Tailored Stories

Chiara Aina
University of Zurich

Research Question

I study the problem of manipulating a boundedly rational agent by controlling her interpretation of signals she is about to receive

Is it possible to persuade others only by providing interpretations of future events?

- Not only possible, but it can also lead the receiver to hold inconsistent beliefs across events
  - Allowing for multiple stories, I provide a disciplined relaxation of the Bayes-plausibility constraint
- Persuasion is generally limited and it depends on the initial beliefs

Set Up

- States $\omega \in \Omega$ with common prior on Ω: $\mu_0 \in int(\Delta(\Omega))$ & Signals $s \in S$
- **Model** m: map assigning to each state a distribution of signals conditional on that state

$$(\pi^m(s|\omega))_{s\in S,\omega\in\Omega} \in [\Delta(S)]^{\Omega}$$

- Adopting model m, an agent forms beliefs conditional on signal s via Bayes rule

$$\mu_s^m = (\mu_s^m(\omega))_{\omega\in\Omega} \in \Delta(\Omega)$$

- **Vector of posterior beliefs**: array of posterior distributions conditional on each signal

$$\boldsymbol{\mu}^m = (\mu_s^m)_{s\in S} \in [\Delta(\Omega)]^S$$

- A vector of posterior beliefs $\boldsymbol{\mu}$ is **Bayes-consistent** if the prior μ_0 is a convex combination of the posterior across signals $(\mu_s)_{s\in S}$ [Bayes-plausibility ⇒ Bayes-consistency]
- **Equivalent representation** between models and Bayes-consistent vectors of posteriors

Properties

- **Fit of a model** m **conditional on the signal** s: how likely a model fits the observed data

$$\Pr^m(s) = \sum_{\omega\in\Omega} \mu_0(\omega)\,\pi^m(s|\omega)$$

There is a multiplicity of models that induce the same posterior conditional on a signal with different levels of fit

- **Movement for** μ_s **in state** ω: how much the target posterior is far from the prior in a state

$$\delta(\omega;\mu_s) = \frac{\mu_s(\omega)}{\mu_0(\omega)}$$

- Maximal movement for μ_s: $\bar{\delta}(\mu_s) = \max_{\omega\in\Omega} \delta(\omega;\mu_s)$

Lemma: *a model m inducing μ_s conditional on signal s has fit* $\Pr^m(s) \in \left[0, \bar{\delta}(\mu_s)^{-1}\right]$

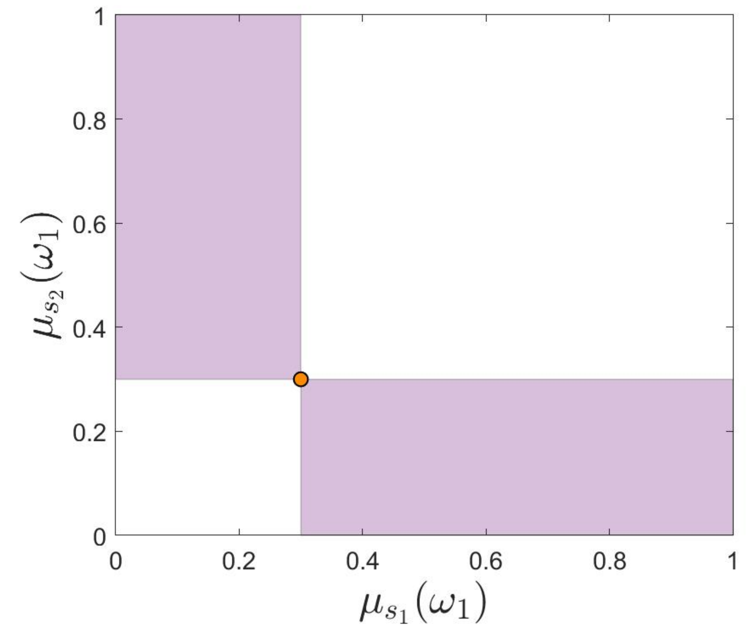

Bayes-consistent vectors of posteriors
Notes: orange point = prior 30%; every point in the purple area corresponds to a model.

Ex-ante Model Persuasion

Receiver

- The receiver does not know the state but she has observed a signal realization
- She needs a model to interpret the signal and update her priors
- The sender communicates a set of models $M \subseteq \mathcal{M}$
  $|M|$ is not greater than the number of models that the receiver is willing to consider

1. Model Adoption $\quad \tilde{m}_s \in \arg\max_{m\in M} \Pr^m(s)$ [Maximum likelihood selection]
2. Action Choice $\quad a^*(\mu_s) \in \arg\max_{a\in A} \mathbb{E}_{\mu_s^{\tilde{m}_s}}[U^R(a,\omega)]$

Tie breaking rule: if indifferent, adopt the model/action maximizing the sender's expected utility

