ideal action $\theta + M_R$, which differs from the sender’s preferred action unless their preferences are aligned.

We call this tension the sender’s *persuasion temptation*. To understand it better, it is helpful to decompose the sender’s utility into two components.

Proposition 2. *Suppose the receiver’s action rule α is a translation with shift $\sigma(\alpha)$. Then the sender’s expected utility from any translation invariant communication scheme π_S is*

$$-\underbrace{\mathbb{E} [(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)))^2]}_{\text{information loss}} - \underbrace{[\sigma(\alpha) - (M_S - \beta(\pi_S))]^2}_{\text{disagreement loss}}. \quad (5)$$

The first component of this decomposition, $\mathbb{E}[(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)))^2]$, is the variance of the communicated anecdote as $\theta_S(\vec{x}, y) + \beta(\pi_S) = \mathbb{E}[\pi_S(\vec{x}, y)]$ by the definition of bias. We interpret this variance as the inherent *information loss* of the communication scheme because this loss would be incurred even if the sender sends this anecdote to herself. The second term captures the disagreement between the receiver’s shift $\sigma(\alpha)$ and the sender’s preferred shift for the receiver which equals $M_S - \beta(\pi_S)$.⁸

3.1 Most Informative Communication Schemes

We can better understand the tension in minimizing both the information loss and the disagreement loss by first characterizing communication schemes that minimize *only* the information loss. This is the most informative communication scheme for a sender who is choosing the action herself based on a single recalled anecdote. For example, a sender who can only keep a single anecdote in memory for her future self solves this problem.


It turns out that most informative communication schemes are always targeting schemes that have to satisfy a simple necessary condition, namely that the offset r^* of the targeting scheme has to equal its bias $\beta(r^*)$ which we call the “balanced-offset” condition.

Theorem 3. *Any translation invariant communication scheme that minimizes the sender’s information loss is a targeting scheme that satisfies the balanced-offset condition $r^* = \beta(r^*)$. Moreover, such a targeting scheme exists.*

The proofs of this section are relegated to Appendix C. We can show that the balanced-offset condition implies that the most informative targeting scheme will have the sender choose an offset close to the global maximum of the density distribution. We now develop this intuition for the special case of foresight (where the sender knows the state of the world). Section 6 extends it to the no-foresight case.

⁸The sender would like the receiver to choose an action at offset M_S after compensating for the bias $\beta(\pi_S)$ in the communication scheme.

Figure 4: Examples of most informative offset r^* under foresight



Green shading shows symmetric density of closest anecdote to offset r^* .

Lemma 1. *For symmetric anecdote distribution F with a peak at 0 the most informative targeting scheme for any number of anecdotes n under foresight is the mean scheme.*

Figure 4a shows the normal anecdote distribution $F \sim N(0, 1)$. Because of symmetry, the closest anecdote is equally likely to lie to the left or right of the offset $r = 0$ (indicated by green shading) and therefore the targeting scheme has bias $\beta(0) = 0$. Moreover, no offset $r < 0$ can implement a most informative targeting scheme because the closest anecdote is more likely to the right of the offset than to the left and therefore $\beta(r) > r$. Analogously, no offset $r > 0$ can implement such a communication scheme either.


Lemma 1 is intuitive because it places the most informative offset at the point where the density of anecdote distribution is the highest. Anecdotes are most densely distributed around the peak and the closest anecdote to $r^* = 0$ is therefore a better predictor of the state of the world compared to the closest anecdote to any other offset $r \neq 0$. This argument carries over to the case of symmetric single-peaked distributions that are not centered at 0. For example, Figure 4b shows an normal anecdote distribution centered around 2 which shifts the most informative offset to the new peak $r^* = 2$.

The general intuition that the most informative offset is close to the maxima of the density distribution even holds when the density has more than one mode and is not symmetric.

Lemma 2. *Assume the sender has foresight, the anecdote density f is uniformly continuous on its domain and fix any $\epsilon > 0$. There exists an \bar{n} such that for any number of anecdotes $n > \bar{n}$ the most informative offset r^* is at most ϵ away from one of the global maxima of the anecdote density.*

Figure 5a shows an example of a bimodal distribution where the anecdote is with equal probability drawn from either the normal distributions $N(-2, 1)$ or $N(2, 1)$. The density of this bimodal distribution has two global maxima at -2 and 2 . When the sender has foresight and 5 anecdotes to choose from there are three offsets that satisfy the balanced-offset condition: $r^* = \pm 1.73$ and

Figure 5: Most informative offset r^* under foresight for bimodal normal mixture distribution



$$F \sim B \cdot N(-2, 1) + (1 - B) \cdot N(2, 1) \text{ where } B \text{ is Bernoulli with } p = \frac{1}{2}.$$

$r^* = 0$. The latter is a local minimum but the former offsets implement the two possible most informative targeting schemes. For $n = 10$ the two most informative targeting schemes at $r^* = \pm 1.98$ are even closer to the peaks of the density as predicted by Lemma 2 (see Figure 5b).

3.2 Persuasion Temptation

We now return to the original problem and allow the sender to take the disagreement loss into account. Note, that if the receiver chooses his best response $\sigma(\alpha) = M_R - \beta(r^*)$ to the sender's strategy (Theorem 2) then the disagreement loss equals $(M_R - M_S)^2$. In the special case where the sender and receiver share the same personal preferences (e.g. $M_R = M_S$) the disagreement loss is zero and therefore a balanced-offset targeting scheme is also an equilibrium of the sender-receiver game as we will show in Section 4. This is intuitive because the most informative targeting scheme solves exactly the problem where the sender is constrained to send her future self a single anecdote.

However, the balanced-offset targeting scheme is no longer an equilibrium when sender and receiver are misaligned. In that case, the sender can reduce her disagreement loss by slightly deviating to a targeting scheme with a bias of $\beta(r^*) + \delta$ ($|\delta|$ is small), which moves the receiver's action towards the sender's preferred action. For example, in the motivating example of Figure 1 she will choose $\delta < 0$. This reduces the disagreement loss to $(|M_R - M_S| - |\delta|)^2$ which is an $O(|\delta|)$ improvement to the sender's utility. The information loss increases but this effect is only $O(\delta^2)$ because we deviated from the targeting scheme that minimized the information loss.⁹ Consequently, a small δ -shift in the sender's communication bias is a strictly profitable deviation. Of course, the receiver will respond by adjusting the shift in her action rule by $-\delta$. In return, the sender will adjust her bias by another δ leading to *partial* unraveling. There is no full unraveling towards

⁹If we write the information loss as a function $L(\beta)$ of the bias β and assume it is twice differentiable then we can use the Taylor expansion around β^* and write $L(\beta^* + \delta) = L(\beta^*) + \frac{1}{2}L''(\beta^* + \zeta\delta)\delta^2$ for some $0 \leq \zeta \leq 1$ because $L'(\beta^*) = 0$ (as L is minimized at β^*). Therefore, if L'' is bounded the change in the information loss is $O(\delta^2)$.

always reporting the minimum (or maximum) anecdote because at some point the increase in the information loss term is no longer second-order.

Note, that in any equilibrium, the disagreement loss is always $(M_R - M_S)^2$ due to Theorem 2 despite the sender’s best efforts to reduce it. The rational receiver can always undo any bias that is added by the sender. However, the sender’s *persuasion temptation* increases her information loss. Communicating with a non-aligned receiver versus an aligned receiver therefore has two costs to the sender. First of all, there is always the disagreement loss $(M_R - M_S)^2$. Second, there is also an *additional* information loss because the temptation to persuade the receiver induces the sender to choose a suboptimal communication scheme. This is an example of *informational homophily* from the sender’s perspective: given a choice between a non-aligned and an aligned receiver the sender always prefers to communicate with the aligned receiver.

The next result shows that this type of homophily is mutual in our model.

Proposition 3. *Let π be a translation invariant communication scheme and let α_π be the best response of the receiver. Then the utility of the receiver is equal to the negative variance of the anecdote $\pi(\vec{x}, y)$,*

$$- \mathbb{E}[(\pi(\vec{x}, y) - (\theta + \beta(\pi)))^2]. \quad (6)$$

Hence, the receiver’s loss is exactly equal to the sender’s (expected) information loss. The receiver suffers no disagreement loss because she can in equilibrium undo any bias imposed by the sender. Therefore, the receiver always prefers to listen to an aligned sender (without persuasion temptation) rather than a non-aligned sender. Moreover, sender and receiver utility differ only by the constant disagreement loss $(M_R - M_S)^2$.

4 Communication Equilibria, Polarized Communication and Informational Homophily


In this section we show that translation invariant equilibria in our sender-receiver game always exist and are characterized by a simple condition that generalizes the “balanced-offset” condition of most informative targeting schemes from Section 3. We then show that these equilibria generically give rise to *polarized communication* in the sense that whenever the sender is more left-leaning (right-leaning) than the receiver she will select anecdotes with a left (right) offset. Polarization in return decreases utility for both the sender and the receiver and therefore implies *informational homophily*.

4.1 Equilibrium Characterization and Existence

We start by defining a *translation invariant equilibrium* which is simply a PBE where both the sender and the receiver use translation invariant strategies.

Definition 5. *A perfect Bayesian equilibrium (π, α) is translation invariant if π is a translation invariant communication scheme and α is a translation.*

Figure 6: Equilibrium offset for normal distribution $N(0, 1)$ and $n = 10$ anecdotes



Theorem 1 tells us that in such translation invariant PBE the sender will use a targeting scheme with offset r that satisfies

$$r = M_S - \sigma(\alpha). \quad (7)$$

This is also shown graphically in our motivating example from Figure 1. Moreover, theorem 2 pins down the shift $\sigma(\alpha)$:

$$\sigma(\alpha) = M_R - \beta(r). \quad (8)$$

Here, $\beta(r)$ denotes the bias $\beta(\pi)$ of targeting scheme π with offset r .

Combining these two conditions gives us a necessary condition for a translation invariant PBE:

$$\begin{aligned} r &= M_S - \underbrace{(M_R - \beta(r))}_{\sigma(\alpha)} \\ \beta(r) - r &= \underbrace{M_R - M_S}_{\Delta}. \end{aligned} \quad (9)$$

Note that this condition reduces to the balanced-offset condition when the preferences of sender and receiver are aligned ($\Delta = 0$).

The above condition is not only necessary but also a sufficient condition for PBE.

Theorem 4. *A pair (π, α) is a translation invariant equilibrium if and only if there exists some value $r \in \mathbb{R}$ such that π is the targeting scheme with offset r , α is a translation with shift $M_R - \beta(r)$, and the bias satisfies*

$$\beta(r) - r = M_R - M_S = \Delta.$$

All the proofs in this section can be found in Appendix D.

The distance between the bias $\beta(r)$ of the targeting scheme and its offset r is therefore the key relationship to understand the emergence of communication equilibria in our game. Figure 6a shows how this distance varies with r in the simple motivating example of Figure 1 where the sender

has foresight and access to $n = 10$ normally distributed anecdotes around the state of the world θ ($F \sim N(0, 1)$). For negative offset $r = 0$ we observe that $\beta(r) > r$ because the closest anecdote to r is more likely to be on the right of r (see Figure 1). Moreover, the distance $\beta(r) - r$ increases as we decrease r because the targeting scheme is approaching the minimum scheme and the bias $\beta(r)$ converges to the expected offset of the minimum anecdote (in this example, the minimum of $n = 10$ anecdotes). This ensures that for any $\Delta > 0$ there will be some $r^* < 0$ that satisfies $\beta(r^*) - r^* = \Delta$. For example, for $M_R - M_S = 2$ we obtain $r^* = -3.56$.

We can show that this intuition generalizes for a broad set of anecdote distributions.

Theorem 5. *For any n , M_S and M_R , if $\mathbb{E}_{\epsilon \sim F} |\epsilon|$ is bounded then a translation invariant PBE exists.*

We restrict attention from now on to anecdote distributions that satisfies this property.

We highlight the main technical ideas of the formal proof provided in Appendix D. Given a communication scheme π_r with offset r , let $z = \pi_r(\vec{x}, y) - (\theta_S(\vec{x}, y) + r)$ denote the distance between the target $\theta_S(\vec{x}, y) + r$ and the closest anecdote (out of n total anecdotes), where z is positive if the closest anecdote is larger and negative if the closest anecdote is smaller. Let $H(r)$ be the expected value of z , given offset r . By definition, we have that $\beta(r) = r + H(r)$. Theorem 4 now implies that to show that a PBE exists, it suffices to show that there exists a value of r such that $H(r) = M_R - M_S$, i.e., $H(r)$ is an onto function in its codomain $(-\infty, +\infty)$.

We show this in two steps. In the first step, we establish that $H(r)$ tends to $+\infty$ and $-\infty$, respectively, as $r \rightarrow -\infty$ and $r \rightarrow +\infty$. At a high level, this is due to that fact that boundedness of $\mathbb{E}[|x - \theta|]$ also implies that $\mathbb{E}[\max_i x_i - \theta_S(\vec{x}, y)]$ and $\mathbb{E}[\min_i x_i - \theta_S(\vec{x}, y)]$ are also bounded (for a fixed n). Therefore, $\beta(r)$ remains bounded even as $|r| \rightarrow \infty$. This proves that $H(r)$ tends to $+\infty$ and $-\infty$, respectively, as $r \rightarrow -\infty$ and $r \rightarrow +\infty$.

Had $H(r)$ been a continuous function, the first step would suffice to prove that $H(r)$ is an onto function taking all values in $(-\infty, +\infty)$. The second step handles discontinuous $H(r)$ by establishing that at any point that $H(r)$ is discontinuous, the left limit of $H(r)$ is smaller than the right limit. To see why, suppose H is not continuous at r_0 . For each possible realization of (θ, \vec{x}, y) , either z is continuous at r_0 or it is not. If not, this means that $\theta_S(\vec{x}, y) + r_0$ is precisely halfway between two anecdotes in \vec{x} , say with absolute distance $d > 0$ to each, in which case the limit of z from below is $-d$ (distance to the anecdote to the left) and the limit of z from above is d (distance to the anecdote to the right). Integrating over all realizations, we conclude that the one-sided limits of H exist and $\lim_{r \rightarrow r_0^-} H(r) < \lim_{r \rightarrow r_0^+} H(r)$.

Together these two steps show that $H(r)$ can take any value in the range $(-\infty, +\infty)$ and therefore a translation invariant PBE always exists

We can use the same logic to characterize the translation invariant equilibrium of the broadcasting extension of our model that we introduced in Section 2.3.

Lemma 3. *A translation invariant equilibrium in the broadcast model satisfies $\beta(r) - r = \overline{M}_R - M_S$ where $\overline{M}_r = \int M_F dG$ is the mean personal preference across receiver types. Moreover, for any n such an equilibrium exists if $\mathbb{E}_{\epsilon \sim F} |\epsilon|$ is bounded.*

Consequently, the sender's incentive to persuade is dictated by the misalignment $\Delta = \overline{M}_R - M_S$ between this representative receiver and the sender.

4.2 Polarized Communication with Non-aligned Preferences

We now show that communication becomes more *polarized* the more misaligned the preferences of sender and receiver are. Before formally defining polarization, we examine the case where sender and receiver are perfectly aligned ($\Delta = 0$).

Corollary 1. *When sender and receiver preferences are aligned ($\Delta = 0$) the balanced-offset condition holds and the equilibrium targeting scheme minimizes both the sender's and receiver's information loss.*

This follows immediately from Theorems 4 and 3. In our example of a normal anecdote distribution $N(0, 1)$ this implies that the sender chooses the mean scheme ($r^* = 0$) where she sends an anecdote closest to her posterior mean.

In principle there can be many balanced-offset equilibria: we will focus on the minimal and maximal balanced-offset equilibria, \underline{r} and \bar{r} :¹⁰

$$\begin{aligned}\underline{r} &= \min \{r | H(r) = 0\} \\ \bar{r} &= \max \{r | H(r) = 0\}\end{aligned}\tag{10}$$

We then define the *least polarized* communication equilibrium $r(\Delta)$:¹¹

$$r(\Delta) = \begin{cases} \max \{r | H(r) = \Delta, r < \underline{r}\} & \text{for } \Delta > 0 \\ \min \{r | H(r) = \Delta, r > \bar{r}\} & \text{for } \Delta < 0 \end{cases}\tag{11}$$


For example, if the sender is more left-leaning than the receiver ($\Delta > 0$) then $r(\Delta)$ is the equilibrium with offset $r(\Delta)$ that is closest to *the left of* the minimal balanced-offset equilibrium \underline{r} . For sufficiently large Δ all possible equilibria must be to the left of \underline{r} . However, for small Δ it is possible that there are equilibria in the interval $[\underline{r}, \bar{r}]$. For symmetric single-peaked distributions such as the normal distribution we know that $\underline{r} = \bar{r} = 0$ and therefore the least polarized equilibrium for $\Delta > 0$ are the ones with minimal absolute offset.

Figure 6b shows the unique equilibrium offset for the normal distribution as a function of Δ . The graph suggests that communication becomes increasingly polarized the more misaligned the preferences of sender and receiver are. The next result shows this observation holds generally.

¹⁰We know that \underline{r} (resp., \bar{r}) exists because $H(r) \rightarrow \infty$ as $r \rightarrow -\infty$ (resp., $H(r) \rightarrow -\infty$ as $r \rightarrow \infty$), and at any point of discontinuity r_0 we have $\lim_{r \rightarrow r_0^-} H(r) \leq H(r_0) \leq \lim_{r \rightarrow r_0^+} H(r)$. To see why this implies existence of \underline{r} , define $r_0 = \inf \{r : H(r) \geq 0\}$; then it suffices to show that $H(r_0) = 0$. If H is continuous at r_0 then this is immediate, but if H is discontinuous at r_0 then $\lim_{r \rightarrow r_0^+} H(r) > \lim_{r \rightarrow r_0^-} H(r) \geq 0$ in violation of the definition of r_0 .

¹¹Likewise, the least polarized equilibrium exists because $H(r) \rightarrow \infty$ as $r \rightarrow -\infty$ and $H(r) \rightarrow -\infty$ as $r \rightarrow \infty$, and at any point of discontinuity r_0 we have $\lim_{r \rightarrow r_0^-} H(r) \leq H(r_0) \leq \lim_{r \rightarrow r_0^+} H(r)$.

Figure 7: Intuition for Proposition 4 ($r(\Delta)$ decreases)



Proposition 4. *Assume some fixed anecdote distribution F and a fixed number of anecdotes n . The least polarized equilibrium $r(\Delta)$ decreases with Δ . Therefore, the sender’s communication becomes more left-polarized (right-polarized) the more her preferences lie to the left of (right of) the receiver’s preferences.*

Figure 7 illustrates the intuition behind the proof for $\Delta > 0$. It shows the graph of $H(r) = \beta(r) - r$ and the minimal balanced-offset equilibrium at \underline{r} for some non-single-peaked distribution. When the difference, Δ , in personal preferences equals Δ_1 , there are three possible equilibria with the least polarized one at $r(\Delta_1)$. For $\Delta = \Delta_2$ there are two equilibria with the least polarized one at $r(\Delta_2) < r(\Delta_1)$. For any $\Delta > \Delta_2$ (such as Δ_3) there is only a single equilibrium with $r(\Delta) < r(\Delta_2)$. Therefore, the function $r(\Delta)$ has a discontinuity at Δ_2 but is decreasing for any $\Delta > 0$.

As the misalignment between sender and receiver becomes extreme, the sender reports essentially only the minimum or maximum anecdote.

Corollary 2. *Assume some fixed anecdote distribution F and a fixed number of anecdotes n . As $\Delta \rightarrow \infty$ ($\Delta \rightarrow -\infty$) the sender’s communication scheme at any translation invariant equilibrium converges to the minimum (maximum) scheme.*

4.3 Informational Homophily

We know from Theorem 3 that all targeting schemes that minimize the sender’s information loss (as well as the receiver’s utility) have to satisfy the balanced-offset condition. Therefore, any equilibrium with $\Delta \neq 0$ has to have strictly lower utility for the sender and weakly lower utility for the receiver compared to the most informative communication scheme.¹²

Intuitively, one might expect that more polarized communication in response to greater misalignment in sender and receiver preferences ($|\Delta| \uparrow$) decreases utility. We show that this holds under foresight.

¹²The utility is strictly lower for the sender because the information loss is weakly lower and the disagreement loss Δ^2 is strictly negative.

Proposition 5. *Assume some fixed anecdote distribution F , a fixed number of anecdotes n and foresight. Both the sender and receiver utility of the least polarized equilibrium decreases with $|\Delta|$.*

The proof establishes that the variance of the sent anecdote in the least polarized equilibrium increases with $|\Delta|$. The result then follows immediately from Propositions 2 and 3.

Proposition 5 implies that anecdotal communication gives rise to *informational homophily*. Both the sender and the receiver prefer to communicate with more like-minded partners because communication is less polarized and therefore more precise. Greater polarization does not “fool” a rational receiver who corrects for the larger bias in the sender’s anecdote, but it comes at the cost of precision which hurts both the sender and receiver.

5 Experts and the Curse of Informedness

In the previous section we characterized communication and welfare when varying the misalignment Δ between sender and receiver preferences but keeping the anecdote distribution and the number of anecdotes constant. We showed that greater misalignment induces greater polarization which in return reduces welfare. Therefore, both sender and receiver will tend to seek out like-minded communication partners (keeping everything else equal).

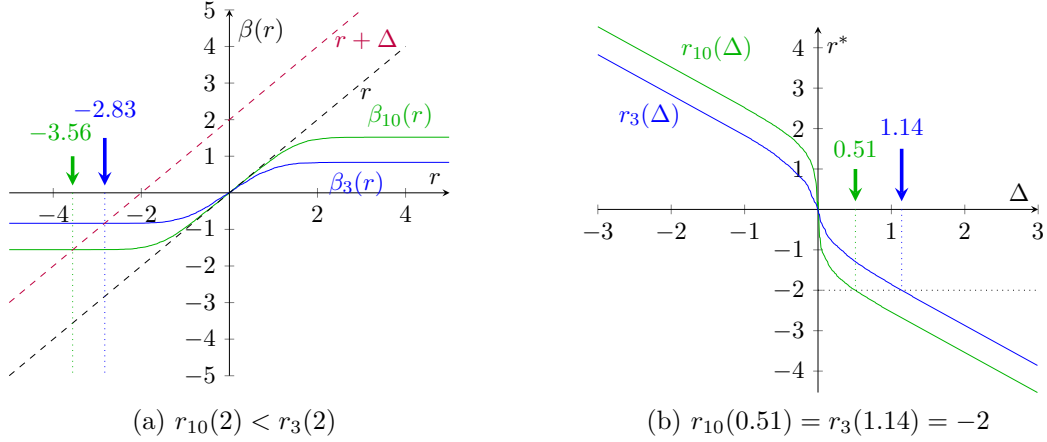
Another natural question is whether communication becomes more or less efficient when the sender is better informed (more of an “expert”). Formally, we are fixing both the anecdote distribution and the preference misalignment Δ in this section and vary the number of anecdotes n . Somewhat surprisingly, increasing n can make communication *less* efficient. We first show that experts use more extreme anecdotes to communicate *unless* the preferences are completely aligned ($\Delta = 0$). Intuitively, the information loss for any fixed offset r decreases with n and therefore the sender has a greater incentive to choose an even more biased anecdote. However, the impact on welfare is ambiguous because extreme anecdotes can also become more informative about the state of the world as n increases. We show that for an (appropriately defined) “light-tailed” distributions (such as the normal distribution) experts provide more precise anecdotes despite the polarization. On the other hand, for “heavy-tailed” distributions the precision of anecdotes sent by experts decreases with n . Intuitively, such experts have access to too many “outlier anecdotes” which makes it difficult for the rational receiver to learn about the state of the world. We call this phenomenon the *curse of informedness*.

5.1 Polarized Experts

Throughout this section we assume that the sender has foresight. We denote the bias of a targeting scheme with offset r and n anecdotes with $\beta_n(r)$.

Figure 8a illustrates the construction of the (unique) targeting equilibrium when $\Delta = 2$ and the anecdote distribution is normally distributed ($F \sim N(0, 1)$) for $n = 3$ and $n = 10$. We observe in this example that increasing the number of anecdotes makes the sender’s communication more polarized as the sender’s target moves from $r_3(2) = -2.83$ to $r_{10}(2) = -3.56$. This observation

Figure 8: Equilibrium offset $r_n(\Delta)$ for normal anecdote distribution



holds generally. For the next result, we assume that the anecdote distribution F is continuous and strictly positive over its domain $[r_{\min}, r_{\max}]$ (where we allow $r_{\min} = -\infty$ and/or $r_{\max} = \infty$).

