Forecasted Learning

EEA 2024

Paul Richter & Raquel Lorenzo

Universitat Pompeu Fabra

paul.richter@upf.edu

raquel.lorenzo@upf.edu

What is the paper about?

Decision-makers often learn in an environment where they:

- Choose between different biases/distorted sources of information (e.g. news they read)
- Know they'll receive info beyond their control in the future (e.g. news others share)

What is the paper about?

Decision-makers often learn in an environment where they:

- Choose between different biases/distorted sources of information (e.g. news they read)
- Know they'll receive info beyond their control in the future (e.g. news others share)

How does the expectation of additional info affect individual learning decisions?

How does news sharing affect agents' choice of media bias?

Results Preview

The expectation of additional information affects the optimal choice of bias

Results Preview

The expectation of additional information affects the optimal choice of bias

This Project:

- Focuses on rational Bayesian DMs and independent sources (Independence conditional on the state)
- Rationalizes phenomenon such as Opposite-Biased-Learning and Coordination
- Characterizes optimal choice of bias with regard to binary signals
- Applies results to strategic setting with social interaction and information sharing

The Model

- A Bayesian decision-maker (DM) chooses $a \in \{L, R\}$
- Unknown state of the world $\theta \in \{L, R\}$
- Objective: choose the correct action, *support the correct policy*

$$u(a, \theta) = \begin{cases} 1 & a = \theta \\ 0 & a \neq \theta \end{cases}$$

The Model

- A Bayesian decision-maker (DM) chooses $a \in \{L, R\}$
- Unknown state of the world $\theta \in \{L, R\}$
- Objective: choose the correct action, *support the correct policy*

$$u(a,\theta) = \begin{cases} 1 & a = \theta \\ 0 & a \neq \theta \end{cases}$$

- Prior belief $p_0 \in [0,1]$ is the probability that $\theta = R$
- Before choosing a, DM obtains info from two sources (ind. signals):
  - $\rightarrow$ A chosen source σ_C, *news she reads*
  - $\rightarrow$ An exogenous source σ_E, *news shared by others*

The Timing

T0. Fix a state θ, a prior p_0 and an exogenous source σ_E

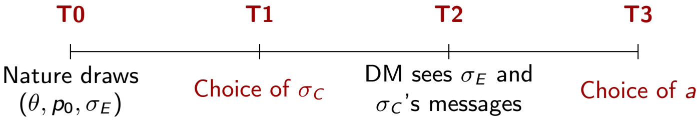

The Timing

T0. Fix a state θ, a prior p_0 and an exogenous source σ_E

T1. Knowing σ_E but not its message, the DM chooses a source σ_C

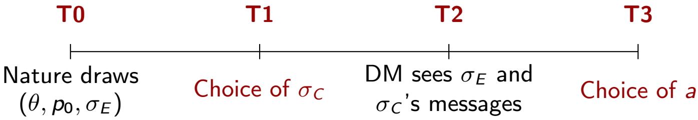

The Timing

T0. Fix a state θ, a prior p_0 and an exogenous source σ_E

T1. Knowing σ_E but not its message, the DM chooses a source σ_C

T2. Messages from σ_C, σ_E are realized and observed by the DM

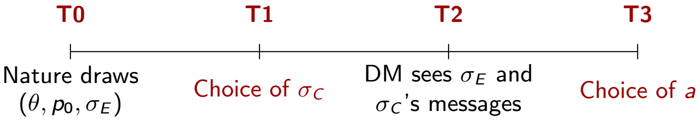

The Timing

T0. Fix a state θ, a prior p_0 and an exogenous source σ_E

T1. Knowing σ_E but not its message, the DM chooses a source σ_C

T2. Messages from σ_C, σ_E are realized and observed by the DM

T3. DM chooses an action a and obtains utility $u(a, \theta)$

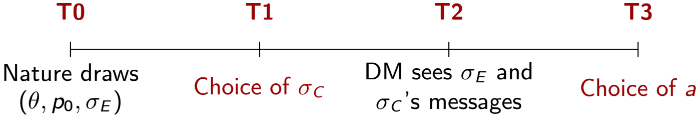

The Sources – A Binary Signal Structure

Sources (signals) can send two possible messages (realizations): l or r

A source σ is characterized by the prob. of each message conditional on θ

- Prob. of message l in state L: $\mathbb{P}(l|L)$
- Prob. of message r in state R: $\mathbb{P}(r|R)$

