Some issues in accuracy-based epistemology
or what philosophers have learned from weather forecasters

The notion of partial belief has seen wide usage in philosophy of recent decades. Since our beliefs seem to come in varying strength, it is quite natural to associate a credal state of a subject with a vector of real numbers from the [0,1] segment, with 1 signifying certainty that the given proposition is true, 0 - certainty that it is false, and the intermediate numbers corresponding to the possible non-extreme levels of strength of belief in its truth. How should we compare and evaluate different credal states? Since one of the key epistemic norms is that one should seek truth, it would seem natural that the way to do that is to somehow calculate the distance of the given credal state from the vector of 0s and 1s corresponding to the true state of the world; in other words, we need a *scoring function* which calculates the *inaccuracy* of a credal state given a valuation on the considered propositions.

It turns out that there are many *prima facie* reasonable scoring functions, and the most popular among philosophers is the so called Brier score, proposed back in the 50s as a way of evaluating weather forecasters so that their proper, "honest" behaviour would be promoted: that is, so that the probabilities of (say) rain they announce would be equal to what in their hearts they believe them to be. (They should expect lying about their beliefs to be less profitable then telling the truth.) I will describe the class of scoring functions which meet this requirement, which are called *proper scoring functions*, and which are considered to be appropriate for measuring epistemic inaccuracy.

Probabilism is the thesis that one's credal state should satisfy the classical probability axioms. There have been many arguments for probabilism, some of which (e.g. those referring to Dutch Books) have lost their appeal in recent years. It turns out that the notion of accuracy can be used in a new type of argument for probabilism and various different epistemic norms. I will briefly describe and illustrate the general argumentative strategy, due mainly to de Finetti and J.M. Joyce.

I will also briefly comment on a recent argument by Easwaran and Fitelson that the goal of minimizing inaccuracy should be expected to conflict with other well-argued-for epistemic goals, such as respecting the evidence.

The last two notions I will describe are that of *calibration* and *refinement*. A rain-forecaster is *perfectly calibrated* when in the class of days on which she announces the probability of rain to be x, the proportion of rainy days to the whole of that class is x. (E.g. it rains on 80% of the days on which she announces that the probability of rain is 80%.) Now, we should expect that calibration is some measure of a credal state's quality, but it definitely is not the whole story. Suppose that it rained on 150 out of 300 days. Forecaster A has been announcing "The probability of rain is 50% today" throughout the whole period. Forecaster B said "It will rain today" on precisely the days it would rain, and said "It won't rain today" on precisely the days it wouldn't. Both A and B are perfectly calibrated, but it's clear B's predictions have been of more practical use. We will say that B is more *refined* than A (I will define this in the talk). I will mention a recent result by Pettigrew on the relationship between accuracy, calibration and refinement; in the context of proper scoring rules the notions are intimately connected.

This is meant to be an introductory talk, with the key notions presented 'from the ground up', the goal of which is to provide students with some information about a new and active field in epistemology.