Communicating with Anecdotes

Nika Haghtalab[1], Nicole Immorlica[2], Brendan Lucier[2], Markus Mobius[2], and Divyarthi Mohan[3]

[1]University of California, Berkeley
[2]Microsoft Research
[3]Tel Aviv University

Abstract

We study a communication game between a sender and receiver where the sender has access to a set of informative signals about a state of the world. The sender chooses one of her signals and communicates it to the receiver. We call this an "anecdote". The receiver takes an action, yielding a utility for both players. Sender and receiver both care about the state of the world but are also influenced by a personal preference so that their ideal actions differ. We characterize perfect Bayesian equilibria when the sender cannot commit to a particular communication scheme. In this setting the sender faces "persuasion temptation": she is tempted to select a more biased anecdote to influence the receiver's action. Anecdotes are still informative to the receiver but persuasion comes at the cost of precision. This gives rise to "informational homophily" where the receiver prefers to listen to like-minded senders because they provide higher-precision signals. In particular, we show that a sender with access to many anecdotes will essentially send the minimum or maximum anecdote even though with high probability she has access to an anecdote close to the state of the world that would almost perfectly reveal it to the receiver. In contrast to the classic Crawford-Sobel model, full revelation is a knife-edge equilibrium and even small differences in personal preferences will induce highly polarized communication and a loss in utility for any equilibrium. We show that for fat-tailed anecdote distributions the receiver might even prefer to talk to poorly informed senders with aligned preferences rather than a knowledgeable expert whose preferences may differ from her own because the expert's knowledge also gives her likely access to highly biased anecdotes. We also show that under commitment differences in personal preferences no longer affect communication and the sender will generally report the most representative anecdote closest to the posterior mean for common distributions.

1 Introduction

Economists usually assume that people learn about the world by updating the parameters of some underlying model as new evidence arrives. Such models can be efficiently communicated to others: for example, abstracts of academic papers might summarize main results in the form of model parameters such as the "elasticity of demand" in a certain industry or an overall toxicity score of a new radiation treatment. This type of communication is very natural for modeling learning amongst experts who have already agreed on a common set of models that provide a "language" for their field of study.

However, communication between experts can be incomprehensible to non-experts who have no understanding of such models. Such agents instead often rely on *anecdotal evidence*. Consider, for example, an investor who is trying to decide how much money to invest in a given mutual fund. An optimal investment decision depends on the state of the future economy. If the economy is growing quickly, the investor would like to invest a lot of money; if it is growing slowly or shrinking, she may want to invest less. A typical investor lacks expertise in economic analysis and so relies on information from other actors, such as politicians, newspapers or financial analysts. This information is often provided in the form of anecdotes, say the percent increase in the number of jobs in a given sector this quarter. In fact, newspaper articles often simply report a selection of related facts on a topic.

In this paper, we model these situations as a communication game between a sender and a receiver. The receiver (e.g., investor in above example) takes an action. Both the sender (e.g., politician) and receiver are impacted by this action and how it relates to the state of the world (for example, the true state of the economy). But they also have *personal preferences* such that their ideal action will differ even if they have access to the same information. These personal preferences are shifts relative to the state of the world. For example, the sender may be "left-leaning" and prefer an action a bit to the left of the state of the world whereas the receiver is "right-leaning."

In our model, only the sender observes informative signals about the state of the world which we refer to as anecdotes. She can select one of these signals to send to the receiver. Importantly, we assume throughout that the communication of anecdotes is always truthful: the sender cannot make up "fake news", for example. As the sender cares about the action of the receiver, she faces a *persuasion temptation*. For example, if the politician would like to persuade the investor that the economy is booming so that he invests his money in the stock market, the politician might select a more positive anecdote. The sender must balance this temptation against the potential *information loss* incurred by sending unrepresentative anecdotes. If her communication carries very little information about the state of the world, the receiver's action will be poorly correlated with the state of the world which in turn hurts the sender.

We will analyze the behavior of the sender and the receiver at equilibrium to study the efficiency and bias of shared anecdotes under this anecdotal communication framework. Our main result is that, when the sender is unable to commit to a communication strategy, the tension between persuasion temptation and information loss leads to a partial unraveling in the communicated anecdote: it is biased to a degree determined by the misalignment in personal preferences. This has implications for homophily: a receiver may prefer an aligned sender with few anecdotes to an unaligned one with many. In contrast, when the sender can commit to her communication strategy, she will send the most representative anecdote.

More precisely, we consider a Bayesian sender and receiver. We note that a Bayesian receiver will anticipate that the sender might communicate a biased anecdote and properly account for it. We characterize sender and receiver's best responses in Theorems 1 and 2 and show that a perfect Bayesian equilibrium exists (see Theorem 4).

In equilibrium, the sender will not be able to change the mean posterior beliefs of the receiver: the receiver will *always* learn *something* because anecdotes are always truthful. However, this does not imply that the sender's attempts to persuade the receiver are without costs. In Theorem 3, we characterize perfect Bayesian equilibria and show that there is monotone mapping between the gap in personal preferences and the bias of the signal the sender sends. This shows that even for small differences in personal preferences between sender and receiver the anecdotes that are communicated can be highly biased. In these cases, the sender will tend to choose anecdotes from the tail of the anecdote distribution. While this does not succeed in persuading the receiver to take a biased action, it does destroy precision: anecdotes in the tail are more thinly distributed and hence reveal information that has high variance about the true state of the world.

This bias in equilibrium gives rise to "informational homophily". In these cases, receivers may prefer to communicate with senders with similar personal preferences to eliminate the loss of information caused by sender's temptation to persuade. In fact, we show in Proposition 3 that a receiver might sometimes prefer to talk to a sender with access to just a few signals (or even just one) compared to a well-informed expert who has access to a vast number of anecdotes but has a different personal preference from the receiver. This will be the case for fat-tailed distributions (such as the Laplace distribution) because the sender can more likely access extreme anecdotes which are less informative to the receiver, even after they have been debiased. This insight can explain why receivers might not seek out or listen to experts with different backgrounds.

A sender may be able to *commit* to a communication scheme in settings where her reputation precedes her, such as a reputable newspaper. When the sender can *commit* to a communication scheme, we see a different type of behavior in the equilibrium. In these cases, senders will send the most informative signal no matter the gap between the personal preferences. As shown in Theorem 5, this signal also minimizes the cost of the receiver. When anecdotes follow a single-peaked distribution, such as the normal or Laplace distribution, the most informative anecdotes are those close to the peak, which (under some mild technical conditions) the sender approximates through her posterior mean.

Our paper is organized as follows. In Section 3, we introduce our model. In Section 4, we characterize the best-response behavior. In Section 5, we characterize sender's and receiver's utility. In Section 6, we characterize the perfect Bayesian equilibria and discuss receiver's choice between expert and non-expert senders. In Section 7, we consider settings where the sender can commit and characterize optimal commitments.

2 Related Work

Our model is related to two strands of literature in economics, namely papers on *framing* and papers on *selection*.

The framing literature allows the sender to send any type of signal. There are subcategories: the Bayesian persuasion literature assumes that the sender can commit to a particular signaling scheme while the "cheap-talk" literature assumes no such commitment. Kamenica and Gentzkow [2011] introduce the Bayesian persuasion model where the sender commits to sending a signal that is consistent with her information in a Bayesian sense. In that model, as in ours, there is a state of the world distributed according to a common prior. The sender commits to a signaling scheme, mapping observations about the state of the world to an arbitrary signal. She then observes the state and transmits the corresponding signal to the receiver. The receiver then picks an action. The sender's payoff is a function of the receiver's action, and so the sender wishes to "persuade" the receiver to take particular actions. Kamenica and Gentzkow [2011] provide a characterization of the optimal signaling scheme. Our work approaches a similar question for a constrained signaling problem, where the sender's signaling scheme is restricted to take the form of sending one of a collection of anecdotes. This constraint imposes friction that limits the sender's ability to persuade, and indeed we find that the optimal choice of the sender under these restrictions will be to communicate as informatively as possible about the state of the world. Crawford and Sobel [1982] introduced the "cheap-talk" model where the sender has no commitment power. The sender can again choose an arbitrary signaling scheme, but is not able to commit to the signaling strategy in advance. Crawford and Sobel [1982] show that despite the lack of commitment, a non-trivial amount of information can be communicated at equilibrium, and moreover such equilibria take the form of sending a coarsening of the signal available to the sender.

Given our restriction that the sender must choose from a set of available anecdotes, our work fits into the literature on selection. The literature on voluntary disclosure was introduced by Grossman and Hart [1980], Grossman [1981], and Milgrom [1981]. These papers consider the setting of a seller who can choose whether to disclose information about a product to a buyer, and can make this choice based on the information itself. Similar to our model of communication via anecdotes, the seller in these papers cannot arbitrarily distort information about the product, but rather simply choose whether or not to reveal it. Importantly, the seller cannot necessarily commit to their revelation strategy in advance. The main result is that in every sequential equilibrium, the seller fully discloses her information. This so-called unraveling is driven by the fact that the seller can not commit to a signaling scheme.