	$m = l$	m=r
$\theta = L$	$\mathbb{P}(l|L)$	$1 - \mathbb{P}(l|L)$
$\theta = R$	$1 - \mathbb{P}(r|R)$	$\mathbb{P}(r|R)$

where $\mathbb{P}(l|L) > 1 - \mathbb{P}(r|R)$

The Sources – Assumptions for this Talk

Two sources of information: σ_L and σ_R

- DM chooses between the two $\sigma_C \in \{\sigma_L, \sigma_R\}$
- On top, expects $\sigma_E \in \{\sigma_L, \sigma_R\}$, *everyone can access the same media*

The Sources – Assumptions for this Talk

Two sources of information: σ_L and σ_R

- DM chooses between the two $\sigma_C \in \{\sigma_L, \sigma_R\}$
- On top, expects $\sigma_E \in \{\sigma_L, \sigma_R\}$, *everyone can access the same media*

Each source **biased** towards one state: σ_L Left-biased and σ_R Right-biased

Definition

For every source σ characterized by $\mathbb{P}(l|L)$, $\mathbb{P}(r|R)$ as above:

- the source is *Left-biased* if $\mathbb{P}(l|L) > \mathbb{P}(r|R)$
- the source is *Right-biased* if $\mathbb{P}(r|R) > \mathbb{P}(l|L)$

The Sources – Assumptions for this Talk

The two sources, σ_L and σ_R, are **symmetric**

Definition

Two sources σ^L, σ^R are *symmetric* if:

$\mathbb{P}(m_L = l|L) = \mathbb{P}(m_R = r|R)$ and $\mathbb{P}(m_L = r|R) = \mathbb{P}(m_R = l|L)$

The Sources – Assumptions for this Talk

The two sources, σ_L and σ_R, are **symmetric**

Definition

Two sources σ^L, σ^R are *symmetric* if:

$\mathbb{P}(m_L = l|L) = \mathbb{P}(m_R = r|R)$ and $\mathbb{P}(m_L = r|R) = \mathbb{P}(m_R = l|L)$

Messages from σ_C, σ_E are **independent** conditional on θ

- If DM (mis)matches σ_C with σ_E is not due to correlation
- Note, σ_C and σ_E can send different messages even if $\sigma_C = \sigma_E$