Proposition 6. *Assume that the sender has foresight and that the anecdote distribution F is fixed, continuous and strictly positive over its domain $[r_{\min}, r_{\max}]$ (where r_{\min} can be $-\infty$ and/or r_{\max} can be ∞). Suppose also that the personal preferences of receiver and sender differ by a fixed $\Delta \neq 0$. Then the communication at equilibrium becomes more polarized as n increases, in the following sense: there exist sequences $R_{\min}(n)$ and $R_{\max}(n)$ with $R_{\min}(n) \rightarrow r_{\min}$ and $R_{\max}(n) \rightarrow r_{\max}$ such that for any equilibrium offset r_n^* (i.e., $\beta_n(r_n^*) - r_n^* = \Delta$) we have that either $r_n^* \leq R_{\min}(n)$ or $r_n^* \geq R_{\max}(n)$ for all n .*

All the proofs for this section are in Appendix E. The intuition for this proposition is shown in Figure 8a. The bias $\beta_n(r)$ of the targeting scheme gets closer to the diagonal as n increases: the more anecdotes the sender has access to, the easier it is to find an anecdote close to the offset r . Therefore, the information loss for the sender from communicating in a more polarized fashion becomes smaller as n increases and hence the sender's persuasion temptation will induce her to choose a more extreme targeting scheme.

Figure 8b demonstrates that even a small misalignment in personal preferences can produce a high level of polarization. When the sender has access to $n = 3$ anecdotes the differences in personal preferences have to be at least $\Delta \geq 1.14$ for the sender to target an anecdote two standard deviations away from the mean anecdote. A more informed sender with access to $n = 10$ anecdotes will exhibit that degree of polarization already for $\Delta \geq 0.51$. Proposition 6 therefore implies that the sender's informedness and polarization are complements.

5.2 The Curse of Informedness

Does polarized communication by experts decrease the sender's utility (and by Proposition 3 the receiver's utility)? In Section 4, we observed greater polarization and lower utility when fixing the

Table 1: Examples of strong heavy-tailed, strong light-tailed and knife-edge distributions

Distribution	probability density function $f(\epsilon)$	Type
Pareto	$f(\epsilon) = \frac{1}{(\epsilon +1)^3}$	strong heavy tails ($\alpha = 0$)
Gaussian	$f(\epsilon) = \frac{1}{2\sqrt{\pi}} \exp(-\frac{1}{2}\epsilon^2)$	strong light tails ($\alpha = 2$)
Laplace	$f(\epsilon) = \frac{1}{2} \exp(- \epsilon)$	knife-edge ($\alpha = 1$)

number of anecdotes n and increasing the sender/receiver alignment ($|\Delta| \uparrow$) (see Proposition 5).

It is not obvious that a similar result should hold when we fix the sender/receiver alignment Δ and let the sender become more informed ($n \uparrow$). While communication becomes more polarized (which decreases utility for fixed n) there are also more anecdotes to choose from, so even the more extreme anecdotes may be more informative to the receiver. The net impact on both the sender’s and receiver’s utility therefore seems ambiguous and depends on how informative extreme anecdotes are about the state of the world.

We next show that the net utility impact depends on the tails of the anecdote distribution. We will show that if F is sufficiently “heavy-tailed” then polarization will make the expert’s signal increasingly uninformative so that the receiver would prefer to talk to *any less-informed sender* (even highly misaligned ones). We call this the “curse of informedness”. On the other hand, if F is “light-tailed” then the expert’s anecdotes become increasingly informative as n increases despite the polarization.

Definition 6. *We say that the anecdote density has strong heavy tails if its hazard rate satisfies*

$$\frac{f(x)}{1 - F(x)} = u(x)x^{\alpha-1} \tag{12}$$

for a positive function $u(x)$ which is bounded from above by some constant \bar{U} and satisfies $x^\beta \frac{u'(x)}{u(x)} \rightarrow 0$ for some $\beta > 0$ and $0 \leq \alpha < 1$ and $|x| > \underline{x}$. The anecdote distribution has a strong light tails if

$$\frac{f(x)}{1 - F(x)} = u(x)x^{\alpha-1} \tag{13}$$

for a positive function $u(x)$ which is bounded from below by some constant \underline{U} with $x^\beta \frac{u'(x)}{u(x)} \rightarrow 0$ for some $\beta > 0$ and $\alpha > 1$, $|x| > \underline{x}$.

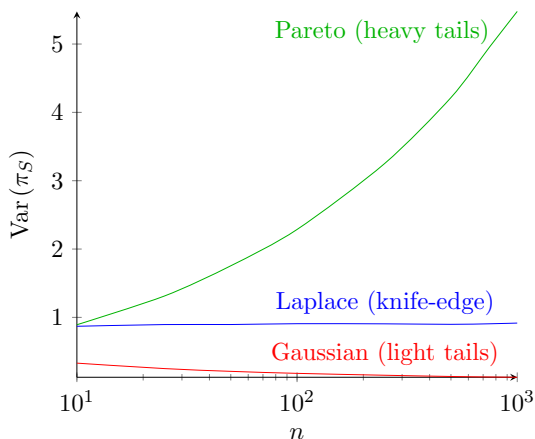
Distributions with strong heavy (light) tails have tails that decline at a less than (more than) exponential rate everywhere. For example, Table 1 defines the strong heavy-tailed Pareto distribution ($\alpha = 0$), the strong light-tailed Gaussian distribution ($\alpha = 2$) and the Laplace distribution which is a knife-edge case with its exponential tails ($\alpha = 1$).

We can now state the main result.

Theorem 6. Assume that the sender has foresight, the anecdote distribution F is fixed, continuous, and strictly positive over its domain, and the personal preferences of receiver and sender differ by a fixed $\Delta \neq 0$. If F is strong-heavy-tailed then the variance of the anecdotes sent at any equilibrium converges to ∞ as $n \rightarrow \infty$. If instead the anecdote distribution is strong-light-tailed then the variance of the anecdote sent at any equilibrium converges to 0 as $n \rightarrow \infty$.

Figure 9 plots the variance of the sent anecdote for $n \in [10, 1000]$ for the Pareto, Gaussian and Laplace distributions from Table 1. Consistent with Theorem 6, the variance stays constant around 0.9 for the knife-edge Laplace distributions, converges to 0 for the light-tailed Gaussian and diverges for the heavy-tailed Pareto distribution.

Figure 9: Variance of sent anecdote for Pareto, Gaussian and Laplace distributions for $\Delta = 1$



An immediate consequence is that for heavy-tailed distributions greater expertise lowers both the sender’s and receiver’s utility unless both of them are perfectly aligned (this follows from Proposition 3). Therefore, a receiver might prefer to listen to more-aligned but less informed sender rather than a less-aligned expert because the expert finds it tempting to report easily accessible extreme anecdotes. With light-tailed distributions, such extreme anecdotes are still highly informative about the state of the world and therefore the receiver will generally prefer to listen to experts even if their preferences are not aligned with his own.

Importantly, receivers in our model do not distrust experts when anecdotes are heavy-tailed because they think experts tell them lies. Rather they do not trust experts with different preferences because they correctly anticipate that such experts will tell them unrepresentative data points which are too noisy to infer the state of the world.

The curse of informedness also affects the incentives of senders to *acquire* information in the first place. When the anecdote distribution is heavy-tailed and senders expect to be matched with unaligned receivers they have an incentive to *remain less informed*. Being too well informed makes them less credible sources of information for receivers (unless the receiver is perfectly aligned).

To connect these results to the real-world, we can look at the dramatic decline in trust of scientists between 2019 and 2024. The COVID-19 pandemic gave a lot of visibility to experts who

advised the public on how to reduce infections, deal with school closings and how to manage the economy. Academic experts in the US tend to be more left-leaning than the rest of the population.¹³ Therefore, the personal preference M_S of the majority of experts will lie to the left of the average voter’s personal preferences ($\Delta = \overline{M}_R - M_S > 0$). The high degree of uncertainty during the pandemic might also serve as a good example of a heavy-tailed anecdote distribution. Hence, the “curse of informedness” (and Lemma 4.1) implies that the variance in experts’ messages was potentially large during the pandemic. Consistent with this story, 2023 Pew Research study documents that the share of US adults who “have a great deal or fair amount of confidence in scientists to act in the public’s best interests” declined from 87% to 73% between April 2020 and October 2023 with strong level of trust falling from 39% to 23% (Kennedy and Tyson, 2023). At the same time, distrust more than doubled from 12% to 27%. While the decline was more pronounced among Republicans, it decreased for both Democrats and Republicans.

6 Communication with Commitment

Up to now we have assumed that the sender cannot commit to a communication scheme. However, sometimes commitment is possible. For example, a sender might have a reputation for a particular type of reporting or a reputable newspaper might commit to always selecting an unbiased set of facts for their articles (and might be punished by readers if they are later found out to have deviated from this communication scheme). Other examples of commitment include the interpretation of our model as a behavioral game played between current and future self, where current self decides which anecdotes to save to memory so that future self makes the best possible decisions.

We formally defined the commitment equilibrium as a modification of Definition 1 in Section 2.4 (see Equation 4). As before, Theorems 1 and 2 imply that if the sender uses a translation invariant communication scheme π , then the receiver’s best response will be a translation and vice versa. Hence, we again focus on translation invariant commitment equilibria. The following theorem characterizes the set of such equilibria.

Theorem 7. *The sender’s optimal translation invariant commitment π_S is the most informative communication scheme that minimizes only the information loss. In particular, it is a targeting scheme that satisfies the balanced-offset condition $r^* = \beta(r^*)$ and the receiver’s response is a translation with shift $M_R - r^*$.*

The proof is provided in Appendix F. Intuitively, commitment removes the sender’s persuasion temptation: since she cannot deceive the receiver, she will ignore the unavoidable disagreement loss and simply minimize the information loss. She therefore will implement the socially optimal communication scheme subject to the constraint that she can only send a single anecdote. Hence,

¹³A study of voter registration of 7,243 faculty at 40 leading U.S. universities in the fields of Economics, History, Journalism/Communications, Law, and Psychology found that 3,623 were registered Democrats and 314 Republicans with an overall ratio of 11.5:1 while the ratio in the overall population is close to 1 (Langbert et al., 2016).

the sender solves the same problem as in Section 3 and behaves as if her personal preferences are aligned with the receiver's.

6.1 Characterizing Commitment Equilibria

Our characterization of most informative communication schemes in Section 3.1 immediately provides a characterization of the translation invariant commitment equilibria when the sender has *foresight* in the following cases:

1. When n is large, the sender's optimal communication scheme is a targeting scheme with a target close to the global maximum density, by Lemma 2.
2. When the anecdote distribution F is symmetric and single-peaked at 0 then the sender's optimal targeting scheme is the mean scheme by Lemma 1.

Intuitively, one might expect that the mean scheme is also an equilibrium *without* foresight. Even though the sender does not know the state of the world, her estimate of the posterior mean will concentrate around the true state (for large n) and we therefore would expect the sender to select an anecdote close to the posterior mean since the density of the anecdote distribution is also maximized around the true state. We show that this intuition is correct for a class of distribution which we call *well-behaved*.

Definition 7. *We say that the anecdote distribution F is well-behaved if the following holds.*

1. *The distribution is strictly single-peaked and symmetric with finite variance.*
2. *Let $g(x) = f'(x)/f(x)$. That is, $g(x) = \frac{d \log f(x)}{dx}$. We assume that $|g'(x)| \leq c_1$ for all x , and some constant $c_1 > 0$. That is, $|g(x)| \leq c_1|x| + c_2$.¹⁴*
3. *F has exponential tails. That is, there is a constant $Q > 0$, such that for $x > Q$, we have $1 - F(x) \leq c_3 \exp(-|x|)$ for a constant $c_3 > 0$, and $x < -Q$ we have $F(x) \leq c_3 \exp(-|x|)$.*

For example, the normal distribution $F \sim N(0, 1)$ and the Laplace distribution with density $f(\epsilon) = \frac{1}{2} \exp(-|\epsilon|)$ are well-behaved. Recall that $f(\epsilon)$ is the density of the signal distribution at $\theta + \epsilon$, an offset of ϵ from the true state of the world.

Theorem 8. *For any well-behaved anecdote distribution, in the no-foresight setting, the unbiased communication scheme that sends the closest signal to $\theta_S(\vec{x}, y)$ strictly dominates any biased signaling scheme for sufficiently large n and is optimal among all unbiased communication schemes.*

The proof is delegated to a separate online appendix. Intuitively, the sender would like to send *precisely* the posterior mean to the receiver. However, since she is constrained to sending a signal she has to contend with the second best which is to send the signal closest to the posterior mean. When we interpret our model as a model of memory where the current self communicates with her future self by storing a single anecdote in memory we can think of the anecdote closest to the posterior mean as the “most representative anecdote”.

¹⁴Note that, for $x < 1$ we bound $|g(x)| \leq c_2$ and otherwise we can bound $|g(x)| \leq c_1$.

6.2 Biased Commitment Equilibria under No-Foresight

We have just shown that for single-peaked and symmetric anecdote distributions the mean communication scheme is optimal under foresight as well as under no-foresight when n is large. However, unbiased communication ($r^* = 0$) is not necessarily optimal for small n under no-foresight even when the anecdote distribution is single-peaked and symmetric.

The simplest example is the case where the sender has $n = 2$ anecdotes.

Proposition 7. *Suppose $n = 2$, signals are drawn from a symmetric distribution around θ , and the sender has no foresight. Then the optimal communication scheme is either the minimum scheme or the maximum scheme.*

The proofs of this subsection appear in Appendix F. The optimal scheme still has to satisfy the balanced-offset condition of Theorem 7 such that $\beta(r^*) = r^*$. However, there are three distinct candidate targeting schemes: the mean scheme ($r^* = 0$), the minimum scheme ($r^* < 0$), and the maximum scheme ($r^* > 0$), where the offset (and equivalently the bias) is zero, negative, and positive, respectively.

The mean scheme incurs a higher level of information loss compared to the minimum and maximum schemes due to a distinctive aspect of the two-signal setting under no-foresight: both anecdotes have the same absolute distance to the posterior mean $\theta_S(\vec{x}, y)$ or empirical mean $\frac{x_1+x_2}{2}$. Consequently, the information loss of a communication scheme primarily hinges on the uncertainty regarding whether the posterior mean lies to the left or right of the communicated signal. This uncertainty is minimized by having the mean always lying to the left (or the right), i.e., consistently sending the minimum (or the maximum, respectively).

The biasedness of the optimal communication scheme is not an artifact of the two-signal setting as demonstrated in the following example.

Proposition 8. *Suppose $n = 3$, signals are drawn from a uniform distribution around θ , and the sender has no foresight. Then at every commitment equilibrium the sender uses a communication scheme with non-zero bias.*

The idea behind Proposition 8 is that, conditional on the value of the sender's posterior mean, the conditional density over signal realizations is not necessarily single-peaked. For uniform distributions and $n = 3$, the correlation between the posterior mean and the minimum and maximum signals is stronger than the correlation between the posterior mean and the middle signal. One can therefore communicate more information about the posterior mean through a biased communication scheme ($r^* \neq 0$) that sometimes returns the minimum signal (or, by symmetry, sometimes returns the maximum signal).

Why does bias help in this example? Recall that there is intrinsic error in the sender's posterior mean. This variance is unavoidable. But it can introduce correlation with particular anecdotes. This correlation can be used to help minimize the variance between the posterior mean and the anecdote passed to the receiver. This is why, for the uniform case, it is helpful to bias toward

more extreme signals: even though they are not more informative than the moderate signals when it comes to the true state of the world, they are more informative with respect to the sender’s posterior mean. The interplay between these two sources of errors therefore introduces an incentive for the sender to systematically bias their communication.

However, for large n the correlation between the posterior mean and the closest signal to the posterior mean dissipates and therefore the unbiased mean scheme re-emerges as the optimal communication scheme under no-foresight as shown in Theorem 8.

7 Conclusion

We introduce a model of strategic communication where the sender communicates raw data points (anecdotes) that are informative about an underlying state of the world. While this type of anecdotal communication is less efficient than communicating posterior beliefs about the state directly, it does not require that the sender and receiver have a shared representation of the state of the world. Therefore, our model can help us understand how newspapers convey news to their diverse readership by assembling articles from various quotes and factual statements, or how politicians rely on examples and anecdotes to connect with their voters. We have shown four main results. First of all, our model naturally gives rise to polarization as the sender balances the need to inform the receiver with the temptation to persuade him which makes her target biased anecdotes. Second, polarization increases, and welfare decreases with the misalignment between sender’s and receiver’s preferences. Third, polarization is amplified when experts communicate because they have access to more outlier anecdotes. Fourth, polarization will hurt both the sender and the receiver when the anecdote distribution is heavy-tailed giving rise to the curse of informedness.

Stories as “Carriers of Beliefs”. On a conceptual level, our model fits within a paradigm where stories or anecdotes provide the basis for both communicating as well as storing beliefs. While we mainly focus on the former interpretation in this paper, we also explained that our model can be viewed as the current self (sender) communicating with a future self (receiver). This can be expanded into a theory of *optimal storage of stories* for a decision-maker who has limited memory and can only store information about the state of the world through representative data points rather than posterior beliefs. We believe that this model of storing “beliefs” is appropriate in many situations where we need to learn about the world without being an expert. For instance, when describing notable traits of a colleague, friend, or acquaintance – like kindness, intelligence, or creativity – we could illustrate them through specific anecdotes. These anecdotes might showcase instances when the individual demonstrated kindness, solved a challenging problem, or devised an innovative solution, rather than relying solely on formal metrics such as like altruism or IQ scores.

Extensions. We hope that our model can serve as a starting point for future extensions. For example, our model assumes that the sender can only send a single anecdote. This makes our analysis particularly simple because we can focus on translation invariant equilibria. However,

senders can and do communicate more than one anecdote. For example, a 500-word news article contains on average 10 to 20 sentences and 5 to 10 facts (Kiss et al., 2024). How can one adapt our analysis? We can gain some intuition by looking at the special case where sender and receiver are aligned ($\Delta = 0$), the sender has foresight and the anecdote distribution is symmetric with a single peak at 0. If we focus on some small enough interval of width δ around the state of the world where the anecdote density does not change much, we can think of δn anecdotes that are approximately uniformly distributed over this interval. If the sender can only send one anecdote she will send the closest one to the state of the world which is an expected distance $O(1/n)$ away from the true state and the (quadratic) information loss is therefore $O(1/n^2)$. With two anecdotes, the sender could combine any of approximately $\delta n/2$ anecdotes to the left and right of the true state and choose among these $O(n^2)$ combinations the one whose average is closest to the true state. We conjecture that the expected distance of this best average is now only $O(1/n^2)$ away from the true state and the information loss is therefore only $O(1/n^4)$. This heuristic argument suggests that the information loss from anecdotal communication declines rapidly with the number of anecdotes that the sender can transmit. The analysis in Sections 4 and 5 then suggests that polarization becomes even more extreme when sender and receiver are no longer aligned which might mitigate or even reverse any reductions in the information loss that are achievable in the optimal case.

References

- Ambrus, Attila and Aaron Kolb**, “On defining ex ante payoffs in games with diffuse prior,” *Economic Theory*, 2021, 72, 445–472.
- Cambon, Sarah Chaney**, “November Employment Report Shows U.S. Economy Added 263,000 Jobs,” <https://www.wsj.com/articles/november-jobs-report-unemployment-rate-economy-growth-2022-11669944597> 2022. Accessed: 2022-12-28.
- Crawford, Vincent P and Joel Sobel**, “Strategic information transmission,” *Econometrica: Journal of the Econometric Society*, 1982, pp. 1431–1451.
- Dye, Ronald A**, “Disclosure of nonproprietary information,” *Journal of accounting research*, 1985, pp. 123–145.
- Dziuda, Wioletta**, “Strategic argumentation,” *Journal of Economic Theory*, 2011, 146 (4), 1362–1397.
- Frenkel, Sheera, Adam Satariano, and Ryan Mac**, “Meta Lays Off More Than 11,000 Employees,” <https://www.nytimes.com/2022/11/09/technology/meta-layoffs-facebook.html> 2022. Accessed: 2022-12-28.
- Glazer, Jacob and Ariel Rubinstein**, “A study in the pragmatics of persuasion: a game theoretical approach,” *New perspectives on games and interaction*, 2008, pp. 121–140.

- Grossman, Sanford J**, “The informational role of warranties and private disclosure about product quality,” *The Journal of Law and Economics*, 1981, *24* (3), 461–483.
- **and Oliver D Hart**, “Disclosure laws and takeover bids,” *The Journal of Finance*, 1980, *35* (2), 323–334.
- Hart, Sergiu, Ilan Kremer, and Motty Perry**, “Evidence games: Truth and commitment,” *American Economic Review*, 2017, *107* (3), 690–713.
- Jung, Woon-Oh and Young K Kwon**, “Disclosure when the market is unsure of information endowment of managers,” *Journal of Accounting research*, 1988, pp. 146–153.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian persuasion,” *American Economic Review*, 2011, *101* (6), 2590–2615.
- Kartik, Navin**, “Strategic communication with lying costs,” *The Review of Economic Studies*, 2009, *76* (4), 1359–1395.
- , **Marco Ottaviani, and Francesco Squintani**, “Credulity, lies, and costly talk,” *Journal of Economic theory*, 2007, *134* (1), 93–116.
- Kennedy, Brian and Alex Tyson**, “American Trends Panel Wave 135 (September 25–October 1, 2023),” Technical Report, Pew Research 2023.
- Kiss, Oliver, Markus Mobius, Tanya Rosenblat, and David Rothschild**, “Assembling News Like Legos,” Technical Report, Microsoft Research 2024.
- Langbert, Mitchell, Anthony J. Quain, and Daniel Klein**, “Faculty Voter Registration in Economics, History, Journalism, Law, and Psychology,” *Econ Journal Watch*, September 2016, *13* (3), 422–451.
- Lin, Xiao and Ce Liu**, “Credible Persuasion,” *arXiv preprint arXiv:2205.03495*, 2022.
- Lipnowski, Elliot, Doron Ravid, and Denis Shishkin**, “Persuasion via weak institutions,” *Journal of Political Economy*, 2022, *130* (10), 2705–2730.
- Martini, Giorgio**, “Multidimensional Disclosure,” Technical Report 2018.
- Milgrom, Paul**, “What the seller won’t tell you: Persuasion and disclosure in markets,” *Journal of Economic Perspectives*, 2008, *22* (2), 115–131.
- Milgrom, Paul R**, “Good news and bad news: Representation theorems and applications,” *The Bell Journal of Economics*, 1981, pp. 380–391.
- Min, Daehong**, “Bayesian persuasion under partial commitment,” *Economic Theory*, 2021, *72* (3), 743–764.

Needelman, Joshua, “We Revisited People Who Left Their Jobs Last Year. Are They (Still) Happy?,” <https://www.nytimes.com/2023/01/16/business/quitting-jobs-happiness.html> 2023. Accessed: 2023-1-16.

Nguyen, Anh and Teck Yong Tan, “Bayesian persuasion with costly messages,” *Journal of Economic Theory*, 2021, *193*, 105212.

Sher, Itai, “Credibility and determinism in a game of persuasion,” *Games and Economic Behavior*, 2011, *71* (2), 409–419.

Zhang, Kun, “Withholding Verifiable Information,” *arXiv preprint arXiv:2206.09918*, 2022.

A Diffuse Prior

In this section we discuss our assumption of a diffuse prior and its implications on the posterior of the agents. Throughout, we restrict attention to translation invariant sender and receiver strategies that we introduce in Section 2.2. We start by showing that sender and receiver beliefs *after* sending an anecdote are well defined if the prior on the state of the world θ is diffuse. We then show that players' ex-ante payoffs *before* a message is sent are also well-defined.

A.1 Posterior Distributions of the Sender and Receiver

In statistics, it is common to use an *improper prior* as uninformative priors. The simplest way to formalize a diffuse prior that reveals no information about θ is to consider the density to be a constant $\mu(\theta) = c$ for all $\theta \in \mathbb{R}$. It is important to note that while μ is not a proper probability distribution (since $\int_{\theta} \mu(\theta) d\theta = \infty$), it is still possible that the posterior formed can still be proper and well defined¹⁵.

Claim A.1. *Given a diffuse prior and $\vec{x} \sim \theta + F$, $y \sim \theta + G$, the posterior distribution of the sender conditioned on \vec{x}, y , is a proper distribution. Moreover, when F and G are symmetric, the posterior mean $\theta_S(\vec{x}, y)$ is an unbiased estimator of θ .*

Similarly, given a diffuse prior, a translation invariant communication scheme π , and a signal $x = \pi(\vec{x}, y)$. The posterior distribution of the receiver conditioned on π, x , is a proper distribution. Further, the posterior mean of the receiver is $\theta_R(\pi, x) = x - \beta(\pi)$.

We prove this claim below.

Sender's posterior distribution. Recall that the anecdotes x_1, \dots, x_n are drawn independently from $\theta + F$. Thus the pdf of an anecdote given θ is $f(x - \theta)$. Similarly, y is drawn from $\theta + G$ and hence the pdf of y given θ is $g(y - \theta)$.

We first observe that in the foresight case, when $G = 0$, the sender's posterior is a point mass at y .

