Learning and Persuading with Anecdotes

Nika Haghtalab, UC Berkeley

Communication and Opinion Formation

Our actions are informed by complex statistical beliefs about the world:

- Vote for politician X?
  - →What policies would she support? Has she been ethical? Is she trustworthy? Is she bipartisan?
- Get the COVID vaccine today?
  - → Is it effective? Is it safe? Am I high priority enough?

Our information comes from varying degrees of complexity:

- High quality detailed statistical analysis: scientific papers, investigative journalism, survey many pieces of information.
- Retelling of experiences: Tweets, FB, most news pieces, about a single activity or view of the whole

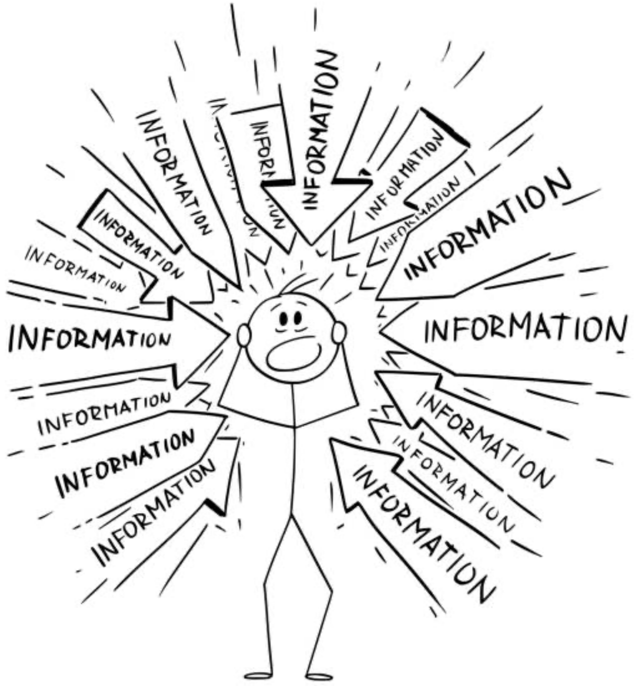

Generalization VS Communication

Account for the difference in generalization and communication:

- Generalization: Beliefs learned from many pieces of information
- Communication: Stories we retell to justify our beliefs or persuade others

Machine learning models use different abstractions

→Information and communication: Individual samples or data pieces

→Beliefs:

- Complex functions that describe your actions in any one scenario.
- Posterior distributions that describe your belief about what led to the state of the world.

Not claiming that machine learning and human learning are the same!

Learning and Persuading with Anecdotes

Joint work with

Nicole Immorlica
Microsoft Research

Brendan Lucier
Microsoft Research

Markus Mobius
Microsoft Research

Divyarthi Mohan
Princeton University

Learning and Persuading with Anecdotes

Consider an environment where communication is retelling of an anecdote.

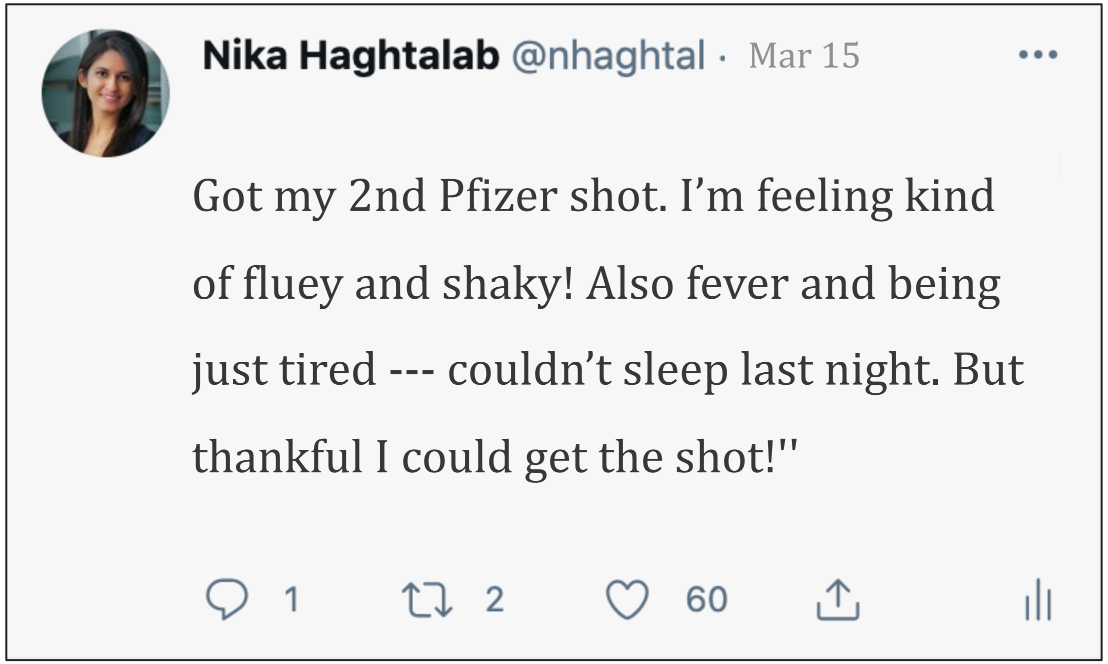

Nika Haghtalab @nhaghtal · Mar 15

Got my 2nd Pfizer shot. I’m feeling kind of fluey and shaky! Also fever and being just tired --- couldn’t sleep last night. But thankful I could get the shot!''

Anecdote: One person’s account

… More than 3% of recipients of Moderna COVID-19 vaccine develop severe temporary side effects including fatigue (9.7%), myalgia (8.9%), arthralgia (5.2%), headache (4.5%), pain (4.1%) …

published scientific study

Summarized statistics of many accounts

Anecdote: A person shares one of k actual observations. Can’t make up stories.
Persuasion: Share the anecdote that gets your listeners to take actions you like.

Questions

What do rational communication and learning with anecdotes look like?

Communicating anecdotes is less efficient.

Does restriction to anecdote introduce **bias in communication**?

Does restriction to anecdotes contribute to **belief polarization**?