Milgrom [1981] further analyze a setting where the seller is constrained in the amount of information she can reveal. Namely, she has access to a set of data points about her product (akin to our anecdotes) and can only reveal a fixed number of them (e.g., just one as in our setting). He then shows that the seller always reveals the most favorable information about the quality of the product. We see a similar unraveling in our setting with an unobserved signaling scheme – the sender ends up sending an extreme signal, but not necessarily the most extreme signal (due to the structure of payoffs which differs from that of Milgrom [1981]).

This contrasts to the setting of an observed signaling scheme, which can be interpreted as having no unraveling (the sender sends the most informative signal).

There are other settings with partial unraveling. Dye [1985] and Jung and Kwon [1988] show unraveling breaks down when the receiver is uncertain whether the sender is informed. Martini [2018] shows unraveling breaks down when information is multi-dimensional. Ali et al. [2020] describe a disclosure game in which senders (buyers in their setting) gain by coarsening their information. Fishman and Hagerty [1990] consider a setting, similar to ours, in which the sender is restricted in the number of signals she can disclose and study the optimal amount of discretion a designer should permit the sender.

In our paper, as in much of the literature, the unraveling, or partial unraveling, is driven by a lack of the power to commit on the part of the sender prior to observing the signals. A recent work also studies the power of the receiver to commit to a mapping from received information to sender payoff as in mechanism design (the difference being that the sender is restricted to voluntary disclosure strategies). Hart et al. [2017] define a disclosure game in which the equilibrium outcome without commitment coincides with the optimal outcome with commitment.

A related line of work studies the power of commitment for the *receiver* in persuasion games with hard information. Glazer and Rubinstein [2004] consider a game of persuasion where the sender can communicate via cheap talk about a multi-dimensional state of the world, after which the receiver can verify one of the features before taking a binary action. They show that commitment does not provide additional power to the receiver: the optimal action choice rule is implementable at ex post equilibrium. Closer to our model, Glazer and Rubinstein [2008] show that the same result holds without cheap talk and partial verification. Instead, each state of the world is associated with a subset of feasible signals, and the sender always sends exactly one signal before the receiver takes an action. They again find that the receiver gains no additional benefit from being able to commit. Sher [2011] extends this latter model to non-binary actions, and shows that the same result holds under a concavity relationship between receiver and sender utilities. 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, and we find that even a small desire to persuade can lead to a significant loss of communication fidelity (and hence welfare) in the absence of commitment.

More generally, this paper complements a long line of work on communication and disclosures games (e.g., Milgrom and Roberts [1986], Verrecchia [2001], Di Tillio et al. [2021], Dziuda [2011], Wolinsky [2003], Chen [2011], Jovanovic [1982], Seidmann and Winter [1997]).

3 Model

We consider a communication game played by two players, a sender and a receiver. The sender has information about a payoff-relevant state of the world $\theta \in \Theta = \mathbb{R}$ drawn from a common prior, say the danger of COVID-19 for instance. The receiver, in turn, chooses a payoff-relevant action $a \in \mathcal{A} = \mathbb{R}$, say when to wear a mask.

Preferences. Players' preferences over their actions depend on the state of the world. For example, if COVID-19 is highly dangerous, both players might desire a higher frequency of mask-wearing. However, their preferences can differ. Perhaps the sender values global health more and personal freedom less than the receiver, thereby preferring a higher frequency of mask-wearing than the receiver in any state of the world. 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 the personal preferences are publicly known and write $\Delta = (M_S - M_R)$ for the known difference in personal preferences. Intuitively, Δ captures the preference misalignment between sender and receiver.[1]

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 noise distribution F over the reals, we model these shareable signals as a set of n samples $x_1, \ldots, x_n$ where each $x_i = \theta + \epsilon_i$ for a noise parameter $\epsilon_i \sim F$ drawn independently. We will write $\vec{x} = (x_1, \ldots, x_n)$ for the profile of samples which we will refer to as *anecdotes* from now on. For example, an anecdote might come in the form of a survey or research paper that the sender has access to. 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 distribution F as well as the number of anecdotes, n, is common knowledge.[2]

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* by an additional signal. Given a distribution G over the reals, the sender has access to a signal $y = \theta + \gamma$ for a noise parameter $\gamma \sim G$. 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 G represents a diffuse prior such as $N(0, \infty)$. In the latter case, signal y reveals no additional information about θ to the

[1] 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.

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

sender. As with $\vec{x}$, we assume y is private knowledge of the sender but that the distribution G is common knowledge.

Communication The sender communicates exactly one anecdote in $\vec{x}$ to the receiver influencing the receiver's action. While the anecdote is communicated honestly, the sender can cherry-pick from among the set of anecdotes she has access to. Note that the sender's side information y cannot be communicated, only an anecdote. While our anecdotes can not be manipulated or falsified, they can be selected in a biased manner, as is common in much 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.

Equilibrium. 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 $(\vec{x}, y)$. A strategy for the sender in our game is a communication scheme $\pi\colon \mathbb{R}^n \times \mathbb{R} \to \mathbb{R}$ 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 $i \in [n]$. The receiver then selects action a after observing the revealed signal.

Formally, the timing of our game is as follows. In round 0, nature chooses state θ, anecdotes $x_1, \ldots, x_n$, and signal y. Anecdotes $\vec{x}$ and signal y are visible to the sender but not the receiver. In round 1, the sender selects one of the anecdotes as described above; this choice is observed by the receiver. In round 2, the receiver selects an action. Payoffs are then realized as described above. Since the receiver does not observe the choice of nature, a strategy for the receiver is an action rule $\alpha\colon \mathbb{R} \to \mathbb{R}$ that maps the sender's chosen anecdote to a choice of action. 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.

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

1. *For each x, action $\alpha^*(x)$ maximizes expected receiver utility given distribution $B(x)$ over θ, i.e.,*
$$\alpha^*(x) \in \arg\max_a \{ \mathop{\mathbb{E}}_{\theta \sim B(x)} [u_R(a, \theta)] \}.$$

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 $\pi^*(\vec{x}, y) = x$.*

3. *For each $\vec{x}$ and y, $\pi^*(\vec{x}, y)$ maximizes sender utility given α^*, i.e.,*
$$\pi^*(\vec{x}, y) \in \arg\max_{x_i \in \vec{x}} \{ \mathop{\mathbb{E}}_{\theta} [u_S(\alpha^*(x_i), \theta) \mid (\vec{x}, y)] \}.$$

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

In Section 7 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, e.g., in journalism the public generally understands the bias of various newspapers. In this case an equilibrium does not need to meet requirement (3), 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

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

Diffuse Prior. For tractability, we will make an assumption about the primitives of our model. Our assumption states that the prior over the state of the world, θ, is diffuse. 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 emphasize that the exact form of the prior over θ is not important as long as it is a diffuse prior that reveals no information about θ, i.e., its density is almost uniform everywhere. This assumption simplifies our reasoning about sender and receiver posterior means to show that these means are unbiased estimators of θ. This assumption also emphasizes the interesting extreme where the receiver, absent communication from the sender, has absolutely no knowledge about the state of the world. We formalize this assumption and its implications further in Appendix A.

4 Best-Response Characterizations

Communication serves multiple purposes: it both transmits information and influences actions. When preferences are not aligned, these purposes can be at odds. To transmit as much information as possible, a sender intuitively wants to send the most accurate anecdote, i.e., one that minimizes the variance of the receiver's posterior. But if the sender wants to influence the receiver's action, she might wish to send a slightly biased anecdote, one that pulls the receiver's action towards her own personal preference. As we will see, in her best-response, the sender balances between these objectives by targeting a particular bias in her communication. The receiver, in turn, chooses an action by debiasing the communication.

To formalize these results, it is useful to introduce the notion of a translation-invariant communication scheme. 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 $(x_1 + \delta, x_2 + \delta, \ldots, x_n + \delta)$.

Definition 2. *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 $\delta \in \mathbb{R}$.*

Not all communication schemes are translation invariant. For example, the 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 irrational and the maximum otherwise. However, in our world with a diffuse prior, where specific numbers and their properties have no meaning, most "natural" communication schemes are translation invariant. For example, sending the minimum anecdote or the median anecdote or the anecdote closest to the posterior mean of the sender are all translation invariant.[3]

We can define a similar notion for the receiver's action rule.

Definition 3. *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 $\sigma(\alpha)$. Action rules that are translations are also "natural" in our setting. Such rules correspond to a receiver who simply believes the anecdote she hears is representative of the state, albeit potentially with a shift. 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.[4]

To prove that this intuitive form of action rule is in fact a best response to some communication schemes, we first define the *bias* of a communication scheme. Given a translation-invariant communication scheme π, we'll say the bias of π, $\beta(\pi)$, is equal to $\mathbb{E}_{\theta,\vec{x},y}[\pi(\vec{x}, y) - \theta]$. For example, if the sender always sends the signal closest to her posterior mean and the signal distribution is symmetric then the bias is equal to 0. On the other hand, if the sender always selects the minimum anecdote then the bias is the expected distance of the minimum anecdote from the state θ.

