

I will present a recent project titled “Empirical Bayes for Compound Adaptive Experiments”. The paper is too preliminary to circulate. But here is an abstract:

We investigate Empirical Bayes analysis in the context of compound adaptive experiments, where the arm distribution in each experiment follows a normal distribution with an unknown mean parameter that we aim to estimate. There are two primary approaches to EB estimation: g-modeling, which estimates the prior by maximizing the marginal likelihood, and f-modeling, which directly computes posterior means from the sample distribution of observations. We establish that g-modeling remains a valid EB procedure even when it incorrectly assumes exogenous data collection; it holds regardless of the sampling algorithm used and the endogeneity of sample sizes. One can apply standard g-modeling methods by treating the data as if it were exogenously sampled, and restrict attention to only the sample mean of the data. Remarkably, we show that the risk guarantees established for g-modeling with i.i.d data remain valid for adaptively generated data, with no need for prior knowledge of the sampling algorithm, even when it varies across experiments. In contrast, the f-modeling approach results in biased estimates. We validate the robustness of the g-modeling approach through simulations involving commonly used adaptive algorithms and illustrate its applicability using a real-world dataset comprising multiple sequential experiments.