Fast value iteration: an application of Legendre-Fenchel duality to a class of deterministic dynamic programming problems in discrete time

Carpio, Ronaldo

Kamihigashi, Takashi

(Citation)
Journal of Difference Equations and Applications,26(2):209-222

(Issue Date)
2020-02-01

(Resource Type)
journal article

(Version)
Version of Record

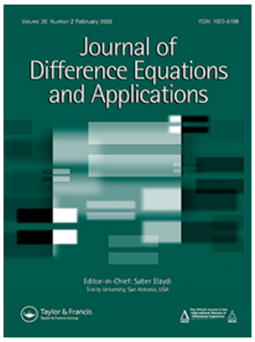

Taylor & Francis
Taylor & Francis Group

Journal of Difference Equations and Applications

ISSN: 1023-6198 (Print) 1563-5120 (Online) Journal homepage: https://www.tandfonline.com/loi/gdea20

Fast value iteration: an application of Legendre-Fenchel duality to a class of deterministic dynamic programming problems in discrete time

Ronaldo Carpio & Takashi Kamihigashi

To cite this article: Ronaldo Carpio & Takashi Kamihigashi (2020) Fast value iteration: an application of Legendre-Fenchel duality to a class of deterministic dynamic programming problems in discrete time, Journal of Difference Equations and Applications, 26:2, 209-222, DOI: 10.1080/10236198.2020.1713770

To link to this article: https://doi.org/10.1080/10236198.2020.1713770

Published online: 31 Jan 2020.

JOURNAL OF DIFFERENCE EQUATIONS AND APPLICATIONS
2020, VOL. 26, NO. 2, 209–222
https://doi.org/10.1080/10236198.2020.1713770

Fast value iteration: an application of Legendre-Fenchel duality to a class of deterministic dynamic programming problems in discrete time*

Ronaldo Carpio[a] and Takashi Kamihigashi[b]

[a]School of Business and Finance, University of International Business and Economics, Beijing, People's Republic of China; [b]RIEB and CCSS, Kobe University, Kobe, Japan

ABSTRACT
We propose an algorithm, which we call 'Fast Value Iteration' (FVI), to compute the value function of a deterministic infinite-horizon dynamic programming problem in discrete time. FVI is an efficient algorithm applicable to a class of multidimensional dynamic programming problems with concave return (or convex cost) functions and linear constraints. In this algorithm, a sequence of functions is generated starting from the zero function by repeatedly applying a simple algebraic rule involving the Legendre-Fenchel transform of the return function. The resulting sequence is guaranteed to converge, and the Legendre-Fenchel transform of the limiting function coincides with the value function.

ARTICLE HISTORY
Received 30 October 2018
Accepted 16 December 2019

KEYWORDS
Dynamic programming; Legendre-Fenchel transform; Bellman operator; convex analysis

2010 MATHEMATICS SUBJECT CLASSIFICATIONS
37N40; 46N10; 90C39

1. Introduction

It has been known since Bellman and Karush [3–6] that Legendre-Fenchel (LF) duality [9] can be utilized to solve finite-horizon dynamic programming (DP) problems in discrete time. Although there have been subsequent applications of Legendre-Fenchel (LF) duality to DP (e.g. [2,8,17,18,20]), to our knowledge there has been no serious attempt to exploit Legendre-Fenchel (LF) duality to develop an algorithm to solve infinite-horizon DP problems.

In this paper we propose an algorithm, which we call 'Fast Value Iteration' (FVI), to compute the value function of a deterministic infinite-horizon DP problem in discrete time. FVI is an efficient algorithm applicable to a class of multidimensional DP problems with concave return functions (or convex cost functions) and linear constraints.

