

Making and Evaluating Point Forecasts

Tilmann Gneiting
University of Heidelberg

Vortrag im Rahmen des Institutskolloquiums

8. Februar 2012, 17:15 Uhr

Seminarraum, Ludwigstraße 33 I

Single-valued point forecasts continue to be issued and used in almost all realms of science and society. Typically, competing point forecasters or forecasting procedures are compared and assessed by means of an error measure or scoring function, such as the absolute error or the squared error. However, this common practice can lead to grossly misguided inferences, unless the scoring function and the forecasting task are carefully matched.

Effective point forecasting requires that the scoring function be specified a priori, or that the forecaster receives a directive in the form of a statistical functional, such as the mean or a quantile of the predictive distribution. If the scoring function is specified a priori, the forecaster can issue an optimal point forecast, namely, the Bayes rule. If the forecaster receives a directive in the form of a functional, it is critical that the scoring function be consistent for it, in the sense that the expected score is minimized when following the directive.

A functional is elicitable if there exists a scoring function that is strictly consistent for it. Expectations, ratios of expectations and quantiles are elicitable. For example, a scoring function is consistent for the mean functional if and only if it is a Bregman function. It is consistent for a quantile if and only if it is generalized piecewise linear. Not all functionals are elicitable; for instance, conditional value-at-risk is not, despite its popularity in quantitative finance.