



Approximate dynamic programming with post-decision states as a solution method for dynamic economic models



Isaiah Hull*

Swedish Riksbank, Research Division, SE-103 37 Stockholm, Sweden

ARTICLE INFO

Article history:

Received 11 September 2013

Received in revised form

21 September 2014

Accepted 27 March 2015

Available online 7 April 2015

JEL classification:

C60

C61

C63

D52

Keywords:

Numerical solutions

Approximations

Heterogeneous agents

Nonlinear numerical solutions

Dynamic programming

ABSTRACT

I introduce and evaluate a new stochastic simulation method for dynamic economic models. It is based on recent work in the operations research and engineering literatures (Van Roy et al., 1997; Powell, 2007; Bertsekas, 2011), but also had an early application in economics (Wright and Williams, 1982, 1984). The baseline method involves rewriting the household's dynamic program in terms of post-decision states. This makes it possible to choose controls optimally without computing an expectation. I add a subroutine to the original algorithm that updates the values of states not visited frequently on the simulation path; and adopt a stochastic stepsize that efficiently weights information. Finally, I modify the algorithm to exploit GPU computing.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

I introduce and evaluate a new stochastic simulation method. It is based on dynamic programming, rather than the parameterization of expectational equations; and is similar to value function iteration (VFI), but does not require agents to explicitly compute an expectation over future state values when choosing controls. The baseline method, referred to as approximate dynamic programming with post-decision states (ADP-POST), has been shown to substantially reduce convergence times in high dimensional engineering and operations research problems. I use this method to solve dynamic economic problems, focusing primarily on DSGE model applications. I start by demonstrating how the algorithm can be modified to yield performance gains over common global solution methods in representative agent models. I then show how these gains grow when the approach is modified and used to solve a model with many heterogeneous agents and a state space with 12,005,000 nodes.

The algorithms presented in this paper are built around a post-decision state version of the Bellman equation, rather than the pre-decision state version. This approach makes it possible to solve two classes of models that are typically computationally intractable: (1) models with many continuous state variables; and (2) models with several correlated

* Tel.: +46 76 589 06 61; fax: +46 8 0821 05 31.

E-mail address: isaiah.hull@riksbank.se

exogenous processes. However, the formal tests will focus on the former, since the latter primarily presents a programming – and not convergence time – challenge.

The method used in this paper is based on a neuro-dynamic programming application in [Van Roy et al. \(1997\)](#). There, it was used to solve a retail inventory management problem with 33 state variables. Standard dynamic programming methods, such as value function iteration (infinite horizon) and backwards recursion (finite horizon), could not feasibly solve problems with such large state spaces. Even if all of the variables were discretized into the coarsest possible grid, it would still contain 2^{33} (8.59 billion) nodes.¹ In comparison, the neuro-dynamic programming algorithm they constructed around post-decision states was not only solvable, but delivered decision rules that yielded a 10% cost reduction over the standard, heuristic method.

Within economics, [Wright and Williams \(1982, 1984\)](#) first exploited a solution method based around post-decision states in two applications. [Judd \(1998\)](#) provided an expanded overview of the approach and explained how to construct the post-decision state Bellman equation. Similarly, [Sutton and Barto \(1998\)](#) introduced the technique to the reinforcement learning literature, but used it primarily as a strategy for absorbing the action space into the state space.² [Carroll \(2006\)](#) also made use of post-decision states, but to avoid root-finding operations.

[Powell \(2007\)](#) developed the method introduced by [Van Roy et al. \(1997\)](#) into a set of solution algorithms for high dimensionality optimization problems. However, all of the algorithms in [Powell \(2007\)](#) were designed for a partial equilibrium context in operations research; and most focused on inventory management applications. In problems of this variety, an agent faces random demands for a product and must choose how much inventory to hold. Using a post-decision state version of the Bellman equation works particularly well for this class of problems because the post-decision state is simply the inventory carried into the period, less the realization of demand. Outside of portfolio allocation problems, dynamic programming problems in economics do not typically take this form, and subsequently will not be able to achieve computational gains by collapsing the action space into the state space. For this reason, it is difficult to generalize what has been found in other literatures about the performance of the baseline algorithm to dynamic economic models.