The best response of a receiver to a translation-invariant communication scheme simply shifts the received anecdote by the bias to obtain an unbiased estimate of the state (and then additionally by the receiver's personal preference) – hence his action rule is a translation.

Theorem 1. *For any translation-invariant communication scheme π, the best response of the receiver to π is a translation with shift $M_R - \beta(\pi)$.*

Proof. Let $x = \pi(\vec{x})$ and let the belief distribution be $B(x) = D_{\pi,x}$. Note by definition of bias, $x - \beta(\pi)$ is an unbiased estimator of θ. By Assumptions 1, the receiver's posterior mean about the state of the world is simply equal to the value of the unbiased estimator

[3] As we will see in Section 6, there are perfect Bayesian equilibria in which the sender selects a translation-invariant communication scheme.

[4] Similarly, we will see in Section 6, there are perfect Bayesian equilibria in which the receiver's action rule is a translation.

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

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

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. □

Note that Theorem 1 characterizes the receiver's best response among *all* action rules to a translation-invariant communication scheme. In other words, the best response to a translation-invariant scheme is in fact itself a translation.

With the receiver's action rule in hand, we can similarly ask about the sender's best response to an action rule that is a translation. We start by introducing a particular communication scheme where the sender selects a signal that is closest to a shift r from her posterior mean. We call these *targeting schemes.* Formally, write θ_S for the random variable corresponding to the sender's posterior mean of θ, and $\theta_S(\vec{x},y)$ for the realization of this posterior mean given $(\vec{x},y)$.

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 that $\theta_S(\vec{x}+\delta,y+\delta) = \theta_S(\vec{x},y)+\delta$, formalized in Appendix A. Hence a targeting scheme is translation invariant.

We can now explicitly calculate the best response of a sender to an action rule that is a translation. It is the targeting scheme whose offset negates the shift of the action rule.

Theorem 2. *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 $M_S-\sigma(\alpha)$.*

Proof. 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

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

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. □

We again note that Theorem 2 characterizes the sender's best-response among all (not necessarily translation-invariant) communication schemes. That is, the best response to an action rule that is a translation is translation-invariant.

5 Utility Calculations

Before proving the existence of translation-invariant equilibria in Section 6 we first explore the sender and receiver utility for such strategy profiles. This allows us to make welfare statements about these equilibria.

Sender Utility. First we show that the sender's utility is driven by two components: the information loss of her scheme and the persuasion temptation she faces. Suppose that the sender is using communication scheme π_S, and the receiver's action rule α is a best response to communication scheme π_R (which may not be equal to π_S). Then we can write the sender's expected utility as a function of the variance of her chosen anecdote, the preference misalignment, and the difference between the true and perceived bias.

Proposition 1. *Suppose the receiver's action rule is a translation that is a best response to a translation-invariant communication scheme π_R. Then the sender's expected utility from any communication scheme π_S (not necessarily translation-invariant) is*

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

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

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

Let $w = \pi_S(\vec{x},y) - (\theta_S(\vec{x},y) + \beta(\pi_S))$ and $z = (M_R - M_S) + (\beta(\pi_S) - \beta(\pi_R))$ 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. □

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 can interpret this variance as the *information loss* of the communication scheme. Holding the second component fixed, the sender prefers to minimize this term or, equivalently, maximize the amount of information she communicates. Such a scheme will have some fixed bias, e.g., for Gaussian noise distributions, the minimum variance anecdote is the one closest to the posterior mean which has zero bias. However, this bias impacts the second component of the sender's utility too, $[(M_R - M_S) + (\beta(\pi_S) - \beta(\pi_R))]^2$. This component can be thought of as the *persuasion temptation.* When the receiver's perception of the sender's bias $\beta(\pi_R)$ is fixed, the sender best-responds by choosing a slightly more biased scheme, namely one with $\beta(\pi_S) = \beta(\pi_R) + (M_S - M_R)$, in an attempt to persuade the receiver to take an action closer to her own personal preference. In a thought experiment, if we hold the first component fixed and imagine allowing the sender and receiver to successively best-respond to one-another, we see that there would be complete unraveling to a scheme with the maximum bias. Hence it is the tension between the information loss and persuasion temptation that results in partial unraveling.

Receiver Utility With this interpretation in hand, we can show that the receiver utility is driven by the information loss alone, i.e., the variance of the anecdote chosen by the sender.

Proposition 2. *Let π be a translation invariant scheme, and let α be the best response of the receiver. Then the loss of the receiver is the variance of the anecdote $\pi(\vec{x}, y)$,*

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

Proof. From Theorem 1 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 scheme π. Thus, we have receiver's loss (for any fixed θ) equals

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

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

□

Note that this theorem implies that, if the sender can commit to her strategy and we impose translation invariance on the sender, the receiver's action rule will be a translation. We will explore this further in Section 7.

6 Perfect Bayesian Equilibria

We now characterize and prove the existence of translation-invariant equilibria. We will show that equilibria must satisfy a particular fixed-point equation that pins down the relationship between the preference misalignment and the bias of the sender's communication scheme. We then argue this fixed-point equation has a solution which is, in some cases, unique – in particular for strictly single-peaked and symmetric anecdote distributions. This allows us to explore comparative statics such as which environments create more biased communication or higher utility.

6.1 Characterization and Existence

To start, we note an implication of best-response Theorems 1 and 2 is that if we can find a pair (π, α) that are mutual best-responses among the class of strategies described (i.e., targeting schemes and translations, respectively), then they must form a perfect Bayesian equilibrium of our communication game. We call such an equilibrium *translation-invariant*. Translation-invariance is a particularly appealing equilibrium-selection criterion in our model with a diffuse prior as the communicated anecdote and resulting action of such equilibria relative to θ are independent of the value of θ in expectation.

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

We can use the best-response Theorems 1 and 2 to characterize translation-invariant equilibria. To do so, we need to understand the bias of a targeting scheme or, equivalently, the target that will result in a given bias.

Definition 6. *Given a bias $\delta \in \mathbb{R}$, we define $r(\delta) \in \mathbb{R}$ to be the value such that the targeting scheme with offset $r(\delta)$ has bias δ, if such a value exists.*

We note that the offset and bias of a targeting scheme are not necessarily equal due to the varying density of the anecdote distribution. It turns out the difference between the offset and bias is the key desiderata for the sender. In particular, this difference must equal the preference misalignment in any perfect Bayesian equilibrium.

Theorem 3. *A pair (π, α) is a translation invariant perfect Bayesian equilibrium if and only if there exists some value $\delta \in \mathbb{R}$ such that π is the targeting scheme with offset $r(\delta)$, α is a translation with shift $(M_R - \delta)$, and*

$$r(\delta) - \delta = M_S - M_R.$$

Proof. Suppose (π, α) is a translation invariant PBE. Then α must be a translation. Define δ so that $M_R - \delta$ is the shift of α. Then by Theorem 2 we know that π is a targeting scheme with offset $M_S - (M_R - \delta) = \delta + (M_S - M_R)$. Moreover, by Theorem 1 we must have that $\sigma(\pi) = M_R - \beta(\pi)$. Since $\sigma(\pi) = M_R - \delta$ we conclude that $\beta(\pi) = \delta$. Since π has offset $\delta + (M_S - M_R)$, we conclude from the definition of $r(\delta)$ that

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

as required.

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

□

It remains to show that such equilibria exist, i.e., that the fixed-point equation has a solution. Note the right-hand side is simply the preference misalignment and is hence fixed by the primitives of the model. To gain intuition for why the scheme can be selected to satisfy this equation for an arbitrary constant, consider a Gaussian noise distribution for anecdotes. Since this is zero-mean, when the offset r is zero, the bias of the scheme δ is also zero, so the left-hand side can be set to 0. To show it can take an arbitrary positive value, consider shifting the offset to the right (i.e., making the target more positive). The bias of the scheme will become positive as well, but as the density of anecdotes is decreasing, the bias will be smaller than the target. In fact, the decreasing density implies that the *difference* between offset and bias must grow as the offset grows. Hence the left-hand side can take any positive value (a similar argument shows it can take any negative value as well). The following theorem formalizes this intuition and generalizes it to a broad class of noise distributions.

Theorem 4. *For any n, M_S and M_R, if $\mathbb{E}_{x\sim F}[|x-\theta|]$ is bounded then a translation invariant PBE exists.*

Proof. Fix n and $M_S - M_R$. Let's start with finding a condition that pins down the offset $r(\delta)$ of a PBE from Theorem 3. Given an offset r, denote by z 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 $(\theta, \vec{x}, y)$.

We can now calculate the expected bias δ of a targeted communication scheme with offset r:

$$\delta = r + \int_{-\infty}^{\infty} z h(z; r) dz \tag{1}$$

Write $H(r) = \int_{-\infty}^{\infty} z h(z; r) dz$. Theorem 3 now implies that to show that a PBE exists, it suffices to show that there exists a value of r such that

$$H(r) = -(M_S - M_R). \tag{2}$$

We will show (2) in two steps. First, we will show that $H(r) \to \infty$ as $r \to -\infty$ and $H(r) \to -\infty$ as $r \to \infty$. We know $\mathbb{E}_{x\sim F}[|(x-\theta)|]$ is bounded by assumption; say $\mathbb{E}_{x\sim F}[|(x -$

$\theta)|] < c_0$. Then

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

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

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

So for any $Z > 0$, we have that $H(r) \leq -Z$ for all sufficiently large r, and hence $H(r) \to -\infty$ as $r \to \infty$. A symmetric argument[5] shows that $H(r) \to \infty$ as $r \to -\infty$.

