Binary Scoring Rules that Incentivize Precision

Eric Neyman George Noarov Matt Weinberg

Motivation

Consider forecasting the weather via a sequence of computationally costly weather simulations. Forecaster/*expert* runs simulations and reports back the forecast for tomorrow. We want to devise a reward scheme that:

- Rewards forecaster, by comparing forecast and realized weather;
- Incentivizes honest forecasts — i.e. elicits forecaster's true belief;
- Incentivizes the forecaster to apply *maximum* effort — i.e. to run the simulation many times in order to obtain a *precise* forecast.

We consider reward schemes based on *proper scoring rules*: a classic way to reward the forecaster's prediction while enforcing truthfulness. However,

Not all scoring rules equally encourage the forecaster's precision.

We build a framework for ranking proper scoring rules by their incentivization properties, and find explicit optimally incentivizing proper scores.

Model

Reward scheme

Consider predicting p, the probability of rain tomorrow. Today, p is drawn uniformly from $[0,1]$. *Expert* gets a coin with bias p, each flip priced at $c > 0$. Today, she flips the coin as often as desired, and submits forecast $q \in [0,1]$.

Tomorrow, a *symmetric binary scoring rule* $s : (0,1) \to \mathbb{R}$ rewards the expert's prediction q: The expert receives $s(q)$ if it rains, and $s(1-q)$ if it does not rain.

The scoring rule s is required to be *proper* and *normalized*.

- s is *proper*, that is, expert is incentivized to report honestly:

  For all p, expert's expected reward $p\,s(q)+(1-p)\,s(q)$ is maximized at $q = p$.
- s is *normalized*, i.e. satisfies two conditions:
  - A completely uninformed expert gets reward 0 — that is, $s(\frac{1}{2}) = 0$.
  - A perfect expert gets expected reward 1 — that is, $\int_0^1 (x\,s(x) + (1-x)\,s(1-x))\,dx = 1$.

Information acquisition

- The expert is Bayesian, and starts off with a uniform prior on p.
- Expert's initial prediction is $q_0 = \frac{1}{2}$, the mean of her prior.
- After each sample, she updates her belief and forms prediction via Laplace Rule of Succession: if n flips and h heads, prediction is $q_n = \frac{h+1}{n+2}$.
- Expert dynamically reevaluates whether to keep sampling after each flip.

Decision-making: Locally and globally adaptive experts

- *Locally adaptive* expert myopically stops flipping as soon as per-flip cost c exceeds ex-ante expected reward gain from flipping one more time.
- *Globally adaptive* expert keeps flipping until it is not part of her globally optimal strategy for the future.

Main Result: An Incentivization Index

Consider a locally or globally adaptive expert. Consider any per-flip cost $c > 0$, and a proper normalized symmetric binary scoring rule s.

Choose $\ell \geq 1$, and let $\mathrm{Error}_c^\ell(s) := \mathbb{E}\left[|p-q|^\ell\right]$ be the expected ℓ^{th} power error of expert's prediction q when she is rewarded via scoring rule s and pays c per flip. The expectation is over randomness in p and expert's flips.

Let $R(x) := x\,s(x) + (1-x)\,s(1-x)$ be the ex-ante expected reward for truthful predictions x.

Define the ℓ^{th}-order **Incentivization Index** of scoring rule s by

$$\mathrm{Ind}^\ell(s) := \int_0^1 \left(\frac{x(1-x)}{R''(x)}\right)^{\frac{\ell}{4}} dx.$$

The Incentivization Index **characterizes** the expected prediction error:

$$\lim_{c \to 0} c^{-\ell/4} \cdot \mathrm{Error}_c^\ell(s) = \mu_\ell \cdot 2^{\ell/4} \cdot \mathrm{Ind}^\ell(s),$$

where μ_ℓ is the ℓ^{th} moment of a standard Gaussian.

Interpretation:

For two scoring rules s_1, s_2, if $\mathrm{Ind}^\ell(s_1) < \mathrm{Ind}^\ell(s_2)$ then s_1 gives *better* expected ℓ^{th} power error than s_2, for all small enough costs c.

Simulation results:

Even though the above characterization is shown asymptotically as $c \to 0$, our simulations show that the Incentivization Index is trustworthy even for nonvanishing/"practical" values of c.

Intuition behind the Index

Suppose a locally adaptive expert has made n flips and gotten h heads. Her expected ex-ante payoff gain from flipping once more is:

$$\underbrace{\frac{h+1}{n+2} R\left(\frac{h+2}{n+3}\right) + \frac{n-h+1}{n+2} R\left(\frac{h+1}{n+3}\right)}_{\text{expected reward after another flip}} - \underbrace{R\left(\frac{h+1}{n+2}\right)}_{\text{exp. reward now}} \approx R''(p) \cdot \frac{p(1-p)}{2n^2},$$

via Taylor for c small enough that she flips long enough until $\frac{h+1}{n+2} \approx p$.

