The value of knowing drivers' opportunity cost in Ride Sharing systems

Ran Snitkovsky

March 1

Abstract

Consider a ride sharing platform, and a large population of strategic potential drivers, heterogeneous in terms of their opportunity costs, who choose whether or not to work for that platform. The platform is endowed with knowledge about the different drivers' opportunity costs. How can the platform implement a matching policy that uses this knowledge in order to improve system efficiency? Can such improvement be quantified? In this work we introduce an analytically-tractable mean field (fluid) model that accounts for the dynamic nature of drivers' spatial location, revenue, and availability status. Based on this model we compare drivers' equilibrium participation under two different matching policies. Our analysis leads to improvement bounds on the equilibrium performance: We show that a policy which utilizes knowledge about drivers' opportunity costs can perform up to two times better than a policy that does not do so, in terms of the number of drivers it attracts and in terms of the rate of matches it produces. We demonstrate by simulation that the mean field model provides an accurate approximation for a corresponding (stochastic) discrete model, in which the discussed improvement is observed empirically.