Next 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 $(\theta,\vec{x},y)$, 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 $\lim_{r\to r_0^-} H(r) < \lim_{r\to r_0^+} H(r)$.

Now we are ready to prove (2). Since $H(r) \to -\infty$ as $r \to \infty$, 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 $H(r) \to \infty$ as $r \to -\infty$.

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 $\lim_{r\to r_2^+} H(r) \leq -(M_S - M_R)$. So since H is discontinuous at r_2, we know (from our analysis of the directionality of discontinuities of H)

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

that $\lim_{r \to r_2^-} H(r) < -(M_S - M_R)$. 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 (2). □

6.2 Comparative Statics

Having pinned down the equilibria, we can now make qualitative statements about its structure. In particular, we can reason about *which* anecdote is selected by the communication scheme for aligned or extreme misaligned preferences. This then has implications for the utilities of the sender and receiver in different economic environments.

Throughout the remainder of this section, we make the additional assumption that the anecdote distribution F is strictly single-peaked and symmetric.

We first consider an environment in which the sender and receiver are completely aligned, i.e., $M_S = M_R$. In this case, Theorem 3 shows that, at equilibrium, the sender must send the closest signal to some target point, with the additional property that the expectation of what is sent is precisely equal to that target point. By symmetry, this property always holds for the unbiased signaling scheme, and hence unbiased communication is supported at equilibrium.

Corollary 1. *Assume that the anecdote distribution F is single-peaked and symmetric. If $M_S = M_R$ then there is an unbiased translation invariant PBE, with the sender always selecting the anecdote closest to $\theta_S(\vec{x}, y)$.*[6]

Proof. Since $M_S = M_R$, Theorem 3 implies that the equilibrium targeting scheme has offset r and bias δ where $r = \delta$. From (1), this means that $\int_{-\infty}^{\infty} z h(z; r) dz = 0$. By symmetry, this occurs when $r = 0$. □

In the other extreme, as the preferences of the sender and receiver diverge, the persuasion temptation of the sender will eventually overwhelm her disutility from information loss. Thus, fixing the number of anecdotes that the sender knows, a sufficiently extreme sender will then always send an extreme anecdote. A far-left sender will send the minimum anecdote; a far-right sender will send the maximum one. The following corollary formalizes this.

Corollary 2. *Assume that the anecdote distribution F is single-peaked and symmetric. Then the sender's communication scheme at any translation-invariant equilibrium converges to the minimum scheme as $M_S - M_R \to -\infty$ (in the sense that it is a targeting scheme with offset that approaches $-\infty$) and converges to the maximum scheme as $M_S - M_R \to \infty$ (in the sense that the offset approaches ∞).*

[6] It may be tempting to conjecture that if the sender and receiver are fully aligned, this unbiased scheme is also the communication scheme that both the sender and receiver prefer. After all, the calculations in Section 5 show that both sender and receiver prefer schemes that minimize the variance of the selected anecdote, and the persuasion temptation of the sender disappears under aligned preferences. However, even for single-peaked and symmetric distributions, this is not always the case: it may be possible to reduce variance by introducing bias into the communication scheme. We present such an example in Section 7.

Proof. Consider $M_S < M_R$ (the other case is analogous). For single-peaked symmetric distributions we have $r < \delta < 0$ and $\delta - r = -(M_S - M_R)$. This implies $r < (M_S - M_R)$. Hence, $r \to -\infty$ as $(M_S - M_R) \to -\infty$. □

Even when the sender sends heavily biased anecdotes, the receiver will debias them at equilibrium. However, this biasing and debiasing will tend to increase the variance of the received signal, as the sender chooses to communicate anecdotes that are further in the extreme tail of the signal distribution. The ultimate impact on utility is ambiguous and can depend on the signal distribution, as we explore in Section 6.3.

6.3 Choosing Between Senders

Our comparative statics allow us to compare the informativeness of signals received from different possible senders. This raises the natural question of what type of sender a receiver should listen to for advice. For example, a receiver could choose a very informed sender with drastically different personal preference, or a much less informed sender with whom her preferences are aligned. Which communication provides her with better information? While the more informed sender has more anecdotes to pull from and therefore is in principle better equipped to advise the receiver, she also selects anecdotes in a more biased way which reduces precision.

As it turns out, the choice of option that results in higher receiver utility depends on the distribution from which anecdotes are drawn. We'll first show that if the anecdotes are drawn from a Gaussian distribution, then it is preferable to communicate with a sufficiently informed sender, no matter how great the difference in personal preferences. Intuitively, this is because even though the more informed sender's choice of anecdote will be heavily biased, the information content is still very high after the receiver debiases the anecdote. On the other hand, this outcome is not universal: if anecdotes are drawn from a Laplace distribution, then no matter how informed the more-informed sender is, there is a level of difference in viewpoint such that the receiver would prefer to communicate with the less-informed sender.

The following proposition makes this comparison more precise.

Proposition 3. *Consider a receiver with personal preference M_R who can choose between two senders:*

a. A poorly-informed sender with access to n_0 anecdotes but who shares the same personal preference as the receiver, or

b. a well-informed expert with access to $n > n_0$ anecdotes and personal preference $M_S \neq M_R$.

If the expert's anecdotes are drawn from a normal distribution, then for any n_0 there exists n_1 such that if $n > n_1$, the receiver will always prefer to talk to the well-informed expert regardless of M_S. In contrast, if the expert's anecdotes are drawn from a Laplace distribution then for any $n_0 \geq 2$ and any choice of $n > n_0$, there exists $\Delta_0 > 0$ such that if $|M_S - M_R| > \Delta_0$ the receiver will prefer to talk to the poorly-informed sender.

Proof. To prove this theorem, we compare the receiver's loss for the following two schemes and situations: 1) Sending the minimum/maximum signal out of $n \to \infty$ signals, denoted by signal M_n 2) Sending the closest out of n_0 signals to θ, denoted by C_{n_0}. By Proposition 2 together with Corollary 2 and Corollary 1, it is enough to compare the variance of M_n to the variance of C_{n_0}.

Calculations for the Gaussian distribution. For the case of Gaussian distribution, let M_n be the minimum of n i.i.d. Gaussian variables with $\mu = 0$ and $\sigma = 1$. The Fisher–Tippett–Gnedenko (FTG) extreme value theorem Fisher and Tippett [1928], Gnedenko [1943] then implies that

$$\lim_{n\to\infty} \Pr\left[M_n \leq y\right] = \exp(-\exp(a_n(y - b_n))),$$

for $a_n = \sqrt{2\ln n}$ and $b_n = \sqrt{2\ln n} - \frac{\ln\ln n + \ln(4\pi)}{2\sqrt{2\ln n}}$ (from Lemma 1.2.1 Bovier [2005]). This implies that M_n is distributed as a Gumbel distribution with variance $O\left(\frac{\pi^2 \ln(n)}{3}\right)$ and mean $\sqrt{2\ln n} - \Theta(\frac{\ln\ln n}{\sqrt{\ln n}})$. In particular, the variance of the minimum signal tends to 0 as $n \to \infty$.

However, for any constant n_0, the variance of C_{n_0} will be some constant $\Omega(1)$. Thus, for sufficiently large n, the variance of M_n will be strictly less than the variance of C_{n_0}.

Calculations for the Laplace distribution Let M_n be the minimum of n i.i.d. Laplace variable with $\mu = 0$ and $\beta = 1$. Again using Fisher–Tippett–Gnedenko (FTG) extreme value theorem, we have that

$$\lim_{n\to\infty} \Pr\left[M_n \leq y\right] = \exp(-\exp(a_n(y - b_n))),$$

where $a_n = 1$ and $b_n = \ln \frac{n}{2}$, see Appendix C. This implies that for $n \to \infty$, M_n is distributed as a Gumbel distribution with variance $\pi^2/6$ and mean $\ln \frac{n}{2} + \gamma$, where $\gamma \approx 0.5$ is the Euler constant.[7]

Next, we calculate the variance of C_n, i.e., the signal that is closest to 0 among n i.i.d Laplace variables with $\mu = 0$ and $b = 1$. This is equivalent to the variance of $Z_n = \min_i\{z_1, \ldots, z_n\}$, where $z_i = |x_i|$ and z_i is an exponential distribution with pdf $\exp(-z)$. It is well known that the minimum of n exponential variables, Z_n, is an exponential variable with pdf $n\exp(-nz)$. Therefore, the variance of C_n

$$2\int_0^\infty n\exp(-nz)z^2 \, dz = \frac{2}{n^2}.$$

Given these calculations, we note that for any $n_0 > 1$, $Var(C_{n_0}) < Var(M_n)$ for the Laplace distribution. Therefore, the receiver will have strictly higher utility by choosing the poorly informed sender. ☐

[7] Mean and variance of the Gumbel distribution with CDF $\exp\left(-\exp\left(-\frac{x-\mu}{\beta}\right)\right)$ are $\mu + \beta\gamma$ and $\pi^2\beta^2/6$, respectively.