Model – Actions

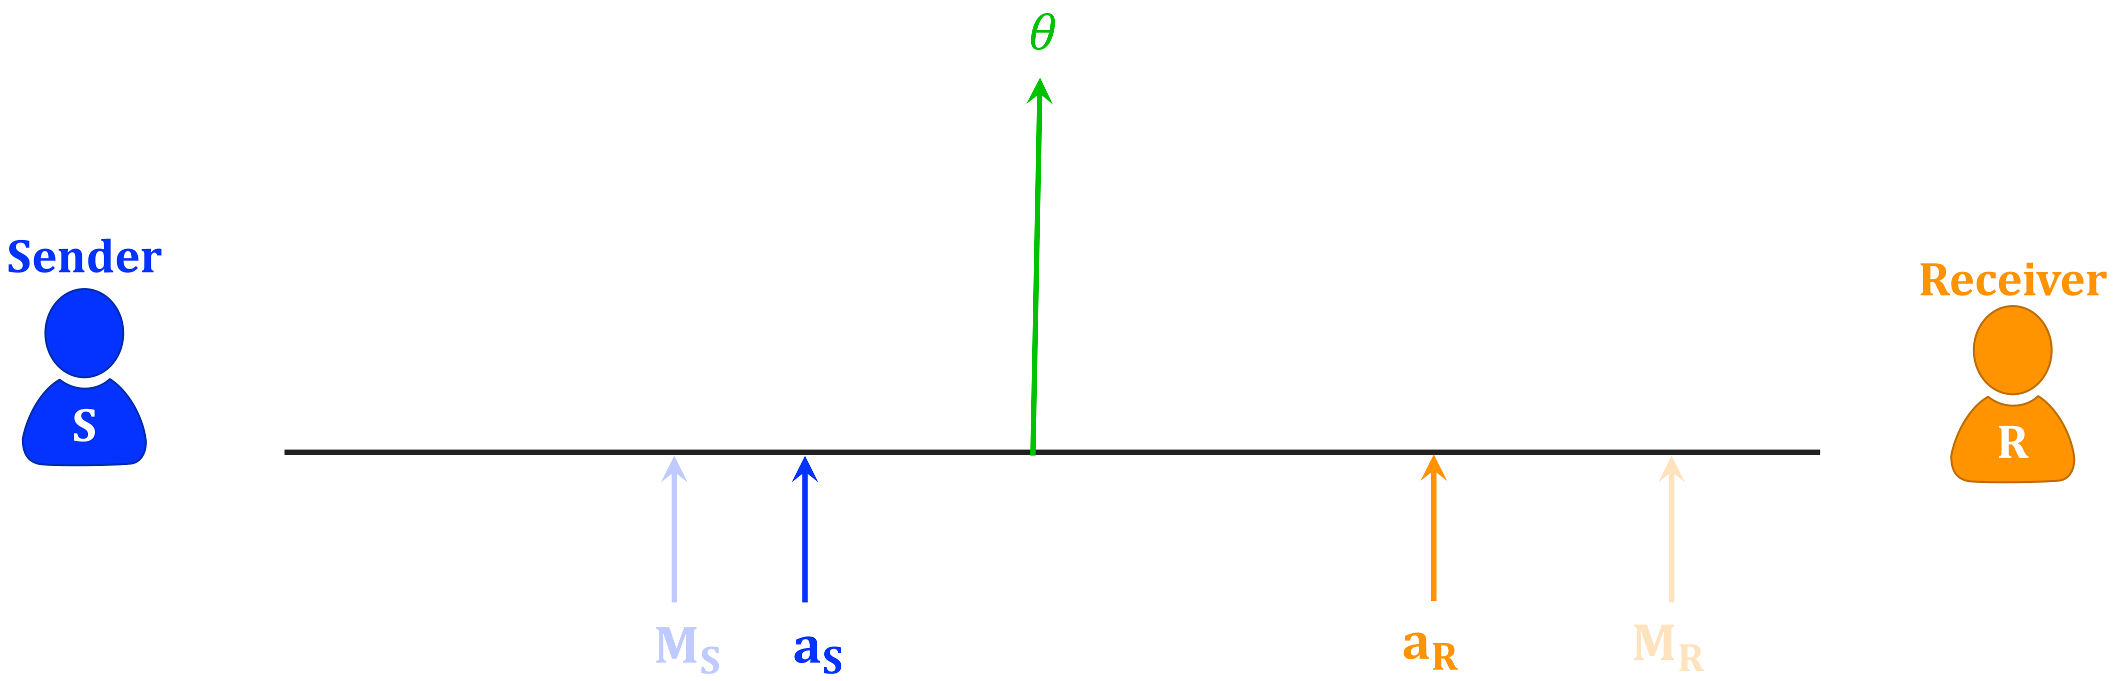

Moral stances $\mathbf{M_S}$, $\mathbf{M_R}$: Actions that would have been taken in absence of any information about the world.

Actions $\mathbf{a_S}$, $\mathbf{a_R}$: Actions taken if the state of the world θ were known.

$$\begin{aligned}\mathbf{a}_i &= \operatorname{argmin}_a \mathbb{E}[\alpha(a-\theta)^2 + (1-\alpha)(a-\mathbf{M}_i)^2] \\ &= \alpha\theta + (1-\alpha)\mathbf{M}_i\end{aligned}$$