For any G that is a proper distribution, we see that the posterior of θ given y is proper.

$$\begin{aligned} \mu(\theta|y) &= \frac{g(y - \theta)\mu(\theta)}{\int_{\hat{\theta}} g(y - \hat{\theta})\mu(\hat{\theta})d\hat{\theta}} \\ &= \frac{g(y - \theta)}{\int_{\hat{\theta}} g(y - \hat{\theta})d\hat{\theta}} \\ &= \frac{g(y - \theta)}{\int_{\gamma} g(\gamma)d\gamma} = g(y - \theta) \end{aligned}$$

The first equality is just the definition of a posterior, and the second equality holds since μ is the diffuse prior with $\mu(\theta) = c$ for all θ . The third equality does a change of variables to $\gamma = y - \hat{\theta}$.

¹⁵Also, in a more Frequentist view, μ can be thought of as the likelihood function to capture the absence of data.

Finally, the last step follows because G is a proper distribution. Hence $\mu(\theta|y)$ is a proper posterior distribution.

Recall that $\mu(\theta|\vec{x}, y) = \frac{\hat{f}(\vec{x}|y, \theta)\mu(\theta|y)}{\int_{\hat{\theta}} \hat{f}(\vec{x}|y, \hat{\theta})\mu(\hat{\theta}|y)d\hat{\theta}}$, where $\hat{f}(\vec{x}|\theta, y)$ is the conditional pdf of \vec{x} given θ, y . That is, we can use $\mu(\theta|y)$ as a prior. Since $\mu(\theta|y)$ is a proper distribution, the posterior $\mu(\theta|\vec{x}, y)$ is also proper.

For the non-foresight case, when G is diffuse, we can use a similar argument as above to first compute the posterior given x_1 and y . We get,

$$\begin{aligned}\mu(\theta|x_1, y) &= \frac{f(x_1 - \theta)g(y - \theta)\mu(\theta)}{\int_{\hat{\theta}} f(x_1 - \hat{\theta})g(y - \hat{\theta})\mu(\hat{\theta})d\hat{\theta}} \\ &= \frac{f(x_1 - \theta)}{\int_{\epsilon} f(\epsilon)d\epsilon} \\ &= f(x_1 - \theta)\end{aligned}$$

This is again by noting that $\mu(\theta) = c$, $g(y - \theta) = c'$, and doing a change of variable to $\epsilon = x_1 - \hat{\theta}$. Thus $\mu(\theta|x_1, y)$ is a proper posterior distribution because f is proper distribution. Now using this as a prior, we get that $\mu(\theta|\vec{x}, y)$ is a proper posterior.

Sender's posterior mean. We observe that, given a diffuse prior and symmetric anecdote distributions F, G , $\mu(\theta|\vec{x}, y) = \mu(-\theta|-\vec{x}, -y)$. With this, it is easy to see that, for $\theta = 0$, $\mathbb{E}_{\vec{x}, y}[\theta_S(\vec{x}, y)|\theta = 0] = 0$. Moreover, we show below that $\theta_S(\vec{x}, y) = \theta_S(\vec{x} + t, y + t) - t$, and hence $\mathbb{E}_{\vec{x}, y}[\theta_S(\vec{x}, y)|\theta] = \theta$. Thus, the sender's posterior mean is an unbiased estimator of θ .

We will now show that $\mu(\theta|\vec{x}, y) = \mu(\theta + t|\vec{x} + t, y + t)$, and this would imply,

$$\theta_S(\vec{x}, y) = \int_{\theta} \theta \cdot \mu(\theta|\vec{x}, y)d\theta = \int_{\theta} \theta \cdot \mu(\theta + t|\vec{x} + t, y + t)d\theta = \theta_S(\vec{x} + t, y + t) - t.$$

We have $\mu(\theta|\vec{x}, y) = \mu(\theta + t|\vec{x} + t, y + t)$ because,

$$\begin{aligned}\mu(\theta|\vec{x}, y) &= \frac{\prod_i f(x_i - \theta) \cdot g(y - \theta)\mu(\theta)}{\int_{\hat{\theta}} \prod_i f(x_i - \hat{\theta}) \cdot g(y - \hat{\theta})\mu(\hat{\theta})d\hat{\theta}} \\ &= \frac{\prod_i f(x_i + t - \theta - t) \cdot g(y + t - \theta - t)\mu(\theta + t)}{\int_{\hat{\theta}} \prod_i f(x_i + t - \hat{\theta} - t) \cdot g(y - \hat{\theta})\mu(\hat{\theta})d\hat{\theta}} \\ &= \mu(\theta + t|\vec{x} + t, y + t)\end{aligned}$$

Receiver's posterior distribution. We show that the receiver's posterior distribution given a translation invariant π and $x = \pi(\vec{x}, y)$ is a proper distribution. Let $h_{\pi}(x|\theta)$ be the pdf of the signal sent given π and θ . Observe that, by definition of translation invariant, $\pi(\vec{x} - \theta, y - \theta) = \pi(\vec{x}, y) - \theta$. Therefore, $h_{\pi}(x|\hat{\theta}) = h_{\pi}(x - \hat{\theta}|0)$. Note that, $h_{\pi}(\cdot|0)$ only depends on π, F , and G .

$$\begin{aligned}
\mu(\theta|\pi, x) &= \frac{h_\pi(x|\theta)\mu(\theta)}{\int_{\hat{\theta}} h_\pi(x|\hat{\theta})\mu(\hat{\theta})d\hat{\theta}} \\
&= \frac{h_\pi(x - \theta|0)}{\int_{\hat{\theta}} h_\pi(x - \hat{\theta}|0)d\hat{\theta}} \\
&= \frac{h_\pi(x - \theta|\theta = 0)}{\int_{\epsilon} h_\pi(\epsilon|0)d\epsilon} = h_\pi(x - \theta|0)
\end{aligned}$$

Receiver's posterior mean. Given a translation invariant π , for any θ , recall that the bias $\beta(\pi) = \int_x (x - \theta)h_\pi(x|\theta) = \int_z zh_\pi(z|0)$. Hence we get that the posterior mean of the sender $\theta_R(\pi, x)$ is

$$\begin{aligned}
\int_{\hat{\theta}} \hat{\theta}\mu(\hat{\theta}|\pi, x)d\hat{\theta} &= \int_{\hat{\theta}} \hat{\theta}h_\pi(x - \hat{\theta}|0)d\hat{\theta} \\
&= \int_z (x - z)h_\pi(z|0)dz \\
&= x - \int_z zh_\pi(z|0)dz \\
&= x - \beta(\pi)
\end{aligned}$$

Thus for any θ , translation invariant π , we get $\mathbb{E}_x[\theta_R(\pi, x)|\theta] = \mathbb{E}_x[x - \beta(\pi)|\theta] = \theta$ (follows directly from the definition of $\beta(\pi)$). Therefore, the receiver's posterior mean $x - \beta(\pi)$ is an unbiased estimator of θ .

A.2 Ex Ante Payoffs

When describing payoffs and equilibria in our communication game, it is also important to understand the players' beliefs about payoffs *before* any messages are received. Concerns about the well-definedness of ex ante payoffs in signaling games with diffuse priors were also noted recently by Ambrus and Kolb (2021). They provide sufficient conditions under which such a game's ex ante payoffs are well-defined for a class of strategies that include translation invariant strategies. In this section we describe their conditions and results and show that they are satisfied by our model under some modifications that are without loss of generality for translation invariant strategies. This will justify our use of expectations over θ when analyzing payoffs and beliefs, despite θ being drawn from an improper prior.

We now describe the setting of Ambrus and Kolb (2021) in the context of our model. There is a state of the world θ and N players. Each player receives a collection of signals, where each signal s_{ij} observed by player i is drawn independently as $\theta + \hat{s}_{ij}$ where $\hat{s}_{ij} \sim F_{ij}$ admits positive density.¹⁶ Each player i then chooses an action a_i from a set $A_i \subseteq \mathbb{R}$ of allowable actions that

¹⁶The model of Ambrus and Kolb (2021) assumes at most one signal per player, but their analysis holds without

can depend on the set of signals received. Each player i then receives a payoff u_i that depends on θ and the profile \vec{a} of actions chosen. It is assumed that payoffs are translation invariant, in the sense that $u_i(\theta + \delta, \vec{a} + \delta) = u_i(\theta, \vec{a})$ for all $\theta \in \mathbb{R}$, $\vec{a} \in \mathbb{R}^n$, and $\delta \in \mathbb{R}$. A strategy μ_i is a (possibly randomized) mapping from received signals to an action. We write $u_i(\theta, \vec{\mu})$ for the expected utility obtained under strategy profile $\vec{\mu}$ where the expectation is over the realization of signals and any randomness in $\vec{\mu}$.

Condition A.1 (Compactness and Translation Invariance). *For each player, given any set of signals received, the set of potential actions is a compact set. Moreover, whenever all signals are offset by a real number $\delta \neq 0$, then the set of potential actions is likewise offset by δ .*

Condition A.2 (Bounded Interim Payoffs). *For each player i there is some finite constant C_i such that, for any realization of θ , the expectation (over signal realizations) of $\max_{\vec{a}} |u_i(\theta, \vec{a})|$ is at most C_i , where the maximum is over action profiles that are allowable given the realized signals.*

Condition A.3 (Irreducibility). *For each player i there is no pair of distinct strategies μ_i, μ'_i such that the expected payoffs of μ_i and μ'_i are identical for all θ and all translation invariant strategies of the other players.*

Roughly speaking, Ambrus and Kolb (2021) show that if Conditions A.1, A.2, and A.3 are satisfied, then ex ante expected payoffs under any profile of translation invariant strategies is well-defined and equal to the expected payoffs conditional on any fixed value of θ . To state their result more formally, we first require some definitions.

Definition 8 (Diffusing Sequence). *Let λ denote the Lebesgue measure. We say a sequence $(P_m)_{m \in \mathbb{N}}$ of Borel probability measures on \mathbb{R} is diffusing if for any set W with $\lambda(W) \in (0, \infty)$ and any $\eta > 0$, there exists $M \in \mathbb{N}$ such that for all $m > M$,*

- $P_m(W) > 0$ and
- for all measurable $Y \subseteq W$, $\left| \frac{P_m(Y)}{P_m(W)} - \frac{\lambda(Y)}{\lambda(W)} \right| < \eta$.

Roughly speaking, a diffusing sequence is an infinite sequence of probability distributions that, in the limit, tends toward an improper distribution that behaves as a diffuse prior. For example, the sequence of distributions $(N(0, m))$ (i.e., P_m is a mean-zero Gaussian with variance m) is a diffusing sequence.

Definition 9 (Admissibility). *A class of strategy profiles is said to be admissible if for any profile $\vec{\mu} = (\mu_1, \dots, \mu_N)$ of strategies in the class, there exists a vector $u^* \in \mathbb{R}^N$ such that for any diffusing sequence (P_m) and every player i , $\lim_{m \rightarrow \infty} \mathbb{E}_{\theta \sim P_m} [u_i(\theta, \vec{\mu})] = u_i^*$.*

We are now ready to state the result of Ambrus and Kolb (2021).

change if one player can receive multiple independent signals.

Theorem A.1 ((Ambrus and Kolb, 2021)). *Suppose the game satisfies Conditions A.1, A.2, and A.3. Then the class of profiles of translation invariant strategies is admissible.*¹⁷

We now show how to apply Theorem A.1 to our signaling game. As the theorem applies to translation invariant strategies, we will focus attention on translation invariant strategies in our game. Recall that the sender's strategy is a communication scheme π , and a translation invariant strategy of the receiver is a translation action rule which is characterized by an offset $\sigma \in \mathbb{R}$. We can therefore map our game into the setting of Ambrus and Kolb (2021) as follows: the sender receives signals (\vec{x}, y) and selects from the set $\{x_1, \dots, x_n\}$. The receiver is not given any signals and chooses an offset σ . If the sender chooses x and the receiver chooses offset σ , they both receive payoffs as defined in our original game when the receiver takes action $x + \sigma$. Note that this formulation is a simultaneous-move game and that payoffs are indeed translation invariant with respect to θ and the received signals.

We would like to apply Theorem A.1 to this reformulation of our signaling game. However, we note that Conditions A.1 and A.2 do not directly apply in our model since the action space of the receiver is unbounded. Indeed, the receiver in our game can choose any offset $\sigma \in \mathbb{R}$ (an unbounded set). Moreover, the resulting cost to the sender and receiver can be arbitrarily high for extreme choices of σ . However, as we now show, it is without loss to restrict attention to a bounded space of receiver actions.

Lemma A.1. *Given F , G , and n , there exists $C > 0$ such that any choice of offset $\sigma > C$ is strictly dominated by $\sigma = C$, and any choice of offset $\sigma < -C$ is strictly dominated by $\sigma = -C$.*

Proof. Over the space of translation invariant communication schemes for the sender, the one that induces the (pointwise) maximal posterior over θ for any given message x is the one that always returns the minimum anecdote $\min\{x_1, \dots, x_n\}$. Note that $\mathbb{E}[\min\{x_1, \dots, x_n\}] \geq \mathbb{E}[\theta - \sum_i |x_i - \theta|] \geq \theta - nc$ where $c = \mathbb{E}_{\epsilon \sim F}[|\epsilon|]$ is a constant depending on F . In particular, this means that $E[D_{\pi, x}] \leq x + nc$ for every choice of translation invariant π . Thus, for any realization of x , a receiver action of $a > x + nc + M_R$ is strictly dominated by action $a = x + nc + M_R$. A symmetric argument, for the maximum scheme, shows that a receiver action of $a < x - nc + M_R$ is strictly dominated by $a = x + nc + M_R$. Taking $C = nc + |M_R|$ therefore satisfies the conditions of the lemma. \square

Lemma A.1 shows that when we restrict the sender to translation invariant communication schemes, it is without loss of generality to restrict the receiver to choosing actions that lie in a compact set centered on the received signal. As the sender is also mechanically restricted to choosing (a distribution over) one of the received anecdotes as an action, which is compact given the anecdote realization and invariant to translations, we conclude that Condition A.1 is satisfied

¹⁷In fact, they show a more general result: a weaker condition than translation invariance, in which the strategies tend toward a translation invariant limit strategy in the limit as θ tends to $\pm\infty$, is necessary and sufficient for the ex ante payoffs to be well-defined and determined by payoffs in the limit as θ tends to $\pm\infty$. We will only use the sufficient condition of full translation invariance.

under this restriction. Likewise, under the restriction (without loss) that the receiver chooses an action from this bounded set, the sender and receiver’s expected utilities are likewise bounded uniformly over the choice of θ (with the extremal value occurring when the sender chooses the maximum anecdote, and the receiver always takes the maximum allowable action $x + C$ given the sender’s message). Thus Condition A.2 is satisfied as well.

Finally, we note that Condition A.3 is satisfied as well. Indeed, any two distinct strategies (i.e., communication schemes) of the sender, say π and π' with $\pi(\vec{x}) \neq \pi'(\vec{x})$, would result in different outcomes (and hence payoffs, if we fix θ) on anecdote realization \vec{x} when the receiver chooses the identity action rule $\alpha(x) = x$ (i.e., offset $\sigma = 0$). On the other hand, any two distinct choices of offset $\sigma \neq \sigma'$ would result in different payoffs to the sender and receiver for any given choice x of the sender.

We conclude that Conditions A.1, A.2, and A.3 all hold after we eliminate strictly dominated strategies from the space of translation invariant strategies. We can therefore conclude from Theorem A.1 that any profile of undominated translation invariant strategies (α, π) is *admissible*. They induce expected payoffs equal to the limit of payoffs (for those strategies) under any diffusing sequence of priors over θ . The notion of expected payoffs of given strategies with respect to θ is therefore well-defined, equal to the expected payoffs under *any* realization of θ , and consistent with our posterior belief calculations.

B Proofs for Section 2

Proof of Theorem 1. Recall that the sender wishes to maximize $u_S(a, \theta) = -(a - (\theta + M_S))^2$. Since $a = \pi(\vec{x}, y) + \sigma(\alpha)$ from the definition of shift, the sender’s goal is to choose π so that $\pi(\vec{x}, y)$ maximizes

$$-\mathbb{E}_{\theta}[(\pi(\vec{x}, y) + \sigma(\alpha) - (\theta + M_S))^2 \mid (\vec{x}, y)].$$

This is the expectation of a quadratic loss. Using bias-variance decomposition and the fact that $\theta_S(\vec{x}, y) = \mathbb{E}_{\theta}[\theta \mid (\vec{x}, y)]$ is an unbiased estimator of θ and the variance of $\theta_S(\vec{x}, y)$ is a constant that is independent of $\pi(\cdot)$, this goal is achieved by choosing $\pi(\vec{x}, y)$ to maximize

$$-(\pi(\vec{x}, y) + \sigma(\alpha) - (\theta_S(\vec{x}, y) + M_S))^2$$

for each \vec{x} and y . For any realization of \vec{x} and y , this expression is maximized by setting $\pi(\vec{x}, y)$ as close as possible to $\theta_S(\vec{x}, y) + M_S - \sigma(\alpha)$. Since the only constraint on $\pi(\vec{x}, y)$ is that it be chosen from the profile of anecdotes \vec{x} , the result follows. \square

Proof of Theorem 2. Let $x = \pi(\vec{x})$ and let the belief distribution be $B(x) = D_{\pi, x}$. Given any translation invariant π , the posterior distribution $D_{\pi, x}$ is a proper distribution (see Claim A.1 in Appendix A for more details). Since bias equals the expected difference between the communicated anecdote and θ , we have that $x - \beta(\pi)$ is an unbiased estimate of θ . Since we work with diffuse priors, the receiver’s posterior mean about the state of the world is simply equal to the value of

the unbiased estimator (formalized in Claim A.1 in Appendix A). The receiver wishes to maximize

$$\begin{aligned}\mathbb{E}_{\theta \sim D_{\pi,x}} [u_R(\alpha(x), \theta)] &= \mathbb{E}_{\theta \sim D_{\pi,x}} [-(\alpha(x) - (\theta + M_R))^2] \\ &= -(\alpha(x) - (x - \beta(\pi) + M_R))^2 - \mathbb{E}_{\theta \sim D_{\pi,x}} [(\theta - x + \beta(\pi))^2]\end{aligned}$$

where the second equality is the bias-variance decomposition and follows because $x - \beta(\pi)$ is an unbiased estimator of θ . This maximum is achieved for $\alpha(x) = x - \beta(\pi) + M_R$, i.e., a translation of $M_R - \beta(\pi)$ as claimed. \square

Proof of Proposition 2. The proof follows by direct manipulations of the sender's expected utility function.

$$\begin{aligned}\mathbb{E}_{\theta, \vec{x}, y} [u_S(\alpha(\pi_S(\vec{x}, y)), \theta)] &= \mathbb{E}_{\theta, \vec{x}, y} [((\alpha(\pi_S(\vec{x}, y)) - (\theta + M_S))^2)] \\ &= \mathbb{E}_{\vec{x}, y} [((\alpha(\pi_S(\vec{x}, y)) - (\theta_S(\vec{x}, y) + M_S))^2)] \\ &= \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) + \sigma(\alpha) - (\theta_S(\vec{x}, y) + M_S))^2] \\ &= \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) + \sigma(\alpha) - (\theta_S(\vec{x}, y) + M_S) + \beta(\pi_S) - \beta(\pi_S))^2] \\ &= \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)) + \sigma(\alpha) - (M_S - \beta(\pi_S)))^2]\end{aligned}$$

Let $w = \pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S))$ and $z = \sigma(\alpha) - (M_S - \beta(\pi_S))$ so that the above expectation is $\mathbb{E}[(w + z)^2] = \mathbb{E}[w^2 + z^2 + 2wz]$. Note

$$\mathbb{E}[wz] = z \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)))] = z \mathbb{E}_{\theta, \vec{x}, y} [(\pi_S(\vec{x}, y) - (\theta + \beta(\pi_S)))]$$

where the second inequality follows because $\theta_S(\vec{x}, y)$ is a valid posterior mean for θ . But the right-hand side is zero by definition of bias. Therefore $\mathbb{E}[wz] = 0$ and hence the claim follows. \square

Proof of Proposition 3. From Theorem 2 we see that $\alpha(x) = x + M_R - \beta(\pi)$, where the sender sends signal $x = \pi(\vec{x}, y)$ and $\beta(\pi)$ is the bias of the communication scheme π . Thus, we have receiver's loss (for any fixed θ) equals

$$\begin{aligned}\mathbb{E}_{\vec{x}, y} [(\pi(\vec{x}, y) + M_R - \beta(\pi) - \theta - M_R)^2] &= \mathbb{E}_{\vec{x}, y} [(\pi(\vec{x}, y) - \theta - \beta(\pi))^2] \\ &= \mathbb{E}_{\vec{x}, y} [(\pi(\vec{x}, y) - \mathbb{E}_{\vec{x}, y} [\pi(\vec{x}, y)])^2]\end{aligned}$$

since by definition of bias $\mathbb{E}_{\vec{x} \sim F_\theta, y \sim G_\theta} [\pi(\vec{x}, y)] = \beta(\pi) + \theta$. \square

C Proofs for Section 3

Proof of Proposition 2. The proof follows by direct manipulations of the sender's expected utility function.

$$\begin{aligned}
\mathbb{E}_{\theta, \vec{x}, y} [u_S(\alpha(\pi_S(\vec{x}, y)), \theta)] &= \mathbb{E}_{\theta, \vec{x}, y} [((\alpha(\pi_S(\vec{x}, y)) - (\theta + M_S))^2)] \\
&= \mathbb{E}_{\vec{x}, y} [((\alpha(\pi_S(\vec{x}, y)) - (\theta_S(\vec{x}, y) + M_S))^2)] \\
&= \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) + \sigma(\alpha) - (\theta_S(\vec{x}, y) + M_S))^2] \\
&= \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) + \sigma(\alpha) - (\theta_S(\vec{x}, y) + M_S) + \beta(\pi_S) - \beta(\pi_S))^2] \\
&= \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)) + \sigma(\alpha) - (M_S - \beta(\pi_S)))^2]
\end{aligned}$$

Let $w = \pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S))$ and $z = \sigma(\alpha) - (M_S - \beta(\pi_S))$ so that the above expectation is $\mathbb{E}[(w + z)^2] = \mathbb{E}[w^2 + z^2 + 2wz]$. Note

$$\mathbb{E}[wz] = z \mathbb{E}_{\vec{x}, y} [(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)))] = z \mathbb{E}_{\theta, \vec{x}, y} [(\pi_S(\vec{x}, y) - (\theta + \beta(\pi_S)))]$$

where the second inequality follows because $\theta_S(\vec{x}, y)$ is a valid posterior mean for θ . But the right-hand side is zero by definition of bias. Therefore $\mathbb{E}[wz] = 0$ and hence the claim follows. \square

Proof of Theorem 3. We prove the Theorem through contradiction. Assume, that π^* is a most informative communication scheme with bias β^* and utility $U(\pi^*)$ which is not a targeting scheme satisfying $r = \beta^*$.

Consider the targeting scheme π^T with offset $r = \beta^*$ and associated bias β^T . We know that $\beta^T \neq \beta^*$ - otherwise, we would have found a targeting with bias β^* which minimizes the Euclidean distance to $\beta^* + \theta_S(\vec{x}, y)$ and therefore $U(\pi^T) > U(\pi^*)$ because $\pi^* \neq \pi^T$. Therefore, π^* would not be most informative which is a contradiction.

We now calculate the sender's utility when using targeting scheme π^T :

$$\begin{aligned}
U(\pi^T) &= -\mathbb{E} \left[(\pi^T(\vec{x}, y) - \beta^T - \theta_S(\vec{x}, y))^2 \right] \\
&= -\mathbb{E} \left[(\pi^T(\vec{x}, y) - \beta^* - \theta_S(\vec{x}, y) + \beta^* - \beta^T)^2 \right] \\
&= -\mathbb{E} \left[(\pi^T(\vec{x}, y) - \beta^* - \theta_S(\vec{x}, y))^2 \right] - 2(\beta^* - \beta^T) \underbrace{\mathbb{E} [\hat{\pi}(\vec{x}, y) - \beta^* - \theta_S(\vec{x}, y)]}_{\beta^T - \beta^*} - (\beta^* - \beta^T)^2 \\
&= -\underbrace{\mathbb{E} \left[(\pi^T(\vec{x}, y) - \beta^* - \theta_S(\vec{x}, y))^2 \right]}_{< (\pi^*(\vec{x}, y) - \beta^* - \theta_S(\vec{x}, y))^2} + (\beta^* - \beta^T)^2 > U(\pi^*) \tag{14}
\end{aligned}$$

But this implies that π^* is not a most informative communication scheme which is a contradiction.

□

Proof of Lemma 1. First, it is easy to see that the balanced-offset condition is satisfied at $r^* = 0$: the closest anecdote is equally likely to be on the left or right because of symmetry.