7 Commitment

So far we have assumed that the sender has no ability to 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 select 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 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 have formally defined the commitment equilibrium in Section 3. Optimal commitment is weakly advantageous to the sender, given that it is a relaxation of the perfect Bayesian equilibrium concept. In this section, we consider the implications of commitment on both the sender and receiver.

As before, Proposition 2 implies that if the sender uses a translation invariant communication scheme π, then the receiver's best response α_π will be a translation. Hence, we can again focus on translation-invariant commitment equilibria. The following theorem characterizes the set of such equilibria.

Theorem 5. *The sender's optimal commitment π_S is a signaling scheme that minimizes the variance of the signal sent. That is,*

$$\pi_S \in \text{argmin}_\pi \mathop{\mathbb{E}}_{\theta,\vec{x},y}[(\pi(\vec{x},y) - \mathbb{E}[\pi(x,y)])^2]$$

Moreover, this minimizes the receiver's loss.

Proof. For any committed communication scheme π_S, the receiver's response is a best response α_{π_S}. For translation invariant schemes, the best response α_{π_S} is responding to $\pi_R = \pi_S$. Applying Proposition 1 with $\pi_R = \pi_S$, we have that the sender's expected loss from any given translation-invariant communication scheme π_S is

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

Minimizing this loss therefore corresponds to choosing π_S that minimizes $\mathbb{E}[(\pi_S(\vec{x},y) - (\theta_S(\vec{x},y) + \beta(\pi_S)))^2]$. But as discussed immediately following Proposition 1, this is precisely the variance of the communicated anecdote, as claimed. Moreover, by Proposition 2 this is also the receiver's loss. Therefore, the optimal commitment communication also minimizes the receiver cost. □

Theorem 5 implies that under commitment the sender behaves as if her personal preference is aligned with the receiver's personal preference. Commitment removes the sender's persuasion temptation: she therefore will implement the socially optimal communication subject to the constraint that she can only send a single anecdote.

7.1 Single-peaked and Symmetric Anecdote Distributions

In this section we describe the variance-minimizing signaling scheme for strictly single-peaked and symmetric anecdote distributions. Intuitively, the sender will choose anecdotes where the anecdote distribution has maximal density. For single-peaked distributions this will be exactly the neighborhood around the state of the world θ. This suggests that the sender should want to send signals that are close to the posterior mean θ_S, and hence will want to select a communication scheme with low bias. Theorem 6 below formalizes this insight 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 noise 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, $|g(x)| \leq c_1|x| + c_2$* [8]
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 6. *For any well-behaved signal distribution, 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 signaling 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 do 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".

7.1.1 Overview of Theorem 6

For the proof of Theorem 6 we bound the losses from the biased and unbiased signaling schemes and show that the unbiased scheme dominates.

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

Proposition 4. *Given any well-behaved noise distribution F, the unbiased targeting communication scheme with $\delta = r(\delta) = 0$ which selects the closest signal to the sender's posterior mean $\theta_S(\vec{x}, y)$ has signaling loss:*

$$\frac{1}{2n^2 f(0)^2} + o\left(\frac{1}{n^2}\right) \tag{3}$$

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

$$\frac{1}{2n^2 f(\delta(\pi))^2} + o\left(\frac{1}{n^2}\right) \tag{4}$$

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

In the rest of the section we provide an overview of the proof of Proposition 4 . The full proofs are relegated to Appendix D.

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, $\mathbb{E}_{\theta,\vec{x},y}[(\pi(\vec{x}, y) - \theta_S(\vec{x}, y) - \beta(\pi))]$, of any translation invariant signaling scheme with bias $\beta(\pi) = \delta$ is at least $\mathbb{E}_{\theta,\vec{x},y}[X_\delta^2]$.

Optimal unbiased scheme. Since the bias of the scheme that sends signal closest to the posterior mean is itself 0, this is the optimal amongst all unbiased signaling 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 will see in Section 7.2, 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 $\theta_S(\vec{x})$. 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 $\mathbb{E}[X_0^2]$ 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 $\theta_S(\vec{x})$, 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 $\theta_S(\vec{x})$; 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 each[9]. See Figure 1.

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 $\theta_S(\vec{x}) \in C$. We formally show this in Lemmas 11 and 12 [10].

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 Appendix D for the proof, and Lemma 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 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 $\theta_S(\vec{x})$ 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 4.

7.2 Discussion

Unbiased communication schemes are not necessarily optimal if the distribution is not single-peaked and symmetric or with small n in the absence of foresight.

The following example demonstrates that a biased signaling scheme can be optimal even with single-peaked and symmetric distributions when n is small.

Proposition 5. *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 signaling scheme with non-zero bias.*

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

[10] For these lemmas we assume that the density function θ_S is sufficiently "nice" in C. Refer to Section D.2 for details about this assumption, and why we can make this assumption without loss of generality.

Figure 1: Intuition for proof of theorem 6

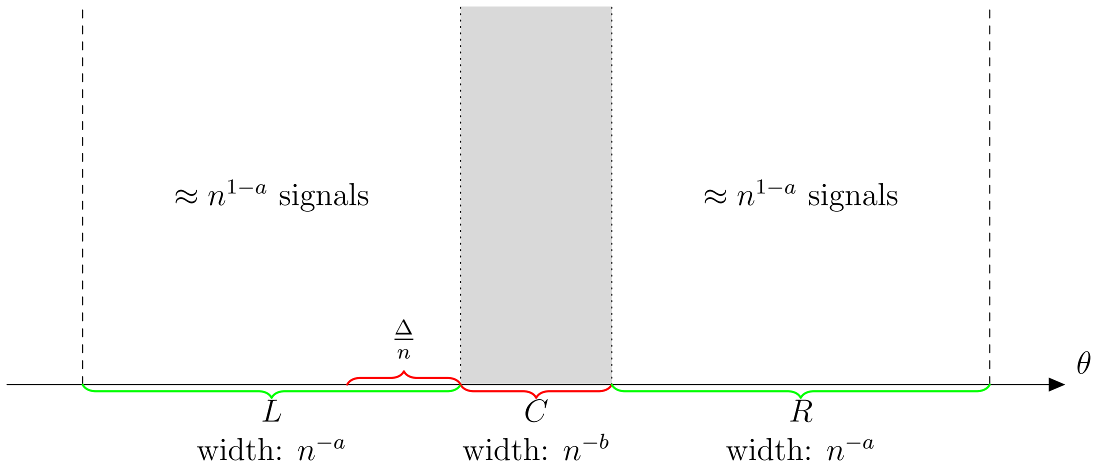

We assume $a < 1$ and $1 < b < 2a - \frac{1}{2}$ which ensures that the intervals L and R contain many signals but the collective influence of these signals on the posterior mean is $O(n^{\frac{1}{2}-2a})$ and hence smaller than $\frac{1}{n}$. It also implies that if we consider a posterior that is contained 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.

The idea behind Proposition 5 is that, conditional on the value of the sender's posterior mean, the conditional density over signal realizations is not necessarily single-peaked. For the setting of Proposition 5, the optimal unbiased scheme is precisely the one that always returns the middle signal. However, for uniform distributions, 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 that sometimes returns the minimum signal (or, by symmetry, sometimes returns the maximum signal). The full proof appears in Appendix B.

Why does bias help? Recall that there is intrinsic error in the sender's posterior mean. This variance is unavoidable. But it can introduce correlation with particular samples. This correlation can be used to help minimize variance between the posterior mean and the signal 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 incentive for the sender to systematically bias their communication.

References

S Nageeb Ali, Greg Lewis, and Shoshana Vasserman. Voluntary disclosure and personalized pricing. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 537–538, 2020.

Anton Bovier. Extreme values of random processes. *Lecture Notes Technische Universität Berlin*, 2005.

Ying Chen. Perturbed communication games with honest senders and naive receivers. *Journal of Economic Theory*, 146(2):401–424, 2011.

Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

Alfredo Di Tillio, Marco Ottaviani, and Peter Norman Sørensen. Strategic sample selection. *Econometrica*, 89(2):911–953, 2021.

Ronald A Dye. Disclosure of nonproprietary information. *Journal of accounting research*, pages 123–145, 1985.

Wioletta Dziuda. Strategic argumentation. *Journal of Economic Theory*, 146(4):1362–1397, 2011.

R. A. Fisher and L. H. C. Tippett. Limiting forms of the Frequency distribution of the largest or smallest member of a sample. *Mathematical Proceedings of the Cambridge Philosophical Society*, 24(2):180–190, 1928. doi: 10.1017/S0305004100015681.

Michael J Fishman and Kathleen M Hagerty. The optimal amount of discretion to allow in disclosure. *The Quarterly Journal of Economics*, 105(2):427–444, 1990.

