Department of Decision Sciences
Statistics Seminar

Partially Observable Risk-Sensitive Markov Decision Processes

Nicole Bäuerle
Karlsruher Institut für Technologie (KIT)

Thursday, 11th October 2018
12:30 pm Room 3-E4-SR03 Via Roentgen 1 Milano

Abstract

We consider the problem of minimizing a certainty equivalent of the total or discounted cost over a finite and an infinite time horizon which is generated by a Partially Observable Markov Decision Process (POMDP). In contrast to a risk-neutral decision maker this optimization criterion takes the variability of the cost into account. It contains as a special case the classical risk-sensitive optimization criterion with an exponential utility. We show that this optimization problem can be solved by embedding the problem into a completely observable MDP with extended state space and give conditions under which an optimal policy exists. The state space has to be extended by the joint conditional distribution of current unobserved state and accumulated cost. In case of an exponential utility, the problem simplifies considerably and we rediscover what in previous literature has been named information vector. However, since we do not use any change of measure techniques here, our approach is simpler. A small numerical example, namely the classical repeated casino game with unknown success probability is considered to illustrate the influence of the certainty equivalent and its parameters.

The talk is based on joint work with Ulrich Rieder.

Department of Decision Sciences

Via Röntgen 1 - 20136Milano

Tel. 02 5836.5632
Fax 02 5836.5630

Bocconi