Calibrated Forecasting and Persuasion

Atulya Jain [1] Vianney Perchet [2]

[1]HEC Paris

[2]ENSAE, Criteo AI Lab

Motivation

- Probability forecasts are used widely to provide information.
- Decision makers (DMs) follow recommendation if they are credible.
- Use statistical tests (or learning algorithms) to determine credibility.
- Strategic interaction between rational agent and AI agent.

Calibration test

- The expert passes the calibration test, if for any forecast $f \in [0,1]$, the realized proportions of rainy days when she announced f is close to f.

$$|f - \frac{\sum_{t=1}^{T} \mathbf{1}_{\{f_t=f\}}\omega_t}{\sum_{t=1}^{T} \mathbf{1}_{\{f_t=f\}}}| \leq \epsilon_T \quad \forall f \qquad (Finite)$$

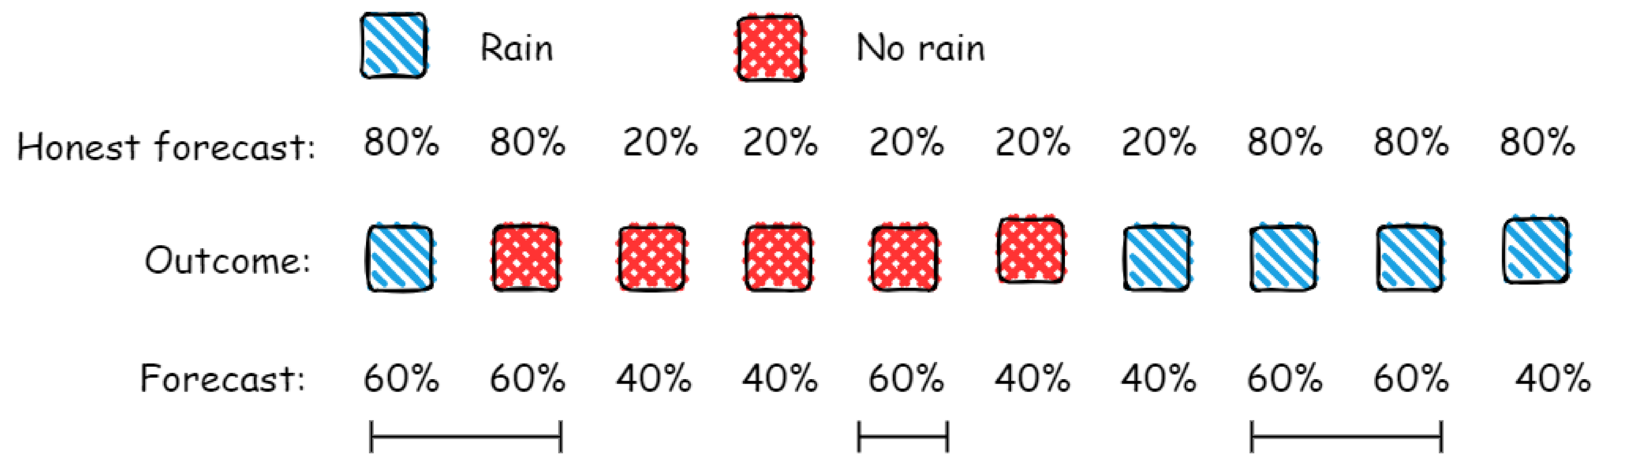

$$f = \lim_{T\to\infty} \frac{\sum_{t=1}^{T} \mathbf{1}_{\{f_t=f\}}\omega_t}{\sum_{t=1}^{T} \mathbf{1}_{\{f_t=f\}}} \quad \forall f \qquad (\star)$$

Forecast = Realized Outcomes

Dynamic forecasting game

- Expert knows the stochastic process: $\mu \in \Delta\{0,1\}^{\infty}$:
- At stage t, expert knows true probability p_t and sends forecast f_t.
- The DM performs the calibration test.
  - Pass: play acc. to forecast: $a_t = \hat{a}(f_t)$, where $\hat{a}(f_t)$: DM's best action given belief f_t.
  - Fail: play punishment action that results in loss $c >> 0$ for expert.
- Outcome ω_t is observed and players obtain payoff $U_E(\omega_t, a_t)$ and $U_{DM}(\omega_t, a_t)$.
- Expert's Goal: Maximize long-run average payoff while passing the calibration test:

$$\liminf_{T\to\infty} \frac{\sum_{t=1}^{T} U_E(\omega_t, a_t)}{T}$$

such that $(\star)$ holds.

