Truthful Elicitation of Imprecise Forecasts

Anurag Singh[1] **Siu Lun Chau**[2] **Krikamol Muandet**[1]

[1]CISPA Helmholtz Center for Information Security, Germany [2]Nanyang Technological University, Singapore

Summary. Truthful elicitation of imprecise probability (IP) is a long-standing open problem, with prior work showing the impossibility of a deterministic scoring rule for truthful elicitation of credal sets. This poster presents our recent attempt [1], where we propose a novel *randomised* strictly proper IP scoring rule. We argue that additional communication of the choice of aggregation function between forecasters and decision-makers (DMs) can enable credal set elicitation. However, this communication allows for DM's manipulation, jeopardizing truthfulness. To allow truthful elicitation, we propose that the DM shares a *distribution* over aggregations. This table shows the current taxonomy of elicitation in IP.

Forecaster's Belief:	precise	imprecise (credal set)		
Communication:	N/A	N/A	ρ	$p(\rho)$
Scoring Rule:	Strictly-Proper	Impossible	Proper	Strictly-Proper

Setup. In our work, we consider an agent selecting an input x from a finite set $\mathcal{X} = \{x_1, \dots, x_n\}$, where the utility $u(x, o)$ depends on the chosen input x and outcome $o \in \mathcal{O}$. For a forecaster, $\mathcal{X} := \Delta(\mathcal{O})$ or $2^{\Delta(\mathcal{O})}$ and $u(x, o) := s(p, o)$, highlighting the decision-making aspect of elicitation. For the DM, $\mathcal{X} := \mathcal{A} = \{a_1, \dots, a_m\}$, i.e. the action space and $u(x, o) := u_{\mathrm{DM}}(a, o)$.

Definition 1 (Aggregation Function)**.** Given a credal set $\mathcal{Q}$, an aggregation function ρ is a map that combines utilities via a positive linear combination $\rho\left(\{\mathbb{E}_q[u(x,o)]\}_{q\in\mathcal{Q}}\right) = \int_{q\in\mathcal{Q}} \boldsymbol{w}(q)\mathbb{E}_q[u(x,o)]\,dq$ where $\boldsymbol{w}(q) \geq 0$ for all $q \in \mathcal{Q}$.

This allows us to parameterise the IP scoring rule $s_\rho : 2^{\Delta(\mathcal{O})} \to \mathbb{R}$ with the aggregation rule ρ, which we implement by generalising the tailored scoring rules to accommodate the aggregation rule, i.e., $s_\rho(\mathcal{Q}, o) = ku(a^*_{\rho,\mathcal{Q}}, o) + c$ where $k, c \in \mathbb{R}_{\geq 0}$ and $a^*_{\rho,\mathcal{Q}} = \arg\max_{a\in\mathcal{A}} \rho(\{\mathbb{E}_q[u_{DM}(a,o)]\}_{q\in\mathcal{Q}})$

Connection to Social Choice Theory. When interpreting IP as a "collective" report of precise probabilities, a social choice perspective naturally emerges for the downstream DM. This perspective applies even to a single-agent forecaster. Following Arrow's axiomatisation, we outline three desirable properties of any aggregation ρ: Pareto Efficiency (PE), Independence from Irrelevant Alternatives (IIA), and Non-Dictatorship (ND). This allows us to interpret the prior impossibility results on truthful elicitation of credal sets with existing impossibility results in social choice. Mainly Arrow's impossibility results, and to relate the truthful elicitation of credal sets to results on strategy-proof voting systems.

Figure 1. *Reporting the true belief uniquely maximizes the expected score.*

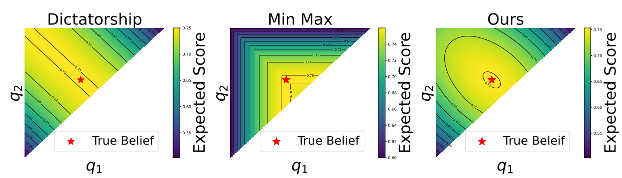

Characterisation of the strictly proper scoring rule. The DM shares a distribution $p(\rho)$ for truthful elicitation. Then the expected utility for forecast $\mathcal{Q}$ with belief $\mathcal{P}$ is

$$V^{\mathcal{P}}_\theta(\mathcal{Q}) := \mathbb{E}_{\rho\sim\theta}[V^{\mathcal{P}}_\rho(\mathcal{Q})] = \mathbb{E}_{\theta\sim p(\rho)}[\rho[\{\mathbb{E}_p[s_\rho(\mathcal{Q}, o)]\}_{p\in\mathcal{P}}]].$$

A scoring rule is said to be strictly proper if for all $\mathcal{P}$, $V^{\mathcal{P}}_\theta(\mathcal{P}) > V^{\mathcal{P}}_\theta(\mathcal{Q})$ for all $\mathcal{Q}$ such that $\mathcal{P} \neq \mathcal{Q}$.

Theorem 1. *(Strictly proper IP tailored scoring rules) An IP scoring rule s_θ is strictly proper if $p(\rho)$ is a distribution with full support for the class of linear aggregations of ρ. Then for any $k_\rho, c_\rho \in \mathbb{R}_{\geq 0}$ and an arbitrary function $\Pi : 2^{\Delta(\mathcal{O})} \to \mathbb{R}$, the score is defined as*

$$s_\theta(\mathcal{Q}, o)(\rho) = \begin{cases} k_\rho u_{DM}(a^*_{\rho,\mathcal{Q}}, o) + c_\rho & \text{if } p(\rho) > 0 \\ \Pi_o(\mathcal{Q}) & \text{if } p(\rho) = 0 \end{cases}.$$

Simulation. Figure 1 shows the results of simulations with a binary outcome (e.g., chance of rain tomorrow) for the true imprecise belief $[0.4, 0.6]$. We evaluate the scoring rule s_ρ with ρ being dictatorship or min-max. We compare these to our randomised strictly proper IP scoring rule s_θ. The forecaster reports an interval $\mathcal{Q} := [q_1, q_2]$ denoting the lower and upper probabilities. For our implementation, we consider $\mathcal{A} = [0, 1]$ and $u(a, o) := (o - a)^2$.

References

[1] Anurag Singh, Siu Lun Chau, and Krikamol Muandet. "Truthful Elicitation of Imprecise Forecasts". In: *arXiv preprint arXiv:2503.16395* (2025).