Optimality of Own-Biased Learning w/o Additional Info

The expected value of each source *in isolation* depends on the prior belief

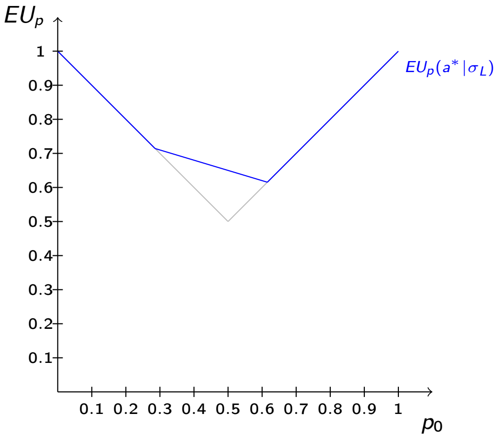

Optimality of Own-Biased Learning w/o Additional Info

The expected value of each source *in isolation* depends on the prior belief

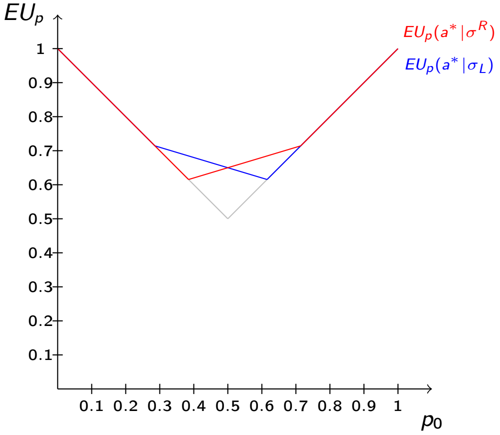

Optimality of Own-Biased Learning w/o Additional Info

The difference in expected value of sources also depends on the prior belief

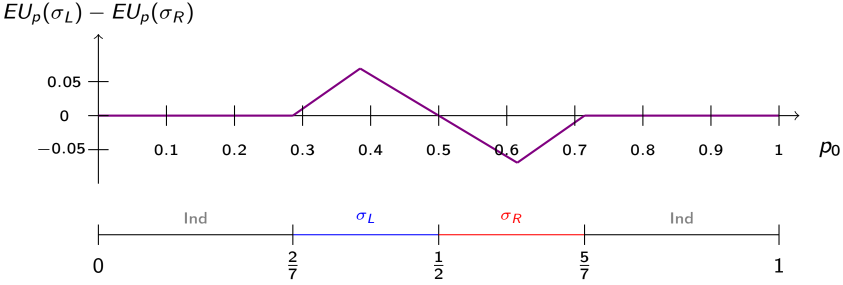

- Own-biased learning is weakly optimal
- Strictly for sufficiently uncertain DMs

Solving for the Optimal Strategy w/ Additional Info

Optimal strategy can be found:

1. Comparing expected values at each interim posterior belief, p_{m_E}
2. Computing weighted averages using Martingale property

Solving for the Optimal Strategy w/ Additional Info

Optimal strategy can be found:

1. Comparing expected values at each interim posterior belief, p_{m_E}
2. Computing weighted averages using Martingale property

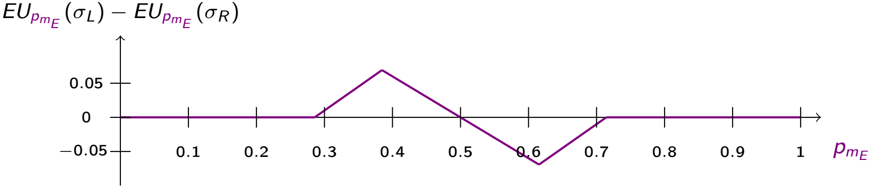

Solving for the Optimal Strategy w/ Additional Info

Optimal strategy can be found:

1. Comparing expected values at each interim posterior belief, p_{m_E}
2. Computing weighted averages using Martingale property

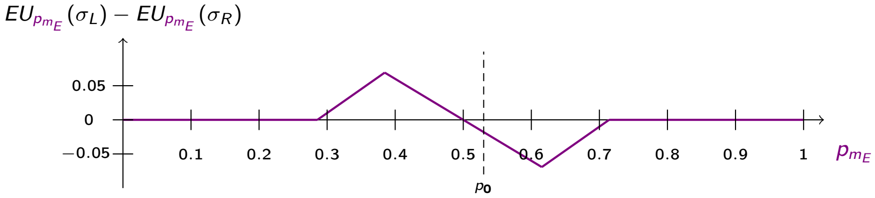

Solving for the Optimal Strategy w/ Additional Info

Optimal strategy can be found:

1. Comparing expected values at each interim posterior belief, p_{m_E}
2. Computing weighted averages using Martingale property

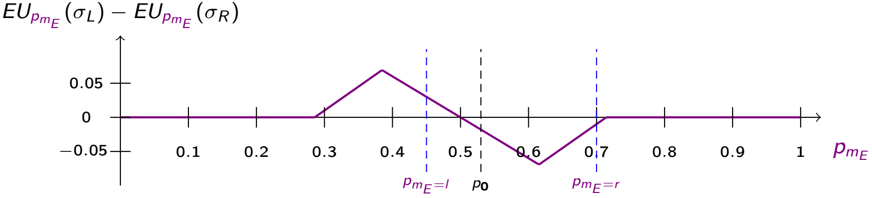

Solving for the Optimal Strategy w/ Additional Info

Optimal strategy can be found:

