

Multi-armed bandits with uncertainty.

Samuel Cohen, *University of Oxford*.

Abstract

Making good decisions based on estimates is difficult, but of clear importance in many applications. This is particularly the case when the decisions made will affect the information available in the future. Formally, this means that the filtration of our problem is not fixed in advance, but depends on the control used. We will consider the 'simplest' problem of this type, a multi-armed bandit problem, while taking account of uncertainty aversion. We will see that an extension of the classical Gittins' index approach is possible in this framework, despite many dynamic consistency issues.