

Abstract

In the economics literature there are two dominant approaches for solving models with optimal experimentation (also called active learning). The first approach is based on the value function and the second on an approximation method. In principle the value function approach is the preferred method. However, it suffers from the *curse of dimensionality* and is only applicable to small problems with a limited number of policy variables. The approximation method allows for a computationally larger class of models, but may produce results that deviate from the optimal solution. Our simulations indicate that when the effects of learning are limited, the differences may be small. However, when there is sufficient scope for learning, the value function solution is more aggressive in the use of the policy variable.