Dynamic Non-Bayesian Persuasion

Masanori Kobayashi | UCLA Economics
masakobayashi@g.ucla.edu

1 Research question and motivation

Standard Bayesian persuasion: The sender sends a signal only once.

Real-life applications (e.g., ads):

- Information often arrives sequentially
- Receivers may update in a non-Bayesian way

With a non-Bayesian receiver,

1. **Can the sender gain by splitting information into multiple stages?**
2. **If so, when does timing matter?**

2 Environment and updating

Persuasion environment: (Θ, p, A, u, v)

- Θ: finite state space
- p: prior over Θ
- A: compact action space
- $u, v : A \times \Theta \to \mathbb{R}$: sender and receiver utilities

Updating rules

- Sender: Bayesian
- Receiver: **systematic distortion**

An updating rule is a systematic distortion if there exists a family of prior-dependent functions $D := (D_p)_{p \in \Delta(\Theta)}$, called a **distortion rule**, such that the receiver's posterior is

$$\boxed{D_p(q), \quad \text{where } q = \mu_B(\sigma, p)(\cdot \mid s)},$$

Here, $\mu_B(\sigma, p)(\cdot \mid s)$ is the Bayesian posterior induced by experiment $\sigma : \Theta \to \Delta(S)$, prior p, and signal $s \in S$.

3 Static and dynamic persuasion

The receiver acts only at the end.

STATIC: one update

$$\sigma, s \longrightarrow \boxed{D_p(q)} \longrightarrow a.$$

DYNAMIC: update twice

$$\sigma^1, s_1 \to \boxed{D_p(q)}$$

$$\to \sigma^2_{s_1}, s_2 \to \boxed{D_{D_p(q)}(\mu_B(\sigma^2_{s_1}, D_p(q))(\cdot \mid s_2))} \to a.$$

4 Main result

The key property: divisibility

SIMULTANEOUS

process (s_1, s_2) jointly
apply the rule once
final belief r^{sim}

DIVISIBLE
if $r^{\mathrm{sim}} = r^{\mathrm{seq}}$

SEQUENTIAL

update after s_1
use the distorted belief as the next prior
update after s_2
final belief r^{seq}

THEOREM

An updating rule is divisible

$$\Longleftrightarrow$$

Static and dynamic persuasion give the sender the same value in every environment.

Regularity conditions: D_p is a continuous injection, does not shrink posterior support, and its response to fully revealing signals is prior-independent.

Why?

- If the rule is divisible, the sender can always replicate the sequential outcome with a single experiment.
- If the rule is not divisible, one can always construct an environment in which the sender strictly prefers sequential persuasion by a separating-hyperplane argument.

5 A revealing example

Grether's (1980) α–β rule weights base rates and likelihoods geometrically:

$$\mu(\theta \mid s) = \frac{p(\theta)^\alpha \sigma_\theta(s)^\beta}{\sum_{\theta'} p(\theta')^\alpha \sigma_{\theta'}(s)^\beta}.$$

Here α controls the weight on the prior (base rates), while β controls the weight on new signal likelihoods. Bayes' rule is $\alpha = \beta = 1$.

Divisible if and only if $\alpha = 1$.

6 Why dynamics can help

Judge–prosecutor example: $\Theta = A = \{0, 1\}$, $p = 1/3$, $u(a, \theta) = \mathbf{1}\{a = \theta\}$, $v(a, \theta) = \mathbf{1}\{a = 1\}$

- The judge convicts ($a = 1$) whenever the posterior is at least 1/2.

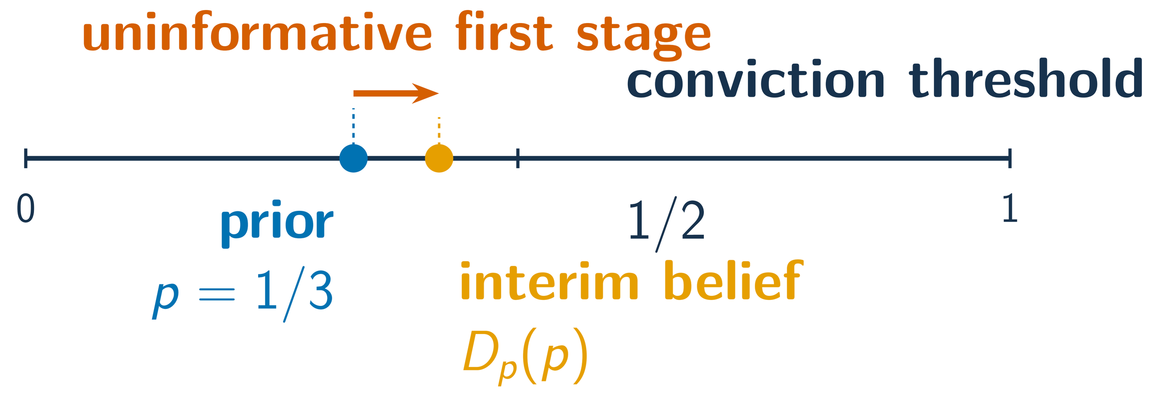

With **base-rate neglect** ($\alpha < 1$), even an uninformative first stage moves beliefs toward 1/2. The second persuasion stage therefore starts closer to the sender's target:

$$\boxed{V^{\mathrm{dynamic}} > V^{\mathrm{static}}.}$$

7 Take-away

- The static persuasion model is tractable, and it is without loss precisely when the updating rule is divisible.
- If an analyst wants to study biases represented by updating rules that fail divisibility, dynamic persuasion needs to be considered.
- **Example:** Recency bias modeled as base-rate neglect ($\alpha < 1$) in Grether's (1980) α–β rule.

Grether (1980); Kamenica–Gentzkow (2011); Cripps (2018); de Clippel–Zhang (2022).

Research page