Welfare Comparisons for Biased Learning

Mira Frick, Ryota Iijima, and Yuhta Ishii

Abstract :
We study robust welfare comparisons of learning biases, i.e., deviations from correct Bayesian updating. Given a true signal distribution, we deem one bias more harmful than another if it yields lower objective expected payoffs in all decision problems. We characterize this ranking in static (one signal) and dynamic (many signals) settings. While the static characterization compares posteriors signal-by-signal, the dynamic characterization employs an "efficiency index" quantifying the speed of belief convergence. Our results yield welfare-founded quantifications of the severity of well-documented biases. Moreover, the static and dynamic rankings can disagree, and "smaller" biases can be worse in dynamic settings.