Sender

- *What does the sender know?* The receiver's preferences, the (common) prior, and the number of models that the receiver is willing to consider
- The sender does not know the state, but he is endowed with a model t
  Used to computer: (i) predictive probabilities of each signal $\Pr^t(s)$, and (ii) posterior conditional on each signal μ_s^t
- **Sender's Value of** μ, calculated over signal and state realizations using model t

$$V(\mu) = \mathbb{E}^t[U^S(a^*(\mu_s),\omega)] = \sum_{s\in S} \Pr^t(s)\,\mathbb{E}_{\mu^t}\left[U^S(a^*(\mu_s),\omega)|s\right]$$

Many Models: Choose the set of models M^* that maximizes his value at $\mu^M = \left(\mu_s^{\tilde{m}_s}\right)_{s\in S}$

$$M^* \in \arg\max_{M\subseteq\mathcal{M}} V(\mu^M) \quad \text{such that} \quad \tilde{m}_s \in \arg\max_{m\in M} \Pr^m(s)$$

One Model: If the receiver considers only one model from the sender, the problem is

$$m^* \in \arg\max_{m\in\mathcal{M}} V(\mu^m)$$

To solve these, it is enough to characterize the set of feasible vector of posterior belief because, from the perspective of the sender, there is a fixed distribution over the signals $\left(\Pr^t(s)\right)_{s\in S}$

Set of Feasible Vectors of Posterior Beliefs: One Model

With only a model, the sender can only induce vectors of posteriors that are Bayes-consistent

Set of Feasible Vectors of Posterior Beliefs: Many Models

With more models, the sender can also induce Bayes-inconsistent vectors of posteriors

Theorem: *a vector of posteriors $\boldsymbol{\mu}$ is feasible if* $\sum_{s\in S} \bar{\delta}(\mu_s)^{-1} \geq 1$

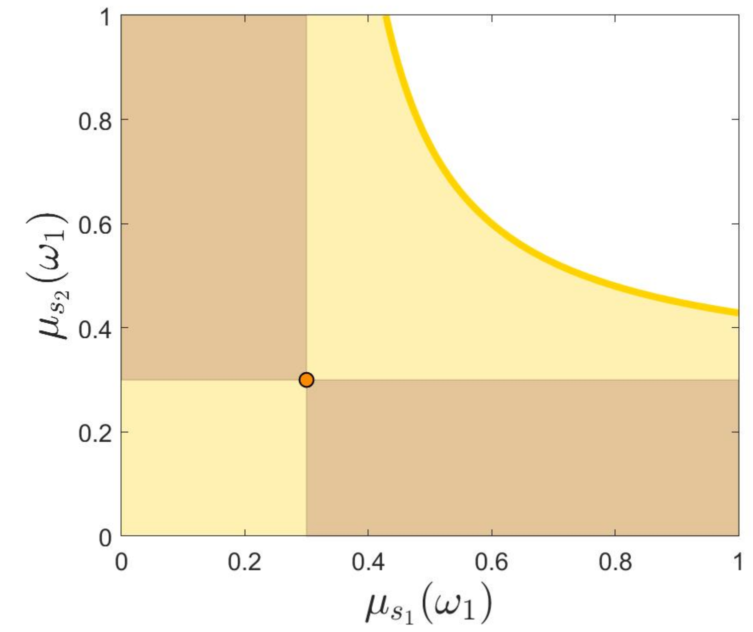

Set of Feasible Vectors of Posteriors
Notes: every vector of posterior in the yellow area is feasible.