Persuasion problem

- States: $\Omega = [0,1]$ and prior distribution of states: $P = (\lambda, p)$
- Expert commits to a signaling policy $G : \Omega \to \Delta S$.
- After signal realization $s \in S$, DM has posterior mean $q_s \in \Delta\Omega$ and plays action $\hat{a}(q_s)$.
- Expert's utility from belief q : $\hat{u}(q) = \sum_{\omega\in\{0,1\}} q(\omega) U_E(\omega, \hat{a}(q))$.
- $Q = (\mu, q)$ is a garbling of $P = (\lambda, p)$ if $\exists$ stochastic matrix G:

$$\lambda G = \mu$$

$$(\lambda p) G = \mu q$$

- Expert's Goal: Find distribution of posteriors that maximizes expected utility:

$$Per(P, \hat{u}) = \max_{Q\in G(P)} \sum_{q\in Supp(Q)} \mu(q)\hat{u}(q).$$

Garbling of honest forecasts

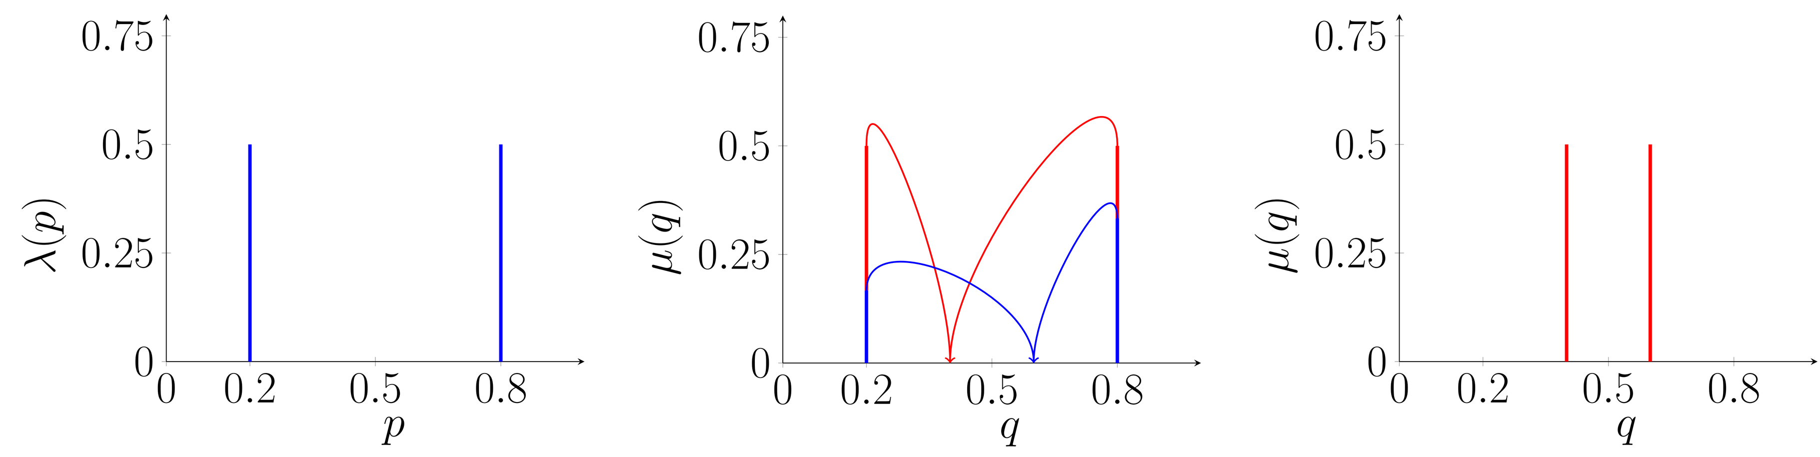

$P = (\lambda, p)$: honest forecasts $\quad G$: garbling $\quad Q = (\mu, q)$: forecasts

Results

Long-run average payoff an expert can attain:

Expert ↓ / Process →	General	Stationary + Ergodic
Informed	$\mathbb{E}_P[\hat{u}]$	$Cav\ \hat{u}\|_C(p)$
Uninformed	$\overline{Co}\ \hat{u}(p)$	$\hat{u}(p)$