Now consider any $r < 0$ and any realization \vec{x} of n anecdotes. We construct a coupled realization \tilde{x} as follows: for any $x_i < r$ we keep the anecdote and $\tilde{x}_i = x_i$. For any $r \leq x_i \leq -r$ we map the anecdote to to $\tilde{x}_i = r$ and for any $x_i > -r$ we map the anecdote to $\tilde{x}_i = x_i - 2r$.

We can see that the coupled distribution is symmetric around r and therefore $\tilde{\beta}(r) = r$. However, because of the coupling all the anecdotes to the right of r are closer to the offset than for the coupled original realization. Therefore, $\tilde{\beta}(r) < \beta(r)$ and hence $r < \beta(r)$. Hence, the targeting scheme with offset r cannot be most informative. □

Proof of Lemma 2. We prove the claim by contradiction. If the lemma does not hold then there is some $\epsilon^* > 0$ and a sequence (n_k) and associated most informative offsets $(r_{n_k}^*)$ that are more than ϵ^* away from any global maximum of the anecdote density.

First, we observe that there has to be some $\delta > 0$ such that $f(r_{n_k}^*) < f_{\max} - \delta$ for all k where f_{\max} is the global maximum reached by the density function. Suppose not, there would be a subsequence $(n_{k'})$ of (n_k) such that $f(r_{n_{k'}}^*) \rightarrow f_{\max}$. Moreover, there has to be a subsequence that lies in some bounded interval - otherwise because of uniform continuity the integral of the density function could not be 1. But this implies by Bolzano-Weierstrass that there is convergent subsequence $(n_{k''})$ of $(n_{k'})$. By continuity, $f(\lim_{k'' \rightarrow \infty} r_{n_{k''}}^*) = f_{\max}$ and therefore $r_{n_{k''}}^*$ eventually is ϵ^* -close to a global maximum which is a contradiction.

We can therefore assume that there exists some $\delta > 0$ such that $f(r_{n_k}^*) < f_{\max} - \delta$ for all k . Because of uniform continuity we can find, for any $\delta' > 0$ an $\epsilon' > 0$ such that $|f(x) - f(r_{n_k}^*)| < \delta'$ for any $x - r_{n_k}^* < \epsilon'$.

The information loss, or equivalently the variance of the closest anecdote to the offset $r_{n_k}^*$, can then be expressed as:

$$\text{Var}(\pi_{r_{n_k}^*}(\vec{x})) = \mathbb{E}[(\pi_{r_{n_k}^*}(\vec{x}) - (\theta + r_{n_k}^*))^2],$$

the expected squared distance of the closest anecdote to the point $\theta + r_{n_k}^*$, because most informative communication schemes satisfy the balanced-offset condition $\beta(r_{n_k}^*) = r_{n_k}^*$. Hence, we have,

$$\begin{aligned} \text{Var}(\pi_{r_{n_k}^*}(\vec{x})) &= \int_0^\infty nz^2 [f(r_{n_k}^* + z) + f(r_{n_k}^* - z)] [1 - F(r_{n_k}^* + z) + F(r_{n_k}^* - z)]^{n-1} dz \\ &\geq \int_0^{\epsilon'} nz^2 2(f(r_{n_k}^*) - \delta') [1 - 2(f(r_{n_k}^*) + \delta')z]^{n-1} dz \\ &= \frac{f(r_{n_k}^*) - \delta'}{f(r_{n_k}^*) + \delta'} \frac{1}{2n^2(f(r_{n_k}^*) + \delta')^2} + o(1/n^2) \\ &> \frac{f(r_{n_k}^*) - \delta'}{f(r_{n_k}^*) + \delta'} \frac{1}{2n^2(f_{\max} - \delta + \delta')^2} + o(1/n^2), \end{aligned} \tag{15}$$

where the first inequality follows from the uniform continuity of f and let $\epsilon' < 1/2(f(r_{n_k}^*) + \delta')$, the

next equality follows simply because the $o(1/n^2)$ terms includes subtractions that are exponentially small in n , and finally the last inequality follows by our assumption that $r_{n_k}^*$ is bounded away from the global maxima.

We can compare this loss to the information loss when choosing a targeting scheme with offset \tilde{r} such that $f(\tilde{r}) = f_{\max}$:

$$\begin{aligned}
\text{Var}(\pi_{\tilde{r}}(\vec{x})) &= \mathbb{E}[(\pi_{\tilde{r}}(\vec{x}) - (\theta + \tilde{r}))^2] - (\beta(\tilde{r}) - \tilde{r})^2 \\
&\leq \int_0^\infty nz^2 [f(\tilde{r} + z) + f(\tilde{r} - z)] [1 - F(\tilde{r} + z) + F(\tilde{r} - z)]^{n-1} dz \\
&\leq \int_0^{\epsilon'} nz^2 2(f(\tilde{r}) + \delta') [1 - 2(f(\tilde{r}) - \delta')z]^{n-1} dz \\
&+ \int_{\epsilon'}^\infty nz^2 [f(\tilde{r} + z) + f(\tilde{r} - z)] [1 - F(\tilde{r} + z) + F(\tilde{r} - z)]^{n-1} dz \\
&\leq \int_0^{\epsilon'} nz^2 2(f(\tilde{r}) + \delta') [1 - 2(f(\tilde{r}) - \delta')z]^{n-1} dz \\
&+ \int_{\epsilon'}^\infty nz^2 [f(\tilde{r} + z) + f(\tilde{r} - z)] [1 - 2\epsilon'(f(\tilde{r}) - \delta')]^{n-1} dz \\
&= \frac{f_{\max} + \delta'}{f_{\max} - \delta'} \frac{1}{2n^2(f_{\max} - \delta')^2} + o(1/n^2)
\end{aligned} \tag{16}$$

For any $r_{n_k}^*$ we can choose δ' small enough such that $\text{Var}(\pi_{r_{n_k}^*}(\vec{x})) > \text{Var}(\pi_{\tilde{r}}(\vec{x}))$ for sufficiently large n_k . Therefore, $r_{n_k}^*$ is not a most informative offset for some k as it is dominated by a targeting scheme with offset \tilde{r} and lower information loss. This is a contradiction. \square

Proof of Proposition 3. From Theorem 2 we see that $\alpha(x) = x + M_R - \beta(\pi)$, where the sender sends signal $x = \pi(\vec{x}, y)$ and $\beta(\pi)$ is the bias of the communication scheme π . Thus, we have receiver's loss (for any fixed θ) equals

$$\begin{aligned}
\mathbb{E}_{\vec{x}, y} [(\pi(\vec{x}, y) + M_R - \beta(\pi) - \theta - M_R)^2] &= \mathbb{E}_{\vec{x}, y} [(\pi(\vec{x}, y) - \theta - \beta(\pi))^2] \\
&= \mathbb{E}_{\vec{x}, y} [(\pi(\vec{x}, y) - \mathbb{E}_{\vec{x}, y}[\pi(\vec{x}, y)])^2]
\end{aligned}$$

since by definition of bias $\mathbb{E}_{\vec{x} \sim F_\theta, y \sim G_\theta} [\pi(\vec{x}, y)] = \beta(\pi) + \theta$. \square

D Proofs for Section 4

Proof of Theorem 4. Suppose (π, α) is a translation invariant PBE. Then α must be a translation, with shift $\sigma(\alpha)$. Then by Theorem 1 we know that π is a targeting scheme with offset $r = M_S - \sigma(\alpha)$. Moreover, by Theorem 2 we must have that $\sigma(\alpha) = M_R - \beta(\pi)$. Since by definition $\beta(r) = \beta(\pi)$,

we conclude that

$$r = M_S - (M_R - \beta(r)) = \beta(r) + (M_S - M_R)$$

as required.

The other direction follows immediately from Theorems 2 and 1, because π and α are best responses to each other. Note that, in Theorem 2, $\alpha(x)$ maximizes the receiver utility given the belief distribution $D_{\pi,x}$. Hence, we get that (π, α) with belief distribution $B(x) = D_{\pi,x}$ is a translation invariant PBE. \square

Proof of Theorem 5. Fix n and $M_S - M_R$. Let's start with finding a condition that pins down the offset r of a PBE from Theorem 4. Given a communication scheme π_r with offset r , let $z = \pi_r(\vec{x}, y) - (\theta_S(\vec{x}, y) + r)$ denote the distance between the target $\theta_S(\vec{x}, y) + r$ and the closest anecdote (out of n total anecdotes), where z is positive if the closest anecdote is larger and negative if the closest anecdote is smaller. Write $h(z; r)$ for the density of z given r , over all randomness in (θ, \vec{x}, y) .

We can now calculate the expected bias $\beta(r)$ of a targeting communication scheme with offset r :

$$\beta(r) = r + \int_{-\infty}^{\infty} zh(z; r)dz \quad (17)$$

Write $H(r) = \int_{-\infty}^{\infty} zh(z; r)dz$. Theorem 4 now implies that to show that a PBE exists, it suffices to show that there exists a value of r such that

$$H(r) = -(M_S - M_R). \quad (18)$$

We will show that (18) holds, i.e., $H(r)$ is an onto function in its codomain $(-\infty, +\infty)$, in two steps. In the first step, we establish that $H(r)$ tends to $+\infty$ and $-\infty$, respectively as $r \rightarrow -\infty$ and $r \rightarrow +\infty$. Had $H(r)$ been a continuous function, this step would suffice to prove (18). The second step handles discontinuous $H(r)$ by establishing that at any point that $H(r)$ is discontinuous, the left limit of $H(r)$ is smaller than the right limit. Together these steps show that $H(r)$ can take any value.

For the first step, we will show that $H(r) \rightarrow \infty$ as $r \rightarrow -\infty$ and $H(r) \rightarrow -\infty$ as $r \rightarrow \infty$. We know $\mathbb{E}_{\epsilon \sim F}[|\epsilon|]$ is bounded by assumption; say $\mathbb{E}[|x - \theta|] < c_0$. Then

$$\begin{aligned} \mathbb{E}_{\theta, \vec{x}, y} [\max_i x_i - \theta_S(\vec{x}, y)] &= \mathbb{E}_{\theta, \vec{x}, y} [\max_i x_i - \theta + \theta - \theta_S(\vec{x}, y)] \\ &= \mathbb{E}_{\theta, \vec{x}, y} [\max_i x_i - \theta] + E_{\theta, \vec{x}, y}[\theta - \theta_S(\vec{x}, y)] \\ &= \mathbb{E}_{\theta, \vec{x}} [\max_i x_i - \theta] + 0 \\ &\leq \sum_i E_{\theta, x_i} [|x_i - \theta|] \\ &\leq nc_0 \end{aligned}$$

where the second equality is linearity of expectation and the third equality follows because $\theta_S(\vec{x}, y)$ is a valid posterior mean.

Now choose any $Z > 0$ and suppose $r \geq nc_0 + Z$. Then

$$\begin{aligned} \int_{-\infty}^{\infty} zh(z; r)dz &= \mathbb{E}_{\theta, \vec{x}, y} [(\arg \min_{x_i \in \vec{x}} |x_i - r - \theta_S(\vec{x}, y)|) - r - \theta_S(\vec{x}, y)] \\ &\leq \mathbb{E}_{\theta, \vec{x}, y} [\max_i x_i - r - \theta_S(\vec{x}, y)] \\ &\leq -Z. \end{aligned}$$

So for any $Z > 0$, we have that $H(r) \leq -Z$ for all sufficiently large r , and hence $H(r) \rightarrow -\infty$ as $r \rightarrow \infty$. A symmetric argument¹⁸ shows that $H(r) \rightarrow \infty$ as $r \rightarrow -\infty$.

In the second step we show that, roughly speaking, if $H(r)$ is discontinuous at some r then the one-sided limits still exist, and the limit from above will be strictly greater than the limit from below. To see why, suppose H is not continuous at r_0 . For each possible realization of (θ, \vec{x}, y) , either z is continuous at r_0 or it is not. If not, this means that $\theta_S(\vec{x}, y) + r_0$ is precisely halfway between two anecdotes in \vec{x} , say with absolute distance $d > 0$ to each, in which case the limit of z from below is $-d$ (distance to the anecdote to the left) and the limit of z from above is d (distance to the anecdote to the right). Integrating over all realizations, we conclude that the one-sided limits of H exist, $\lim_{r \rightarrow r_0^-} H(r) < \lim_{r \rightarrow r_0^+} H(r)$, and moreover $\lim_{r \rightarrow r_0^-} H(r) \leq H(r_0) \leq \lim_{r \rightarrow r_0^+} H(r)$.

Now we are ready to prove (18). Since $H(r) \rightarrow -\infty$ as $r \rightarrow \infty$, there must exist some finite r_1 such that $H(r_1) < -(M_S - M_R)$. Choose $r_2 \leq r_1$ to be the infimum over all r' such that $H(r) \leq -(M_S - M_R)$ for all $r \in (r', r_1]$. That is, $(r_2, r_1]$ is a maximal (on the left) interval on which $H(r) \leq -(M_S - M_R)$. Note that r_2 must be finite, since $H(r) \rightarrow \infty$ as $r \rightarrow -\infty$.

Suppose for contradiction that $H(r_2) \neq -(M_S - M_R)$. It must then be that H is discontinuous at r_2 , as otherwise there is an open ball around r_2 on which H is either less than or greater than $-(M_S - M_R)$, but either way this contradicts the definition of r_2 .

From the definition of r_2 we have that $\lim_{r \rightarrow r_2^+} H(r) \leq -(M_S - M_R)$. So since H is discontinuous at r_2 , we know (from our analysis of the directionality of discontinuities of H) that $\lim_{r \rightarrow r_2^-} H(r) < -(M_S - M_R)$. But this then means that there exists some $\epsilon > 0$ such that $H(r) < -(M_S - M_R)$ for all $r \in (r_2 - \epsilon, r_2)$, contradicting our choice of r_2 .

We conclude that $H(r_2) = -(M_S - M_R)$, so r_2 is the desired value of r proving (18). \square

Proof Proposition 4. Suppose $\Delta = M_R - M_S > 0$ (the other case is analogous). Recall that $r(\Delta) = \max_r \{r : H(r) = \Delta\} < 0$ is the maximum offset of any equilibrium targeting scheme. Assume towards contradiction that there exists $\Delta' > \Delta$ such that $r(\Delta') > r(\Delta)$. Recall that, by definition, $r(\Delta') < \underline{r}$. We can therefore let $r^* = \inf_{r \in [r'(\Delta), \underline{r}]}\{r : H(r) \leq \Delta\}$. This infimum is well-defined since $H(\underline{r}) = 0 < \Delta$.

We claim that $H(r^*) = \Delta$. This is because, from the proof of Theorem 5, we know that one-sided limits of H exist at r^* and $\lim_{r' \rightarrow r^{*-}} H(r) \leq H(r^*) \leq \lim_{r' \rightarrow r^{*+}} H(r)$. But from the definition

¹⁸By taking $r = -nc_0 - Z$ and observing $\mathbb{E}[\min_i -\theta] \geq -nc_0$.

of r^* we have that (a) $\Delta \leq \lim_{r' \rightarrow r^{*+}} H(r)$, and (b) $H(r') > \Delta$ for all $r' \in [r(\Delta), r^*]$. So the only way to satisfy $\lim_{r' \rightarrow r^{*-}} H(r) \leq \lim_{r' \rightarrow r^{*+}} H(r)$ is to have $\lim_{r' \rightarrow r^{*-}} H(r) = \lim_{r' \rightarrow r^{*+}} H(r) = \Delta$ and hence $H(r^*) = \Delta$.

We have now reached the desired contradiction, since $H(r^*) = \Delta$ and $r^* > r'(\Delta) > r(\Delta)$, contradicting the definition of $r(\Delta)$. \square

Proof of Corollary 2. Consider $\Delta = M_R - M_S > 0$ (the other case is analogous). We then have $r < \beta(r) < \underline{r}$ and $\beta(r) - r = \Delta$. This implies $r < \underline{r} - \Delta$. Hence, $r \rightarrow -\infty$ as $\Delta \rightarrow \infty$. \square

Proof of Proposition 5. Suppose $\Delta < 0$ such that $r(\Delta) \geq \bar{r}$. Let $\pi_r(\vec{x}) = \arg \min_{x_i \in \vec{x}} |x_i - (\theta + r)|$ be the closest anecdote to the offset $\theta + r$ and let $z = \pi_r(\vec{x}) - (\theta + r)$ denote the corresponding distance. The bias $\beta(r) \in (0, r)$ then equals $\beta(r) = \mathbb{E}[\pi_r(\vec{x})] - \theta$. The variance of $\pi_r(\vec{x})$ is given by

$$\begin{aligned} \text{Var}[\pi_r(\vec{x})] &= \mathbb{E}[(\pi_r(\vec{x}) - (\theta + \beta(r)))^2] \\ &= \mathbb{E}[(z - (\beta(r) - r))^2] \\ &= \mathbb{E}[z^2] - (\beta(r) - r)^2 \quad (\text{Since } \mathbb{E}[z] = (\beta(r) - r)) \\ &= n \int_0^\infty z^2 (f(z+r) + f(r-z))(1 - F(r+z) + F(r-z))^{n-1} dz - (\beta(r) - r)^2 \end{aligned}$$

Let $P(r, z) = 1 - F(r+z) + F(r-z)$ be the probability that an anecdote does not lie in the interval $(\theta + r - z, \theta + r + z)$. Hence we can re-write the above equality as,

$$\begin{aligned} \text{Var}[\pi_r(\vec{x})] &= - \int_0^\infty z^2 \left(\frac{\partial P(r, z)^n}{\partial z} \right) dz - (\beta(r) - r)^2 \\ &= -z^2 H(r, z)^n \Big|_{z=0} + \int_0^\infty 2z P(r, z)^n dz - (\beta(r) - r)^2 \\ &= \int_0^\infty 2z P(r, z)^n dz - (\beta(r) - r)^2 \end{aligned}$$

The last equality follows because $z^2 P(r, z)^n = 0$ at $z = 0$ and we have $\lim_{z \rightarrow \infty} z^2 \cdot P(r, z) = 0$ for $n \geq 2$. To see why, note that, for any $z > 2|r|$, observe that if $|x - \theta| \leq z/2$ then $|x - (\theta + r)| \leq z$. By Chebyshev's inequality we have

$$P(r, z) \leq \Pr_{x \sim F} [|x - \theta| \geq z/2] \leq 4\sigma^2/z^2,$$

where θ is the mean of the anecdote distribution and σ^2 is its (finite) variance. Hence for $n \geq 2$ we have $\lim_{z \rightarrow \infty} z^2 P(r, z)^n = 0$. We therefore conclude $z^2 P(r, z)^n \Big|_{z=0}^\infty = 0$.

This implies that,

$$\begin{aligned}
\frac{\partial \text{Var}[\pi_r(\vec{x})]}{\partial r} &= \int_0^\infty 2nz \frac{\partial P(r, z)}{\partial r} P(r, z)^{n-1} dz - 2(\beta(r) - r)(\beta'(r) - 1) \\
&= -2 \mathbb{E}[z] - 2(\beta(r) - r)(\beta'(r) - 1) \\
&= 2(r - \beta(r)) + 2(r - \beta(r))(\beta'(r) - 1) \\
&\geq 0 \quad (\text{Since } \beta'(r) \geq 0)
\end{aligned}$$

where the equalities follows because the expected value of z is $\beta(r) - r$ and the last inequality follows because the bias increases with r . Hence, for $r \geq \bar{r}$, the variance increases with r . A similar argument follows for $r < \underline{r}$. □

E Proofs for Section 5

Proof of Proposition 6. We prove this result by contradiction. Assume that the result does not hold and hence there is some $\Delta \neq 0$, some distribution F , and constants $R_L > r_{\min}$ and $R_M < r_{\max}$ and a subsequence (n_k) and associated equilibrium offsets $r^*(n_k)$ such that $R_L \leq r^*(n_k) \leq R_M$.

Because of continuity there is some $f_{\min}, f_{\max} > 0$ such that $f_{\min} \leq f(x) \leq f_{\max}$ for any $x \in [R_L, R_M]$. We can then write for any $r \in [R_L, R_M]$:

$$\beta_n(r) - r = \int_0^\infty nz [f(r+z) - f(r-z)] [1 - F(r+z) + F(r-z)]^{n-1} dz \quad (19)$$

Fix some $0 < \epsilon' < \frac{1}{2}f_{\min}$ and we get:

$$\begin{aligned}
-\int_0^{\epsilon'} nz f_{\max} [1 - 2zf_{\min}]^{n-1} dz + o(1/n^2) &\leq \beta_n(r) - r \\
&\leq \int_0^{\epsilon'} nz f_{\max} [1 - 2zf_{\min}]^{n-1} dz + o(1/n^2) \quad (20)
\end{aligned}$$

This allows us to bound $\beta(r) - r$ as follows:

$$|\beta_n(r) - r| \leq \frac{f_{\max}}{2f_{\min}} \frac{1}{n} + o(1/n) \quad (21)$$

Therefore, for sufficiently large n_k we have $\beta_{n_k}(r_{n_k}^*) - r_{n_k}^* \neq \Delta$ and therefore $r_{n_k}^*$ cannot be an equilibrium offset. □

Proof of Theorem 6 (heavy tails). The sent anecdote is y away from offset $r > 0$ and has density $h(y)$:

$$h(y) = nf(r+y) [1 - (F(r+|y|) - F(r-|y|))]^{n-1} \quad (22)$$

We prove the claim for heavy tails through contradiction: assume that $\text{Var}_n(y) \not\rightarrow \infty$. This

implies that there exists some $M > 0$ and a subsequence (n_k) with associated offsets r_{n_k} such that $\text{Var}_{n_k}(y) < M$.

We next introduce two symmetric and nested intervals around the offset as well as a formula for the probability that y falls into such an interval as well as conditional expectation of y in this case.

Growing interval I_k . We define $I_k = [r_{n_k} - b_k, r_{n_k} + b_k]$ where $b_k = r_{n_k}^\tau$ for some $0 < \tau < 1$ (where τ will be specified later). This is a symmetric interval around the offset r_{n_k} which gradually increases in width with r_{n_k} . Importantly, it doesn't increase linearly with r_{n_k} but only at rate $r_{n_k}^\tau$ - this ensures that the width of the interval becomes small relative to the offset.

Nested interval \hat{I}_k . We next use Chebyshev's inequality which tells us that for any fixed $\epsilon > 0$:

$$\text{Prob} \left(|y + \Delta| \geq \sqrt{\frac{M}{\epsilon}} \right) \leq \epsilon \quad (23)$$

This implies that we can ensure that with probability at least ϵ the closest anecdote is in the interval $\hat{I}_{n_k} = \left[r_{n_k} - \sqrt{\frac{M}{\epsilon}} - |\Delta|, r_{n_k} + \sqrt{\frac{M}{\epsilon}} + |\Delta| \right]$ for all n_k . Note that for sufficiently large k we have $\hat{I}_k \subset I_k$.

Formula for $P(y \in I)$. Consider any symmetric interval $I = [r_{n_k} - b, r_{n_k} + b]$ around the offset (such as I_k or \hat{I}_k).

Then the probability $P(y \in I)$ can be expressed as:

$$\begin{aligned} P(y \in I) &= \int_{-b}^0 n f(r+y) [1 - (F(r+|y|) - F(r-|y|))]^{n-1} dy \\ &+ \int_0^b n f(r+y) [1 - (F(r+|y|) - F(r-|y|))]^{n-1} dy \\ &= \int_0^b n [f(r+y) + f(r-y)] [1 - (F(r+y) - F(r-y))]^{n-1} dy \\ &= 1 - [1 - (F(r+b) - F(r-b))]^n \end{aligned} \quad (24)$$

Formula for $E(y|y \in I)$. We can use an analogous derivation to show:

$$E(y|y \in I) = \int_0^b y n [f(r+y) - f(r-y)] [1 - (F(r+y) - F(r-y))]^{n-1} dy \quad (25)$$

The rest of the proof for heavy tails proceeds in three steps:

Step 1: We show that the density f of the anecdote distribution is essentially constant over the interval I_k . This allows us to simplify equation 24 for any $I = [r_{n_k} - b, r_{n_k} + b] \subset I_k$ as follows:

$$P(y \in I) \approx 1 - [1 - 2f(r)b]^n \quad (26)$$

Step 2: We then show that $P(y \in \hat{I}_k) \geq 1 - \epsilon$ implies $P(y \notin I_k) < \epsilon^{Dr_{n_k}^\tau}$ for some constant D .

Step 3: We can now finally show that $E(y) \rightarrow 0$ which is a contradiction since $E(y) = \Delta$. Hence our initial assumption that $Var_n(y) \not\rightarrow \infty$ was false.