She will flip while $R''(p) \cdot \frac{p(1-p)}{2n^2} \gtrsim c$, thus making $n \approx \sqrt{\frac{p(1-p)R''(p)}{2c}}$ flips.

Her final estimate $q \sim$ Bin $\left(p, \sqrt{\frac{p(1-p)R''(p)}{2c}}\right) \approx N\left(p, \sqrt[4]{\frac{2p(1-p)c}{R''(p)}}\right)$. So:

$$\mathrm{Error}_c^2(s) \approx \int_0^1 \mathrm{Var}\left(N\left(p, \sqrt[4]{\frac{2p(1-p)c}{R''(p)}}\right)\right) dp = \int_0^1 \sqrt{\frac{2p(1-p)c}{R''(p)}} dp = \sqrt{c} \cdot \mathrm{Ind}^2(s).$$

Optimal Scoring Rules: Closed Form and Performance

For every $\ell \geq 1$, the unique (up to normalization) scoring rule that optimizes the Incentivization Index has a closed form! For $\ell \geq 1$, the optimal score is:

$$s_{\ell,\mathsf{OPT}}(x) := \begin{cases} \kappa_\ell \int_{1/2}^x \left(t^{\ell-8}(1-t)^{2\ell+4}\right)^{1/(\ell+4)} dt, & \text{if } x \leq 1/2, \\ \kappa_\ell \int_{1/2}^x \left(t^{\ell}(1-t)^{2\ell-4}\right)^{1/(\ell+4)} dt, & \text{if } x \geq 1/2. \end{cases}$$

Here, κ_ℓ is a normalization constant. Some special cases: for $\ell = 2$, we have for $x \in [1/2, 1)$ that

$$s_{2,\mathsf{OPT}}(x) = \frac{3}{5}\kappa_2 \left(x^{5/3} - 0.5^{5/3}\right),$$

and as $\ell \to \infty$, the optimal rule pointwise converges, for $x \in (0,1)$, to

$$s_{\infty,\mathsf{OPT}}(x) := \frac{5}{9}(48x^4 - 128x^3 + 96x^2 - 11).$$

How good are classical scores (log $s_{\log}(x) = \ln x$, Brier $s_{\text{quad}}(x) = -(1-x)^2$, spherical $s_{\text{sph}}(x) = -x/\sqrt{x^2+(1-x)^2}$, "hs" $s_{\mathsf{hs}}(x) = -\sqrt{(1-x)/x}$) at incentivizing precision?

Scoring Rule	1	2	4	8	16	32	64	128	256	512
	Value of ℓ									
hs	0.990	0.997	0.992	0.979	0.962	0.947	0.935	0.927	0.922	0.919
Logarithmic	0.970	0.990	0.998	0.993	0.982	0.969	0.959	0.951	0.946	0.943
Quadratic	0.905	0.946	0.979	0.996	0.999	0.995	0.989	0.984	0.980	0.978
Spherical	0.853	0.899	0.940	0.968	0.984	0.992	0.995	0.995	0.995	0.994
OPT (l = 1)	1.000	0.993	0.971	0.938	0.905	0.877	0.856	0.842	0.833	0.827
OPT (l = 2)	0.992	1.000	0.991	0.969	0.941	0.915	0.896	0.882	0.873	0.868
OPT (l = 4)	0.966	0.991	1.000	0.992	0.973	0.953	0.936	0.924	0.916	0.910
OPT (l = 8)	0.925	0.964	0.991	1.000	0.994	0.981	0.969	0.958	0.951	0.946
OPT (l = 16)	0.885	0.931	0.971	0.994	1.000	0.996	0.989	0.981	0.976	0.972
OPT (l = 32)	0.854	0.903	0.949	0.980	0.996	1.000	0.998	0.994	0.990	0.987
OPT (l = 64)	0.835	0.885	0.932	0.967	0.988	0.998	1.000	0.999	0.997	0.995
OPT (l = 128)	0.824	0.874	0.921	0.958	0.981	0.994	0.999	1.000	0.999	0.998
OPT (l = 256)	0.818	0.868	0.915	0.952	0.976	0.990	0.997	0.999	1.000	1.000
OPT (l = 512)	0.815	0.864	0.912	0.949	0.973	0.987	0.995	0.998	1.000	1.000
OPT (l -> Infinity)	0.812	0.861	0.908	0.945	0.970	0.984	0.992	0.996	0.998	0.999

Figure 1: For each score s and power ℓ, this table shows normalized ratio $\sqrt[\ell]{\mathrm{Ind}^\ell(s_{\ell,\mathsf{OPT}})/\mathrm{Ind}^\ell(s)}$.

Future Directions

- *Extension to non-uniform priors.* What if a-priori, the probability of rain tomorrow is non-uniform?
- *Extension to other metrics of precision.* We aim to minimize the ℓ^{th} power distance between prediction q and true probability p. What about optimizing the Bregman divergence between p and q?
- *Other models of costly information acquisition.*
- *Other structures of effort levels.* In our work, forecaster has countably many effort levels.