Title: The Counterfactual Scoring Rule: Designing Belief-Elicitation Questions (joint with Ruoyu Wang and Xinhan Zhang)

Abstract: Belief elicitation is widely used in experimental economics, but incentivized reports can distort subjects' primary decisions by interacting with main-task incentives. We characterize when belief statistics can be elicited without distortion in experiments with arbitrary task-payoff structures. We propose the Counterfactual Scoring Rule (CSR), a novel mechanism that nondistortionarily elicits any belief statistic by decomposing it into a component spanned by the task payoff and a residual component recovered through supplemental, action-independent statistics. The rank of the residual matrix serves as a tight question budget: the minimum number of elicitation questions required. When supplemental questions are unavailable, we fully characterize robust nondistortionary elicitation -- truthful reporting and task-optimal choice remain optimal under small misspecifications -- by a joint alignment condition: the target statistics are an affine transformation of the task payoff with a common, action-independent shift. The characterization builds on Kirchhoff's law applied to graphs induced by co-optimal actions, which are prone to locally fragile incentives.

You can also find an older version here:
https://arxiv.org/abs/2602.10474