1. Comparing expected values at each interim posterior belief, p_{m_E}
2. Computing weighted averages using Martingale property

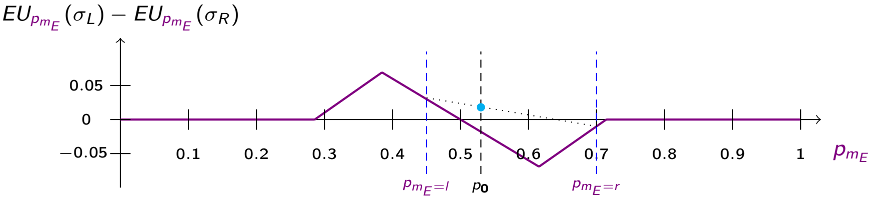

Solving for the Optimal Strategy w/ Additional Info

Optimal strategy can be found:

1. Comparing expected values at each interim posterior belief, p_{m_E}
2. Computing weighted averages using Martingale property

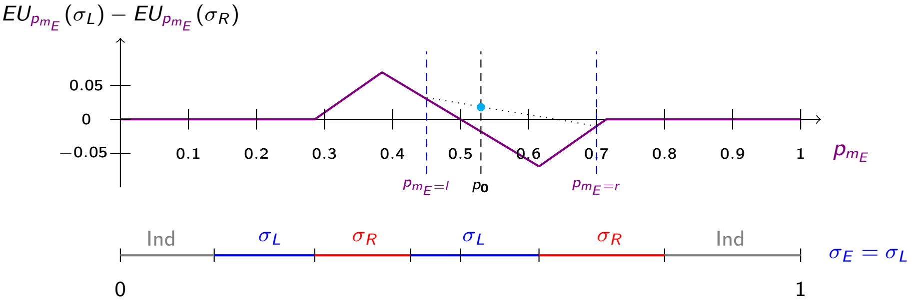

Example of a very certain Left-biased DM who chooses own-biased learning

Solving for the Optimal Strategy w/ Additional Info

Optimal strategy can be found:

1. Comparing expected values at each interim posterior belief, p_{m_E}
2. Computing weighted averages using Martingale property

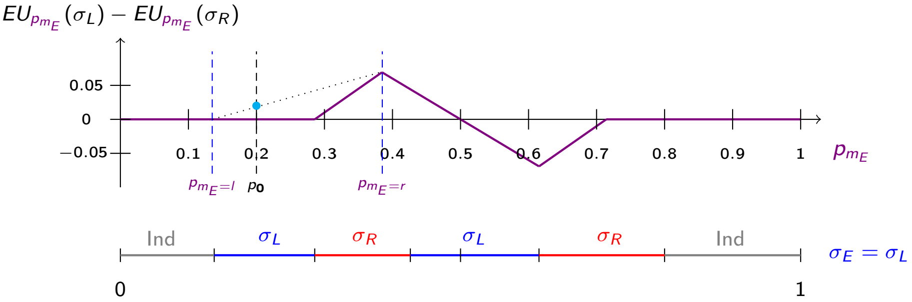

Example of a Right-biased DM who chooses opposite-biased learning

Structure of the Optimal Strategy

CS Thresholds

Proposition

For a given σ_L, σ_R and $\sigma_E \in \{\sigma_L, \sigma_R\}$, there exist five thresholds

$0 \leq p_1 < p_2 \leq p_3 \leq p_4 < p_5 \leq 1$ *(all strict when $\mathbb{P}_L(l|L), \mathbb{P}_L(r|R) \neq 1$) s.t.*

i) *If $p_0 \in [0, p_1] \cup [p_5, 1]$, i DM is **indifferent** between any choice of σ_C*

ii) *If $p_0 \in [p_1, p_2] \cup [p_4, p_5]$, DM chooses **own-biased learning***

iii) *If $p_0 \in [p_2, \min\{p_3, \frac{1}{2}\}] \cup [\max\{p_3, \frac{1}{2}\}, p_4]$, DM chooses **opp-biased learning***

iv) *If $p_0 \in [\min\{p_3, \frac{1}{2}\}, \max\{p_3, \frac{1}{2}\}]$, DM chooses to **match** σ_E*