Optional action, minimizes squared loss
to moral stances and state of the world

Model – Observations

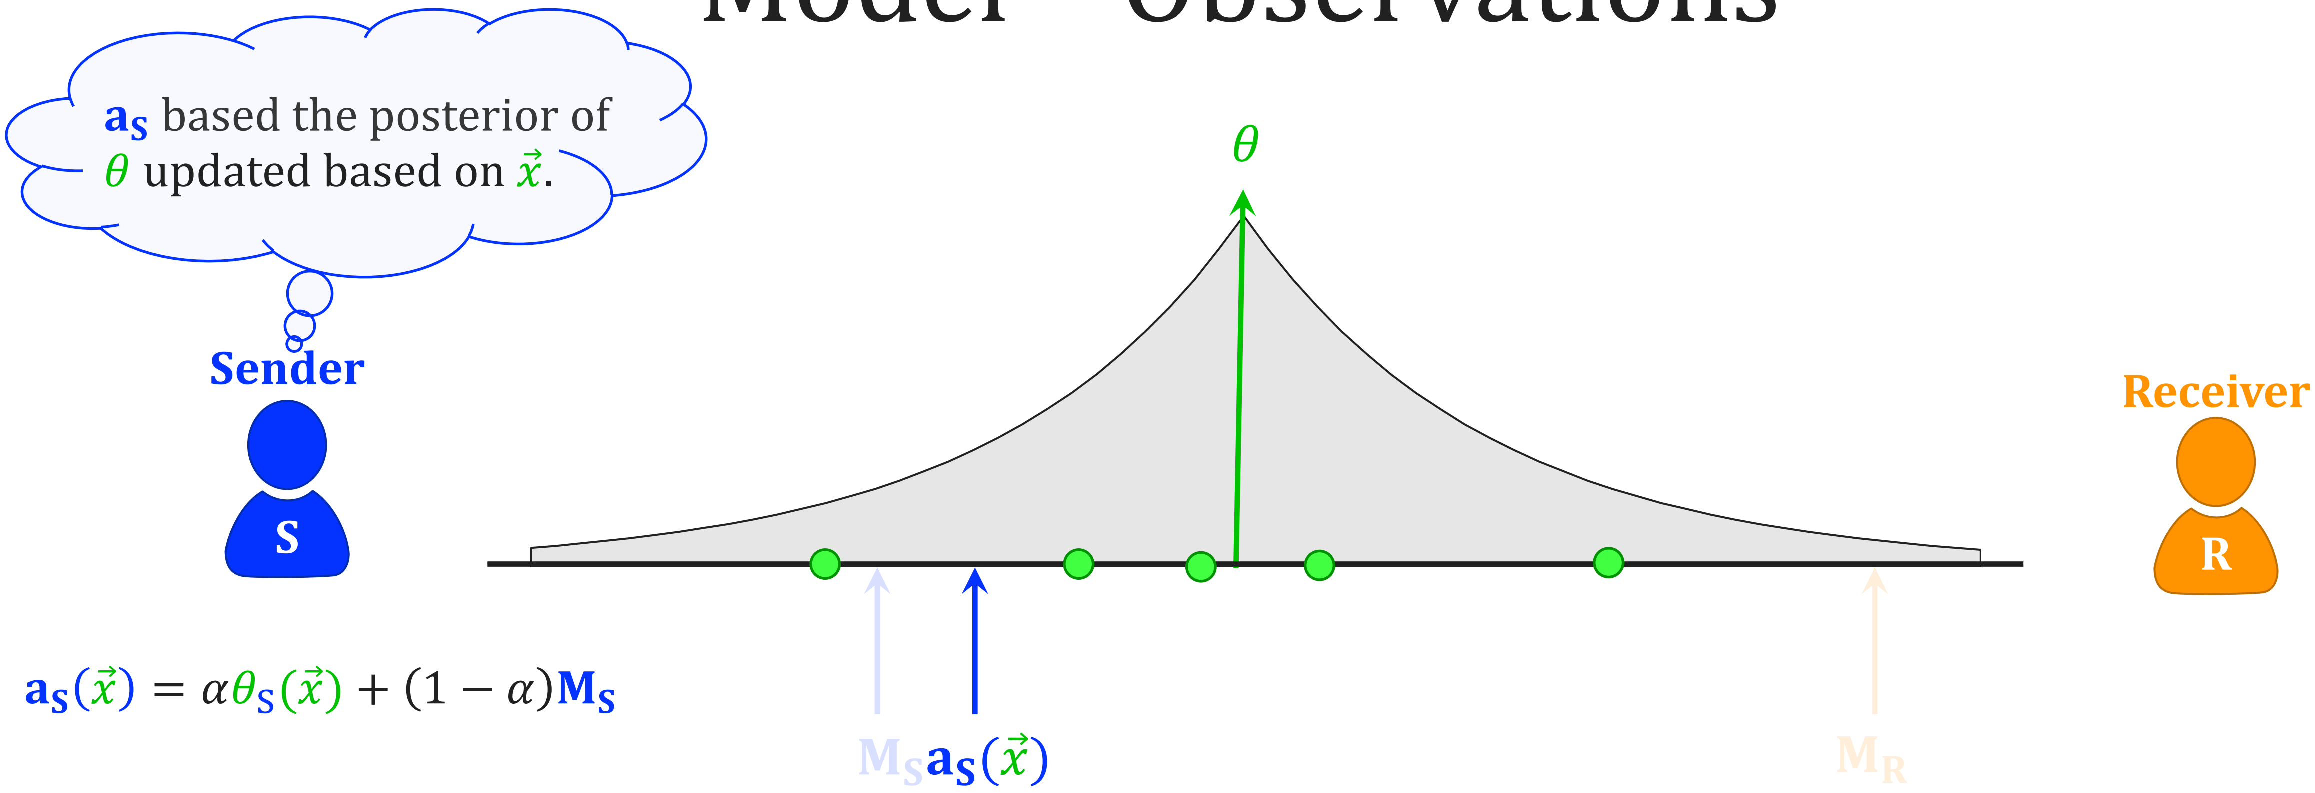

$\mathbf{a_S}$ based the posterior of θ updated based on $\vec{x}$.

$\mathbf{a_S}(\vec{x}) = \alpha\theta_S(\vec{x}) + (1-\alpha)\mathbf{M_S}$

Neither players know the state of world θ. Diffuse Prior: θ equally likely anywhere in $\mathbb{R}$.

Sender observes $\vec{x} = x_1, \ldots, x_k$ i.i.d from a distribution parameterized by θ
→ Single peaked at θ, symmetric, known pdf, f, given θ.

Model – Sender

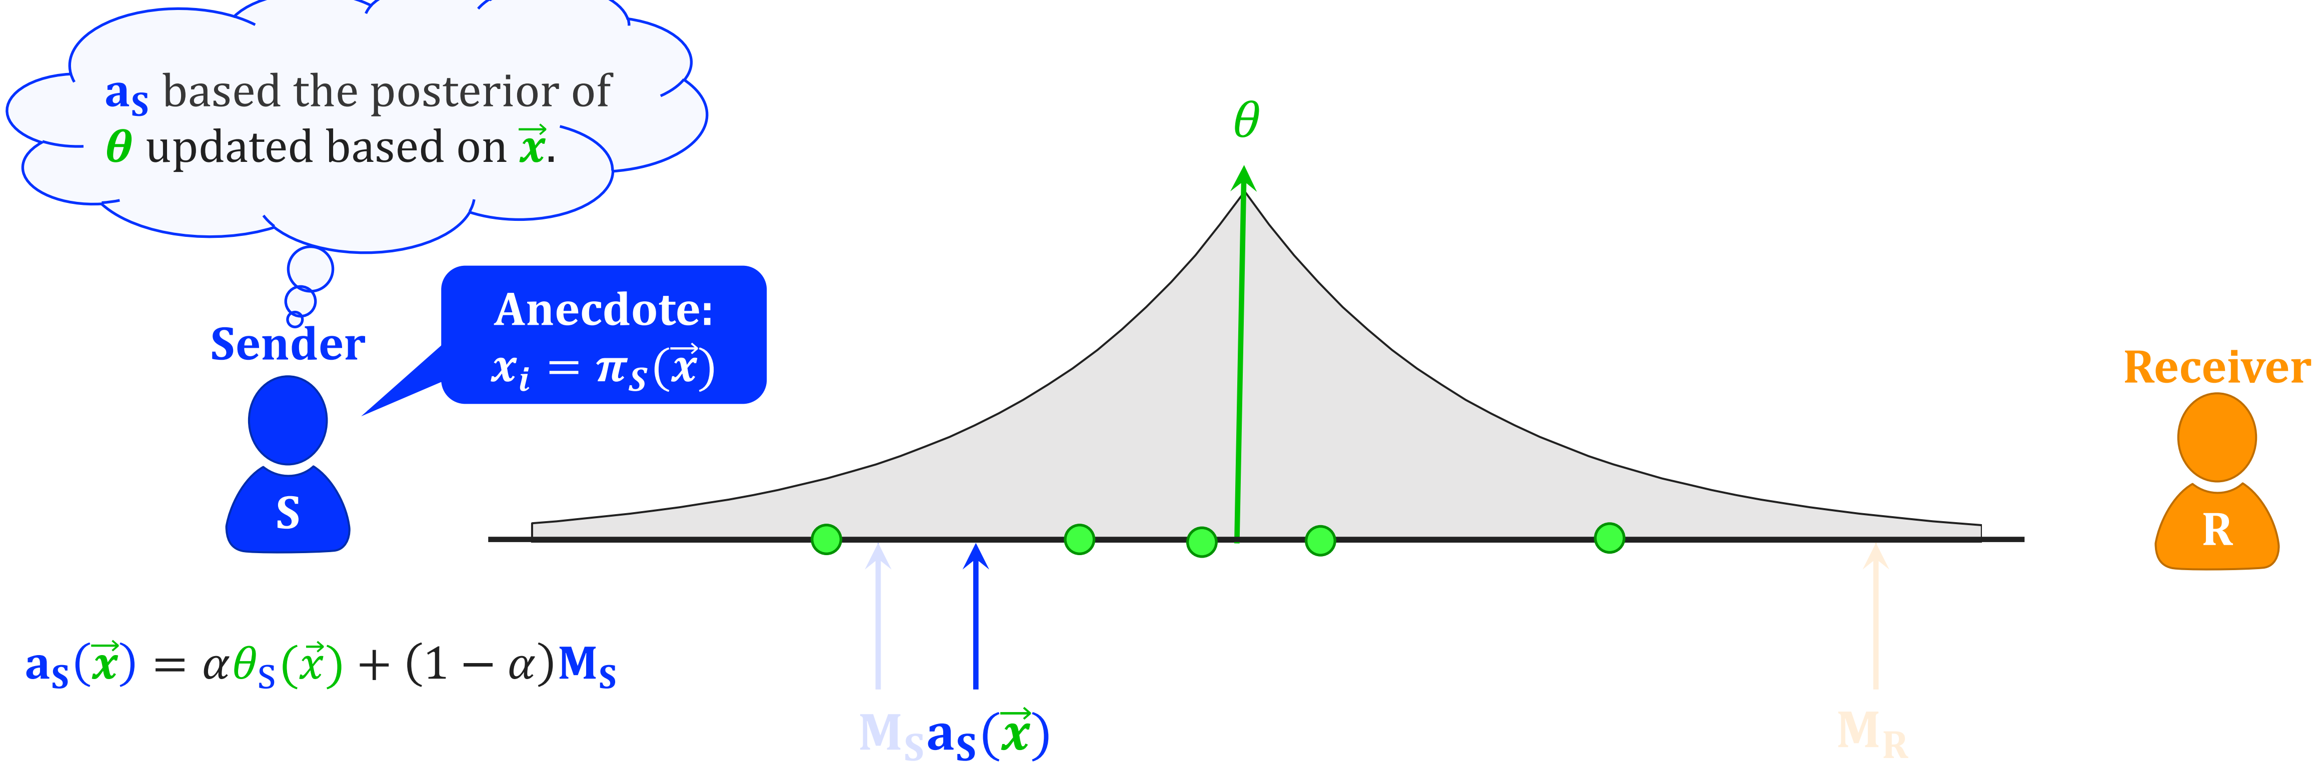

$\mathbf{a}_S$ based the posterior of $\boldsymbol{\theta}$ updated based on $\vec{\boldsymbol{x}}$.

Sender

Anecdote: $x_i = \pi_S(\vec{x})$

$$\mathbf{a}_S(\vec{x}) = \alpha\theta_S(\vec{x}) + (1-\alpha)\mathbf{M}_S$$

Neither players know the state of world θ. Diffuse Prior: θ equally likely anywhere in $\mathbb{R}$.

Sender observes $\vec{x} = x_1, \dots, x_n$ i.i.d from a distribution parameterized by θ

→ Single peaked at θ, symmetric, known pdf, f, given θ.

