A polyhedral approximation approach to concave numerical dynamic programming

Kenichi Fukushima [a,*], Yuichiro Waki [b]

[a] *University of Wisconsin-Madison, 1180 Observatory Drive, Madison, WI 53706, United States*
[b] *University of Queensland, Australia*

ARTICLE INFO

ABSTRACT

This paper introduces a numerical method for solving concave continuous state dynamic programming problems which is based on a pair of polyhedral approximations of concave functions. The method is globally convergent and produces computable upper and lower bounds on the value function which can in theory be made arbitrarily tight. This is true regardless of the pattern of binding constraints, the smoothness of model primitives, and the dimensionality and rectangularity of the state space. We illustrate the method's performance using an optimal firm management problem subject to credit constraints and partial investment irreversibilities.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

This paper concerns continuous state numerical dynamic programming problems in which the return and constraint functions are continuous and concave. Such problems arise frequently in economics, often in inner loops for algorithms that solve much harder problems. There is therefore a desire for solution methods that are reliable, precise, and efficient.

Existing methods with the broadest applicability and greatest reliability are those based on value iterations, which if implemented exactly will generate a sequence of functions that converges to the value function from any starting point thanks to its contraction mapping property. When the state variables are continuous, however, an exact implementation of the procedure is infeasible as it requires storing infinitely many numbers in memory and solving infinitely many optimization problems per iteration. In practice one therefore approximately implements the procedure in one way or another.

There are two standard approaches here. The first is to discretize the state space, that is, simply replace the original state space with one that is finite. This approach is numerically stable—the iterations are guaranteed to converge because their contraction mapping property is preserved—but generally slow. The second approach is to compute the updated function values on a finite grid and then interpolate those values, either exactly or approximately, to generate a function to be used as input in the next iteration. This approach is often faster than the first but is generally less reliable as most interpolation methods break the contraction property of the iterations and can thereby cause non-convergence (see, e.g., Judd, 1998, p. 438). For problems with one-dimensional state spaces there is a satisfactory solution to the latter problem based on shape

---

* Corresponding author. Tel.: +1 608 262 8789.
  E-mail address: fukushima@wisc.edu (K. Fukushima).

preserving splines (Judd and Solnick, 1994). However comparable techniques remain relatively scarce for problems with multi-dimensional state spaces. In particular, currently known techniques (cf. Gordon, 1995; Stachurski, 2008), when applied to concave problems, generally introduce non-concavities which make it difficult to solve the optimization problems reliably and efficiently.

In addition to confronting users with this difficult tradeoff, existing methods are also limited in their ability to tell precisely how accurate the computed solution is. It is now common practice to address this issue by checking if certain necessary conditions for optimality—such as intertemporal Euler equations—hold with high accuracy. There are conditions under which such tests are known to have sound theoretical foundations (Santos, 2000); however it is not straightforward to adapt them to problems with occasionally binding constraints and/or other sources of non-smoothness.

The purpose of this paper is to introduce a method based on a pair of polyhedral approximations of concave functions which improves upon existing methods along these dimensions. In particular, the method is globally convergent, preserves concavity of the problem, and produces computable upper and lower bounds on the value function which can in theory be made arbitrarily tight. Furthermore, these properties hold true regardless of the pattern of binding constraints, the smoothness of model primitives, and the dimensionality and rectangularity of the state space.

The aforementioned features of our method make it particularly well suited for solving in a robust manner problems with occasionally binding constraints, non-differentiabilities, or multi-dimensional state spaces that may be non-rectangular. One such problem is the optimal firm management problem with credit constraints and partial investment irreversibilities in Khan and Thomas (2011). We use this as an example to test the practical performance of our method and find it to be reasonably efficient.

Our method consists of two components and each has important predecessors in the literature. The first component, which produces lower bounds on the value function, is close to a method based on piecewise affine interpolations analyzed by Santos and Vigo-Aguiar (1998) and the lottery based method of Phelan and Townsend (1991) for solving dynamic contracting problems. The second component, which produces upper bounds on the value function, is similar in some ways to what Nishimura and Stachurski (2009) used to analyze a model of primary commodity markets. Both components also fall into a broad class of methods outlined by Gordon (1995) and Stachurski (2008) for which convergence is guaranteed. Our method is also closely related to Judd et al.'s (2003) method for solving repeated games, and can in fact be viewed as its adaptation to dynamic programming problems. As far as we know, however, no paper has combined these strands in the literature into a general purpose method of the kind that we develop here.

The main limitation of our approach is that it works only with concave problems, and this constraint sometimes does bind in practice. While it is usually possible to get around it by introducing lotteries or other randomization devices, doing so may or may not be reasonable depending on the application.

2. Setup

We focus throughout on a general infinite horizon dynamic programming problem as treated in Stokey et al. (1989, Chapter 9). Our terminology on convex analysis follows Rockafellar (1970).

The decision problem is described by the following elements. The endogenous state variable x and control variable y (which becomes the next period's endogenous state) both belong to $X \subset \mathbb{R}^n$, which we take to be a polytope (i.e., the convex hull of a finite set of points). The random shocks z follow a time homogeneous Markov chain with finite state space Z, and the probability of transiting from state z to z' is $\pi(z'|z)$. The discount factor is $\beta \in (0,1)$. The return function $r : X \times X \times Z \rightarrow \mathbb{R}$ is continuous and concave in its first $2n$ arguments, and we let $r_{\min} := \min r(X \times X \times Z)$ and $r_{\max} := \max r(X \times X \times Z)$. The set of feasible controls at state $(x,z) \in X \times Z$ is $\Gamma(x,z) = \{y \in X : h(x,y,z) \geq 0\}$, where $h : X \times X \times Z \rightarrow \mathbb{R}^m$. Each component of h is continuous and concave in its first $2n$ arguments and $\Gamma(x,z)$ is non-empty for any $(x,z) \in X \times Z$. For later reference we let $P := \mathbb{R}^n$ and define $r^* : P \times P \times Z \rightarrow \mathbb{R}$ as

$$r^*(p,q,z) = \min_{(x,y) \in X \times X} \{p \cdot x + q \cdot y - r(x,y,z) : h(x,y,z) \geq 0\}.$$

Associated with this problem is the operator $\mathbb{T}$ which maps $v : X \times Z \rightarrow \mathbb{R}$ to $\mathbb{T}v : X \times Z \rightarrow \mathbb{R}$, given by

$$\mathbb{T}v(x,z) = \sup_{y \in \Gamma(x,z)} \{r(x,y,z) + \beta \mathbb{E}v(y,z)\} \tag{1}$$

where the operator $\mathbb{E}$ maps $v : X \times Z \rightarrow \mathbb{R}$ to $\mathbb{E}v : X \times Z \rightarrow \mathbb{R}$, given by

$$\mathbb{E}v(y,z) = \sum_{z' \in Z} v(y,z')\pi(z'|z).$$

As is well known, $\mathbb{T}$ is a monotone β-contraction on $B(X \times Z)$, the Banach space of bounded real valued functions on $X \times Z$ equipped with the supremum norm $\|\cdot\|$, and has the value function V as its unique fixed point in $B(X \times Z)$. Standard arguments imply that V is continuous and concave.