Step 1. The density $h(y)$ of the sent anecdote includes the term

$$1 - (F(r + y) - F(r - y)) \quad (27)$$

for $y > 0$. Using the mean-value theorem we get

$$1 - (F(r + y) - F(r - y)) = 1 - y(f(r + \zeta) + f(z - \zeta)) \quad (28)$$

for some $\zeta \in [0, y]$. Since f has strong heavy tails we know that (for some constant $C > 0$):

$$\begin{aligned} F(x) &= 1 - C \exp\left(-\int_{\underline{x}}^x u(t)t^{\alpha-1} dt\right) \\ f(x) &= Cu(x)x^{\alpha-1} \exp\left(-\int_{\underline{x}}^x u(t)t^{\alpha-1} dt\right) \end{aligned} \quad (29)$$

We can now express $f(r + \zeta)$ (and analogously $f(r - \zeta)$) as

$$\begin{aligned} f(r + \zeta) &= Cu(r + \zeta)(r + \zeta)^{\alpha-1} \exp\left(-\int_{\underline{x}}^{r+\zeta} u(t)t^{\alpha-1} dt\right) \\ &= Cu(r + \zeta)r^{\alpha-1}\left(1 + \frac{\zeta}{r}\right)^{\alpha-1} \exp\left(-\int_{\underline{x}}^r u(t)t^{\alpha-1} dt\right) \exp\left(-\int_r^{r+\zeta} u(t)t^{\alpha-1} dt\right) \\ &= f(r) \frac{u(r + \zeta)}{u(r)} \left(1 + \frac{\zeta}{r}\right)^{\alpha-1} \exp(-\zeta u(r + \zeta')(r + \zeta')^{\alpha-1}) \end{aligned} \quad (30)$$

where $\zeta' \in [0, \zeta]$. Next we use the mean-value theorem again:

$$\begin{aligned} \frac{u(r + \zeta)}{u(r)} &= \exp(\ln(u(r + \zeta)) - \ln(u(r))) \\ &= \exp\left(\zeta \frac{u'(r + \zeta'')}{u(r + \zeta'')}\right) \text{ for some } \zeta'' \in [0, \zeta] \end{aligned} \quad (31)$$

We obtain:

$$f(r + \zeta) = f(r) \exp\left(\zeta \frac{u'(r + \zeta'')}{u(r + \zeta'')}\right) \left(1 + \frac{\zeta}{r}\right)^{\alpha-1} \exp(-\zeta u(r + \zeta')r^{\alpha-1} \left(1 + \frac{\zeta'}{r}\right)^{\alpha-1}) \quad (32)$$

We observe:

$$\begin{aligned}
\left| \zeta \frac{u'(r + \zeta'')}{u(r + \zeta'')} \right| &\leq \left| (r + \zeta'')^\beta \frac{u'(r + \zeta'')}{u(r + \zeta'')} \right| \rightarrow 0 \quad \text{for } \tau < \beta \text{ and suff. large } r \\
\frac{\zeta}{r} &\rightarrow 0 \\
\frac{\zeta'}{r} &\rightarrow 0 \\
\zeta r^{\alpha-1} &\leq r^\tau r^{\alpha-1} \rightarrow 0 \quad \text{for } 0 < \tau < 1 - \alpha
\end{aligned} \tag{33}$$

Therefore, as long as $0 < \tau < \min(1 - \alpha, \beta)$ we can ensure that for any small $\vartheta > 0$ we have for sufficiently large n_k :

$$\frac{f(r + \zeta)}{f(r)} \in [1 - \vartheta, 1 + \vartheta] \tag{34}$$

Similarly, we can derive

$$\frac{f(r - \zeta)}{f(r)} \in [1 - \vartheta, 1 + \vartheta] \tag{35}$$

This allows us to bound equation 28:

$$1 - 2yf(r)(1 + \vartheta) < 1 - (F(r + y) - F(r - y)) < 1 - 2yf(r)(1 - \vartheta) \tag{36}$$

More generally, we can bound formula 24 for $P(y \in I)$ for any nested interval $I = [r_{n_k} - b, r_{n_k} + b]$ in I_k for all k as:

$$1 - [1 - 2bf(r)(1 - \vartheta)]^n < P(y \in I) < 1 - [1 - 2bf(r)(1 + \vartheta)]^n \tag{37}$$

Once fixing $\vartheta > 0$ this bound holds for all sufficiently large n_k . This completes Step 1 of the argument.

Step 2. Recall, that Chebychev's inequality ensured us that the sent anecdote lies in the \hat{I}_k interval with probability of at least $1 - \epsilon$:

$$P(y \in \hat{I}_k) \geq 1 - \epsilon$$

By using inequality 37 from Step 1 and setting $\hat{b} = \sqrt{\frac{M}{\epsilon}} + |\Delta|$ we obtain:

$$\begin{aligned}
1 - [1 - 2\hat{b}f(r_{n_k})(1 + \vartheta)]^n &> 1 - \epsilon \\
[1 - 2\hat{b}f(r_{n_k})(1 + \vartheta)]^n &< \epsilon
\end{aligned} \tag{38}$$

We also can use inequality 37 to bound $P(y \notin I_k)$:

$$P(y \notin I_k) < [1 - 2r_{n_k}^\tau f(r_{n_k})(1 - \vartheta)]^n \tag{39}$$

We now introduce a helper lemma.

Lemma E.1. Assume $A > 1$ and x such that $Ax < 1$. Then the following holds:

$$(1 - Ax)^n \leq (1 - x)^{An} \quad (40)$$

Proof: The claim is equivalent to:

$$\ln(1 - Ax) \leq A \ln(1 - x) \quad (41)$$

Consider the function $g(x) = A \ln(1 - x) - \ln(1 - Ax)$. Note that $g(0) = 0$. We also obtain:

$$g'(x) = -\frac{A}{1-x} + \frac{A}{1-Ax} = \frac{A(A-1)x}{(1-x)(1-Ax)} \quad (42)$$

Hence, we have $g'(x) < 0$ for $x < 0$ and $g'(x) > 0$ for $x > 0$ which implies $g(x) \geq 0$ and hence proves the claim.

We can use lemma E.1 to further bound inequality 39:

$$\begin{aligned} P(y \notin I_k) &< \left[1 - \frac{r_{n_k}^\tau (1 - \vartheta)}{\hat{b}(1 + \vartheta)} 2\hat{b}f(r_{n_k})(1 + \vartheta) \right]^n \\ &\leq \left[\left[1 - 2\hat{b}f(r_{n_k})(1 + \vartheta) \right]^n \right]^{\frac{r_{n_k}^\tau (1 - \vartheta)}{\hat{b}(1 + \vartheta)}} \\ &< \epsilon^{Dr_{n_k}^\tau} \end{aligned} \quad (43)$$

where $D = \frac{1 - \vartheta}{\hat{b}(1 + \vartheta)}$. This completes Step 2 of the proof.

Step 3. We can decompose $E(y)$ as follows:

$$\begin{aligned} E(y) &= P(y < r_{n_k} - b_k)E(y|y < r_{n_k} - b_k) + P(y \in I_k)E(y|y \in I_k) + P(y > r_{n_k} + b_k)E(y|y > r_{n_k} + b_k) \\ &\geq P(y \notin I_k)E(y|y < r_{n_k} - b_k) + P(y \in I_k)E(y|y \in I_k) \end{aligned} \quad (44)$$

Denote the expectation of a single anecdote conditional on it having a value $x < 0$ with $\mu_0 < 0$. The expectation of the max anecdote conditional on having a realization less than $r_{n_k} - b_k$ has to be at least μ_0 (since $r > 0$) and hence the expected distance from r can be at most $r - \mu_0$. We can use the bound 43 from Step 2 to obtain:

$$E(y) \geq \underbrace{(r_{n_k} - \mu_0)\epsilon^{Dr_{n_k}^\tau}}_{B_k} + P(y \in I_k)E(y|y \in I_k) \quad (45)$$

It is easy to see that $B_k \rightarrow 0$ as $r_{n_k} \rightarrow \infty$. Therefore, we only have to show that $E(y|y \in I_k) \rightarrow 0$ in order to get a contradiction that $E(y) \neq \Delta$.

We can write $E(y|y \in I_k)$ as:

$$\begin{aligned} E(y|y \in I_k) &= \int_0^{b_k} yn [f(r_{n_k} + y) - f(r_{n_k} - y)] [1 - (F(r_{n_k} + y) - F(r_{n_k} - y))]^{n-1} dy \\ &= \int_0^{b_k} \frac{f(r_{n_k} + y) - f(r_{n_k} - y)}{f(r_{n_k} + y) + f(r_{n_k} - y)} \underbrace{yn [f(r_{n_k} + y) + f(r_{n_k} - y)] [1 - (F(r_{n_k} + y) - F(r_{n_k} - y))]^{n-1}}_{w(y)} dy \end{aligned}$$

We know from Step 1:

$$\frac{-2\vartheta}{2 - 2\vartheta} \leq \frac{f(r_{n_k} + y) - f(r_{n_k} - y)}{f(r_{n_k} + y) + f(r_{n_k} - y)} \leq \frac{2\vartheta}{2 - 2\vartheta} \quad (46)$$

Moreover, we can simplify:

$$\begin{aligned} \int_0^{b_k} w(y) dy &= - \int_0^{b_k} y [(1 - (F(r_{n_k} + y) - F(r_{n_k} - y)))^n]' dy \\ &= -b_k P(y \notin I_k) + \int_0^{b_k} (1 - (F(r_{n_k} + y) - F(r_{n_k} - y)))^n dy \\ &= \underbrace{-b_k P(y \notin I_k)}_{\rightarrow 0} + \underbrace{\int_0^{\hat{b} \frac{1+\vartheta}{1-\vartheta}} (1 - (F(r_{n_k} + y) - F(r_{n_k} - y)))^n dy}_{\text{between 0 and } \hat{b} \frac{1+\vartheta}{1-\vartheta}} + \\ &\quad + \underbrace{\int_{\hat{b} \frac{1+\vartheta}{1-\vartheta}}^{b_k} (1 - (F(r_{n_k} + y) - F(r_{n_k} - y)))^n dy}_{W_k} \end{aligned} \quad (47)$$

For the last term we again use inequality 37:

$$\begin{aligned} (1 - (F(r_{n_k} + y) - F(r_{n_k} - y)))^n &< [1 - 2bf(r_{n_k})(1 - \vartheta)]^n \\ W_k &< \int_{\hat{b} \frac{1+\vartheta}{1-\vartheta}}^{b_k} [1 - 2yf(r_{n_k})(1 - \vartheta)]^n dy \\ &= \int_{\hat{b} \frac{1+\vartheta}{1-\vartheta}}^{b_k} \left[1 - \frac{y(1 - \vartheta)}{\hat{b}(1 + \vartheta)} 2\hat{b}f(r_{n_k})(1 + \vartheta) \right]^n dy \\ &\leq \int_{\hat{b} \frac{1+\vartheta}{1-\vartheta}}^{b_k} \left[[1 - 2\hat{b}f(r_{n_k})(1 + \vartheta)]^n \right]^{\frac{y(1-\vartheta)}{\hat{b}(1+\vartheta)}} dy \quad (\text{lemma E.1}) \\ &< \int_{\hat{b} \frac{1+\vartheta}{1-\vartheta}}^{b_k} \epsilon^{\frac{y(1-\vartheta)}{\hat{b}(1+\vartheta)}} dy \quad (\text{inequality 38}) \\ &< -\frac{\hat{b}(1 + \vartheta)}{\ln(\epsilon)(1 - \vartheta)} \int_1^\infty \exp(-z) dz \\ &= -\frac{\hat{b}(1 + \vartheta)}{\ln(\epsilon)(1 - \vartheta)} \end{aligned} \quad (48)$$

Putting everything together we know that:

$$\int_0^{b_k} w(y)dy < \left(1 - \frac{1}{\ln(\epsilon)}\right) \frac{\hat{b}(1 + \vartheta)}{1 - \vartheta} \quad (49)$$

Therefore we get a bound on the conditional expectation:

$$\frac{-2\vartheta}{2 - 2\vartheta} \left(1 - \frac{1}{\ln(\epsilon)}\right) \frac{\hat{b}(1 + \vartheta)}{1 - \vartheta} \leq E(y|y \in I_k) \leq \frac{2\vartheta}{2 - 2\vartheta} \left(1 - \frac{1}{\ln(\epsilon)}\right) \frac{\hat{b}(1 + \vartheta)}{1 - \vartheta} \quad (50)$$

Since we can choose ϑ as small as we want (for sufficiently large (n_k)) we can deduce that $E(y|y \in I_k) \rightarrow 0$. This completes the argument for Step 3. \square

Proof of Theorem 6 (light tails). Recall by Proposition 5 that, for any n the variance of the anecdote sent in the targeting scheme increases with $|r|$. In particular, the variance of the sent anecdote is at most the variance of anecdote π_∞ , the maximum of the anecdotes available to the sender. It is therefore sufficient to argue that the variance of the maximum anecdote converges to 0 as $n \rightarrow \infty$.

Let us denote the density of the maximum distribution with $h_n(x)$ and the CDF with $H_n(x) = F(x)^n$. The median of this distribution is denoted with \hat{x}_n and is defined as:

$$H_n(\hat{x}_n) = F(\hat{x}_n)^n = \frac{1}{2} \quad (51)$$

Next, we look at the value of $H_n(x)$ for deviations y away from the median:

$$H_n(\hat{x}_n + y) = \left[1 - C \exp\left(-\int_{\underline{x}}^{\hat{x}_n + y} u(t)t^{\alpha-1} dt\right)\right]^n \quad (52)$$

Let's focus on the inner term:

$$\begin{aligned} 1 - C \exp\left(-\int_{\underline{x}}^{\hat{x}_n + y} u(t)t^{\alpha-1} dt\right) &= 1 - C \exp\left(-\int_{\underline{x}}^{\hat{x}_n} u(t)t^{\alpha-1} dt\right) \exp\left(-\int_{\hat{x}_n}^{\hat{x}_n + y} u(t)t^{\alpha-1} dt\right) \\ &= 1 - C \exp\left(-\int_{\underline{x}}^{\hat{x}_n} u(t)t^{\alpha-1} dt\right) \exp(-yu(\zeta)\zeta^{\alpha-1}) \end{aligned} \quad (53)$$

for some $\zeta \in [\hat{x}_n, \hat{x}_n + y]$ (for $y > 0$) and $\zeta \in [\hat{x}_n + y, \hat{x}_n]$ (for $y < 0$) by the mean value theorem.

Next, we calculate the moment $\int_{-\infty}^{\infty} y^k h_n(\hat{x}_n + y) dy$ for $k \geq 1$. We will be interested in the

cases $k = 1$ and $k = 2$ which we need to calculate the variance of the maximum distribution.

$$\begin{aligned}
\int_{-\infty}^{\infty} y^k h_n(\hat{x}_n + y) dy &= \int_{-\infty}^0 y^k h_n(\hat{x}_n + y) dy + \int_0^{\infty} y^k h_n(\hat{x}_n + y) dy \\
&= \int_{-\infty}^0 y^k H'_n(\hat{x}_n + y) dy - \int_0^{\infty} y^k (1 - H_n(\hat{x}_n + y))' dy \\
&= \left[y^k H_n(\hat{x}_n + y) \right]_{-\infty}^0 - \int_{-\infty}^0 k y^{k-1} H_n(\hat{x}_n + y) dy - \left[y^k (1 - H_n(\hat{x}_n + y)) \right]_0^{\infty} \\
&+ \int_0^{\infty} k y^{k-1} (1 - H_n(\hat{x}_n + y)) dy \\
&= \underbrace{\int_0^{\infty} k y^{k-1} (1 - H_n(\hat{x}_n + y)) dy}_{P_n} - \underbrace{\int_{-\infty}^0 k y^{k-1} H_n(\hat{x}_n + y) dy}_{M_n} \tag{54}
\end{aligned}$$

Now we use lemma E.1 for $y < 0$:

$$\begin{aligned}
H_n(\tilde{x}_n + y) &= \left[1 - C \exp \left(- \int_{\underline{x}}^{\tilde{x}_n} u(t) t^{\alpha-1} dt \right) \exp(-y u(\zeta) \zeta^{\alpha-1}) \right]^n \\
&\leq \left\{ \left[1 - C \exp \left(- \int_{\underline{x}}^{\tilde{x}_n} u(t) t^{\alpha-1} dt \right) \right]^n \right\}^{\exp(-y u(\zeta) \zeta^{\alpha-1})} \\
&= \left(\frac{1}{2} \right)^{\exp(-y u(\zeta) \zeta^{\alpha-1})} \tag{55}
\end{aligned}$$

Similarly, we can use the analogous version E.1 for $y > 0$ (which implies $A < 1$):

$$H_n(\tilde{x}_n + y) \geq \left(\frac{1}{2} \right)^{\exp(-y u(\zeta) \zeta^{\alpha-1})} \tag{56}$$

This implies:

$$\begin{aligned}
0 &\leq P_n \leq \int_0^{\infty} k y^{k-1} \left(1 - \left(\frac{1}{2} \right)^{\exp(-y u(\zeta) \zeta^{\alpha-1})} \right) dy \\
0 &\leq |M_n| \leq \int_{-\infty}^0 |k y^{k-1}| \left(\frac{1}{2} \right)^{\exp(-y u(\zeta) \zeta^{\alpha-1})} dy \tag{57}
\end{aligned}$$

It is easy to see that for $\alpha > 1$ the right bounds converge to 0 and therefore $P_n \rightarrow 0$ and $M_n \rightarrow 0$. From this it follows:

$$\lim_{n \rightarrow \infty} E(y^k) = 0 \tag{58}$$

Therefore, we have $E(y) \rightarrow 0$ and $Var(y) \rightarrow 0$. \square

F Proofs for Section 6

Proof of Theorem 7. For any committed translation invariant communication scheme π_S , the receiver's response α_{π_S} is a translation with shift $M_R - \beta(\pi_S)$. Applying Proposition 2 with $\sigma(\alpha_{\pi_S}) = M_R - \beta(\pi_S)$, we have that the sender's expected loss from any given translation invariant communication scheme π_S is

$$\mathbb{E}[(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)))^2] + (M_R - M_S)^2.$$

Therefore, the sender's optimization problem corresponds to choosing π_S that minimizes the information loss $\mathbb{E}[(\pi_S(\vec{x}, y) - (\theta_S(\vec{x}, y) + \beta(\pi_S)))^2]$, which is precisely the variance of the communicated anecdote (by the definition of bias). By Theorem 3 we have that the translation invariant scheme that minimizes the information loss must be a targeting scheme with $r^* = \beta(r^*)$. This immediately implies that the receiver's shift is $\sigma(\alpha_{\pi_S}) = M_R - r^*$. \square

Proof of Proposition 7. We first observe that when $n = 2$, the sender has no foresight (i.e., $y = \perp$) and the anecdote distribution is symmetric at 0, the posterior mean of the sender $\theta_S(x_1, x_2, y) = (x_1 + x_2)/2$. Therefore, both anecdotes have the same absolute distance to the mean.

The information loss of the sender under any translation invariant scheme π is

$$\begin{aligned} \mathbb{E}[(\pi(\vec{x}) - \theta_S(\vec{x}) - \beta(\pi))^2] &= \mathbb{E}[(\pi(\vec{x}) - \theta_S(\vec{x}))^2] - \beta(\pi)^2 \\ &= \mathbb{E}[(x_1 - x_2)^2/4] - \beta(\pi)^2, \end{aligned}$$

because no matter whether π chooses x_1 or x_2 the distance to the average is the same and doesn't depend on π . The unbiased communication scheme clearly has more information loss than any biased communication scheme. In fact, the information loss is minimized by the most biased communication scheme, which means both the maximum scheme and the minimum scheme are most informative (by symmetry). \square

Proof of Proposition 8. Suppose the anecdote distribution is uniform $[0, 1]$. Given three signals $x_1 \leq x_2 \leq x_3$, the sender's posterior is a uniform distribution on $[x_1 - \frac{1}{2}, x_3 - \frac{1}{2}]$.

Optimal Unbiased Scheme. Consider the optimal unbiased scheme, call it π_0 . This communication scheme sends the closest signal to $\theta_S(\vec{x})$. Since θ_S is the midpoint of the interval $[x_1, x_3]$, and since x_2 falls in that interval, the optimal unbiased scheme always sends signal x_2 .

Let's calculate the mean squared error of signal x_2 relative to θ . The CDF of x_2 is given by

$$H(w) = \Pr[x_2 < w] = w^3 + 3w^2(1 - w)$$

for $w \in [0, 1]$, since the first term is the probability that all three samples are less than w , and the second term is the probability that two of the three samples are less than w . Now write

$d = |\theta - x_2| = |1/2 - x_2|$, noting that d is a random variable. Then 1 minus the CDF of d is given by

$$\tilde{H}(z) = \Pr[d > z] = \Pr_{w \sim H}[w < (1/2 - z)] + \Pr_{w \sim H}[w > (1/2 + z)] = 2 \Pr_{w \sim H}[w < (1/2 - z)] = 2H(1/2 - z)$$

for $z < 1/2$, and $\tilde{H}(z) = 0$ for $z \geq 1/2$. Here we used that $\Pr_{w \sim H}[w > (1/2 + z)] = \Pr_{w \sim H}[w < (1/2 - z)]$ by symmetry.

The total loss of communication scheme π_0 is therefore

$$\begin{aligned} \mathbb{E}[d^2] &= \int_0^\infty \Pr[d^2 > z] dz \\ &= \int_0^\infty \Pr[d > \sqrt{z}] dz \\ &= \int_0^\infty 2H(1/2 - \sqrt{z}) dz \\ &= 1/20 \end{aligned}$$

where the final equality is via numerical calculation.

A Better Biased Communication Scheme. We will now build a communication scheme with strictly less loss than π_0 . Write π_r for the targeting scheme with offset r , which by definition returns whichever of the three points is closest to $\theta_S(\vec{x}) + r$. We will eventually choose $r = 1/5$, but for now we'll proceed with general r .

Which point does π_r return? Write x^* for the random variable representing the point that π_r returns. Recall that $\theta_S(\vec{x}) = (x_1 + x_3)/2$, so $\theta_S(\vec{x}) + r$ is always closer to x_3 than x_1 . The distance to point x_3 is $|\theta_S(\vec{x}) + r - x_3| = (x_3 - x_1)/2 + r$, and the distance to point x_2 is $|\theta_S(\vec{x}) + r - x_2| = (x_3 + x_1)/2 + r - x_2$. So the point x_2 will be closest precisely if $x_2 > x_1 + 2r$. To summarize: $x^* = x_2$ if $x_2 > x_1 + 2r$, otherwise $x^* = x_3$.

As before, let's work out the CDF for x^* . What is the probability that $x^* < w$ for some fixed value of $w \in [0, 1]$? If all three points are less than w (which happens with probability w^3) then x^* certainly is. On the other hand, if $x_2 > w$, then certainly $x^* > w$ as well. If $x_2 < w$ and $x_3 > w$ (which happens with probability $3w^2(1 - w)$), then $x^* < w$ only if $x^* = x_2$, which occurs if and only if $x_2 > x_1 + 2r$. The conditional probability of that last event is equivalent to the probability that two random variables, each drawn uniformly from $[0, w]$, are at least distance $2r$ apart from each other. So, we can write the CDF as

$$\begin{aligned} H[w] &= \Pr[x^* < w] \\ &= w^3 + 3w^2(1 - w) \Pr[|x_1 - x_2| > 2r \mid x_2 < w] \\ &= w^3 + 3w^2(1 - w) \cdot 2 \cdot \int_0^{w-2r} \frac{1}{w} \cdot \frac{w - (x + 2r)}{w} dx. \end{aligned}$$

To justify the last equality, consider drawing one point uniformly from $[0, w]$, so with uniform density $\frac{1}{w}$. What is the probability that a second drawn point is at least $2r$ larger? If the first point (call it x) is greater than $w - 2r$ the probability is 0. Otherwise it is $\frac{w-(x+2r)}{w}$. Integrating over x gives the probability of this event. We then double that probability to account for the possibility that the first point drawn is the larger one.

Now write $d = |x^* - r - \theta| = |x^* - (1/2 + r)|$. This will be the distance between the receiver's action and θ (where recall we fixed $\theta = 1/2$), if the receiver shifts the received signal x^* by r . Note that this may not be the optimal action of the receiver, but the optimal action performs at least as well as $\mathbb{E}[d^2]$.

Now 1 minus the CDF of d is given by

$$\tilde{H}(z) = \Pr[d > z] = \begin{cases} H(1/2 + r - z) + 1 - H(1/2 + r + z) & \text{if } 0 < z < \frac{1}{2} - r, \\ H(1/2 + r - z) & \text{if } \frac{1}{2} - r < z < \frac{1}{2} + r, \\ 0 & \text{if } z > \frac{1}{2} + r. \end{cases}$$

Note that unlike the case of π_0 , the fact that $r > 0$ breaks symmetry in the calculation of \tilde{H} . But the reasoning is the same: $d > z$ precisely if either x^* is greater than $1/2 + r + z$ or x^* is less than $1/2 + r - z$.