Jacob Glazer and Ariel Rubinstein. On optimal rules of persuasion. *Econometrica*, 72(6): 1715–1736, 2004.

Jacob Glazer and Ariel Rubinstein. A study in the pragmatics of persuasion: a game theoretical approach. *New perspectives on games and interaction*, pages 121–140, 2008.

B. Gnedenko. Sur la distribution limite du terme maximum d'une serie aleatoire. *Annals of Mathematics*, 44(3):pp. 423–453, 1943. ISSN 0003486X. doi: 10.2307/1968974. URL http://www.jstor.org/stable/1968974.

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

Sanford J Grossman and Oliver D Hart. Disclosure laws and takeover bids. *The Journal of Finance*, 35(2):323–334, 1980.

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

Boyan Jovanovic. Truthful disclosure of information. *The Bell Journal of Economics*, pages 36–44, 1982.

Woon-Oh Jung and Young K Kwon. Disclosure when the market is unsure of information endowment of managers. *Journal of Accounting research*, pages 146–153, 1988.

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

Giorgio Martini. Multidimensional disclosure. 2018.

Paul Milgrom and John Roberts. Relying on the information of interested parties. *The RAND Journal of Economics*, pages 18–32, 1986.

Paul R Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.

Daniel J Seidmann and Eyal Winter. Strategic information transmission with verifiable messages. *Econometrica: Journal of the Econometric Society*, pages 163–169, 1997.

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

Robert E Verrecchia. Essays on disclosure. *Journal of accounting and economics*, 32(1-3): 97–180, 2001.

Asher Wolinsky. Information transmission when the sender's preferences are uncertain. *Games and Economic Behavior*, 42(2):319–326, 2003.

A Diffuse Prior

In this section we discuss our assumption of a diffuse prior and its implications on the posterior of the agents.

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 defined[11].

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

Claim 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 signaling scheme π, and a signal $x = \pi(\vec{x}, y)$. The posterior distribution of the receiver conditioned on π, x, is a proper distribution. Moreover, the posterior mean $\theta_R(\pi, x) = x - \beta(\pi)$ is an unbiased estimator of θ.

We prove the following claim below.

Sender's posterior distribution. Recall that the anecdotes $x_1, \ldots, x_n$ are drawn independently from $\theta + F$. Thus the pdf of an anecdote given θ is $f(x - \theta)$. Similarly, y is drawn from $\gamma + G$ 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 $\gamma = y - \hat{\theta}$. 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 noise 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 in an unbiased estimator of θ.

We will now see 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.

B Three Uniform Signals

In this section we prove Proposition 5, which is that the optimal signaling scheme may be biased even for $n = 3$ signals drawn from a uniform distribution.

Suppose that F is the uniform distribution on $[-1/2, 1/2]$. For notational convenience, it will be convenient to assume (as the analyst) that $\theta = 1/2$, so that all samples are drawn from $[0, 1]$. Note that, by translation invariance and the diffuse prior on θ, this assumption is without loss of generality. We also assume that we are in the setting with no foresight, where signal y reveals no additional information to the sender, so we will tend to drop signal y from the notation.

Given 3 realized signals $x_1 \leq x_2 \leq x_3$, the sender's posterior is a uniform distribution on $[x_3 - 1/2, x_1 + 1/2]$. The posterior mean is therefore $\theta_S(\vec{x}) = \frac{x_1 + x_3}{2}$. In what follows we will think of x_1, x_2, x_3 as random variables corresponding to the least, middle, and largest signal, respectively.

Optimal Unbiased Scheme Consider the optimal unbiased scheme, call it π_0. By Theorem 2, this 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 signaling 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 Scheme We'll now build a 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 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.

C FTG Theorem for Laplace

We first show that the minimum of n variables drawn from a Laplace distribution is in the domain of Gumbel extreme type. From Theorem 1.2.9 of Bovier [2005], we see that it is enough to find a $g(t) > 0$ such that,

$$\lim_{t\to\infty} \frac{1 - F(t + xg(t))}{1 - F(t)} = e^{-x}$$

Recall that for the Laplace distribution with $\mu = 0$ and $\beta = 1$, $F(x) = 1 - \frac{1}{2}e^{-x}$ for $x > 0$. Hence by setting $g(\cdot) \equiv 1$ we get,

$$\lim_{t\to\infty} \frac{1 - F(t + xg(t))}{1 - F(t)} = \lim_{t\to\infty} \frac{e^{-t-xg(t)}}{e^{-t}} = e^{-x}$$

Moreover, to find the required a_n, b_n for the extreme value theorem for Gumbel, it is enough to get γ_n such that,

$$n(1 - F(\gamma_n + xg(\gamma_n))) = e^{-x}$$

Thus, by setting $\gamma = \ln \frac{n}{2}$ we get,

$$n(1 - F(\gamma_n + xg(\gamma_n))) = n \cdot \frac{1}{2} e^{-\ln(n/2)} \cdot e^{-x} = e^{-x},$$

because $F(x) = \frac{1}{2}e^{-x}$ and $g \equiv 1$.

Finally, we set $a_n = 1/g(\gamma_n)$ and $b_n = \gamma_n$, to get $a_n = 1$ and $b_n = \ln \frac{n}{2}$.

D Proof of Proposition 4

We start by proving that the that the signaling loss of the unbiased 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 Appendix D.3 we bound the signaling loss of a biased scheme.

Let $I = [-n^{-\frac{1}{2}+\varepsilon}, n^{-\frac{1}{2}+\varepsilon}] + \theta$. 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 $C \in P$, 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 8), and
3. there are *sufficiently* many signals in $N(C) = L \cup R$.

In Section D.1, we bound the signaling loss contributed by this high probability event. Further, in Section D.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 4), when C is a weak interval (in Lemma 5), or when there are very few signals in $N(C)$ (in Lemma 6).

With this we are ready to bound the signaling loss. Recall that the signaling loss of the unbiased 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,

$$\begin{aligned}
\mathbb{E}[X_0^2] &= \mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\notin I\}] + \sum_{C\in W_I}\mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\in C\}] + \sum_{C\in\mathcal{P}\setminus W_I}\mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\in C\}] \\
&= \sum_{C\in\mathcal{P}\setminus W_I}\mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\in C,K_0\}] \qquad (5)\\
&\quad \underbrace{+\mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\notin I\}] + \sum_{C\in W_I}\mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\in C\}] + \sum_{C\in\mathcal{P}\setminus W_I}\mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\in C,\overline{K_0}\}]}_{\text{rare events}} \qquad (6)\\
&\le \frac{1}{2n^2f(0)^2} + o(1/n^2)
\end{aligned}$$

This is because by Lemma 3 the term in Eq. (5) is $\sum_{C\in\mathcal{P}\setminus W_I}\mathbb{E}[X_0^2\cdot\mathbf{1}\{\theta_S\in C,K_0\}]\le\frac{1}{2n^2f(0)^2}+o(1/n^2)$, and from Lemmas 4, 5, and 6 we see that all terms in Eq. (6) contribute at most $o(1/n^2)$. Thus giving us the required bound on the signaling loss.

D.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 8. *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 $\tau(\overline{\theta})\le c_2n^{1+1/22}e^{-n^{1/22\alpha}}$ for all $\overline{\theta}\in C$. Let $W_I\subset\mathcal{P}$ be the set of all such intervals C.*

By Claim 8 we will see that the probability that $\theta_S(\vec{x},y)\in C$ for some *weak interval* C is negligible $O(n^{-4logn+1})$. Moreover, if C is not *weak*, that is, $\tau(\overline{\theta})\ge c'n^{1/22}n^{-4\log n+1}$ then we get that $\tau(\overline{\theta}')=\tau(\overline{\theta})(1+O(n^{-\frac{1}{22}}))$ for all $\overline{\theta},\overline{\theta}'\in C$ by Claim 7.

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 2. *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,

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

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 $f(x) = f(0)(1 - O(n^{-1/2+\epsilon}))$. ◻

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 3. *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 $0 < \varepsilon < 1$, we have*

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

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 2 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

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

◻

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 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 3 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 $n^{-1/2+\epsilon+a}M$. Therefore by union bound, we get $\Pr[\exists J \in \mathcal{J} : |Y(J) - k_m| \geq \varepsilon_1 k_m] \leq (n^{-1/2+\epsilon+a}M) \cdot \exp\left(-\frac{\varepsilon_1^2 f(0) n^{1-a}}{3M}\right)$ which is at most $O(\exp^{-n^{1/20}})$, for $\varepsilon_1 = O(n^{-1/22})$ and $M = n^{1/100}$. ◻

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

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 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 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 2)
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 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 7)

Thus we get,

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

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 $1 - x \le e^{-x}$ we get,

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

□

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

Lemma 3. *Fix any $C \in \mathcal{P} \setminus W_I$. Then we have,* $\mathbb{E}[X_0^2 \cdot \mathbf{1}(K_0)|\theta_S \in C] \le \frac{1}{2n^2 f(0)^2} + o(1/n^2)$.

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

□

In Section D.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.

D.2 Contribution of Rare Events

In this section we bound the loss from the rare events from Eq. (6).