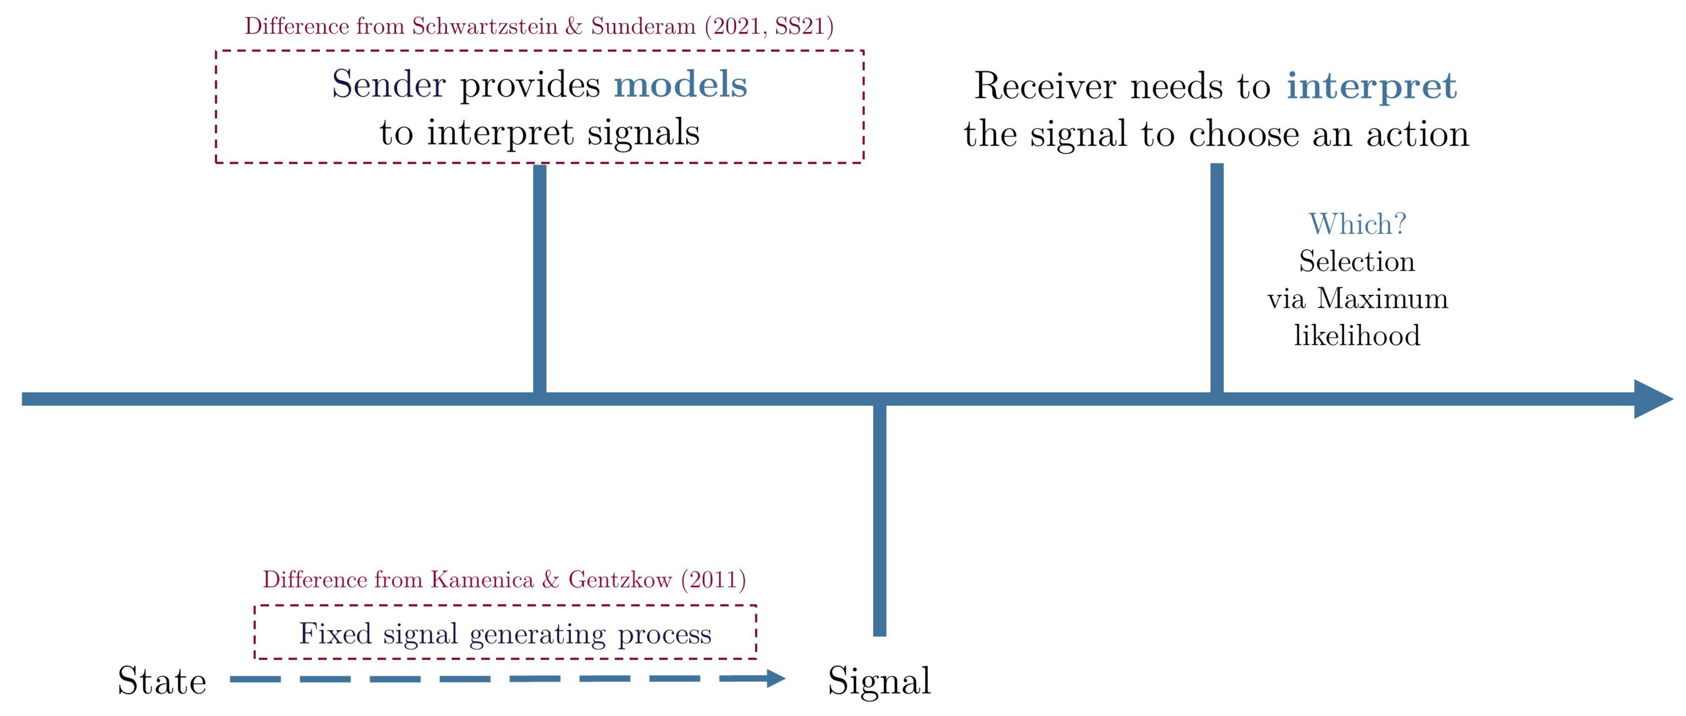

Comparative Statics

- Generally not all vectors of posteriors are feasible, but there are exceptions
- The more signals, the more belief manipulability
- The more uniform priors, the more belief manipulability

Applications

1. **Firehose of Falsehood:** model of Russian propaganda based on a large number of possibly contradictory and mutually inconsistent messages (Paul & Matthew, 2016)
   - With conflicting narratives, belief polarization occurs: there is a threshold in prior such that voters with prior higher (lower) than the threshold would hold extreme high (low) posteriors regardless the election outcome
2. **Finance:** with misaligned incentives an advisor can effectively manipulate investors to invest in his preferred asset
   - Even without knowing investors' relavant information such as past experience, the advisor communicates ad-hoc stories to maximize his return
3. **Nudging:** proposing ad-hoc narratives can be seen as a soft intervention to influence in a not coercive manner choices of an agent with the purpose of increase her welfare
   - Confidence manipulation by a paternalistic planner, via distorting the interpretation of signals, is optimal to influence the agent's behavior in a risky task
4. **Intra-personal Phenomena:** a mechanism through which the individual may distort his beliefs without assuming exogenous parameter of memory loss, inattention, first-impression, etc.
   - In a multi-selves model, an agent has incentives to distort his self-confidence in order to offset his time inconsistent preferences

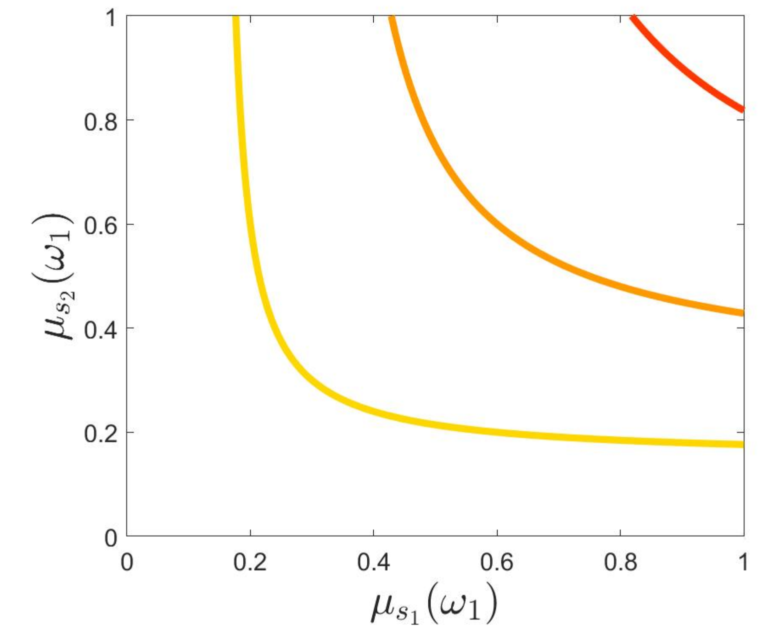

Bayes-consistent vectors of posteriors
Notes: the lighter the color line, the further away from the uniform prior: priors 15%, 30%, 45%.

chiara.aina@econ.uzh.ch