Finally, as before, the total loss of the communication scheme π_r is

$$\begin{aligned} \mathbb{E}[d^2] &= \int_0^\infty \Pr[d^2 > z] dz \\ &= \int_0^\infty \Pr[d > \sqrt{z}] dz \\ &= \int_0^\infty \tilde{H}(\sqrt{z}) dz \end{aligned}$$

For $r = 1/5$, this integral evaluates to approximately 0.036, which is less than $1/20$.

Intuition and Discussion. Why is π_r better than π_0 ? In this case, $\theta_S(\vec{x}) = (x_1 + x_3)/2$, so $\theta_S(\vec{x})$ is highly correlated with x_1 and x_3 and much less correlated with x_2 . This fact is specific to the uniform distribution. By selecting the point closest to $\theta_S(\vec{x}) + 1/5$, we are trading off probability of returning x_2 with probability of returning x_3 . Because of the improved correlation with x_3 , the location of x_3 is more highly concentrated, given $\theta_S(\vec{x})$, than the location of x_2 . So by targeting an “expected” location of x_3 relative to $\theta_S(\vec{x})$ (in this case, $\theta_S(\vec{x}) + 1/5$), we can reduce the variance of the distance to the closest point. \square

Online Appendix: Communicating with Anecdotes

Nika Haghtalab
UC Berkeley

Nicole Immorlica
Microsoft Research

Brendan Lucier
Microsoft Research

Markus Mobius
Microsoft Research

Divyarthi Mohan
Tel Aviv University

April 2024

S-1 Single-peaked and Symmetric Anecdote Distributions

In this online appendix we examine the intuition that the information loss minimizing communication scheme will target close to the global maximum density by considering the symmetric and strictly single-peaked anecdote distribution around θ . Indeed, under foresight, by Lemma 1 we have that the sender optimal targeting scheme is the mean scheme.

In general, this intuition suggests that the sender should want to send signals that are close to the posterior mean θ_S , and hence she will want to select a communication scheme with low bias. Theorem 8 below formalizes this insight for the *no-foresight* setting by showing that unbiased communication is optimal for the sender as long as (a) the single-peaked distribution is “well-behaved” and (b) the sender has access to sufficiently many signals.

We start by defining a well-behaved anecdote distribution.

Definition 7. *We say that the anecdote distribution F is well-behaved if the following holds.*

1. *The distribution is strictly single-peaked and symmetric with finite variance.*
2. *Let $g(x) = f'(x)/f(x)$. That is, $g(x) = \frac{d \log f(x)}{dx}$. We assume that $|g'(x)| \leq c_1$ for all x , and some constant $c_1 > 0$. That is, $|g(x)| \leq c_1|x| + c_2$.¹*

*This work is partially supported by the National Science Foundation under grant CCF-2145898, the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 866132), the Office of Naval Research under grant N00014-24-1-2159, an Alfred P. Sloan fellowship, a Schmidt Sciences AI2050 fellowship, and a Google Research Scholar award. Some parts of this work were done when the authors visited the Simons Institute for the Theory of Computing and when Haghtalab and Mohan were employed at Microsoft Research, New England. We are grateful to Nageeb Ali, Lukas Bolte, Vincent Crawford, Ben Golub, Matt Jackson, Pooya Molavi, Harry Pei, Joel Sobel, Leeat Yariv and seminar participants at University of Chicago, Northwestern, Penn State, Stanford Institute for Theoretical Economics (SITE), Stony Brook Game Theory Festival, and Virtual Seminars in Economic Theory, among others, for helpful comments.

¹Note that, for $x < 1$ we bound $|g(x)| \leq c_2$ and otherwise we can bound $|g(x)| \leq c_1$.

3. F has exponential tails. That is, there is a constant $Q > 0$, such that for $x > Q$, we have $1 - F(x) \leq c_3 \exp(-|x|)$ for a constant $c_3 > 0$, and $x < -Q$ we have $F(x) \leq c_3 \exp(-|x|)$.

For example, the normal distribution $F \sim N(0, 1)$ and the Laplace distribution with density $f(\epsilon) = \frac{1}{2} \exp(-|\epsilon|)$ are well-behaved. Recall that $f(\epsilon)$ is the density of the signal distribution at $\theta + \epsilon$, an offset of ϵ from the true state of the world.

Theorem 8. *For any well-behaved anecdote distribution, in the no-foresight setting, the unbiased communication scheme that sends the closest signal to $\theta_S(\vec{x}, y)$ strictly dominates any biased signaling scheme for sufficiently large n and is optimal among all unbiased communication schemes.*

This result confirms our intuition that the optimal communication scheme sends the signal closest to the posterior mean. Intuitively, the sender would like to send *precisely* the posterior mean to the receiver. However, since she can only send a signal, she has to contend with the second best which is to send the signal closest to the posterior mean. When we interpret our model as a model of memory where the current self communicates with her future self by storing a single anecdote in memory, we can think of the anecdote closest to the posterior mean as the “most representative anecdote”.

S-2 Overview of Theorem 8

For the proof of Theorem 8 we bound the losses from the biased and unbiased communication schemes and show that the unbiased communication scheme dominates.

Proposition 9. *Given any well-behaved anecdote distribution F , the unbiased targeting communication scheme with $r = \beta(r) = 0$ which selects the closest signal to the sender’s posterior mean $\theta_S(\vec{x}, y)$ has signaling loss:*

$$\frac{1}{2n^2 f(0)^2} + o\left(\frac{1}{n^2}\right) \quad (59)$$

In contrast, the biased communication scheme with bias δ has signaling loss:

$$\frac{1}{2n^2 f(\delta(\pi))^2} + o\left(\frac{1}{n^2}\right) \quad (60)$$

These two bounds together imply that the unbiased communication scheme is asymptotically optimal, and the optimal communication scheme is asymptotically unbiased.

Let $X_\delta = \min_i |x_i - \theta_S(\vec{x}, y) - \delta|$ denote the absolute distance of the closest signal to the shift of the posterior mean, $\delta + \theta_S(\vec{x}, y)$. We observe that the signaling loss, $\mathbb{E}_{\theta, \vec{x}, y}[(\pi(\vec{x}, y) - \theta_S(\vec{x}, y) - \beta(\pi))]^2$, of any translation invariant communication scheme with bias $\beta(\pi) = \delta$ is at least $\mathbb{E}_{\theta, \vec{x}, y}[X_\delta^2]$.

Optimal unbiased communication scheme. Since the bias of the communication scheme that sends signal closest to the posterior mean is itself 0, this is the optimal amongst all unbiased communication schemes.

When the sender does not have foresight, the posterior mean $\theta_S(\vec{x}, y)$ depends on the realized signals \vec{x} , and this introduces correlation between the signal realizations and the value of $\theta_S(\vec{x}, y) + \delta$. We therefore cannot model X_δ using independent draws from the signal distribution. Indeed, as we seen in Section 6.2, these correlations can significantly impact $\mathbb{E}[X_\delta^2]$ when the number of signals is small.

Our approach is to argue that as n grows large, the impact of these correlations grows small. Small enough, in fact, that the correlation between $\theta_S(\vec{x}, y) + \delta$ and the signal closest to that point becomes small enough that it is dominated by the statistical noise that would anyway be present if signals were drawn independently of $\theta_S(\vec{x}, y)$. We argue this in three steps.

Step 1: We argue that it suffices to focus on cases where $\theta_S(\vec{x}, y)$ falls within a narrow interval. Let $I = [-n^{-\frac{1}{2}+\varepsilon}, n^{-\frac{1}{2}+\varepsilon}]$ for some $\varepsilon > 0$. Using the law of large numbers, we argue that $\theta_S(\vec{x}, y) \in I$ with all but exponentially small probability (in n). The contribution to $\mathbb{E}[X_\delta^2]$ from events where $\theta_S(\vec{x}, y) \notin I$ is therefore negligible and can be safely ignored. This allows us to assume that $\theta_S(\vec{x}, y) \in I$.

Step 2: To reduce the impact of correlation we won't focus on the exact value of $\theta_S(\vec{x}, y)$, but rather an interval in which it falls. To this end we partition I into subintervals of width n^{-b} , where b is chosen so that any given interval is unlikely to contain a signal. One such subinterval contains the posterior mean $\theta_S(\vec{x}, y)$; call that subinterval C . We then consider longer subintervals L and R to the left and right of C , respectively, of width n^{-a} chosen large enough that we expect many signals to appear in each². See Figure 10.

We bound the impact of correlation by showing that if we condition on the number of signals that appear in L and R , then the actual arrangement of signals within those subintervals (keeping all other signals fixed) has only negligible effect on the posterior mean. Specifically, given any arrangement of the signals within L and R , the probability that the posterior mean falls within C remains large. (See Corollary 5 for more details.)

This implies that there is negligible correlation between the joint density function of a fixed number of k signals in $L \cup R$ and the event that $\theta_S(\vec{x}, y) \in C$. We formally show this in Lemmas S-3.11 and S-3.12³.


Step 3: The analysis in Step 2 is conditional on the number of signals k that fall in $L \cup R$. We now show a concentration result on the distribution of k : with high probability, the number of signals that lie in $L \cup R$ is close to the expected number of signals in the interval $L \cup R$ without any correlation to the event $\theta_S(\vec{x}, y) \in C$. See Section S-3 for the proof, and Lemma S-3.6 for the proof that it suffices to consider only this high-probability event.

Given this concentration result, we can focus on bounding the expected value of X_δ^2 , the squared distance of the signal closest to $\theta_S(\vec{x}, y) + \delta$, given the numbers of signals in L and

²For $\delta \neq 0$, we consider L and R to the left and right of $C + \delta$.

³For these lemmas we assume that the density function θ_S is sufficiently “nice” in C . Refer to Section S-3.2 for details about this assumption, and why we can make this assumption without loss of generality.

Figure 10: Intuition for proof of theorem 8



We assume $a < 1$ and $1 < b < 2a - \frac{1}{2}$ which ensures that the intervals L and R contain many signals but the collective influence of these signals on the posterior mean is $O(n^{\frac{1}{2}-2a})$ and hence smaller than $\frac{1}{n}$. It also implies that if we consider a posterior that is contained in C then a rearrangement of signals in L or R will keep the posterior mean within C with high likelihood *and* the probability of signals drawn from the interval C goes to 0. For example, $a = \frac{4}{5}$ and $b = \frac{12}{11}$ satisfy these conditions.

R . From the analysis in Step 2, we can view these signals as (approximately) independently distributed within L and R . We can therefore bound the expected squared distance between interval $C + \delta$ and the closest signal to interval $C + \delta$ by performing an explicit calculation for independent signals. We still do not know the value of $\theta_S(\vec{x}, y)$ within interval C (and we have not bounded the impact of correlation on that value), but C is sufficiently narrow that this uncertainty has limited impact on $\mathbb{E}[X_\delta^2]$. We conclude that the impact of correlation on $\mathbb{E}[X_\delta^2]$ is absorbed in lower-order terms. This gives us the required results of Proposition 9.

S-3 Proof of Proposition 9

We start by proving that the that the signaling loss of the unbiased communication scheme that sends the anecdote closest to the posterior mean $\theta_S(\vec{x}, y)$ is at most $\frac{1}{2n^2 f(0)} + o(1/n^2)$. Later in Section S-3.3 we bound the signaling loss of a biased communication scheme.

Let $I = [-n^{-\frac{1}{2}+\epsilon}, n^{-\frac{1}{2}+\epsilon}] + \theta$. For the remainder of this section, we fix $\theta = 0$ for brevity, but everything holds for any fixed θ . Let \mathcal{P} be a partition of I into intervals of length n^{-b} . For any $C \in \mathcal{P}$, we define $N(C) = L \cup R$, where L (resp. R) is the neighboring interval of length n^{-a} to the left of C (resp. to the right of C).

We first consider the “high probability event” that the following *desirable* properties hold:

1. $\theta_S(\vec{x}, y) \in I$ and let $C \in \mathcal{P}$ be the interval with $\theta_S(\vec{x}, y)$,
2. C is not *weak* (see Definition 10), and

3. there are *sufficiently* many signals in $N(C) = L \cup R$.

In Section S-3.1, we bound the signaling loss contributed by this high probability event. Further, in Section S-3.2 we bound the loss from the “rare event” that some desirable property does not hold: we bound the loss from the event when $\theta_S \notin I$ (in Lemma S-3.4), when C is a weak interval (in Lemma S-3.5), or when there are very few signals in $N(C)$ (in Lemma S-3.6).

With this we are ready to bound the signaling loss. Recall that the signaling loss of the unbiased communication scheme $\pi(\cdot)$ that sends anecdote closest to $\theta_S(\vec{x}, y)$ is $\mathbb{E}[X_0^2] = \mathbb{E}[(\pi(\vec{x}) - \theta_S(\vec{x}, y))^2]$. Given $\theta_S \in C$, let K_0 be the event that there are *sufficiently* many signals in $N(C)$. We see that,

$$\begin{aligned}
\mathbb{E}[X_0^2] &= \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \notin I\}] + \sum_{C \in W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \in C\}] + \sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \in C\}] \\
&= \sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \in C, K_0\}] \tag{61} \\
&\quad + \underbrace{\mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \notin I\}] + \sum_{C \in W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \in C\}] + \sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \in C, \bar{K}_0\}]}_{\text{rare events}} \tag{62} \\
&\leq \frac{1}{2n^2 f(0)^2} + o(1/n^2)
\end{aligned}$$

This is because by Lemma S-3.3 the term in Eq. (61) is $\sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \in C, K_0\}] \leq \frac{1}{2n^2 f(0)^2} + o(1/n^2)$, and from Lemmas S-3.4, S-3.5, and S-3.6 we see that all terms in Eq. (62) contribute at most $o(1/n^2)$. Thus, giving us the required bound on the signaling loss.

S-3.1 Contribution of the High Probability Event

In this section we explain what the desirable properties are, why they are useful, and bound the signaling loss contributed by the event that these properties hold.

Property (1): $\theta_S(\vec{x}, y) \in I = [-n^{-\frac{1}{2}+\epsilon}, n^{-\frac{1}{2}+\epsilon}]$.

Property (2): Let $C \in \mathcal{P}$ be the interval with $\theta_S(\vec{x}, y)$. We need C to be not *weak*.

We start with the definition of a *weak interval*.

Definition 10. Let $\tau(\cdot)$ be the pdf of the posterior mean $\theta_S(\vec{x}, y)$. We say that an interval $C \in \mathcal{P}$ is *weak* if $\tau(\bar{\theta}) \leq c_2 n^{1+1/22} e^{-n^{1/22\alpha}}$ for all $\bar{\theta} \in C$. Let $W_I \subset \mathcal{P}$ be the set of all such intervals C .

By Claim S-3.7 we will see that the probability that $\theta_S(\vec{x}, y) \in C$ for some *weak interval* C is negligible $O(n^{-4\log n+1})$. Moreover, if C is not *weak*, that is, $\tau(\bar{\theta}) \geq c' n^{1/22} n^{-4\log n+1}$ then we get that $\tau(\bar{\theta}') = \tau(\bar{\theta})(1 + O(n^{-\frac{1}{22}}))$ for all $\bar{\theta}, \bar{\theta}' \in C$ by Claim S-3.6.

Property (3): Next we show that there are *sufficiently* many signals in L and R . We start by proving the following claim that $f(x) = f(0)(1 + O(1/\sqrt{n}))$ for all $x \in I = [-n^{-\frac{1}{2}+\epsilon}, n^{-\frac{1}{2}+\epsilon}]$.

Claim S-3.1. *Given any well-behaved distribution with pdf f , for all $x \in I$, we have $f(x) = f(0)(1 + O(1/\sqrt{n}))$.*

Proof. Without loss of generality, we assume that $x > 0$, since $f(-x) = f(x)$. By mean value theorem we see that $f(x) = f(0) + xf'(\tilde{x})$ for some $\tilde{x} \in [0, x]$. By our assumption on g' we get $|g(x)| \leq cx^m$ for some constants $c > 0$ and $m \geq 0$. This implies $|f'(x)| \leq cx^m f(x)$ for all $x > 0$. Since f is non-increasing in $(0, \infty)$ we see that $f(x) \leq f(\tilde{x}) \leq f(0)$. Moreover, $f'(\tilde{x}) \leq 0$, so we have $f'(\tilde{x}) \geq -cx^m f(x)$. By mean value theorem we have,

$$\begin{aligned} f(x) &= f(0) + xf'(\tilde{x}) \\ &\geq f(0) - xc\tilde{x}^m f(\tilde{x}) && \text{(Since } f'(\tilde{x})/f(\tilde{x}) \geq -c\tilde{x}^m) \\ &\geq f(0)(1 - xc\tilde{x}^m) && \text{(Since } f(0) \geq f(\tilde{x})) \\ &\geq f(0)(1 - cx^{m+1}) && \text{(Since } \tilde{x} \leq x) \end{aligned}$$

Therefore, for all $x \in I$ we have $f(x) \geq f(0)(1 - c(n^{-\frac{1}{2}+\epsilon})^{m+1})$. Note that $m \geq 0$, hence we get $f(x) = f(0)(1 - O(n^{-1/2+\epsilon}))$. \square

Using the above claim that $f(x)$ is approximately $f(0)$ for $x \in I$, we bound the number of signals in a subset $A \subset I$.

Claim S-3.2. *Given any interval $A \subset I$ of length ℓ , the expected number signals in A is $n\ell f(0)(1 - O(n^{-1}))$. Let $Y(A)$ be the number of signals in A . For any $0 < \epsilon < 1$, we have*

$$\Pr[Y(A) \leq (1 - \epsilon)\mathbb{E}[Y(A)]] \leq \exp\left(-\frac{\epsilon^2 f(0)n\ell}{2}\right).$$

Proof. Let $Y_i = 1$ if $x_i \in A$ and 0 otherwise. So we have, $\sum_{i=1}^n Y_i = Y(A)$. By Claim S-3.1 we have $f(x) = f(0)(1 + O(1/\sqrt{n}))$ for all $x \in I$. Therefore, we have $\Pr[x_i \in A] = \int_A f(x)dx = f(0)(1 + O(1/\sqrt{n}))\int_A dx = f(0)(1 - O(1/\sqrt{n}))\ell$. Note that, Y_i are i.i.d. random variables, and $\mathbb{E}[Y(A)] = f(0)(1 - O(1/\sqrt{n}))n\ell$. By using Chernoff bound we get

$$\Pr[Y(A) \leq (1 - \epsilon)\mathbb{E}[Y(A)]] \leq \exp\left(-\frac{\epsilon^2 f(0)n\ell(1 - O(n^{-1/2}))}{2}\right).$$

\square

We partition the interval I into intervals J of length n^{-a}/M for $M = n^{1/22}$. Let \mathcal{J} denote the partition. Note that, the size of \mathcal{J} is $n^{-1/2+\epsilon+a}M$.

Lemma S-3.1. Let $k_m = f(0)n^{1-a}/M$. For each $J \in \mathcal{J}$. Let $Y(J)$ be the number signals in interval J (of length n^{-a}/M). $\Pr[|Y(J) - k_m| \geq \varepsilon_1 k_m] \leq \exp\left(-\frac{\varepsilon_1^2 f(0)n^{1-a}}{3M}\right)$.

Moreover, the probability that there is a $J \in \mathcal{J}$ with $|Y(J) - k_m| \geq n^{-1/20}k_m$ is at most $O(\exp^{-n^{1/22}})$

Proof. By directly invoking Claim S-3.2 on J of length n^{-a}/M we get $\Pr[|Y(J) - k_m| \geq \varepsilon_1 k_m] \leq \exp\left(-\frac{\varepsilon_1^2 f(0)n^{1-a}}{3M}\right)$. Note that, the size of \mathcal{J} is $n^{-1/2+\varepsilon+a}M$. Therefore, by union bound, we get $\Pr[\exists J \in \mathcal{J} : |Y(J) - k_m| \geq \varepsilon_1 k_m] \leq (n^{-1/2+\varepsilon+a}M) \cdot \exp\left(-\frac{\varepsilon_1^2 f(0)n^{1-a}}{3M}\right)$ which is at most $O(\exp^{-n^{1/20}})$, for $\varepsilon_1 = O(n^{-1/22})$ and $M = n^{1/100}$. \square

We will now only focus on the case where all $J \in \mathcal{J}$ has sufficiently many signals, which immediately implies the following Corollary 3. In Lemma S-3.6 we bound the loss of the rare event that this is not the case.

Corollary 3. Let $k_0 = f(0)n^{1-a}(1 - 2/M)$, and $\varepsilon = O(n^{-1/22})$. For all $C \in \mathcal{P}$, let $N(C) = L \cup R$ (of size n^{-a}). Let $Y(C)$ be the number signals in interval $N(C)$. If all $J \in \mathcal{J}$ have $k_m(1 \pm \varepsilon)$ signals, then $Y(C)$ has $k_0(1 \pm \varepsilon)$ signals.

That is, $\Pr[|Y(C) - k_0| \geq \varepsilon k_0] \leq O(\exp^{-n^{1/20}})$.

Proof. Note that any $N(C) \subset I$ of length n^{-a} contains at least $M - 2$ many intervals $J \in \mathcal{J}$. By Lemma S-3.1, we have that all $J \in \mathcal{J}$ has at least $(1 - \varepsilon)f(0)n^{1-a}$ many signals with high probability. Therefore, $N(C)$ contains at least k_0 many signals with probability $1 - O(\exp^{-n^{1/20}})$. Note this is regardless of which $C \in \mathcal{P}$ we are considering. \square

We will now only consider the event where all the desirable properties hold. For each $C \in \mathcal{P} \setminus W_I$, let K_0 denote the event that there are $k_0(1 \pm \varepsilon_1)$ signals in $N(C)$.

Lemma S-3.2. Fix any $C \in \mathcal{P} \setminus W_I$. The distribution of the random variable $X_0 \cdot \mathbf{1}(K_0)$ conditioned on $\theta_S \in C$ is stochastically dominated by the exponential distribution with $\lambda = 2nf(0)$. That is,

$$\Pr[X_{(0)} \cdot \mathbf{1}(K_0) > d | \theta_S \in C] < \exp[-\lambda d].$$

Proof. For $A \subseteq [n]$ let K_A denote the event that $x_i \in N(C)$ iff $i \in A$. We also use x_A to denote $\{x_i\}_{i \in A}$. For all $d > n^{-a} + n^{-b}$ we have $\Pr[X_{(0)} \cdot \mathbf{1}(K_0) > d | \theta_S \in C] = 0$. For all $d < n^{-a}/2 + n^{-b}$, let B_d be the interval of length $2d$ centered around θ_S .

We will use the following results/facts:

1. $\Pr[x_A \notin B_d | K_A, \theta_S \in C] = \Pr[x_A \notin B_d | K_A] \cdot \frac{\Pr[\theta_S \in C | x_A \notin B_d, K_A]}{\Pr[\theta_S \in C | K_A]}$ (by Bayes rule)
2. Since f is near uniform in $N(C)$ we have $\Pr[x_A \notin B_d | K_A] \approx (1 - \frac{d}{n^{-a}})^{|A|}$ (by Claim S-3.1)
3. Since redrawing $x_A \in N(C)$ doesn't change θ_S much, we have $\Pr[\theta_S \in C | x_A \notin B_d, K_A] \leq \Pr[\theta_S \in C' | K_A]$ where $C' = C \pm |A|n^{-1-a}$ (by Lemma S-3.11)

4. $\Pr[K_A|\theta_S \in C] \cdot \frac{\Pr[\theta_S \in C'|K_A]}{\Pr[\theta_S \in C|K_A]} = \frac{\Pr[\theta_S \in C']\Pr[K_A|\theta_S \in C']}{\Pr[\theta_S \in C]}$ (by Bayes rule)
5. Since C is not weak and hence τ is near uniform in C we have, $\frac{\Pr[\theta_S \in C']}{\Pr[\theta_S \in C]} = (1 + o(1/\sqrt{n}))$ (by Claim S-3.6)

Thus, we get,

$$\begin{aligned}
\Pr[X_0 > d \text{ and } K_0|\theta_S \in C] &= \sum_{A:|A|\approx k_0} \Pr[X_0 > d \text{ and } K_A|\theta_S \in C] \\
&= \sum_{A:|A|\approx k_0} \Pr[K_A|\theta_S \in C] \Pr[x_i \notin B_d \forall i \in A|K_A, \theta_S \in C] \\
&= \sum_{A:|A|\approx k_0} \Pr[K_A|\theta_S \in C] \Pr[x_i \notin B_d \forall i \notin A|K_A] \cdot \frac{\Pr[\theta_S \in C|x_A \notin B_d, K_A]}{\Pr[\theta_S \in C|K_A]} \\
&\leq (1 - d/n^{-a})^{k_0} \sum_{A:|A|\approx k_0} \Pr[K_A|\theta_S \in C] \cdot \frac{\Pr[\theta_S \in C|x_A \notin B_d, K_A]}{\Pr[\theta_S \in C|K_A]} \\
&\leq (1 - d/n^{-a})^{k_0} \sum_{A:|A|\approx k_0} \frac{\Pr[\theta_S \in C']\Pr[K_A|\theta_S \in C']}{\Pr[\theta_S \in C]} \\
&\leq (1 - d/n^{-a})^{k_0} (1 + o(1/\sqrt{n}))
\end{aligned}$$

Recall that $k_0 = (1 - \varepsilon_1)f(0)n^{1-a}(1 - 2/M)$. Let $\varepsilon_1 = O(n^{-1/22})$ and $M = n^{1/22}$. Since we can bound $1 - x \leq e^{-x}$ we get,

$$\Pr[X_{(0)} \cdot \mathbf{1}(K_0) > d|\theta_S \in C] \leq \left(1 + O(n^{-1/2})\right) \cdot \exp\{-2nf(0)(1 - O(n^{-1/22}))(1 - O(n^{-1/4}))\} \cdot d$$

□

We finally bound the cost of the event with all the desirable properties.