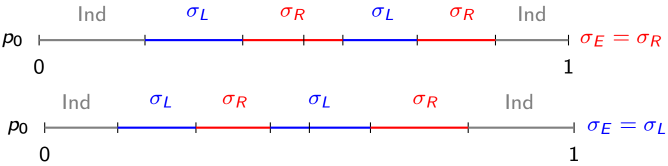

General Features of the Optimal Strategy

Many features extend to any binary exogenous source σ_E, change with informativeness of σ_E

- Extremely certain priors: **Indifference**
- Very certain priors: **Own-BL**
- **Opposite-BL** for moderately certain priors (if σ_E suff. info/bias)
- Uncertain priors: Own-BL, Indifference or **Matching**

Review & Additional Results

Project Key Points

→ Adds **expectation of future, independent information** to standard Bayesian learning model
→ Characterize the preference for source bias/distortion for binary signals.
→ Rationalize opposite-biased learning and consuming varied bias.

Generality of Mechanism ▸ More

→ Additional ind. info can change optimal info choice in more general contexts
→ Sources with more messages, multiple states, other utility functions

Review & Additional Results

Project Key Points

- → Adds **expectation of future, independent information** to standard Bayesian learning model
- → Characterize the preference for source bias/distortion for binary signals.
- → Rationalize opposite-biased learning and consuming varied bias.

Generality of Mechanism ▸ More

- → Additional ind. info can change optimal info choice in more general contexts
- → Sources with more messages, multiple states, other utility functions

Strategic interaction leads to rich predictions ▸ More

- → Game of news choice, where both players want to learn
- → Rationalize (mis)coordinating, (mis)matching news choices..

Thanks for listening!

Related Literature

Individual learning from biased sources for a Bayesian agent:

- **Optimal to consume an own-biased medium:** Calvert, 1985; Suen, 2004; Burke, 2008; Gentzkow and Shapiro, 2006; Meyer, 1991; Mullainathan and Shleifer, 2005; Oliveros and Várdy, 2015.
- **Optimal to multi-home and opposite bias learning:** Che and Mierendorff, 2019; Nikandrova and Pancs, 2018; Mayskaya, 2020; Liang et al., 2022.

(Bayesian) learning from others: Sethi and Yildiz, 2016; Sethi and Yildiz, 2012; Bowen et al., 2021;...

Exogenous manipulation of beliefs:

- **Manipulation of attention:** Gossner et al., 2021; Liang et al., 2022.
- **Persuasion with unknown beliefs:** Dworczak and Pavan, 2022; Laclau, Renou, et al., 2017; Kolotilin et al., 2017.

Complementarity btw signals: Börgers et al., 2013; Brooks et al., 2022

Thresholds of the Optimal Strategy: Likelihood Ratios

$$p_1 = 1 - \frac{\mathbb{P}_L(r|R)\mathbb{P}_E(r|R)}{\mathbb{P}_L(r|L)\mathbb{P}_E(r|L) + \mathbb{P}_L(r|R)\mathbb{P}_E(r|R)}$$

$$p_2 = 1 - \frac{\mathbb{P}_E(r|R)}{\mathbb{P}_E(r|L) + \mathbb{P}_E(r|R)}$$

$$p_3 = \frac{\mathbb{P}_L(r|L)\mathbb{P}_E(l|L) + \mathbb{P}_R(l|L)\mathbb{P}_E(r|L)}{\mathbb{P}_L(r|L)\mathbb{P}_E(l|L) + \mathbb{P}_R(l|L)\mathbb{P}_E(r|L) + \mathbb{P}_L(r|R)\mathbb{P}_E(l|R) + \mathbb{P}_R(l|R)\mathbb{P}_E(r|R)}$$

$$p_4 = \frac{\mathbb{P}_E(l|L)}{\mathbb{P}_E(l|L) + \mathbb{P}_E(l|R)}$$

$$p_5 = \frac{\mathbb{P}_R(l|L)\mathbb{P}_E(l|L)}{\mathbb{P}_R(l|L)\mathbb{P}_E(l|L) + \mathbb{P}_R(l|R)\mathbb{P}_E(l|R)}$$

Back

How does this translate into equilibria?

