Chen Qiu

Econometric Seminar Wednesday, October 23, 2024

Title: Squared Regret Policy Learning

Abstract:

The standard setting of statistical treatment choice assumes a risk neutral planner who evaluates treatment rules by the expected welfare. Yet, the planner may care other features of the welfare distribution than its mean. In this project, we study the estimation of individualized policy rules for a regret averse policy maker with the mean squared regret criterion, a notion advocated by Kitagawa, Lee and Qiu (2022). We focus on the case in which training data gives a rich set of observable characteristics while the assignment policy can depend only on its subset. For example, RCT data may collect sensitive information of individuals but the planner cannot differentiate treatment allocation based on them. For this scenario, we derive three new results that differ drastically from the standard linear regret analyses: first, we show that the population optimal rule is in general fractional due to the within-group heterogeneity of regret; second, we derive the automatic debiased form for the mean squared regret criterion, in which nuisance functions enter nonlinearly and nonsmoothly; third, viewing our debiased criterion as a weighted least square problem, we use a different statistical approach that leads to a new asymptotic upper bound with a convergence rate faster than one over root-n in general (cf. Athey and Wager, 2021). We use the dataset from the National Job Training Partnership Act (JTPA) Study to illustrate the value of our approach.

Link to useful background paper - **Treatment Choice with Nonlinear Regret**