Lemma 4. $\mathbb{E}[X_0^2 \cdot \mathbf{1}\{\theta_S \notin I\}] \le O\left(\exp(-\frac{n^{2\epsilon}A}{2})\right)$ *for some constant $A > 0$.*

Proof. Recall that $\theta_S(\vec{x})$ is the MMSE estimator and $\theta_S - \theta^* \to^d \mathcal{N}(0, C_{\mathcal{I}}/n)$. 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,

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

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

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

Claim 4. *Let T_Q be the event that all $|x_i| > Q$. Then $\mathbb{E}[X_0^2 \mathbf{1}(T_Q \wedge \theta_S \in I)] \le o(1/n^2)$.*

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 $X_0 = \min_i |x_i - \theta_S| \le |x_i| + n^{-1/2+\epsilon}$ for all x_i. Let $t(x) = (|x| + n^{-1/2+\epsilon})^2$. Hence we have,

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

□

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

Lemma 5. *Then* $\sum_{C \in W_I} \mathbb{E}[X_0^2 \mathbf{1}\{\theta_S \in C\}] \leq O\left(Q^2 n^{-4\log n+1}\right) + o(1/n^2)$.

Proof. Let $\overline{T}_Q$ 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,

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

where the last inequality follows from Claim 8 (proved in section D.4) that $\Pr[\theta_S \in W_I] \leq O\left(n^{-4\log n+1}\right)$.

Moreover, by Claim 4 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 6. *Let* $k_0 = f(0)n^{1-a}$, $\varepsilon_1 = O(n^{-1/22})$. *Let* $\overline{K_0}$ *be the event such that* $Y(C) \neq k_0(1 \pm \varepsilon_1)$. *Then,* $\sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_0^2 \cdot \mathbf{1}\{\overline{K_0}$ *and* $\theta_S \in C\}] \leq (Q + n^{-1/2})^2 O\left(\exp\left(-n^{1/20} f(0)\right)\right) + O(\exp^{-Qn}) \leq o(1/n^2)$.

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 4, 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 $\overline{T}_Q$), 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 $\Pr[A_0 \cap \overline{T}_Q \cap \theta_S \in C] \leq \Pr[A_0] \leq \exp\left(-\frac{n^{-1/11} f(0) n^{1-a}}{3M}\right)$.

$$\begin{aligned}\mathbb{E}[X_0^2 \cdot \mathbf{1}\{\overline{K_0} \text{ and } \theta_S \in C\}] &\leq \mathbb{E}[X_0^2 \cdot \mathbf{1}\{A_0 \text{ and } \theta_S \in C\} \cdot \mathbf{1}\{T_Q\}] + \mathbb{E}[X_0^2 \cdot \mathbf{1}\{A_0 \text{ and } \theta_S \in C\} \cdot \mathbf{1}\{\overline{T}_Q\}] \\ &\leq O(\exp^{-Qn}) + (Q + n^{-1/2+\epsilon})^2 \exp\left(-\frac{n^{-1/11} f(0) n^{1-a}}{3M}\right) \\ &\leq o(1/n^2)\end{aligned}$$

□

D.3 Loss of Biased Schemes

In this section we 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 $f(\delta) \geq n^{-1/100}$ [12].

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

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 $C \in P$ 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 8), 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 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,

$$\begin{aligned}
\mathbb{E}[X_\delta^2] &\geq \sum_{C \in \mathcal{P} \setminus W_I} \mathbb{E}[X_\delta^2 \cdot \mathbf{1}\{\theta_S \in C, K_\delta\}] \\
&\geq \sum_{C \in \mathcal{P} \setminus W_I} \Pr[\{\theta_S \in C, K_\delta\}] \cdot \underbrace{\mathbb{E}[X_\delta^2 \,|\theta_S \in C, K_\delta]}_{\text{conditional expectation}} \\
&\geq \frac{1}{2n^2 f(\delta)^2} - o(1/n^2)
\end{aligned} \tag{7}$$

This is because by Lemma 9 the conditional expectation term in Eq. (7) is $\mathbb{E}[X_\delta^2 \,|\theta_S \in C, K] \geq \frac{1}{2n^2 f(\delta)^2}(1 - o(1))$. We note that $\{\vec{x} : \forall J_\delta \in \mathcal{J}_\delta \text{ with suffiencetly few signals}\} \subset K_\delta$ for all $C \in \mathcal{P}$, and by Lemma 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 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 2.

Claim 5. *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 $|x| \le 2(\log n)^2 + n^{-1/2+\epsilon}$. By mean value theorem we have,

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

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

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

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

Claim 6. *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*

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

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 5 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

$$\Pr[Y(A) \ge (1+\varepsilon) \, \mathbb{E}[Y(A)]] \le \exp\left(-\frac{\varepsilon_1^2 f(\delta) n \ell (1 + O(n^{-1/4+\epsilon}))}{3}\right).$$

□

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

Lemma 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 $(1+\varepsilon)k_m$ signals (or less than $(1-\varepsilon)k_m$ is $O(\exp\{\varepsilon^2 f(\delta) n^{1-a}/3M\})$.*

Corollary 4. *Let $k_\delta = f(\delta) n^{1-a}$, let $k^* = (1+\varepsilon_1)k_\delta$, and $k' = (1-\varepsilon_1)k_\delta(1 - 2/M)$. Let $Y(N_\delta)$ be the number signals in interval $N(C_\delta)$ (of length n^{-a}).* $\Pr[(Y(N_\delta) \notin [k', k^*]] \le \exp\left(-\frac{\varepsilon_1^2 f(\delta) n^{1-a}}{3M}\right)$.

Let K_δ denote the event that there are at most $(1+\varepsilon_1)k_\delta$ and at least $(1-\varepsilon_1)k_\delta(1-2/M)$ many signals in $N(C_\delta)$, and there are no signals in C_δ.

Lemma 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 $d < n^{-9/10}$, $\Pr[X_{(\delta)} > d\ |K_\delta, \theta_S \in C] \geq \exp[-\lambda d]\frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta]}{\Pr[\theta_S \in C]}$.*

Proof. Let $k^* = (1+\varepsilon_1)k_\delta$, and $k' = (1-\varepsilon_1)k_\delta(1-2/M)$. For all $k' \leq k \leq k^*$, define $K_\delta(k)$ to be event where $\#x_i \in N(C_\delta) = k$, and there are no signals in C_δ .

For all $d < n^{-a} + n^{-b}$, let E_d denote the event that $\forall x \in L, x \notin B_d$, where $B(d)$ is the interval of length $2d$ centered around $\theta_S + \delta$.

$$
\begin{aligned}
&\Pr[X_{(\delta)} \cdot \mathbf{1}(K_\delta) > d|\theta_S \in C] \\
&= \Pr[\forall x_i \notin B_d, \exists x_1, x_2, \ldots, x_{k'} \in L \text{ and } K_\delta|\theta_S \in C] \\
&= \sum_{k=k'}^{k^*} \Pr[\forall x_i \notin B_d, \exists x_1, x_2, \ldots, x_k \in L \text{ and } K_\delta(k)|\theta_S \in C] \\
&= \sum_{k=k'}^{k^*} \Pr[\exists x_1, x_2, \ldots, x_k \in L \setminus B_d \text{ and } K_\delta(k)|\theta_S \in C] \\
&= \sum_{k=k'}^{k^*} \Pr[K_\delta(k)|\theta_S \in C] \cdot \Pr[\exists x_1, x_2, \ldots, x_k \in L \setminus B_d\ |K_\delta(k), \theta_S \in C] \\
&= \sum_{k=k'}^{k^*} \Pr[K_\delta(k)|\theta_S \in C] \int_{\vec{x}_{[k]} \in L \setminus B_d} f(\vec{x}_{[k]}\ |K_\delta(k), \theta_S \in C) \int_{\tilde{x} \in A(\vec{x}_{[k]}) \setminus (L \cup C)} f(\tilde{x}\ |\vec{x}_{[k]}, \theta_S \in C, K_\delta(k)) d\vec{x} \\
&= \sum_{k=k'}^{k^*} \Pr[K_\delta(k)|\theta_S \in C] \int_{\vec{x}_{[k]} \in L \setminus B_d} f(\vec{x}_{[k]}\ |K_\delta(k), \theta_S \in C) d\vec{x}_{[k]} \\
&= \sum_{k=k'}^{k^*} \Pr[K_\delta(k)|\theta_S \in C] \int_{\vec{x}_{[k]} \in L \setminus B_d} \left(f(\vec{x}_{[k]}|K_\delta(k)) \cdot \frac{\Pr[\theta_S \in C\ |\vec{x}_{[k]}, K_\delta(k)]}{\Pr[\theta_S \in C\ |K_\delta(k)]} \right) d\vec{x}_{[k]} \qquad \text{(Bayes rule)}
\end{aligned}
$$

We will use Lemma 12 and the fact that $\tau(\bar{\theta})$ is approximately constant in $C \notin W_I$.

$$
\begin{aligned}
&\geq \sum_{k=k'}^{k^*} \Pr[K_\delta(k)|\theta_S \in C] \cdot \frac{\Pr[\theta_S \in C \setminus E\ |K_\delta(k)]}{\Pr[\theta_S \in C|K_\delta(k)]} \int_{\vec{x}_{[k]} \in L \setminus B_d} f(\vec{x}_{[k]}|K_\delta(k)) d\vec{x}_{[k]} \qquad \text{(By Lemma 12)} \\
&= \sum_{k=k'}^{k^*} \frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta(k)]}{\Pr[\theta_S \in C]} \Pr[\vec{x}_{[k]} \in L \setminus B_d|K_\delta(k)] \\
&\geq \sum_{k=k'}^{k^*} \frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta(k)]}{\Pr[\theta_S \in C]} \cdot \left(1 - \frac{2d}{n^{-a}}(1 + O(1/n^{1/4}))\right)^k \qquad \text{(By Claim 5)} \\
&\geq \left(1 - \frac{2d}{n^{-a}}(1 + O(1/n^{1/4}))\right)^{k^*} \frac{\Pr[\theta_S \in C \setminus E \text{ and } K_\delta]}{\Pr[\theta_S \in C]}
\end{aligned}
$$

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 $\varepsilon_3 = O(n^{-1/20})$.