Sender **sends one anecdote** $x_i \in \{x_1, \dots, x_k\}$ to the receiver.

→ Using a communication scheme π_S , which might be observable or not by the receiver

Model – Receiver

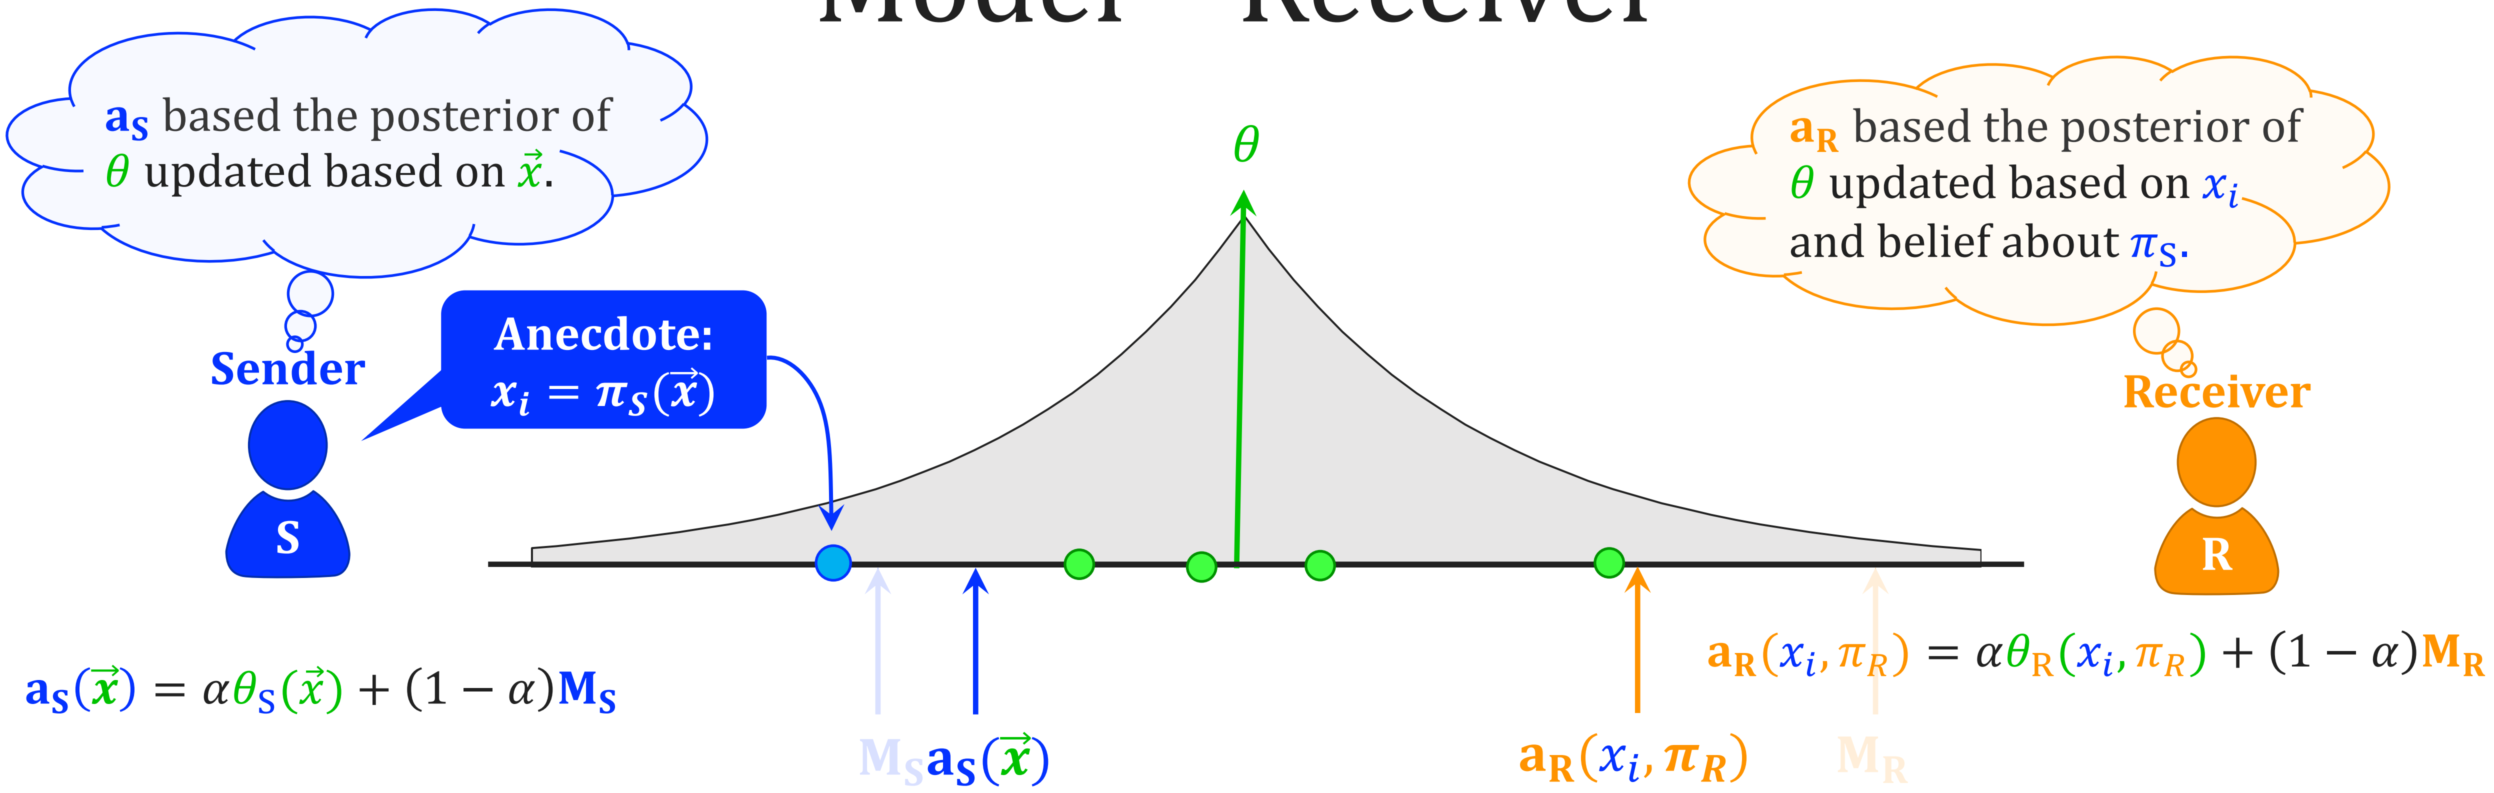

$\mathbf{a}_{\mathrm{S}}$ based the posterior of θ updated based on $\vec{x}$.

Anecdote:
$x_i = \pi_S(\vec{x})$

$\mathbf{a}_{\mathrm{R}}$ based the posterior of θ updated based on x_i and belief about π_S.

$$\mathbf{a}_{\mathrm{S}}(\vec{x}) = \alpha\theta_{\mathrm{S}}(\vec{x}) + (1-\alpha)\mathbf{M}_{\mathrm{S}}$$

$$\mathbf{a}_{\mathrm{R}}(x_i, \pi_R) = \alpha\theta_{\mathrm{R}}(x_i, \pi_R) + (1-\alpha)\mathbf{M}_{\mathrm{R}}$$

Receiver gets anecdote $\pi_S(\vec{x}) = x_i \in \{x_1, \dots, x_k\}$ and has a belief about the communication scheme

Belief π_R about communication scheme:
→ If the communication schme π_S was observed, $\pi_R = \pi_S$, otherwise, we'll consider equilibrium belief.
Receiver's posterior depends on π_R and $\pi_S(\vec{x})$.

