

Testing Decision Makers without Counterfactuals.

Abstract:

A decision-maker (DM) repeatedly makes choices under uncertainty in a bandit environment, where only the realization of the chosen arm is observed. Another competing agent, the adviser (AD), repeatedly provides recommendations, but the realizations of these recommendations are unobserved unless they coincide with the DM's choice. Both agents possess partial information about the arms' realizations. The central question we focus on is whether, in the long run, an outside observer can identify which agent is more informed based solely on the observed decisions, recommendations, and arm realizations.

A test selects one of the agents based on the observed data. We focus primarily on the class of *scoring tests*, which assign a numerical score to each observation and select the agent according to the average score. We study strategic agents whose objective is to be selected by the test. For simultaneous arm choices, we show that there exists a scoring test that successfully identifies the more-informed agent. For sequential arm choices, however, no such scoring test exists.

Finally, we explore the tension between identifying the more-informed agent and maximizing welfare. A DM whose objective is to pass the test may not necessarily make welfare-maximizing decisions. In a binary-arm environment, we show that no scoring test can simultaneously identify the more informed agent and achieve more than half of the welfare attained by welfare-maximizing decisions.