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

□

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

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 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 $\mathbb{E}[X_\delta^2] \geq c_5(\frac{n^{1/100}}{n^2}) >> \frac{1}{2n^2 f(0)^2} \geq \mathbb{E}[X_0^2]$ for a constant $c_5 > 0$.

□

Proof of Proposition 4 (b). By lemma 9, for a sufficiently large n and any $\delta < 2(\log n)^2$ and $f(\delta) > n^{-1/100}$, we have $\mathbb{E}[X_\delta^2] \geq \sum_{C \in \mathcal{P} \setminus W_I} \Pr[\theta_S \in C]\,\mathbb{E}[X_\delta^2 | \theta_S \in C] \geq \sum_{C \in \mathcal{P} \setminus W_I} \frac{1}{2n^2 f(\delta)^2}(1 - o(1))\Pr[\theta_S \in C \setminus E$ 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 $(1 \pm \varepsilon)k_\delta$ many signals in $N(C_\delta)$ and there are no signals in C_δ, that $\Pr[\theta_S \in C \setminus E] \geq \Pr[\theta_S \in C \setminus 2E | Y(N(C_\delta) \cup C_\delta) \in (1 \pm \varepsilon)k_\delta]$, 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 $\Pr[K_\delta] = \Pr[Y(N(C_\delta) \cup C_\delta) \in (1 \pm \varepsilon)k_\delta] \cdot \Pr[Y(C_\delta) = 0 \ | Y(N(C_\delta) \cup C_\delta) \in (1 \pm \varepsilon)k_\delta]$. 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,

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

Let $\mathcal{G}_\delta$ denote the event that all $J_\delta \in \mathcal{J}_\delta$ has $(1\pm\varepsilon)k_\delta/M$ 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 8 (proved in section D.4) we have $\Pr[\theta_S \in W_I] \leq O(n^{-4logn+1})$. By Lemma 7 we have $\Pr[\overline{\mathcal{G}}_\delta] \leq O(\exp\{\varepsilon^2 f(\delta)n^{1-a}/3M\})$. Hence we get $\mathbb{E}[X_\delta^2] \geq \frac{1}{2n^2f(\delta)^2}(1-o(1))$. □

D.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 4.

In Lemma 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 10. *For any signals $\vec{x}$ observed by the sender we have,*

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

where $Var_{\theta\sim D_S(\vec{x})}[\theta]$ 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

$$\theta_S(\vec{x}) = \frac{\int_\theta \theta \prod_i f(x_i - \theta) d\theta}{\int_\theta \prod_i f(x_i - \theta) d\theta} \tag{8}$$

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

$$\begin{aligned}
\frac{\partial \theta_S(\vec{x})}{\partial x_i} &= \frac{\int \theta f'(x_i - \theta) \prod_{j \neq i} f(x_j - \theta) d\theta \int \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2} \\
&\quad - \frac{\int \theta \prod_j f(x_j - \theta) d\theta \int f'(x_i - \theta) \prod_{j \neq i} f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2} \\
&= \frac{\int \theta g(x_i - \theta) \prod_j f(x_j - \theta) d\theta \int \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2} \\
&\quad - \frac{\int \theta \prod_j f(x_j - \theta) d\theta \int g(x_i - \theta) \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2}
\end{aligned} \tag{9}$$

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:

$$\begin{aligned}
\frac{\partial \theta_S(\vec{x})}{\partial x_i} &= \frac{\int \theta \left(g(x_i) - g'(\tilde{x}_i)\theta\right) \prod_j f(x_j - \theta) d\theta \int \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2} \\
&\quad - \frac{\int \theta \prod_j f(x_j - \theta) d\theta \int \left(g(x_i) - g'(\tilde{x}_i)\theta\right) \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2} \\
&= -\frac{\int \theta g'(\tilde{x}_i)\theta \prod_j f(x_j - \theta) d\theta \int \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2} \\
&\quad + \frac{\int \theta \prod_j f(x_j - \theta) d\theta \int g'(\tilde{x}_i)\theta \prod_j f(x_j - \theta) d\theta}{\left(\int \prod_j f(x_j - \theta) d\theta\right)^2}
\end{aligned} \tag{10}$$

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

$$\left|\frac{\partial\theta_S(\vec{x})}{\partial x_i}\right| \leq c_1\left(\frac{\int_\theta \theta^2 \prod_j f(x_j-\theta)d\theta}{\int_\theta \prod_j f(x_j-\theta)d\theta} + \left(\frac{\int_\theta \theta \prod_j f(x_j-\theta)d\theta}{\int_\theta \prod_j f(x_j-\theta)d\theta}\right)^2\right) \tag{11}$$

$$= c_1\left(\mathop{\mathbb{E}}_{\theta\sim D_S(\vec{x})}[\theta^2|\vec{x}] + \mathop{\mathbb{E}}_{\theta\sim D_S(\vec{x})}[\theta|\vec{x}]^2\right)$$

$$= c_1\left(Var_{\theta\sim D_S(\vec{x})}[\theta] + 2\mathop{\mathbb{E}}_{\theta\sim D_S(\vec{x})}[\theta|\vec{x}]^2\right) \tag{12}$$

□

Next we bound the shift in θ_S when rearranging k signals in L (or R) using Lemma 10. We note that $Var_{\theta\sim D_S(\vec{x})}[\theta] = O(1/n)$ 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 $\theta_S : \mathbb{R}^n \to \mathbb{R}$. 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,

$$\begin{aligned}|\theta_S(\vec{y}) - \theta_S(\vec{x})| &= |\nabla\theta_S(\vec{z}) \cdot (\vec{y} - \vec{x})| \\ &\leq \sum_{i=1}^{k} c_1\left(Var_{\theta\sim D_S(\vec{x})}[\theta] + 2\theta_S(\vec{x})^2\right)|y_i - x_i| \quad \text{(By Lemma 10)} \\ &= c_1 O(n^{-1+2\epsilon})\sum_{i=1}^{k}\ell\end{aligned}$$

Since $x_i, y_i \in A$ we bound $|x_i - y_i| \leq \ell$. Further, by bounding $Var_{\theta\sim D_S(\vec{x})}[\theta]$ 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 7. *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 $\overline{\theta} \in C$ we have $\tau(\overline{\theta} + \epsilon') = \tau(\overline{\theta})\left(1 + O(n^{-1/22})\right) + O(e^{-n^{1/22\alpha}})$ for all $0 < \epsilon' \leq 1/n^b$.*

Proof. Fix a posterior mean $\overline{\theta} \in [-n^{-\frac{1}{2}+\epsilon}, n^{-\frac{1}{2}+\epsilon}]$ and consider all the signal draws $X(\overline{\theta}) = \left\{\vec{x} | \overline{\theta}(\vec{x}) = \overline{\theta}\right\}$ that generate this posterior mean. We know that $\tau(\overline{\theta}) = \int_{\vec{x} \in X(\overline{\theta})} \prod_{x_i \in \vec{x}} f(x_i) dx$. Now consider $\overline{\theta} + \epsilon'$. We can couple all the signal realizations in $X(\overline{\theta} + \epsilon')$ and $X(\overline{\theta})$ 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, $\tau(\overline{\theta}+\epsilon') = \int_{\vec{x} \in X(\overline{\theta})} \prod_{x_i \in \vec{x}} f(x_i+\epsilon') dx$.

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

$$\prod_{x_i \in \vec{x}} f(x_i + \epsilon') = \prod_{x_i \in \vec{x}} [f(x_i) + f'(\tilde{x}_i)\epsilon'] \tag{13}$$

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

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

  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 $\Pr_{\vec{x}}[\exists x_i \text{ s.t. } |x_i| > 4(\log n)^2] \leq n e^{-4(\log n)^2} = (n^{-4\log n+1})$.

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

Hence we get

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

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

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

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

Claim 8. $\Pr[\theta_S \in W_I] \leq O(n^{-4\log n+1})$.

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,

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

□

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 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 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,

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

□