Model – Persuasion

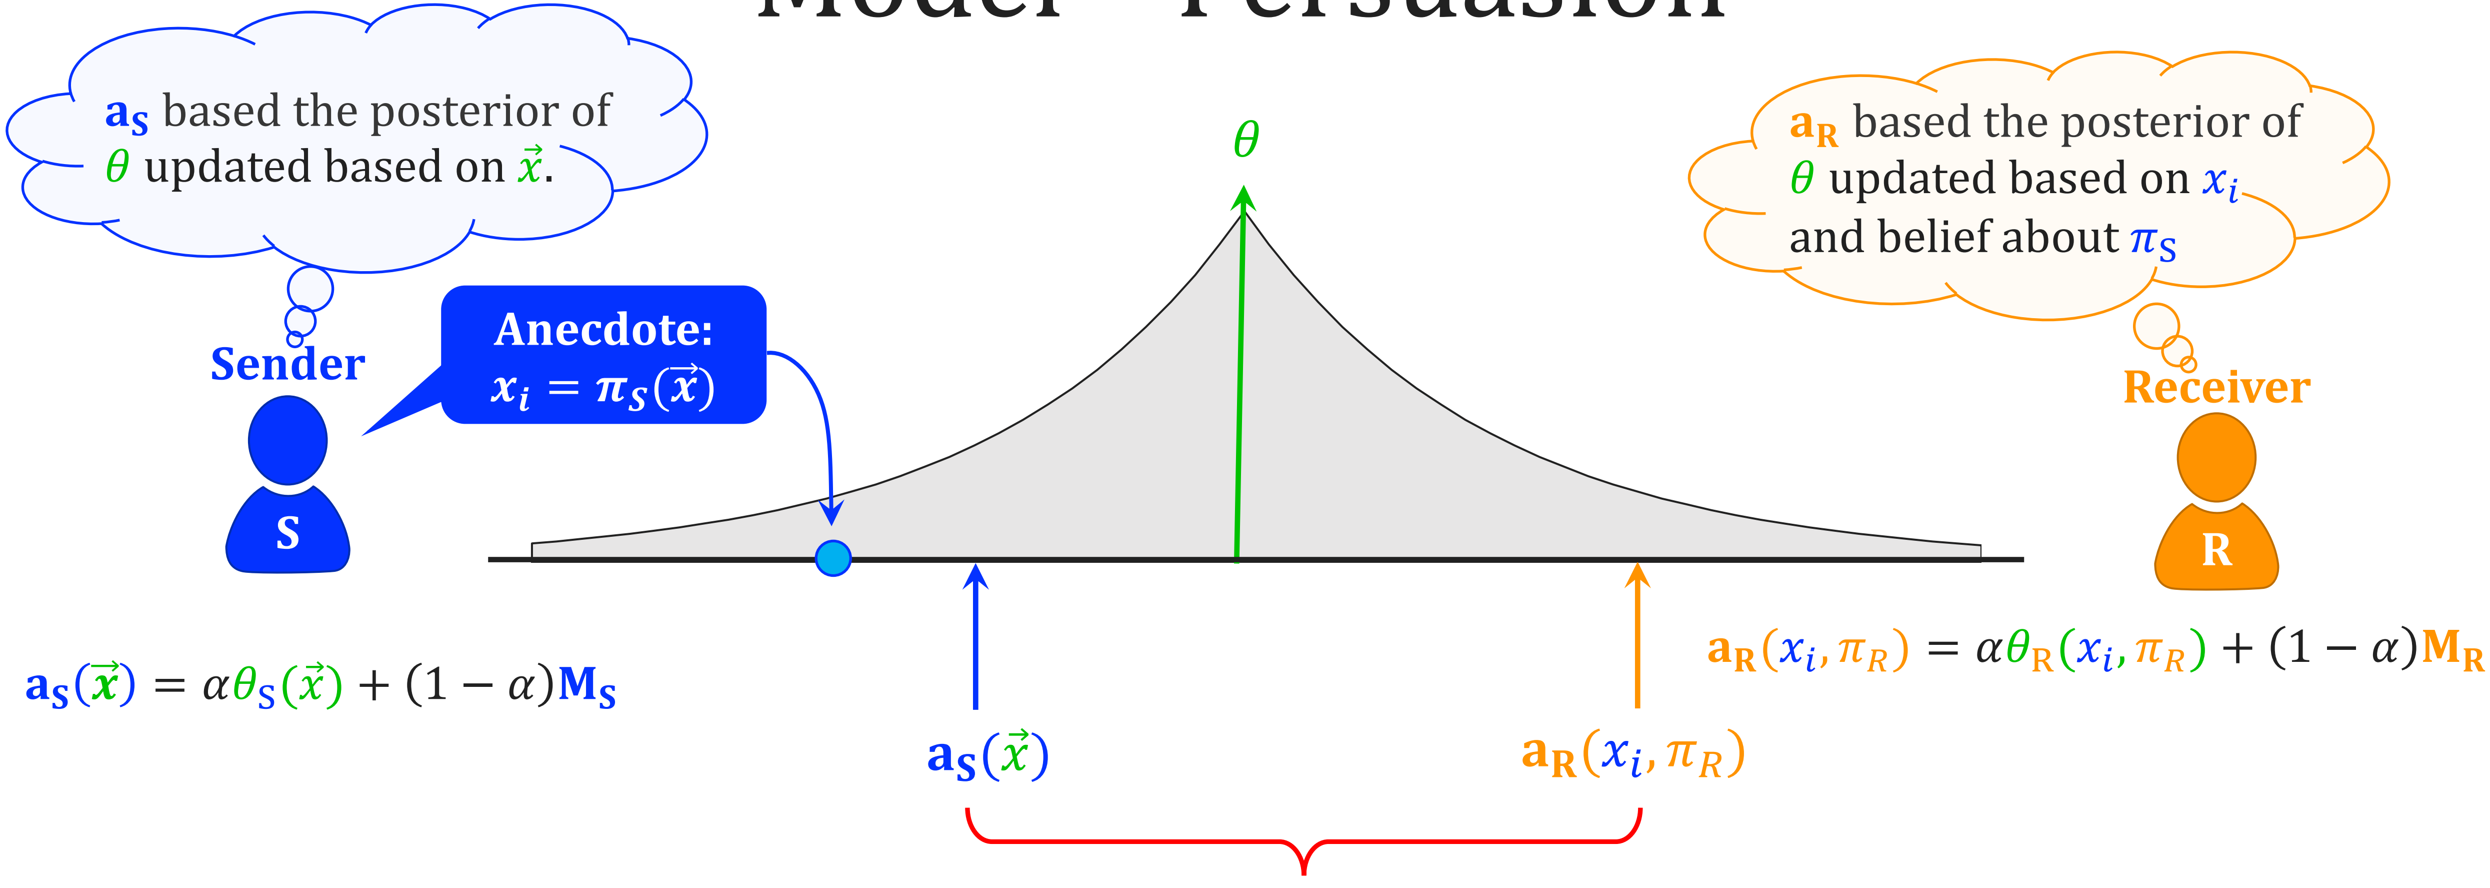

$\mathbf{a}_\mathrm{S}$ based the posterior of θ updated based on $\vec{x}$.

Anecdote:
$x_i = \pi_S(\vec{x})$

$\mathbf{a}_\mathrm{R}$ based the posterior of θ updated based on x_i and belief about π_S

$$\mathbf{a}_\mathrm{S}(\vec{x}) = \alpha\theta_\mathrm{S}(\vec{x}) + (1-\alpha)\mathbf{M}_\mathrm{S}$$

$$\mathbf{a}_\mathrm{R}(x_i, \pi_R) = \alpha\theta_\mathrm{R}(x_i, \pi_R) + (1-\alpha)\mathbf{M}_\mathrm{R}$$

$$\underbrace{\mathbf{a}_\mathrm{S}(\vec{x}) \qquad \mathbf{a}_\mathrm{R}(x_i, \pi_R)}$$

Sender’s Goal

Choose π_S to minimize cost $\mathbb{E}_{\vec{x}}\left[\left(\boldsymbol{a_S}(\vec{x}) - \boldsymbol{a_R}(\pi_S(\vec{x}), \pi_R)\right)^2\right]$

Anecdotes vs. Unrestricted Signals

When do we see biased signaling schemes?

How does the efficiency of signaling schemes change with the number observations?

