Title: Admissibility Among Linear Estimators: The Question, and Some Answers

To estimate the common population mean and variance, we are accustomed to using the sample mean and the sample variance. If the parameter set for these two parameters is restricted in certain ways – e.g., it is known that the population cv is $\leq 10\%$ – there exist other linear estimators of the mean that have uniformly less mean squared error than does the sample mean. Among scalar multiples cs^2 of the sample variance, choices of c less than 1 produce estimators of σ^2 that are better than s^2.

Whether, for an estimator A, there is another that is uniformly better, is a question of *admissibility*. It would be nice to have a set of criteria, a check list, with which one could look at an estimator and tell immediately whether it is admissible.

There is some low-hanging fruit. In some settings, it is easy to prescribe sufficient conditions for a linear estimator to be admissible. And it is fairly simple, in some cases, to specify necessary conditions. But it turns out to be very difficult to shrink the sufficient conditions and expand the necessary conditions until they become both necessary and sufficient.

My objective in this talk is to describe, principally by way of example, the nature of the questions that must be addressed in order to characterize admissible linear estimators in a general linear model.