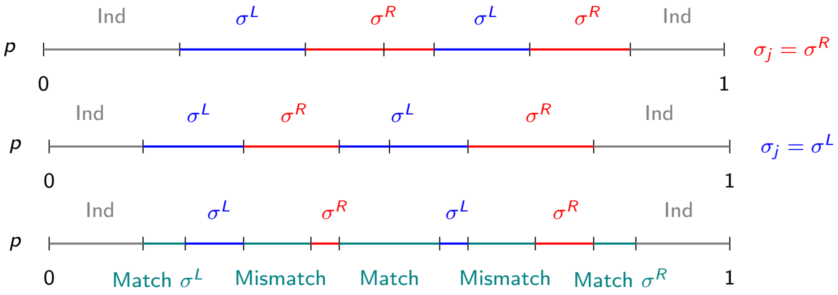

Many different types of equilibria

→ **Dominance-solvable** e.g. dominant strategy and best-reply

→ **Coordination:** both want to match or mismatch

→ **Miscoordination:** one wants to match, one to mismatch

Back

General Importance of Additional Information

Often there exists σ^e that changes the DM's optimal choice of info

$\rightarrow$ If no source of info is better than the other for all priors (in isolation)

$\rightarrow$ Some independent signal can change the DM's optimal info choice

General Importance of Additional Information

Often there exists σ^e that changes the DM's optimal choice of info

$\rightarrow$ If no source of info is better than the other for all priors (in isolation)

$\rightarrow$ Some independent signal can change the DM's optimal info choice

Proposition

For any two sources σ^x, σ^y, and any finite, discrete action, utility space:
If there exists at least one belief, $\hat{p}$, s.t. $EU(\sigma^x|\hat{p}) > EU(\sigma^y|\hat{p})$,
then for all beliefs p there exists $\sigma_e,\ s.t.\ EU(\sigma^x, \sigma_e|p) > EU(\sigma^y, \sigma_e|p)$

Back

Three Types of Optimal Strategy

If σ_E is substantially more informative than σ_L, σ_R

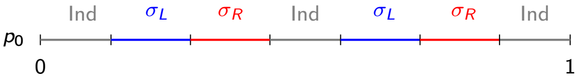

Three Types of Optimal Strategy

If σ_E is substantially more informative than σ_L, σ_R

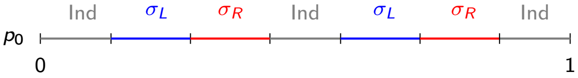

If σ_E is substantially less informative than σ_L, σ_R

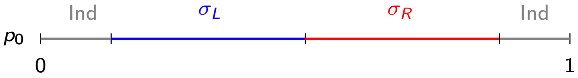

Three Types of Optimal Strategy

If σ_E is substantially more informative than σ_L, σ_R

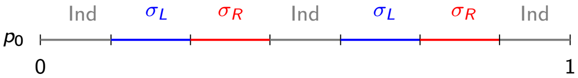

If σ_E is substantially less informative than σ_L, σ_R

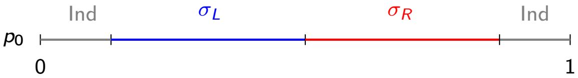

If σ_E is equally or relatively more informative than σ_L, σ_R

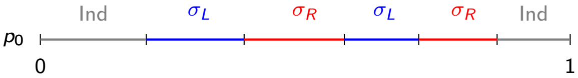

Back