Dynamic Non-Bayesian Persuasion

Masa Kobayashi

August 9, 2026

University of California, Los Angeles

Motivation

Canonical Bayesian Persuasion Model:

- Sender designs and commits to an experiment to influence the receiver's action.
- Applications include advertising (Rayo and Segal, 2010), litigation (Kamenica and Gentzkow, 2011), and others.

Question:

What if the receiver is non-Bayesian/misspecified-Bayesian?

- Evidence: Bhuller and Sigstad (2025), Benjamin (2019), etc.
- Can we simply replace the Bayesian receiver with a non-Bayesian one?
- **Concern**: Real information provision is often sequential rather than all at once.

Results Preview

1. When the receiver is non-Bayesian, the sender's value can depend on whether information is provided all at once or sequentially.
    - One-shot/static persuasion is not generally without loss.
2. The sender is indifferent between static and dynamic persuasion in any environment if and only if the receiver's updating rule is **divisible** (Cripps, 2018).

divisible: given multiple signals, the following two procedures give the same posterior:

1. updating the prior using all signals simultaneously;
2. updating the prior using each signal sequentially.

Model: Environment

Persuasion environment: $\langle \Theta, p, A, u, v \rangle$

- Θ: finite state space
  - p: full-support prior over Θ
- A: compact metrizable action space
- $u, v : A \times \Theta \to \mathbb{R}$: receiver's and sender's utility functions
  - u and v are continuous.

Model: Experiment and Updating Rule

$\sigma = \langle S, (\sigma_\theta)_{\theta \in \Theta} \rangle$: a Blackwell experiment

- S: finite set of signals (fixed)
- $\sigma_\theta \in \Delta(S)$: distribution over S given state θ

Σ: set of all experiments with the fixed signal space S

$\mu : \Sigma \times \Delta(\Theta) \to \Delta(\Theta)^S$: the receiver's updating rule

- μ_B: Bayesian updating.

Assume the sender is Bayesian, but the receiver may not be.

Systematic Distortions

- $\Theta_r := \{\theta \in \Theta : r(\theta) > 0\}$: support of belief r.

Definition (De Clippel and Zhang, 2022)

An updating rule μ is a ***systematic distortion*** *if there exists a* ***distortion rule*** *$D := (D_p)_{p \in \Delta(\Theta)}$ such that $D_p : \Delta(\Theta_p) \to \Delta(\Theta)$ and, for any experiment σ and any signal realization s satisfying $\sum_{\theta \in \Theta} \sigma_\theta(s) p(\theta) > 0$,*

$$\mu(\sigma, p)(\cdot \mid s) = D_p\left(\mu_B(\sigma, p)(\cdot \mid s)\right).$$

Assume the receiver's updating rule is a systematic distortion.

Example: Grether (1980)'s α-β Model

$$\mu(\sigma, p)(\theta \mid s) = \frac{p(\theta)^{\alpha} \sigma_\theta(s)^{\beta}}{\sum_{\theta' \in \Theta} p(\theta')^{\alpha} \sigma_{\theta'}(s)^{\beta}}$$

- $\alpha \in (0, 1]$: $\alpha < 1$ corresponds to base-rate neglect.
- $\beta \in (0, 1]$: $\beta < 1$ corresponds to underreaction to new information.

Distortion rule:

$$D_p(q)(\theta) = \begin{cases} \frac{p(\theta)^{\alpha-\beta} q(\theta)^{\beta}}{\sum_{\theta' \in \Theta_p} p(\theta')^{\alpha-\beta} q(\theta')^{\beta}} & \text{if } \theta \in \Theta_p, \\ 0 & \text{if } \theta \notin \Theta_p, \end{cases}$$

where $q \in \Delta(\Theta_p)$ is the Bayesian posterior.

Static Persuasion: Timeline

1. The sender chooses an experiment σ.
2. After observing σ and the realized signal s, the receiver forms a posterior belief $D_p(\mu_B(\sigma, p)(\cdot \mid s))$.
3. The receiver takes an action $a \in A$ to maximize her expected utility given the posterior belief.

Solution (De Clippel and Zhang, 2022): concavification of the modified indirect utility.

Dynamic Persuasion: Timeline

Parsimonious two-step dynamic persuasion model:

1. The sender chooses a first-step experiment σ^1 and a collection of second-step experiments $\{\sigma^2_{s_1}\}_{s_1 \in S}$.
   - $\sigma^2_{s_1}$ is used after the first-stage signal realization s_1
2. After observing σ^1 and its realization s_1, the receiver updates her belief from p to $D_p(q)$.
   - $q := \mu_B(\sigma^1, p)(\cdot \mid s_1)$
3. After observing $\sigma^2_{s_1}$ and its realization s_2, the receiver updates her belief from $D_p(q)$ to

$$\underbrace{D_{D_p(q)}}_{\text{distortion with “prior” } D_p(q)} \left(\underbrace{\mu_B\left(\sigma^2_{s_1}, D_p(q)\right)(\cdot \mid s_2)}_{\text{Bayesian posterior with “prior” } D_p(q)} \right).$$

4. The receiver chooses an action $a \in A$.

Issue: hard to solve in general.

