

Abstract

Tournaments of heterogeneous candidates can be thought of as probabilistic mechanisms that select high-quality agents. We quantify the efficiency of such selection by the likelihood of selecting the best player, here termed the "predictive power." We study three widely used tournament formats: contests, binary elimination tournaments, and round-robin tournaments. Using a simple model, we demonstrate analytically, and through simulations, how the predictive power of these formats depends on the number of players, noise level, and distribution of players' types. We also present the results of exploratory simulations for two alternative criteria of selection efficiency: the expected ability of the winner and the expected rank of the winner. All three criteria may exhibit unexpected nonmonotonic behavior as functions of the number of players and/or noise level. We discuss the conditions under which different types of behavior should be expected, and their implications for managerial decisions.

<http://dx.doi.org/10.1287/mnsc.1070.0856>