Within operations research, [Powell \(2012\)](#) provides the most recent survey of work that has been done using post-decision state dynamic programming algorithms. Papers such as [Maxwell \(2011\)](#) and [Simao et al. \(2009\)](#) use the post-decision state formulation to solve very high dimensional problems, such as ambulance deployment and large-scale fleet management. Papers such as [Powell and Ryzhov \(2010\)](#) extend the initial algorithm by demonstrating how Bayesian methods can be used to update the value function.

Within economics, techniques such as the Parameterized Expectations Approach (PEA), developed by [Sargent \(1987\)](#), [Marcet \(1988\)](#), and [Den Haan and Marcet \(1990\)](#); and Projection Methods (PM), introduced by [Judd \(1992\)](#) and [McGrattan \(1999\)](#), provide more commonly used strategies for solving optimization problems with large state spaces. More recent work, such as [Maliar and Maliar \(2005\)](#) and [Judd et al. \(2011\)](#), demonstrates how to improve the stability and speed of such approaches. Additionally, [Judd et al. \(2011\)](#) propose methods for reducing the amount of time needed to compute an expectation. The authors accomplish this by precomputing expectations once and outside of the optimization step.³

In this paper, I use the post-decision state Bellman equation as a starting point for a larger algorithm that is designed to solve dynamic economic models. The benefit of constructing the dynamic programming problem around post-decision states is two-fold: first, it makes it possible to avoid computing expectations explicitly⁴; and second, it permits agents to make decisions based exclusively on information gained about the value of states visited on the simulation path, which substantially reduces the dimensionality of the choice problem. In addition to this, I make several modifications to the original algorithm in this paper to improve performance in the context of dynamic economic models.

The first modification adds a Search-Then-Converge (STC) subroutine ([Darken and Moody, 1992](#)) to update the algorithm's stepsize. However, instead of using the iteration count to update the stepsize, as in the original application, I use a function of the frequency with which individual states have been visited. Since nodes near the steady state are visited many times, their step sizes rapidly decline, making additional visits uninformative; however, states near the edges of the ergodic set will be able to retain larger state-specific stepsizes, allowing new visits to remain informative. Note that this modification is particularly important because the stepsize plays a dual role in algorithms based on post-decision states: it not only fosters convergence, but also implicitly computes the expectation over future states.

The second modification updates unvisited states and infrequently visited states near the edges of the ergodic set. This approach first determines a subset of the ergodic set over which the value function or decision rules are well-behaved. Information contained in this region is then used to infer properties of the decision rules over the entire state space. I do this by performing a weighted least squares regression of the decision-rule implied endogenous states on state variables. I use the relative frequencies with which states were visited in the simulations as the weights. This captures the fact that

¹ The approach we consider still requires discretization, but does not suffer as severely from the curse of dimensionality.

² In many dynamic programming applications, the endogenous component of the state space is determined by the realizations of the shock at the start of the period and the controls selected thereafter. In some of these applications, the dimensionality of the problem can be reduced simply by eliminating the choice of controls; and instead considering only the result of those choices (i.e. the end-of-period endogenous states).

³ Similarly, in this paper, I use a family of solution methods that renders the optimization step deterministic; however, it does not precompute expectations. Instead, it implicitly computes a quasi-monte carlo expectation over future states and limits the computation to states visited on the simulation path.

⁴ In fact, the expectation will be computed implicitly during the smoothing step, which adds no additional computational time.

Download English Version:

<https://daneshyari.com/en/article/5098256>

Download Persian Version:

<https://daneshyari.com/article/5098256>

[Daneshyari.com](https://daneshyari.com)