From Proper Scoring Rules to Max-Min Optimal Forecast Aggregation

ERIC NEYMAN, Columbia University
TIM ROUGHGARDEN, Columbia University

This paper forges a strong connection between two seemingly unrelated forecasting problems: incentive-compatible forecast elicitation and forecast aggregation. Proper scoring rules are the well-known solution to the former problem. To each such rule s we associate a corresponding method of aggregation, mapping expert forecasts and expert weights to a "consensus forecast," which we call *quasi-arithmetic (QA) pooling* with respect to s. We justify this correspondence in several ways:

- QA pooling with respect to the two most well-studied scoring rules (quadratic and logarithmic) corresponds to the two most well-studied forecast aggregation methods (linear and logarithmic).
- Given a scoring rule s used for payment, a forecaster agent who sub-contracts several experts, paying them in proportion to their weights, is best off aggregating the experts' reports using QA pooling with respect to s, meaning this strategy maximizes its worst-case profit (over the possible outcomes).
- The score of an aggregator who uses QA pooling is concave in the experts' weights. As a consequence, online gradient descent can be used to learn appropriate expert weights from repeated experiments with low regret.
- The class of all QA pooling methods is characterized by a natural set of axioms (generalizing classical work by Kolmogorov on quasi-arithmetic means).

The full paper is available at https://arxiv.org/abs/2102.07081

ACM Reference Format:
Eric Neyman and Tim Roughgarden. 2021. From Proper Scoring Rules to Max-Min Optimal Forecast Aggregation. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC '21), July 18–23, 2021, Budapest, Hungary*. ACM, New York, NY, USA, 1 page. https://doi.org/10.1145/3465456.3467599

EC '21, July 18–23, 2021, Budapest, Hungary

ACM ISBN 978-1-4503-8554-1/21/07.
https://doi.org/10.1145/3465456.3467599