Lemma S-3.3. *Fix any $C \in \mathcal{P} \setminus W_I$. Then we have, $\mathbb{E}[X_0^2 \cdot \mathbf{1}(K_0)|\theta_S \in C] \leq \frac{1}{2n^2f(0)^2} + o(1/n^2)$.*

Proof. Let $Z(\lambda)$ be the random variable with exponential distribution. We observe that $\mathbb{E}[X_0^2 \cdot \mathbf{1}(K_0)|\theta_S \in C] \leq \mathbb{E}[Z(\lambda)^2](1 + o(n^{-1/2}))$ for $\lambda = 2nf(0)(1 - O(n^{-1/22}))(1 - O(n^{-1/4}))$, because of the stochastic dominance proved above in Lemma S-3.2. Moreover, $\mathbb{E}[Z(\lambda)^2] = \frac{2}{\lambda^2}$. Hence, we get $\mathbb{E}[X_0^2 \cdot \mathbf{1}(K_0)|\theta_S \in C] \leq \frac{1}{2n^2f(0)^2}(1 - O(n^{-1/22}))(1 - O(n^{-1/4})) \leq \frac{1}{2n^2f(0)^2} + o(1/n^2)$.

□

In Section S-3.2 we bound the loss due to the rare events of $\theta_S \notin I$, $C \in W_I$, and $\overline{K_0}$, that is, the number of signals in $N(C)$ is not in $(1 \pm \varepsilon_1)n^{1-a}f(0)(M - 2)/M$. We show that these contribute up to $o(1/n^2)$ loss.

S-3.2 Contribution of Rare Events

In this section we bound the loss from the rare events from Eq. (62).

Lemma S-3.4. $\mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \notin I\}] \leq O\left(\exp\left(-\frac{n^{2\epsilon}A}{2}\right)\right)$ for some constant $A > 0$.

Proof. Recall that $\theta_S(\vec{x})$ is the MMSE estimator and $\theta_S - \theta^* \rightarrow^d \mathcal{N}(0, C_{\mathcal{I}}/n)$. Hence the probability that $\theta_S \notin I$ is at most $\exp(-n^{2\epsilon}A)$ for some $A > 0$. Let $P = \exp(-n^{2\epsilon}A)$. To bound $\mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \notin I\}]$ we see that,

$$\begin{aligned}
\mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \notin I\}] &= \int_0^\infty \Pr[X_0^2 > y \wedge \theta_S \notin I] dy \\
&= \int_0^{1/P^{1/2}} \Pr[\theta_S \notin I] \Pr[X_0^2 > y | \theta_S \notin I] dy + \int_{1/P^{1/2}}^\infty \Pr[X_0^2 > y \wedge \theta_S \notin I] dy \\
&\leq P^{1/2} + 2 \int_{1/P^{1/2}}^\infty \Pr[X_i > \theta_S + \sqrt{y} \wedge \theta_S > 0] dy \quad (\text{For any arbitrary choice of } i) \\
&\leq P^{1/2} + 2 \int_{1/P^{1/2}}^\infty \Pr[X_i > \sqrt{y}] dy \\
&\leq P^{1/2} + 2 \int_{1/P^{1/2}}^\infty (e^{-\sqrt{y}}) dy \quad (\text{Since } \int_x^\infty f(z) dz < e^{-x} \text{ for all } x > Q) \\
&\leq P^{1/2} + 4(1/P^{1/4} + 1) \exp\left(-1/P^{1/4}\right) \\
&= O\left(\exp\left(-\frac{n^{2\epsilon}A}{2}\right)\right)
\end{aligned}$$

Recall that, $P = \exp(-n^{2\epsilon}A)$. So we have $P^{1/2} = O\left(\exp\left(-\frac{n^{2\epsilon}A}{2}\right)\right)$. Since xe^{-x} is $O(e^{-x})$ for x sufficiently large, the term $4(1/P^{1/4} + 1) \exp(-1/P^{1/4}) = O\left(\exp\left(-\exp\left(\frac{n^{2\epsilon}A}{4}\right)\right)\right)$. □

For a well-behaved distribution we have $1 - F(x) \leq c_3 e^{-x}$ for all $x > Q$.

Claim S-3.3. Let T_Q be the event that all $|x_i| > Q$. Then $\mathbb{E}[X_0^2 \mathbf{1}(T_Q \wedge \theta_S \in I)] \leq o(1/n^2)$.

Proof. Since $\theta_S \in I = [-n^{-1/2+\epsilon}, n^{-1/2+\epsilon}]$, and all signals $|x_i| > Q$ are outside I , we have that $X_0 = \min_i |x_i - \theta_S| \leq |x_i| + n^{-1/2+\epsilon}$ for all x_i . Let $t(x) = (|x| + n^{-1/2+\epsilon})^2$. Hence, we have,

$$\begin{aligned}
\mathbb{E}[X_0^2 \mathbf{1}(T_Q \wedge \theta_S \in I)] &\leq \mathbb{E}[(|x_1| + n^{-1/2+\epsilon})^2 \mathbf{1}(T_Q \wedge \theta_S \in I)] \quad (\text{For an arbitrary choice of } i = 1) \\
&= \mathbb{E}[t(x_1) \mathbf{1}(T_Q \wedge \theta_S \in I)] \\
&= \int_{x_1: |x_1| > Q} \cdots \int_{\vec{x}_{-1}} t(x_1) f(x_1) \prod_{i \neq 1} f(x_i) \mathbf{1}(|x_i| > Q) \cdot \mathbf{1}(\theta_S(\vec{x}) \in I) d\vec{x} \\
&\leq \int_{x_1: |x_1| > Q} \cdots \int_{\vec{x}_{-1}} t(x_1) f(x_1) \prod_{i \neq 1} f(x_i) \mathbf{1}(|x_i| > Q) d\vec{x} \\
&\leq \int_{x_1: |x_1| > Q} t(x_1) f(x_1) dx_1 (2 \exp(-Q))^{n-1} \quad (\text{By tail bound Assumption of } f) \\
&\leq \exp\left(-\frac{Q(n-1)}{2}\right) 2 \int_Q^\infty (x + n^{-1/2+\epsilon})^2 f(x) dx \\
&\leq \exp\left(-\frac{Q(n-1)}{2}\right) \cdot Q \cdot O(1) \quad (\text{By tail bound Assumption of } f)
\end{aligned}$$

□

Recall that $W_I \subset \mathcal{P}$ is the set of all intervals C such that $\tau(\bar{\theta}) < c'n^{1/22}n^{-4\log n+1}$ for all $\bar{\theta} \in C$.

Lemma S-3.5. *Then $\sum_{C \in W_I} \mathbb{E}[X_0^2 \mathbf{1}\{\theta_S \in C\}] \leq O(Q^2 n^{-4\log n+1}) + o(1/n^2)$.*

Proof. Let \bar{T}_Q be the event that there is some $|x_i| \leq Q$. Since $\theta_S \in I$ and there is some $|x_i| \leq Q$, we have that $X_0 \leq Q + n^{-\frac{1}{2}+\epsilon}$. Thus,

$$\sum_{C \in W_I} \mathbb{E}[X_0^2 \mathbf{1}\{\theta_S \in C\} \cdot \mathbf{1}\{\bar{T}_Q\}] \leq (Q + n^{-\frac{1}{2}+\epsilon})^2 \Pr[\theta_S \in W_I] \leq (Q + n^{-\frac{1}{2}+\epsilon})^2 O(n^{-4\log n+1})$$

where the last inequality follows from Claim S-3.7 (proved in Section S-3.4) that $\Pr[\theta_S \in W_I] \leq O(n^{-4\log n+1})$.

Moreover, by Claim S-3.3 proved above, we have $\sum_{C \in W_I} \mathbb{E}[X_0^2 \mathbf{1}\{\theta_S \in C\} \cdot \mathbf{1}\{T_Q\}] \leq \mathbb{E}[X_0^2 \mathbf{1}(T_Q \wedge \theta_S \in I)] \leq o(1/n^2)$. Thus, proving the lemma.

□

Lemma S-3.6. *Let $k_0 = f(0)n^{1-a}$, $\varepsilon_1 = O(n^{-1/22})$. Let \bar{K}_0 be the event such that $Y(C) \neq k_0(1 \pm \varepsilon_1)$. Then, $\sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\bar{K}_0 \text{ and } \theta_S \in C\}] \leq (Q + n^{-1/2})^2 O(\exp(-n^{1/20}f(0))) + O(\exp^{-Qn}) \leq o(1/n^2)$.*

Proof. Consider the case where $N(C)$ doesn't have $k_0(1 \pm \varepsilon_1)$ signals. Let A_0 denote that event. By Claim S-3.3, when all $|x_i| > Q$ and $\theta_S \in I$, we bound the expected X_0^2 by $O(\exp^{-Qn})$. If there is even a single $|x_i| \leq Q$ (denoted by the event \bar{T}_Q), then we can bound X_0^2 by $(Q + n^{-1/2+\epsilon})^2$ because $\theta_S \in I = [-n^{-1/2+\epsilon}, n^{-1/2+\epsilon}]$. By corollary 3, we get that $\Pr[A_0 \cap \bar{T}_Q \cap \theta_S \in C] \leq \Pr[A_0] \leq \exp\left(-\frac{n^{-1/11}f(0)n^{1-a}}{3M}\right)$.

$$\begin{aligned}
\mathbb{E}[X_0^2 \cdot \mathbf{1}\{\overline{K}_0 \text{ and } \theta_S \in C\}] &\leq \mathbb{E}[X_0^2 \cdot \mathbf{1}\{A_0 \text{ and } \theta_S \in C\} \cdot \mathbf{1}\{T_Q\}] + \mathbb{E}[X_0^2 \cdot \mathbf{1}\{A_0 \text{ and } \theta_S \in C\} \cdot \mathbf{1}\{\overline{T}_Q\}] \\
&\leq O(\exp^{-Qn}) + (Q + n^{-1/2+\epsilon})^2 \exp\left(-\frac{n^{-1/11}f(0)n^{1-a}}{3M}\right) \\
&\leq o(1/n^2)
\end{aligned}$$

□

S-3.3 Loss of Biased Communication Schemes

In this section we show that for all sufficiently large n and all δ , $\mathbb{E}[X_\delta^2] \geq \frac{1}{2n^2 f(\delta)^2} - o(1/n^2)$. We focus only on $|\delta| \leq 2(\log n)^2$ and δ such that $f(\delta) \geq n^{-1/100}$ ⁴.

Recall \mathcal{P} be a partition of I into intervals of length n^{-b} . For any δ , we denote $C_\delta = C + \delta$ for any $C \in \mathcal{P}$ and $N(C_\delta) = L_\delta \cup R_\delta$, where L_δ (resp. R_δ) is the neighboring interval of length n^{-a} to the left of C_δ (resp. to the right of C_δ).

Similar to the unbiased loss we consider the high probability event where all the following desirable properties hold:

1. $\theta_S(\vec{x}, y) \in I$ and let $C \in \mathcal{P}$ be the interval with $\theta_S(\vec{x}, y)$,
2. C is not *weak* (see Definition 10), and
3. there are *sufficiently few* signals in $N(C_\delta) = L_\delta \cup R_\delta$.

We will show that with high probability all the desirable properties hold.

With this we are ready to bound the signaling loss. Recall that the signaling loss of a biased communication scheme π with $\delta(\pi) = \delta$ is $L(\pi, \delta(\pi)) \geq L(\pi_\delta, \delta) = \alpha^2 \mathbb{E}[X_\delta^2]$. Given $\theta_S \in C$, let K_δ be the event that there are *sufficiently few* signals in $N(C_\delta)$ and there are no signals in C_δ . We see that,

$$\begin{aligned}
\mathbb{E}[X_\delta^2] &\geq \sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_\delta^2 \cdot \mathbf{1}\{\theta_S \in C, K_\delta\}] \\
&\geq \sum_{C \in \mathcal{P} \setminus W_I} \Pr\{\theta_S \in C, K_\delta\} \cdot \underbrace{\mathbb{E}[X_\delta^2 \mid \theta_S \in C, K_\delta]}_{\text{conditional expectation}} \\
&\geq \frac{1}{2n^2 f(\delta)^2} - o(1/n^2)
\end{aligned} \tag{63}$$

This is because by Lemma S-3.9 the conditional expectation term in Eq. (63) is $\mathbb{E}[X_\delta^2 \mid \theta_S \in C, K] \geq \frac{1}{2n^2 f(\delta)^2} (1 - o(1))$. We note that $\{\vec{x} : \forall J_\delta \in \mathcal{J}_\delta \text{ with sufficiently few signals}\} \subset K_\delta$ for

⁴When $f(\delta) \leq n^{-1/100}$ we see that $\mathbb{E}[X_\delta^2] \geq \Omega(\frac{n^{1/100}}{n^2}) \gg \mathbb{E}[X_0^2]$.

all $C \in \mathcal{P}$, and by Lemma S-3.7 we see that $\Pr[\overline{K_\delta}] \leq \Pr[\exists J_\delta \in \mathcal{J}_\delta \text{ with too many signals}] \leq O(\exp\{f(\delta)n^{1/20}\})$. By Lemma S-3.5 we bound the probability of C is weak. Thus, we have,

$$\begin{aligned} \sum_{C \in \mathcal{P} \setminus W_I} \Pr[\{\theta_S \in C, K_\delta\}] &\geq \Pr[\forall J_\delta \in \mathcal{J}_\delta \text{ with sufficiently few signals, and } \theta_S \notin W_I] \\ &\geq 1 - O(\exp\{f(\delta)n^{1/20}\}) - O(\exp\{f(\delta)n^{1/55}\})O(n^{4 \log n - 1}). \end{aligned}$$

We start by showing a more general version of Claim S-3.1.

Claim S-3.4. *Given any well-behaved distribution with pdf f , for all $x \in I + \delta$, we have $f(x) = f(\delta)(1 + O(1/n))$.*

Proof. By mean value theorem we see that $f(x) = f(\delta) + (x - \delta)f'(\tilde{x})$ for some $\tilde{x} \in [\delta, x]$. By our assumption on g' we get $|g(x)| \leq cx^m$ for some constant $c > 0$. This implies $|f'(x)| \leq c|x|^m f(x)$ for all x . Since $|\delta| < 2(\log n)^2$ and $x \in I + \delta$, we have $|x| \leq 2(\log n)^2 + n^{-1/2+\epsilon}$. By mean value theorem we have,

$$\begin{aligned} |f(x) - f(\delta)| &= |(x - \delta)f'(\tilde{x})| \\ &\leq |x - \delta|c|\tilde{x}|^m f(\tilde{x}) \quad (\text{Since } |f'(\tilde{x})| \leq c|\tilde{x}|^m f(\tilde{x})) \\ &\leq c(n^{-1/2+\epsilon})(2(\log n)^2 + n^{-1/2+\epsilon})^m f(\tilde{x}) \\ &\leq c'(n^{-1/2+\epsilon})n^{1/4} f(\tilde{x}) \quad (\text{Since } (\log n)^{2m} = o(n^{1/4})) \\ &\leq f(\delta) + c'n^{-1/4+\epsilon} f(\tilde{x}) \end{aligned}$$

Without loss of generality, we assume that $x, \delta > 0$, since f is symmetric. Thus, we get $f(\tilde{x}_i)$ is between $f(x)$ and $f(\delta)$, because f is single-peaked.

Suppose $f(x) \geq f(\tilde{x}) \geq f(\delta)$, then $f(x) \leq f(\delta) + c'n^{-1/4+\epsilon} f(\tilde{x}) \leq f(\delta) + c'n^{-1/4+\epsilon} f(x)$. So we get $f(x)(1 - c'n^{-1/4+\epsilon}) \leq f(\delta)$. Thus, $f(x) \leq f(\delta) \frac{1}{1 - c'n^{-1/4+\epsilon}} \leq f(\delta)(1 + c''n^{-1/4+\epsilon})$ for some constant $c'' > 0$.

Similarly, if $f(x) \leq f(\tilde{x}) \leq f(\delta)$, then $f(x) \geq f(\delta) - c'n^{-1/4+\epsilon} f(\tilde{x}) \geq f(\delta) - c'n^{-1/4+\epsilon} f(\delta)$. Thus, we get, $f(x) \geq f(\delta)(1 - O(n^{-1/4+\epsilon}))$.

Therefore, for $\delta < 2(\log n)^2$ and all $x \in I + \delta$ we have $f(x) = f(\delta)(1 - O(n^{-\frac{1}{4}+\epsilon}))$. \square

Claim S-3.5. *Given any interval $A \subset I + \delta$ of length ℓ , the expected number signals in A is $n\ell f(\delta)(1 - O(n^{-1/4}))$. Let $Y(A)$ be the number of signals in A . For any $0 < \varepsilon_1 < 1$, we have*

$$\Pr[|Y(A) - \mathbb{E}[Y(A)]| \geq \varepsilon_1 \mathbb{E}[Y(A)]] \leq \exp\left(-\frac{\varepsilon_1^2 f(\delta) n \ell}{3}\right).$$

Proof. Let $Y_i = 1$ if $x_i \in A$ and 0 otherwise. So we have, $\sum_{i=1}^n Y_i = Y(A)$. By Claim S-3.4 we have $f(x) = f(\delta)(1 + O(n^{-1/4+\epsilon}))$ for all $x \in I + \delta$. Therefore, we have $\Pr[x_i \in A] = \int_A f(x) dx = f(\delta)(1 + O(n^{-1/4+\epsilon})) \int_A dx = f(\delta)(1 + O(n^{-1/4+\epsilon}))\ell$. Note that, Y_i are i.i.d. random variables, and

$\mathbb{E}[Y(A)] = f(\delta)(1 + O(n^{-1/4+\epsilon}))n\ell$. By using Chernoff bound we get

$$\Pr[Y(A) \geq (1 + \epsilon)\mathbb{E}[Y(A)]] \leq \exp\left(-\frac{\epsilon_1^2 f(\delta)n\ell(1 + O(n^{-1/4+\epsilon}))}{3}\right).$$

□

We again partition $I + \delta$ into intervals J_δ of size n^{-a}/M . Exactly following Lemma S-3.1 we see that all $J_\delta \subset I + \delta$ have $(1 \pm \epsilon)f(\delta)n^{1-a}(1 + O(n^{-1/4+\epsilon}))$ many signals in L_δ and R_δ .

Lemma S-3.7. *Let $k_m = f(\delta)n^{1-a}/M$. The probability that there is a $J_\delta \in \mathcal{J}_\delta$ with more than $(1 + \epsilon)k_m$ signals (or less than $(1 - \epsilon)k_m$ is $O(\exp\{\epsilon^2 f(\delta)n^{1-a}/3M\})$.*

Corollary 4. *Let $k_\delta = f(\delta)n^{1-a}$, let $k^* = (1 + \epsilon_1)k_\delta$, and $k' = (1 - \epsilon_1)k_\delta(1 - 2/M)$. Let $Y(N_\delta)$ be the number signals in interval $N(C_\delta)$ (of length n^{-a}). $\Pr[(Y(N_\delta) \notin [k', k^*])] \leq \exp\left(-\frac{\epsilon_1^2 f(\delta)n^{1-a}}{3M}\right)$.*

Let K_δ denote the event that there are at most $(1 + \epsilon_1)k_\delta$ and at least $(1 - \epsilon_1)k_\delta(1 - 2/M)$ many signals in $N(C_\delta)$, and there are no signals in C_δ .

Lemma S-3.8. *The distribution of the random variable X_δ conditioned on $K_\delta, \theta_S \in C$ stochastically dominates (up to a factor of $(1 - o(1))$) the exponential distribution with $\lambda = 2nf(\delta)(1 + O(n^{-1/10}))$. That is, for $d < n^{-9/10}$, $\Pr[X_\delta > d \wedge K_\delta | \theta_S \in C] \geq \exp[-\lambda d] \frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta]}{\Pr[\theta_S \in C]}$.*

Proof. Let $k^* = (1 + \epsilon_1)k_\delta$, and $k' = (1 - \epsilon_1)k_\delta(1 - 2/M)$. For all $A \subset [n]$ such that $k' \leq |A| \leq k^*$, define K_A to be event where $x_i \in N(C_\delta)$ iff $i \in A$, and there are no signals in C_δ .

For all $d < n^{-a} + n^{-b}$, let B_d denote the interval of length $2d$ centered around $\theta_S + \delta$.

We will use the following results/facts:

1. $\Pr[x_A \notin B_d | K_A, \theta_S \in C] = \Pr[x_A \notin B_d | K_A] \cdot \frac{\Pr[\theta_S \in C | x_A \notin B_d, K_A]}{\Pr[\theta_S \in C | K_A]}$ (by Bayes rule)
2. Since f is near uniform in $N(C_\delta)$ we have $\Pr[x_A \notin B_d | K_A] \approx (1 - \frac{d}{n^{-a}})^{|A|}$ (by Claim S-3.4)
3. Since redrawing $x_A \in N(C)$ doesn't change θ_S much, we have $\Pr[\theta_S \in C | x_A \notin B_d, K_A] \geq \Pr[\theta_S \in C \setminus E | K_A]$ (by Lemma S-3.11)
4. $\Pr[K_A | \theta_S \in C] \cdot \frac{\Pr[\theta_S \in C \setminus E | K_A]}{\Pr[\theta_S \in C | K_A]} = \frac{\Pr[\theta_S \in C \setminus E \wedge K_A]}{\Pr[\theta_S \in C]}$ (by Bayes rule)

Thus, we get,

$$\begin{aligned}
\Pr[X_\delta > d \cdot \mathbf{1}(K_\delta) | \theta_S \in C] &= \sum_{A: |A| \in (k', k^*)} \Pr[X_0 > d \text{ and } K_A | \theta_S \in C] \\
&= \sum_{A: |A| \in (k', k^*)} \Pr[K_A | \theta_S \in C] \Pr[x_i \notin B_d \forall i \in A | K_A, \theta_S \in C] \\
&= \sum_{A: |A| \in (k', k^*)} \Pr[K_A | \theta_S \in C] \Pr[x_i \notin B_d \forall i \notin A | K_A] \cdot \frac{\Pr[\theta_S \in C | x_A \notin B_d, K_A]}{\Pr[\theta_S \in C | K_A]} \\
&\geq \left(1 - \frac{d}{n^a} (1 + O(n^{-1/4}))\right)^{k^*} \sum_{A: |A| \in (k', k^*)} \Pr[K_A | \theta_S \in C] \cdot \frac{\Pr[\theta_S \in C \setminus E | K_A]}{\Pr[\theta_S \in C | K_A]} \\
&\geq \left(1 - \frac{d}{n^a} (1 + O(n^{-1/4}))\right)^{k^*} \sum_{A: |A| \in (k', k^*)} \frac{\Pr[\theta_S \in C \setminus E \wedge K_A]}{\Pr[\theta_S \in C]} \\
&\geq \left(1 - \frac{d}{n^a} (1 + O(n^{-1/4}))\right)^{k^*} \frac{\Pr[\theta_S \in C \setminus E \wedge K_\delta]}{\Pr[\theta_S \in C]}
\end{aligned}$$

We bound $\left(1 - \frac{2d}{n^a}\right)^{f(\delta)(1+\varepsilon_1)n^{1-a}}$ by observing that $(1-x) \geq e^{-x-x^2}$ for $x < 1/2$. We will choose of $\varepsilon_1 = O(n^{-1/20})$ and consider $d \leq n^{-a-1/20} = n^{-17/20}$, this gives us $\varepsilon_3 = O(n^{-1/20})$.

$$\begin{aligned}
\Pr[X_{(\delta)} \cdot \mathbf{1}(K_\delta) > d | \theta_S \in C] &\geq \left(1 - \frac{2d}{n^a} (1 + O(1/n^{1/4}))\right)^{k^*} \frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta]}{\Pr[\theta_S \in C]} \\
&\geq \frac{\Pr[\theta_S \in C \setminus E \wedge K_\delta]}{\Pr[\theta_S \in C]} \exp(-2dn(1+\varepsilon_2)f(\delta)) \exp(-(2d)^2 n^{1+a} f(\delta)(1+\varepsilon_2)) \\
&\geq \frac{\Pr[\theta_S \in C \setminus E \wedge K_\delta]}{\Pr[\theta_S \in C]} \exp(-2dn(1+\varepsilon_3)f(\delta)) \quad (\text{For } 2d < n^{-a}\varepsilon_3)
\end{aligned}$$

□

Finally, for each $C \in \mathcal{P} \setminus W_I$, we bound the loss $\mathbb{E}[X_\delta^2 \cdot \mathbf{1}\{K_\delta, \theta_S \in C\}]$.

