

“Quantile Average Treatment Effects” (with Michael Lechner)

This paper introduces the Quantile Average Treatment Effect (QATE), a new parameter describing fine-grained causal heterogeneity. The QATE focuses on the 'actionable' part of the causal heterogeneity, which depends on observable characteristics of individual units only. We propose two estimators: a first-stage agnostic estimator, which can be applied with any causal machine learning method, and one specifically tailored to the Modified Causal Forest. A simulation analysis indicates that both estimators are consistent and informative about treatment effect heterogeneity. We illustrate their value in an empirical analysis of heterogeneity in the effect of smoking during pregnancy on birth weight.