

You've Got to be Efficient: Ambiguity, Misspecification and Variational Preferences

Karun Adusumilli

Abstract :

This article introduces a framework for evaluating statistical decisions under both prior ambiguity and likelihood misspecification. We begin with an ambiguity set — a frequentist model that pairs a possibly misspecified likelihood with every possible prior — and uniformly expand it by a Kullback–Leibler radius to accommodate likelihood misspecification. We show that optimal decisions under this framework are equivalent to minimax decisions with an exponentially tilted loss function. Misspecification manifests as an exponential tilting of the loss, while ambiguity corresponds to a search for the least favorable prior. This separation between ambiguity and misspecification enables local asymptotic analysis under global misspecification, achieved by localizing the priors alone. Remarkably, for both estimation and treatment assignment, we show that optimal decisions coincide with those under correct specification, regardless of the degree of misspecification. These results extend to semi-parametric models. As a practical consequence, our findings imply that practitioners should prefer maximum likelihood over the simulated method of moments, and efficient GMM estimators — such as two-step GMM — over diagonally weighted alternatives.