Divisibility

Definition (Informal; Cripps, 2018)
An updating rule is ***divisible*** *if, for any prior and any sequence of signals, the following two procedures yield the same posterior belief:*

- *processing all signals simultaneously;*
- *processing the signals sequentially, using the intermediate posterior as the next prior.*

Main Result

Theorem
Suppose the receiver's updating rule is a systematic distortion that satisfies regularity conditions. Then, TFAE:

1. *The receiver's updating rule is divisible.*
2. *The sender is indifferent between static and dynamic persuasion in every persuasion environment.*

▸ regularity conditions

Application: α-β Model

Grether's α-β rule: with $\alpha, \beta \in (0, 1]$,

$$\mu(\sigma, p)(\theta \mid s) = \frac{p(\theta)^{\alpha} \sigma_\theta(s)^{\beta}}{\sum_{\theta' \in \Theta} p(\theta')^{\alpha} \sigma_{\theta'}(s)^{\beta}}$$

Corollary
Suppose the receiver's belief updating rule is given by Grether's α-β rule. Then the sender is indifferent between static and dynamic persuasion in every environment if and only if $\alpha = 1$.

- The α-β rule is divisible iff $\alpha = 1$

Intuition: judge–prosecutor example in KG

- $\Theta = \{0, 1\}$
  - 0: defendant is innocent, 1: defendant is guilty
  - $p := Pr(\theta = 1) < 1/2$
- $A = \{0, 1\}$
  - 0: acquit, 1: convict
- $u(a, \theta) = \mathbb{1}\{a = \theta\}$
- $v(a, \theta) = a$

Threshold receiver belief: $q' = 1/2$

- $\alpha < 1 \implies$ the receiver's belief drifts towards 1/2 even with no info in the first step.

Take-away

Analyst's modeling tradeoff

- Focusing on geometric underreaction ($\alpha = 1$, $\beta \in (0, 1)$) can justify static persuasion models.
- Base-rate neglect ($\alpha < 1$) makes the timing of information provision payoff-relevant, so one-shot persuasion is generally not without loss.

Contribution and Related Literature

Static Bayesian and non-Bayesian persuasion

- Kamenica and Gentzkow (2011); Rayo and Segal (2010); De Clippel and Zhang (2022)
- **Contribution**: clarifying when the model is WLOG

Concurrent work on dynamic non-Bayesian persuasion

- Azrieli and Das (2025)
    - Receiver with linear underreaction $\implies$ dynamic persuasion can $\uparrow$ sender's payoff
- **Contribution**: complete characterization of updating rules

Conclusion

- Tractable static persuasion framework is without loss if and only if the receiver's updating rule is divisible.
- Non-divisible updating rules require explicit modeling of the timing of information provision, but it is hard to solve in general.

Regular Distortions

Definition

*A distortion rule D is **regular** if it satisfies:*

1. ***Non-decreasing support**: $\Theta_q \subseteq \Theta_{D_p(q)}$ for all $p \in \Delta^\circ(\Theta)$ and $q \in \Delta(\Theta)$.*
2. ***Continuity**: The map $(p, q) \mapsto D_p(q)$ from $\Delta^\circ(\Theta) \times \Delta(\Theta)$ to $\Delta(\Theta)$ is continuous.*
3. ***Injectivity**: D_p is one-to-one for every $p \in \Delta^\circ(\Theta)$.*
4. ***Prior independence under certainty**: For any full-support beliefs p, q,*

$$D_p(\delta_\theta) = D_{D_p(q)}(\delta_\theta) \quad \text{for all } \theta \in \Theta_{D_p(q)}.$$

▸ Return to Theorem

References

Azrieli, Y. and Das, R. (2025). Sequential non-bayesian persuasion. *arXiv preprint arXiv:2508.09464*.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:69–186.

Bhuller, M. and Sigstad, H. (2025). Feedback and learning: The causal effects of reversals on judicial decision-making. *Review of Economic Studies*, 92(4):2359–2397.

Cripps, M. W. (2018). Divisible updating. *Manuscript, UCL*.

De Clippel, G. and Zhang, X. (2022). Non-bayesian persuasion. *Journal of Political Economy*, 130(10):2594–2642.

Grether, D. M. (1980). Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly journal of economics*, 95(3):537–557.

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

Rayo, L. and Segal, I. (2010). Optimal information disclosure. *Journal of political Economy*, 118(5):949–987.