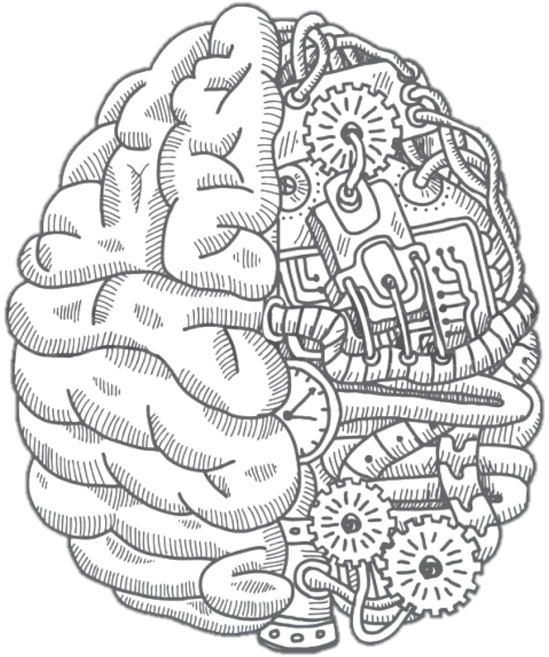

Talk Plan

Understand Sender’s and Receiver’s Perspectives

Optimal scheme for observable π_S

Optimal scheme for non-observable π_S

Talk Plan

Understand Sender's and Receiver's Perspectives

Optimal scheme for observable π_S

Optimal scheme for non-observable π_S

Communication Schemes

Sender uses a communication scheme $\boldsymbol{\pi_S}$ to choose $x_i \in \{x_1, \dots, x_k\}$

Examples: $\boldsymbol{\pi_S}(\vec{\boldsymbol{x}})$

☑ The minimum/maximum signal in $x_1, \dots, x_k$.

☒ The signal closest to 0.

☑ The signal closest to the posterior belief $\theta_S(\vec{x})$.

Translation-Invariant Schemes: Changing the axis doesn't change the scheme's choice.

A useful class of schemes: The signal closest to the posterior belief $\theta_S(\vec{x}) + r$.

→ When $r = 0$ unbiased, when $r \neq 0$ a bias that's between 0 and bias of the min/max signal

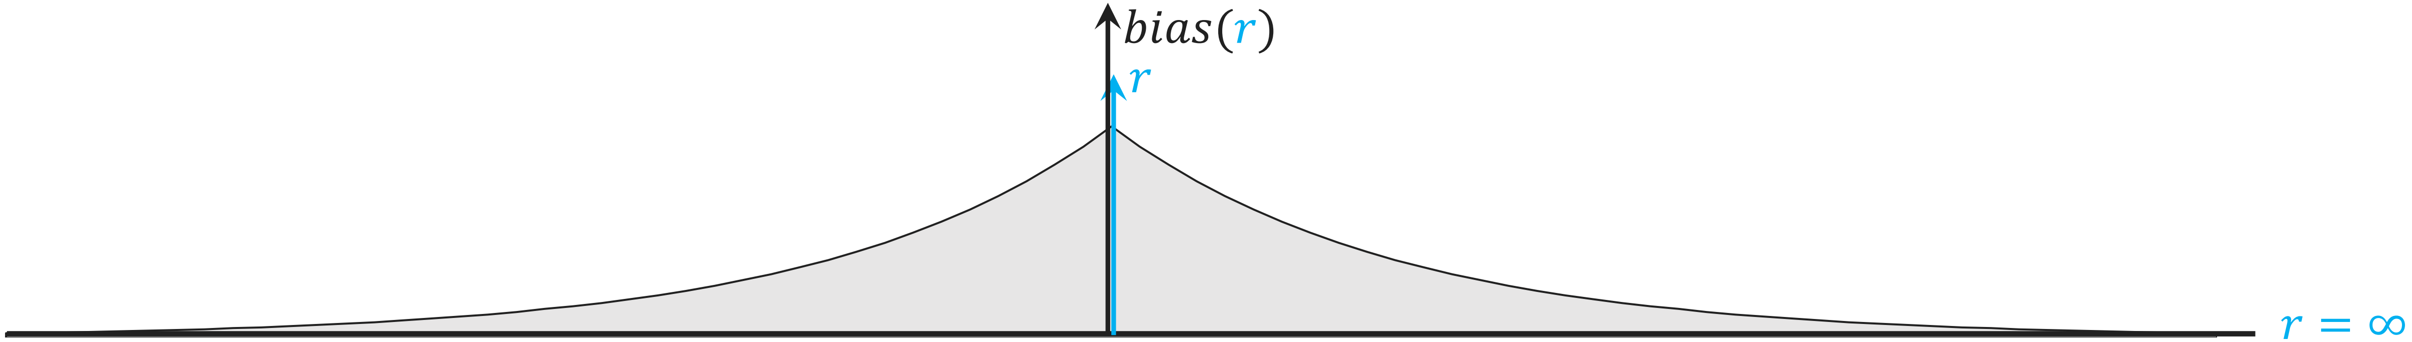

Receiver's Perspective

What does a receiver who gets signal $x_i = \pi_S(\vec{x})$ believe about the state of the world?

→Depends both on x_i and what she scheme she perceives, call it π_R.

→The receiver "undoes" the perceived bias in the communication scheme.

→If $\text{bias}(\pi_R) = \text{bias}(\pi_S)$, receiver's belief $\theta_\text{R}(x_i, \boldsymbol{\pi}_R)$ is unbiased.