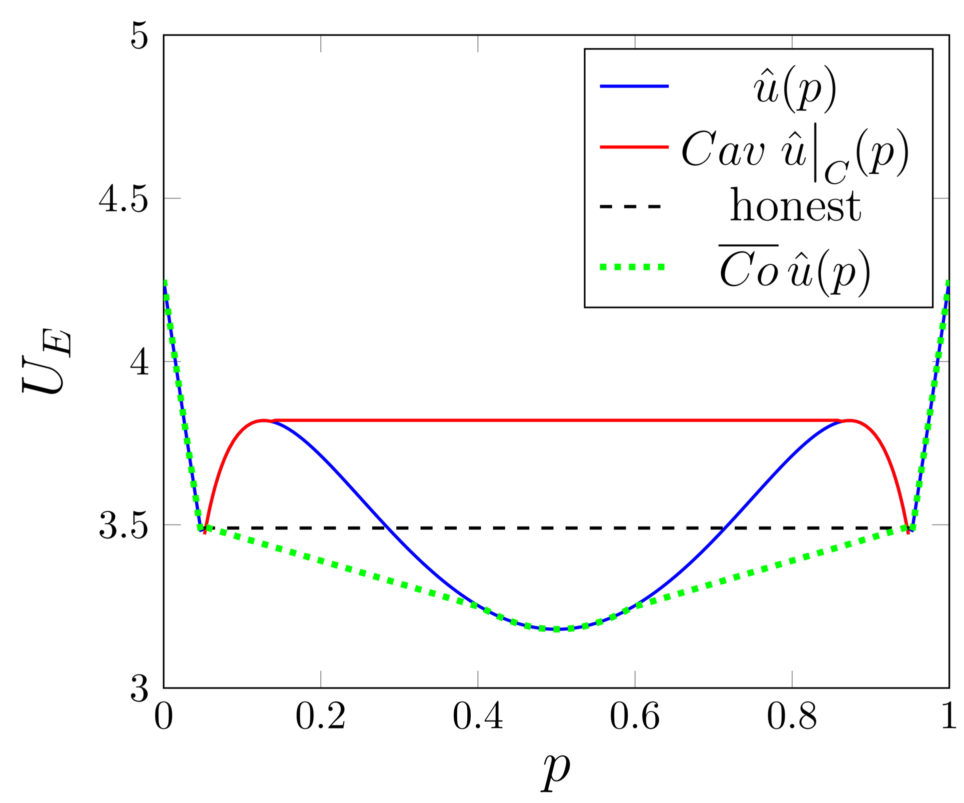

where, $C = [0.05, 0.95]$ is the interval of the honest forecasts and p is the mean probability of rain.

No-regret learning

- A DM has no regret with respect to forecast f if

$$\lim_{T\to\infty} \max_{a\in A} \frac{\sum_{t=1}^{T} \mathbf{1}_{\{f_t=f\}}\big(u_{DM}(\omega_t, a) - u_{DM}(\omega_t, a_t)\big)}{\sum_{t=1}^{T} \mathbf{1}_{\{f_t=f\}}} \leq 0.$$

- The DM's regret measures the difference in payoff he could have gotten and what he got.

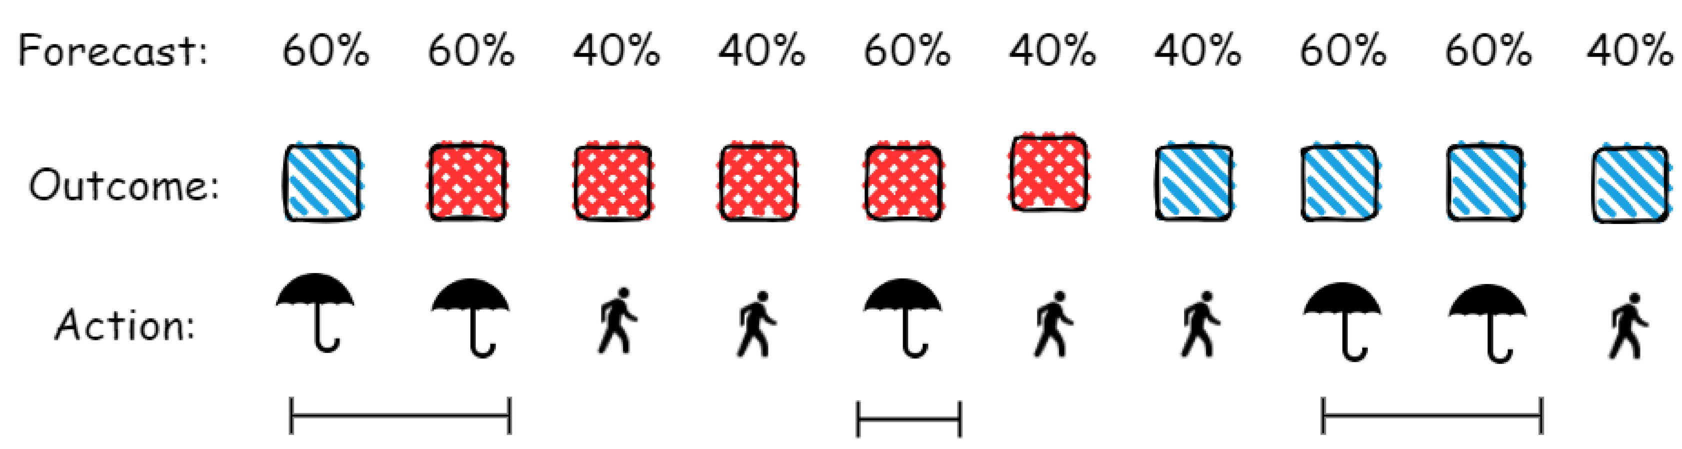

- DM has no regret if he plays acc. to any calibrated forecasting strategy.
- When facing a no-regret learner, expert can guarantee calibration benchmark and in some instances strictly more.

Summary

- We show there is scope for strategic forecasting. Overall, the forecasts need to be accurate but can be less precise than honest forecasts.
- We provide a micro-foundation for the commitment assumption in persuasion models.
- We show a novel connection between calibration test and no-regret learning as heuristics in decision-making.

This research was partially funded by Hi! PARIS Center on Data Analytics and Artificial Intelligence