Lemma S-3.9. *Let K_δ be the event such that $(1-\varepsilon_1)(1-2/M)k_\delta \leq Y(L_\delta) \leq (1+\varepsilon_1)k_\delta$, and no signals in C_δ . Then for all $C \in \mathcal{P} \setminus W_I$,*

1. for $f(\delta) > n^{-1/100}$, we have $\mathbb{E}[X_\delta^2 \cdot \mathbf{1}\{K_\delta, \theta_S \in C\}] \geq \frac{1}{2n^2 f(\delta)^2} (1-o(1)) \Pr[\theta_S \in C \setminus E \text{ and } K_\delta]$,
2. for $f(\delta) \leq n^{-1/100}$, we have $\mathbb{E}[X_\delta^2] \geq c_5 \frac{n^{1/100}}{n^2}$.

Proof. Let $Z(\lambda)$ be the random variable with exponential distribution. We observe that $\mathbb{E}[X_\delta^2 \cdot \mathbf{1}(K_\delta) | \theta_S \in C] \geq \mathbb{E}[Z(\lambda)^2 \cdot \mathbf{1}\{d < (n^{-9/10})^2\}] \cdot \frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta]}{\Pr[\theta_S \in C]}$ for $\lambda = 2nf(\delta)(1+\varepsilon_3)$, because of the stochastic dominance proved above in Lemma S-3.8. Moreover, $\mathbb{E}[Z(\lambda)^2 \cdot \mathbf{1}\{d < (n^{-17/20})^2\}] = \frac{2}{\lambda^2} (1 - O(\lambda n^{-17/20} \exp^{-\lambda n^{-17/20}}))$. Hence, we get $\mathbb{E}[X_\delta^2 \cdot \mathbf{1}(K_\delta) | \theta_S \in C] \geq \frac{1}{2n^2 f(\delta)^2} (1 - o(1)) \frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta]}{\Pr[\theta_S \in C]}$, for $f(\delta) \geq c_4 n^{-1/100}$ for a constant $c_4 > 0$.

We finish by noting that, for sufficiently large n , when $f(\delta) = O(\frac{1}{n^{1/100}})$, we get $\mathbb{E}[X_\delta^2] \geq c_5 \left(\frac{n^{1/100}}{n^2}\right) \gg \frac{1}{2n^2 f(\delta)^2} \geq \mathbb{E}[X_0^2]$ for a constant $c_5 > 0$.

□

Proof of Proposition 9 (b). By lemma S-3.9, for a sufficiently large n and any $\delta < 2(\log n)^2$ and $f(\delta) > n^{-1/100}$, we have $\mathbb{E}[X_\delta^2] \geq \sum_{C \in \mathcal{P} \setminus W_I} \Pr[\theta_S \in C] \mathbb{E}[X_\delta^2 | \theta_S \in C] \geq \sum_{C \in \mathcal{P} \setminus W_I} \frac{1}{2n^2 f(\delta)^2} (1 - o(1)) \Pr[\theta_S \in C \setminus E \text{ and } K_\delta]$.

Given $E \subset C$ as the union of first and last $kn^{-1-a+\epsilon}$ length interval, we define $2E \subset C$ to be the union of the first and last $2kn^{-1-a+\epsilon}$ intervals. Since K_δ is the event that there are $(1 \pm \epsilon)k_\delta$ many signals in $N(C_\delta)$ and there are no signals in C_δ , that $\Pr[\theta_S \in C \setminus E] \geq \Pr[\theta_S \in C \setminus 2E | Y(N(C_\delta) \cup C_\delta) \in (1 \pm \epsilon)k_\delta]$, this is because if $\theta_S \in C \setminus 2E$ and $\vec{x}_{[k]} \in N(C_\delta) \cup C_\delta$ then rearranging the signals $\vec{x}_{[k]}$ by moving the signals in C_δ into $N(C_\delta)$ changes θ_S by at most $kn^{-1-a}f(\delta)$. Moreover, we have $\Pr[K_\delta] = \Pr[Y(N(C_\delta) \cup C_\delta) \in (1 \pm \epsilon)k_\delta] \cdot \Pr[Y(C_\delta) = 0 | Y(N(C_\delta) \cup C_\delta) \in (1 \pm \epsilon)k_\delta]$. The probability that there are no signals in C_δ (of length $n^{-12/11}$) is at least $(1 - O(n^{-1/11}))$. Thus, we get,

$$\mathbb{E}[X_\delta^2] \geq \sum_{C \in \mathcal{P} \setminus W_I} \frac{1}{2n^2 f(\delta)^2} (1 - o(1)) \Pr[\theta_S \in C \setminus 2E \text{ and } Y(N(C_\delta) \cup C_\delta) \in (1 \pm \epsilon)k_\delta] (1 - O(n^{-1/11}))$$

Let \mathcal{G}_δ denote the event that all $J_\delta \in \mathcal{J}_\delta$ has $(1 \pm \epsilon)k_\delta/M$ signals. Thus, we get,

$$\begin{aligned} \mathbb{E}[X_\delta^2] &\geq \sum_{C \in \mathcal{P} \setminus W_I} \frac{1}{2n^2 f(\delta)^2} (1 - o(1)) \Pr[\theta_S \in C \setminus 2E \text{ and } \mathcal{G}_\delta] (1 - O(n^{-1/11})) \\ &= (1 - O(n^{-1/11})) \Pr[\mathcal{G}_\delta \text{ and } \exists C \in \mathcal{P} \setminus W_I \text{ s.t. } \theta_S \in C \setminus 2E] \\ &\geq (1 - O(n^{-1/11})) \left(1 - \Pr[\overline{\mathcal{G}_\delta}] - \Pr[\overline{\exists C \in \mathcal{P} \setminus W_I \text{ s.t. } \theta_S \in C \setminus 2E}] \right) \\ &\geq (1 - O(n^{-1/11})) \left(-\Pr[\overline{\mathcal{G}_\delta}] + \Pr[\exists C \in \mathcal{P} \setminus W_I \text{ s.t. } \theta_S \in C \setminus 2E] \right) \end{aligned}$$

Recall that, $\Pr[\theta_S \in I] \geq (1 - \exp(-n^{2\epsilon}A))$, and by Claim S-3.7 (proved in Section S-3.4) we have $\Pr[\theta_S \in W_I] \leq O(n^{-4\log n + 1})$. By Lemma S-3.7 we have $\Pr[\overline{\mathcal{G}_\delta}] \leq O(\exp\{\epsilon^2 f(\delta) n^{1-a}/3M\})$. Hence, we get $\mathbb{E}[X_\delta^2] \geq \frac{1}{2n^2 f(\delta)^2} (1 - o(1))$. □

S-3.4 Helpful lemmas to bound the correlation between θ_S and the closest signal

In this section we will introduce some helpful lemmas for proving Proposition 9.

In Lemma S-3.10, we characterize the effect of a single signal x on the posterior mean $\theta_S(\vec{x})$. This lemma directly implies Corollary 5, where we show that if for any signals \vec{x} rearranging at most k signals in L (and R) to get \vec{y} guarantees that the new posterior mean is $\theta_S(\vec{y}) \in \theta_S(\vec{x}) \pm O(kn^{-1-a})$.

Lemma S-3.10. *For any signals \vec{x} observed by the sender we have,*

$$\left| \frac{\partial \theta_S(\vec{x})}{\partial x_i} \right| \leq c_1 \text{Var}_{\theta \sim D_S(\vec{x})}[\theta] + 2\theta_S(\vec{x})^2,$$

where $\text{Var}_{\theta \sim D_S(\vec{x})}[\theta]$ is the variance of the sender's posterior distribution $D_S(\vec{x})$.

Proof. Let $h_S(\theta|\vec{x})$ denote the pdf of the sender's posterior distribution, $h(\theta)$ be the (constant) pdf of the diffuse prior. Note that the pdf of a signal x given that the state of the world is θ (denoted by $\hat{f}(x|\theta)$) equals $f(x - \theta)$, where f is the pdf of F . Recall that when the sender observes \vec{x} they update their posterior in a Bayesian way. Hence, we have,

$$\begin{aligned} h_S(\theta|\vec{x}) &= \frac{\prod_i \hat{f}(x_i|\theta)h(\theta)}{\int_{\hat{\theta}} \prod_i \hat{f}(x_i|\hat{\theta})h(\hat{\theta})d\hat{\theta}} \\ &= \frac{\prod_i f(x_i - \theta)}{\int_{\hat{\theta}} \prod_i f(x_i - \hat{\theta})d\hat{\theta}} \end{aligned}$$

and the sender's posterior mean is

$$\theta_S(\vec{x}) = \frac{\int_{\theta} \theta \prod_i f(x_i - \theta)d\theta}{\int_{\theta} \prod_i f(x_i - \theta)d\theta} \quad (64)$$

We want to understand $\frac{\partial \theta_S(\vec{x})}{\partial x_i}$ which is the effect of a single signal x_i on the posterior mean.

We can deduce

$$\begin{aligned} \frac{\partial \theta_S(\vec{x})}{\partial x_i} &= \frac{\int \theta f'(x_i - \theta) \prod_{j \neq i} f(x_j - \theta)d\theta \int \prod_j f(x_j - \theta)d\theta}{\left(\int \prod_j f(x_j - \theta)d\theta \right)^2} \\ &\quad - \frac{\int \theta \prod_j f(x_j - \theta)d\theta \int f'(x_i - \theta) \prod_{j \neq i} f(x_j - \theta)d\theta}{\left(\int \prod_j f(x_j - \theta)d\theta \right)^2} \\ &= \frac{\int \theta g(x_i - \theta) \prod_j f(x_j - \theta)d\theta \int \prod_j f(x_j - \theta)d\theta}{\left(\int \prod_j f(x_j - \theta)d\theta \right)^2} \\ &\quad - \frac{\int \theta \prod_j f(x_j - \theta)d\theta \int g(x_i - \theta) \prod_j f(x_j - \theta)d\theta}{\left(\int \prod_j f(x_j - \theta)d\theta \right)^2} \end{aligned} \quad (65)$$

where $g(y) = \frac{f'(y)}{f(y)}$.

Using the mean value theorem we can write $g(x_i - \theta) = g(x_i) - g'(\tilde{x}_i)\theta$ for $\tilde{x}_i \in [x_i - \theta, x_i]$. We

then obtain:

$$\begin{aligned}
\frac{\partial \theta_S(\vec{x})}{\partial x_i} &= \frac{\int \theta (g(x_i) - g'(\tilde{x}_i)\theta) \prod_j f(x_j - \theta) d\theta \int \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta \right)^2} \\
&\quad - \frac{\int \theta \prod_j f(x_j - \theta) d\theta \int (g(x_i) - g'(\tilde{x}_i)\theta) \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta \right)^2} \\
&= - \frac{\int \theta g'(\tilde{x}_i)\theta \prod_j f(x_j - \theta) d\theta \int \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta \right)^2} \\
&\quad + \frac{\int \theta \prod_j f(x_j - \theta) d\theta \int g'(\tilde{x}_i)\theta \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta \right)^2} \tag{66}
\end{aligned}$$

Next, we note that for a well-behaved distribution, $|g'(\tilde{x})| \leq c_1$. Hence, we can bound, We can then simplify:

$$\left| \frac{\partial \theta_S(\vec{x})}{\partial x_i} \right| \leq c_1 \left(\frac{\int \theta^2 \prod_j f(x_j - \theta) d\theta}{\int \prod_j f(x_j - \theta) d\theta} + \left(\frac{\int \theta \prod_j f(x_j - \theta) d\theta}{\int \prod_j f(x_j - \theta) d\theta} \right)^2 \right) \tag{67}$$

$$\begin{aligned}
&= c_1 \left(\mathbb{E}_{\theta \sim D_S(\vec{x})}[\theta^2 | \vec{x}] + \mathbb{E}_{\theta \sim D_S(\vec{x})}[\theta | \vec{x}]^2 \right) \\
&= c_1 \left(\text{Var}_{\theta \sim D_S(\vec{x})}[\theta] + 2 \mathbb{E}_{\theta \sim D_S(\vec{x})}[\theta | \vec{x}]^2 \right) \tag{68}
\end{aligned}$$

□

Next. we bound the shift in θ_S when rearranging k signals in L (or R) using Lemma S-3.10. We note that $\text{Var}_{\theta \sim D_S(\vec{x})}[\theta] = O(1/n)$ and $\theta_S \in I$.

Corollary 5. *Assume that the posterior mean $\theta_S(\vec{x})$ lies within the interval $C \subset I$. Consider a subset of k signals in a subset A of length ℓ . Any rearrangement of these signals within A changes the posterior mean by $O(k\ell n^{-1+2\epsilon})$.*

Proof. We prove this by using the mean value theorem on the function $\theta_S : \mathbb{R}^n \rightarrow \mathbb{R}$. Given \vec{x} , consider a subset of signals $\vec{x}_{[k]} \in A$. Let \vec{y} be any vector such that $y_i \in A$ for all $i \in [k]$ and $y_i = x_i$ for the rest. By mean value theorem we get for some \vec{z} such that:

$$\theta_S(\vec{y}) = \theta_S(\vec{x}) + \nabla \theta_S(\vec{z}) \cdot (\vec{y} - \vec{x})$$

Note that $|y_i - x_i| \leq \ell$ for all $i \in [k]$ and $y_i - x_i = 0$ otherwise. That is, at most k terms with $|y_i - x_i| \neq 0$. Hence, we get,

$$\begin{aligned}
|\theta_S(\vec{y}) - \theta_S(\vec{x})| &= |\nabla\theta_S(\vec{z}) \cdot (\vec{y} - \vec{x})| \\
&\leq \sum_{i=1}^k c_1 (\text{Var}_{\theta \sim D_S(\vec{x})}[\theta] + 2\theta_S(\vec{x})^2) |y_i - x_i| \quad (\text{By Lemma S-3.10}) \\
&= c_1 O(n^{-1+2\epsilon}) \sum_{i=1}^k \ell
\end{aligned}$$

Since $x_i, y_i \in A$ we bound $|x_i - y_i| \leq \ell$. Further, by bounding $\text{Var}_{\theta \sim D_S(\vec{x})}[\theta]$ by $O(1/n)$, and $\theta_S(\vec{x})$ by $n^{-1/2+\epsilon}$, we get $\theta_S(\vec{y}) = \theta_S(\vec{x}) \pm k\ell O(n^{-1+2\epsilon})$ when rearranging at most k signals in each A . □

Next, we show that the density of the posterior mean is nice in the interval C .

Claim S-3.6. *Assume that the density function f has exponential (or thinner) tails. Let $\tau(\cdot)$ be the density function of the posterior mean. Then for all $\bar{\theta} \in C$ we have $\tau(\bar{\theta} + \epsilon') = \tau(\bar{\theta}) (1 + O(n^{-1/22})) + O(e^{-n^{1/22\alpha}})$ for all $0 < \epsilon' \leq 1/n^b$.*

Proof. Fix a posterior mean $\bar{\theta} \in [-n^{-\frac{1}{2}+\epsilon}, n^{-\frac{1}{2}+\epsilon}]$ and consider all the signal draws $X(\bar{\theta}) = \{\vec{x} | \bar{\theta}(\vec{x}) = \bar{\theta}\}$ that generate this posterior mean. We know that $\tau(\bar{\theta}) = \int_{\vec{x} \in X(\bar{\theta})} \prod_{x_i \in \vec{x}} f(x_i) dx$. Now consider $\bar{\theta} + \epsilon'$. We can couple all the signal realizations in $X(\bar{\theta} + \epsilon')$ and $X(\bar{\theta})$ by considering uniform shifts of the corresponding \vec{x} by ϵ' . This follows from the assumption of a diffuse prior, and we get $\theta_S(\vec{x} + \epsilon') = \theta_S(\vec{x}) + \epsilon'$. That is, $\tau(\bar{\theta} + \epsilon') = \int_{\vec{x} \in X(\bar{\theta})} \prod_{x_i \in \vec{x}} f(x_i + \epsilon') dx$.

Next, consider the probability of observing x versus the coupled signal realizations $x + \epsilon'$.

$$\prod_{x_i \in \vec{x}} f(x_i + \epsilon') = \prod_{x_i \in \vec{x}} [f(x_i) + f'(\tilde{x}_i)\epsilon'] \tag{69}$$

Recall that, by our assumption on g' we have $|f'(x)| \leq c|x|^m f(x)$.

- Note that, for all $|x_i| < 4(\log n)^2$, we have

$$\begin{aligned}
|f(x_i + \epsilon') - f(x_i)| &= |\epsilon' f'(\tilde{x}_i)| \\
&\leq |\epsilon' c(\tilde{x}^m)| f(\tilde{x}) \\
&\leq |\epsilon' c(4 \log n)^{2m}| f(\tilde{x}) \\
&\leq |\epsilon' (n^{-b+1/22})| f(\tilde{x}) \quad (|\epsilon'| \leq n^{-12/11} \text{ and } (\log n)^{2m} = o(n^{1/22}))
\end{aligned}$$

Note that, wlog we can assume that $\text{sign}(x_i + \epsilon) = \text{sign}(x_i)$ because f is symmetric. Thus, we get $f(\tilde{x}_i)$ is between $f(x_i)$ and $f(x_i + \epsilon')$, because f is single-peaked. If $f(x_i) \leq f(\tilde{x}_i) \leq$

$f(x_i + \epsilon')$ we get $0 \leq f(x_i + \epsilon') - f(x_i) \leq c'(n^{-b+1/22})|f(\tilde{x}_i)| \leq c'(n^{-b+\epsilon})|f(x_i + \epsilon')$. Thus, $f(x_i) \leq f(x_i + \epsilon) \leq f(x_i) \left(\frac{1}{1-c'(n^{-b+\epsilon})} \right) \leq f(x_i)(1 + c'n^{-b+\epsilon})$.

Similarly, if $f(x_i) \geq f(\tilde{x}_i) \geq f(x_i + \epsilon)$, then we get $f(x_i) \geq f(x_i + \epsilon) \geq f(x_i) \left(\frac{1}{1+c'(n^{-b+\epsilon})} \right) \geq f(x_i)(1 - c'n^{-b+\epsilon})$.

- Further, by our assumption that f has exponential tails we have $\Pr_{\vec{x}}[\exists x_i \text{ s.t. } |x_i| > 4(\log n)^2] \leq ne^{-4(\log n)^2} = (n^{-4 \log n+1})$.

If $x_i \in [-4(\log n)^2, 4(\log n)^2]$ for all i , then we bound $\prod_{x_i \in \vec{x}} f(x_i + \epsilon') = \prod_i f(x_i)(1 + O(n^{-b+\frac{1}{22}})) = (\prod_i f(x_i))(1 + O(n^{-b+\frac{1}{22}}))^n = (\prod_i f(x_i))(1 + O(n^{1-b+\frac{1}{22}})) = (\prod_i f(x_i))(1 + O(n^{-\frac{1}{22}}))$ for $b = 12/11$

Hence, we get

$$\tau(\bar{\theta} + \epsilon') = \int_{-n^{1/22\alpha}}^{n^{1/22\alpha}} \left(\prod f(x_i) \right) (1 + O(n^{-\frac{1}{22}})) \mathbf{1}\{\vec{x} \in X(\hat{\theta})\} d\vec{x} + O(n^{-4 \log n+1})$$

Therefore, $\tau(\bar{\theta} + \epsilon') = \tau(\bar{\theta})(1 + O(n^{-\frac{1}{22}})) + O(n^{-4 \log n+1})$ for all $0 < \epsilon' < 1/n^b$ and $\bar{\theta} \in C$. \square

Observe that, if $\tau(\bar{\theta}) \geq c'n^{1/22}n^{-4 \log n+1}$ then we can get $\tau(\bar{\theta})(1 + O(n^{-\frac{1}{22}}))$.

Recall that $W_I \subset \mathcal{P}$ is the set of all intervals C such that $\tau(\bar{\theta}) < c'n^{1/22}n^{-4 \log n+1}$ for all $\bar{\theta} \in C$ and some constant $c' > 0$. We show that the total probability mass of these intervals is $O(n^{-4 \log n+1})$.

Claim S-3.7. $\Pr[\theta_S \in W_I] \leq O(n^{-4 \log n+1})$.

Proof. This is simply because there are at most $2n^{b-\frac{1}{2}+\epsilon}$ many intervals in \mathcal{C} (since each interval is of size n^{-b}). Therefore,

$$\Pr[\theta_S \in W_I] \leq \sum_{C \in W_I} \int_C \tau(\bar{\theta}) d\bar{\theta} \leq 2n^{b-\frac{1}{2}+\epsilon} \left(n^{-b} \cdot c'n^{1/22-4 \log n} \right) \leq O(n^{-4 \log n+6/11+\epsilon}).$$

\square

Using Corollary 5 we bound the correlation between the events $\theta_S \in C$ and any realization of k anecdotes in L_δ (and R_δ).

Lemma S-3.11. Fix any $C \in \mathcal{P} \setminus W_I$. For $A \subseteq [n]$, $|A| = k$, let K_A denote the event that $x_i \in L_\delta$ (and R_δ) iff $i \in A$. Let $\vec{s} \in L_\delta^k$ be a set of k (at most $c'n^{1-a}$) signals in L_δ . Then $\Pr[\theta_S \in C | K_A, \vec{x}_A = \vec{s}] \leq \Pr[\theta_S \in C \pm kn^{-1-a} | K_A]$.

Proof. Let $A(\vec{z}) = \{\tilde{z} \in \mathbb{R}^{n-k} : \theta(\vec{z} \cup \tilde{z}) \in C\}$ for any subset of k signals $\vec{z} \in L_\delta^k$. By Corollary 5 we know that by changing \vec{s} to any \vec{z} (in L_δ) for each $\tilde{s} \in A(\vec{s})$ $\theta_S(\tilde{s}, \vec{z}) = \theta_S(\tilde{s}, \vec{s}) + O(kn^{-1-a+\epsilon})$. If $\theta_S(\tilde{s}, \vec{s}) \in C$ then $\theta_S(\tilde{s}, \vec{y}) \in C \pm kn^{-1-a+\epsilon}$.

This implies, $\Pr[\theta_S \in C | K_A, \vec{x}_A = \vec{s}] \leq \Pr[\theta_S \in C \pm kn^{-1-a+\epsilon} | K_A, \vec{x}_A = \vec{y}]$ for all $\vec{z} \in L_\delta^k$. Hence, we get,

$$\Pr[\theta_S \in C | K_A, \vec{x}_A = \vec{s}] \leq \Pr[\theta_S \in C \pm kn^{-1-a} | K_A]$$

□

Similarly, we have a lower bound on $\Pr[\theta_S \in C | K_A, x_A = \vec{s}]$.

Lemma S-3.12. *Fix any $C \in \mathcal{P} \setminus W_I$. For $A \subseteq [n]$, $|A| = k$, let K_A denote the event that $x_i \in L_\delta$ (and R_δ) iff $i \in A$, and no signals in C_δ . Let $\vec{s} \in L_\delta^k$ be a set of k (at most $c'n^{1-a}$) signals in L_δ . Then $\Pr[\theta_S \in C | K_A, x_A = \vec{s}] \geq \Pr[\theta_S \in C \setminus E | K_A]$, where $E \subset C$ as the union of the first and last $kn^{-1-a+\epsilon}$ length sub-interval of C .*

Proof. Let $A(\vec{z}) = \{\tilde{z} \in \mathbb{R}^{n-k} : \theta(\vec{z} \cup \tilde{z}) \in C\}$ for any subset of k signals $\vec{z} \in L_\delta^k$. By Corollary 5 we know that by changing \vec{s} to any \vec{z} (in L_δ) for each $\tilde{s} \in A(\vec{s})$ $\theta_S(\tilde{s}, \vec{z}) = \theta_S(\tilde{s}, \vec{s}) + O(kn^{-1-a+\epsilon})$. If $\theta_S(\tilde{s}, \vec{s}) \in C$ then $\theta_S(\tilde{s}, \vec{y}) \in C \pm kn^{-1-a+\epsilon}$. Thus, $\theta_S(\tilde{s}, \vec{z}) \in C \setminus E$ then $\theta_S(\tilde{s}, \vec{s}) \in C$.

This implies, $\Pr[\theta_S \in C | K_A, \vec{x}_A = \vec{s}] \geq \Pr[\theta_S \in C \setminus E | K_A, \vec{x}_A = \vec{z}]$ for all $\vec{z} \in L_\delta^k$. Hence, we get,

$$\Pr[\theta_S \in C | K_A, \vec{x}_A = \vec{s}] \geq \Pr[\theta_S \in C \setminus E | K_A]$$

□