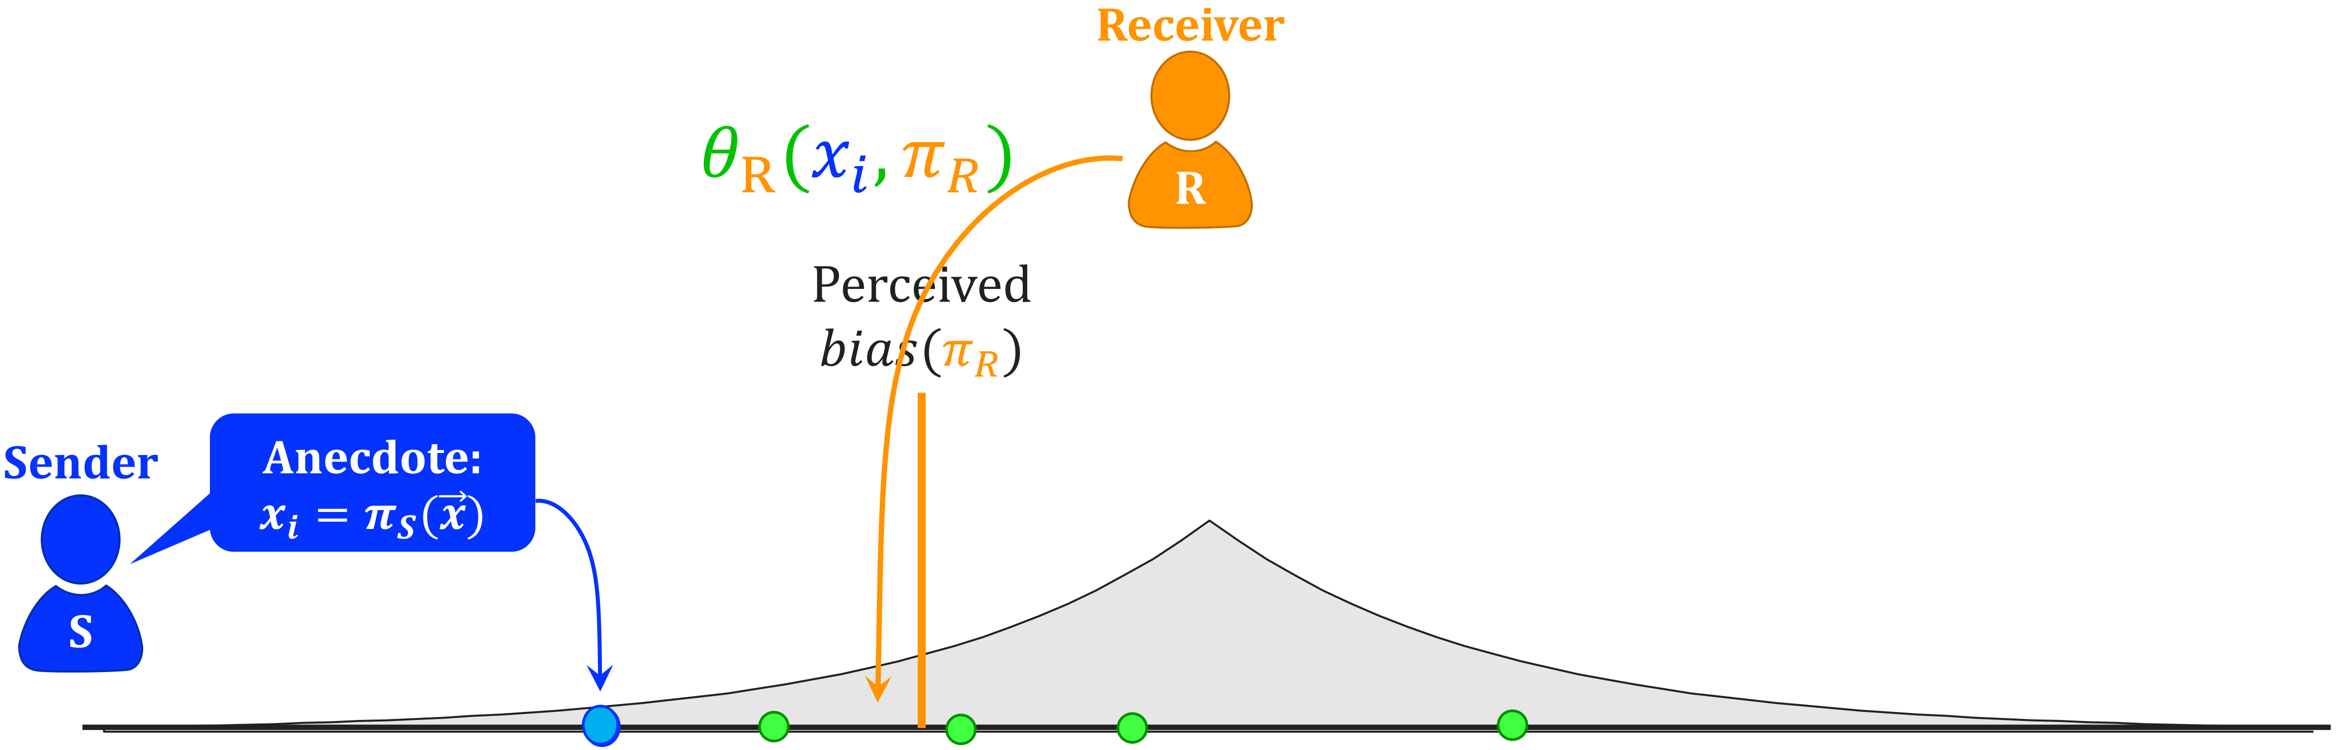

Understanding Sender’s Choices

Sender’s goal of minimization cost $\mathbb{E}_{\vec{x}}\left[\left(\boldsymbol{a}_S(\vec{x}) - \boldsymbol{a}_R(\pi_S(\vec{x}), \pi_R)\right)^2\right]$ takes into account:

① Inability of the sender to express any signal it wants.

② Fundamental gap in moral stances
→ $\boldsymbol{a}_S$ and $\boldsymbol{a}_R$ are attracted to $\mathbf{M_S}$, $\mathbf{M_R}$.

③ Potential Mis-match between the biases of the sender and receiver
→ $\mathrm{bias}(\pi_R)$, $\mathrm{bias}(\pi_S)$

Cost decomposition:

$$\begin{aligned}
\text{Sender}'\text{s Cost} = {} & \alpha^2 \mathbb{E}_{\vec{x}}\left[\left(\theta_S(\vec{x}) - \pi_S(\vec{x}) + bias(\pi_S)\right)^2\right] \quad \text{① Signaling cost} \\
& + (1-\alpha)^2 \, (M_S - M_R)^2 \quad \text{② Fundamental loss in moral stances} \\
& \left.\begin{aligned} & + \alpha^2\left(bias(\pi_R) - bias(\pi_S)\right)^2 \\ & + 2\alpha(1-\alpha)(bias(\pi_R) - bias(\pi_S))(M_S - M_R) \end{aligned}\right\} \text{③ Persuasion Temptation}
\end{aligned}$$

Talk Plan

Understand Sender and Receiver's Perspectives

Optimal scheme for observable π_S

Optimal scheme for non-observable π_S

Observable Communication Scheme

When π_S is observable, $\pi_S = \pi_R$ → Persuasion temptation is 0.

Sender chooses π_S that minimizes $\mathbb{E}_{\vec{x}}\left[\left(\theta_S(\vec{x}) - \pi_S(\vec{x}) + bias(\pi_S)\right)^2\right]$

What's the optimal communication scheme?

$$\text{Sender's Cost} = \alpha^2 \mathbb{E}_{\vec{x}}\left[\left(\theta_S(\vec{x}) - \pi_S(\vec{x}) + bias(\pi_S)\right)^2\right]$$

❶ Signaling cost

$$+ (1-\alpha)^2 \left(M_S - M_R\right)^2$$

❷ Fundamental loss in moral stances

$$\left.\begin{aligned} &+ \alpha^2\left(bias(\pi_R) - bias(\pi_S)\right)^2 \\ &+ 2\alpha(1-\alpha)\left(bias(\pi_R) - bias(\pi_S)\right)\left(M_S - M_R\right) \end{aligned}\right\}$$

❸ Persuasion Temptation

Optimal Communication Scheme

When π_S is observable, $\pi_S = \pi_R$ → Persuasion temptation is 0.

Sender chooses π_S that minimizes $\mathbb{E}_{\vec{x}}\left[\left(\theta_S(\vec{x}) - \pi_S(\vec{x}) + bias(\pi_S)\right)^2\right]$

What's the optimal communication scheme?

Optimal Communication Scheme

If the sender knew the true state of the world θ
→ $\pi_S(\vec{x})$ that's the closest signal to θ would have optimal cost.

Without knowing θ, as # of observations $\theta \rightarrow \infty$, $\theta_S(\vec{x}) \rightarrow \theta$
→ $\pi_S(\vec{x})$ that chooses the **closest signal to $\boldsymbol{\theta_S(\vec{x})}$** has near optimal cost, and is **unbiased**.

Any **biased communication** scheme is **suboptimal.**

Optimal Communication Scheme

When π_S is observable, $\pi_S = \pi_R$ → Persuasion temptation is 0.

Sender chooses π_S that minimizes $\mathbb{E}_{\vec{x}}\left[\left(\theta_S(\vec{x}) - \pi_S(\vec{x}) + bias(\pi_S)\right)^2\right]$

What's the optimal communication scheme?

Optimal Communication Scheme

Choose closest to $\theta_S(\vec{x})$

$$\text{Signaling cost} \leq \frac{\alpha^2}{2k^2 f(0)^2} + o\left(\frac{1}{k^2}\right)$$

Any other π_S

$$\text{Signaling Cost} \geq \frac{\alpha^2}{2k^2 f\left(bias(\pi_S)\right)^2} - o\left(\frac{1}{k^2}\right)$$

Where f is the pdf of the distribution around $\theta = 0$, recall single peaked and symmetric and some additional restrictions.

Talk Plan

- Understand Sender and Receiver's Perspectives
- **Optimal scheme for observable π_S**
  - Unbiased! For large number of observations.
- Optimal scheme for non-observable π_S

Talk Plan

Understand Sender and Receiver's Perspectives

Optimal scheme for observable π_S

Unbiased! For large number of observations.

Optimal scheme for non-observable π_S

Un-observable Communication Scheme

When π_S is not observable → There is temptation to persuade!

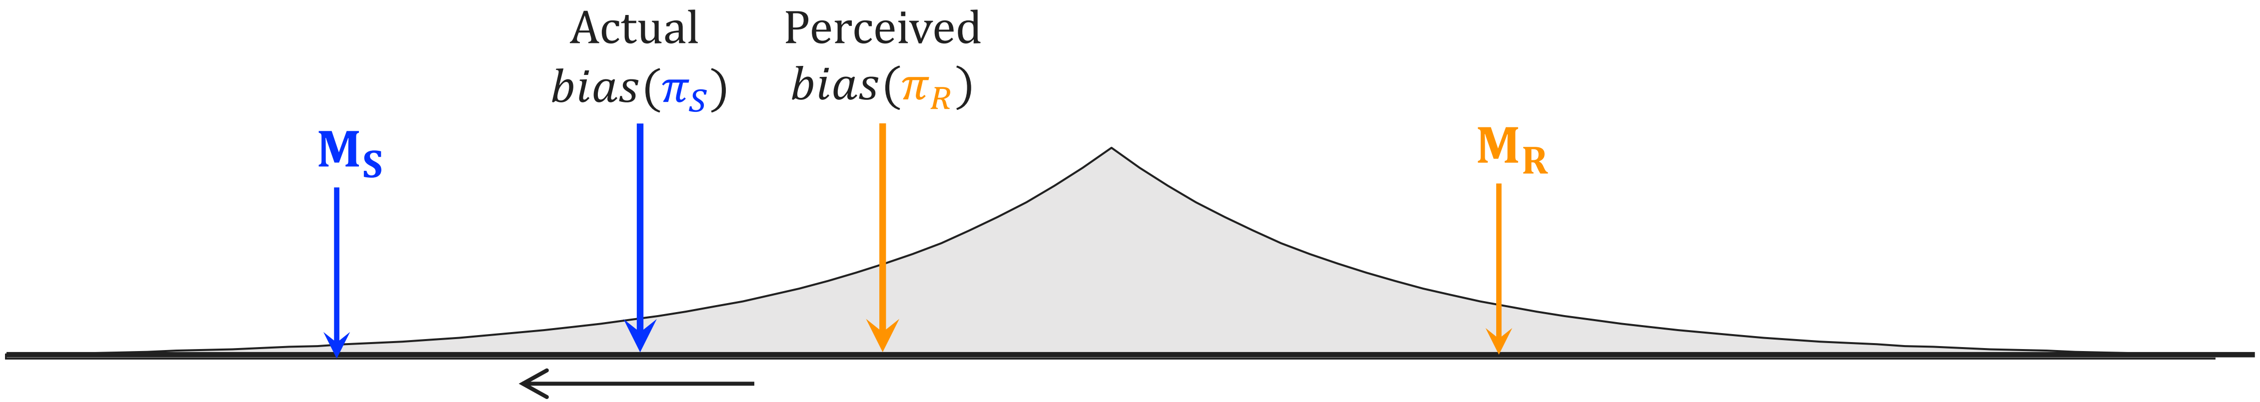

Improves persuasion temptation

Worsens the signaling cost: higher variance when π_S chooses signals farther from the center.

$$\text{Sender's Cost} = \alpha^2 \mathbb{E}_{\vec{x}}\left[\left(\theta_S(\vec{x}) - \pi_S(\vec{x}) + bias(\pi_S)\right)^2\right]$$

❶ Signaling cost

$$+ (1-\alpha)^2\,(M_S - M_R)^2$$

❷ Fundamental loss in moral stances

$$+ \alpha^2\left(bias(\pi_R) - bias(\pi_S)\right)^2$$

$$+2\alpha(1-\alpha)(bias(\pi_R) - bias(\pi_S))(M_S - M_R)$$

❸ Persuasion Temptation

Optimal Un-Observable Communication Scheme

At equilibrium (and with thought exercise of knowing θ):

- Sender's scheme π_S takes the closest signal to $\theta_S(\vec{x}) + r$ for some r, such that

$$\boxed{r - bias(r) = \frac{1-\alpha}{\alpha}(M_S - M_R)}$$

Independent of distribution and # of observations

Implications

1. For any k observation, as $|M_S - M_R| \to \infty$, $|r| \to \infty$
   → π_S converges to taking the most extreme, min/max signal from $\vec{x}$.
2. Similarly, for any $|M_S - M_R|$, as $k \to \infty$, $|r| \to \infty$.

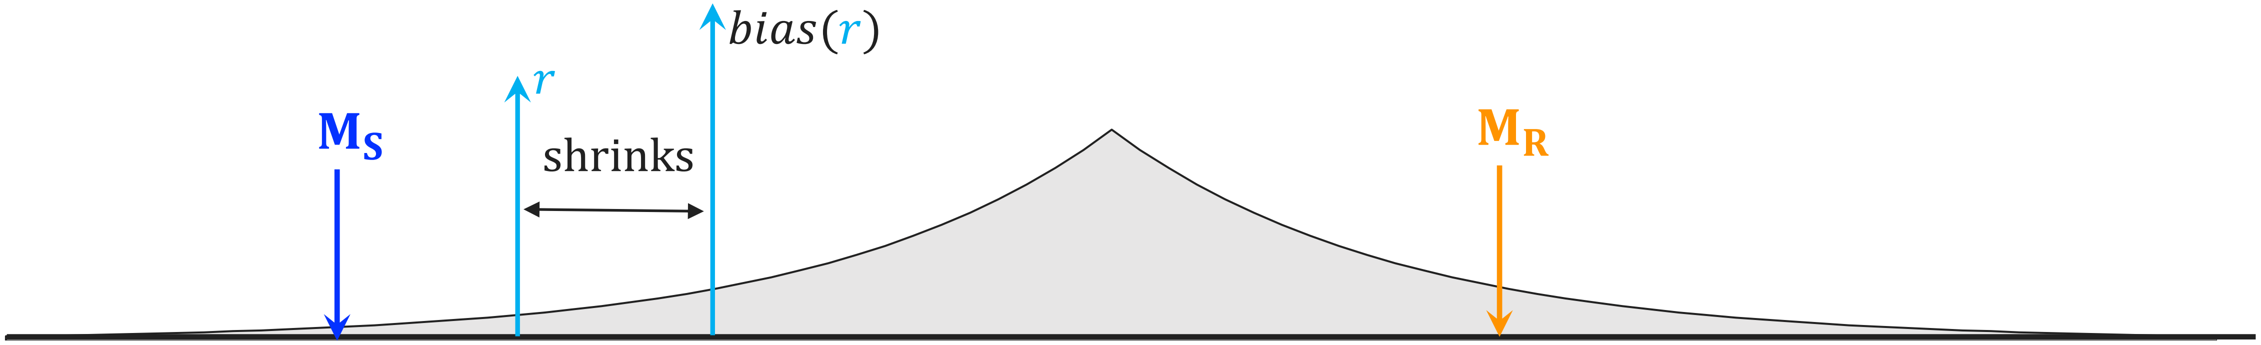

Receiver’s Perspective

Who’d you rather listen to?

- An expert with $k \to \infty$ observations, but with large $|M_S - M_R|$?
- A novice with small $k = 1,2, \ldots$ observations, but with $M_S = M_R$?

Depends on the distribution of observations (extreme value theory)

- Gaussian: The min/max signal has vanishing variance
  → You prefer to listen to the expert
- Laplacian: The min/max signal has a constant variance
  → You’ll choose to listen to the novice.

Homophily caused by the fact that agents communicate in anecdotes.

Talk Plan

Understand Sender and Receiver's Perspectives

Optimal scheme for observable π_S

Unbiased! For large number of observations.

Optimal scheme for non-observable π_S

Biased! For any number of observations.

Talk Plan

Understand Sender and Receiver’s Perspectives

Optimal scheme for observable π_S

Unbiased! For large number of observations.

Optimal scheme for non-observable π_S

Biased! For any number of observations.

Machine Learning and Strategic Behavior

Inspirations from machine learning theory for understanding polarization

Beliefs vs communication (generalization versus samples):

- Posterior distributions that describe your belief about what led to the state of the world.
- Complex **functions** that describe your actions in any one scenario.
  → Haghtalab, Jackson, Procaccia (working paper 2021).

More generally rich interplay between ML and Economics

- Coherent view of strategic behavior and learning behavior

Workshop series on ML in Presence of Strategic Behavior: Alternating between Economics and Computations and and ML conferences (NeurIPS)
Attend and submit!