The FVI algorithm is an implementation of what we call the 'dual Bellman operator', which is a simple algebraic rule involving the Legendre-Fenchel (LF) transform of the return function. A sequence of functions generated by repeated application of the dual Bellman operator is guaranteed to converge, and the Legendre-Fenchel (LF) transform of

CONTACT Takashi Kamihigashi tkamihig@rieb.kobe-u.ac.jp RIEB and CCSS, Kobe University, 2-1 Rokkodai, Nada, Kobe 657-8501, Japan

*Part of this research was conducted while Ronaldo Carpio was visiting RIEB, Kobe University, as a Visiting Researcher in 2015. Earlier versions of this paper were presented at the SCE Conference on Computing in Economics and Finance in Oslo, 2014, and the SAET Conference in Tokyo, 2014.

the limiting function coincides with the value function. Involving no optimization, the dual Bellman operator offers a dramatic computational advantage over standard computational methods such as value iteration and policy iteration (e.g. [22]). We prove that the convergence properties of the iteration of the dual Bellman operator are identical to those of value iteration when applied to a DP problem with a continuous, bounded, concave return function and a linear constraint.

The rest of the paper is organized as follows. In Section 2 we review some basic concepts from convex analysis and show some preliminary results. In Section 3 we present the general DP framework used in our analysis. In Section 4 we apply Legendre-Fenchel (LF) duality to a DP problem with a continuous, bounded, concave return function and a linear constraint. In Section 5 we present our numerical algorithm and compare its performance with that of modified policy iteration. In Section 7 we offer some concluding comments.

2. Preliminaries I: convex analysis

In this section we review some basic concepts from convex analysis and state some well-known results. We also establish some preliminary results.

Let $N \in \mathbb{N}$. Let $\overline{\mathbb{R}}$ denote the extended real line; i.e. $\overline{\mathbb{R}} = [-\infty, \infty]$. For $f : \mathbb{R}^N \to \overline{\mathbb{R}}$, we define $f_*, f^* : \mathbb{R}^N \to \bar{\mathbb{R}}$ by

$$f_*(p) = \inf_{x \in \mathbb{R}^N} \{p^\intercal x - f(x)\}, \quad \forall\, p \in \mathbb{R}^N, \tag{1}$$

$$f^*(p) = \sup_{x \in \mathbb{R}^N} \{p^\intercal x - f(x)\}, \quad \forall\, p \in \mathbb{R}^N, \tag{2}$$

where p and x are $N \times 1$ vectors, and $p^\intercal$ is the transpose of p. The functions f_* and f^* are called the *concave conjugate* and *convex conjugate* of f, respectively.

It follows from (1) and (2) that for any functions $f, g : \mathbb{R}^N_+ \to \overline{\mathbb{R}}$, we have

$$f = -g \;\Rightarrow\; \forall\, p \in \mathbb{R}^N, \quad f_*(p) = -g^*(-p) = -(-f)^*(-p). \tag{3}$$

This allows us to translate any statement about g and g^* to the corresponding statement about $-g$ and $(-g)_*$; this is useful since most results in convex analysis deal with convex functions and convex conjugates. In what follows we focus on concave functions and concave conjugates, and by 'conjugate', we always mean 'concave conjugate'. The *biconjugate* f_{**} of f is defined by

$$f_{**} = (f_*)_*. \tag{4}$$

A concave function $f : \mathbb{R}^N \to \overline{\mathbb{R}}$ is called *proper* if $f(x) < \infty$ for all $x \in \mathbb{R}^N$ and $f(x) > -\infty$ for at least one $x \in \mathbb{R}^N$.[1] The *effective domain* of f is defined as

$$\operatorname{dom} f = \{x \in \mathbb{R}^N : f(x) > -\infty\}. \tag{5}$$

Let F be the set of proper, concave, upper semicontinuous functions from $\mathbb{R}^N$ to $\overline{\mathbb{R}}$. For $f, g : \mathbb{R}^N \to \mathbb{R} \cup \{-\infty\}$, the *sup-convolution* of f and g is defined as

$$(f \# g)(x) = \sup_{y \in \mathbb{R}^N} \{f(y) + g(x - y)\}. \tag{6}$$

The following two lemmas collect the basic properties of conjugates we need later.

Lemma 2.1 (Rockafellar and Wets [23], Theorems 11.1, 11.23): (a) *For any $f \in F$, we have $f_* \in F$ and $f_{**} = f$.*

(b) *For any $f, g \in F$, we have $(f \# g)_* = f_* + g_*$.*

(c) *Let $N' \in \mathbb{N}$ and $u : \mathbb{R}^{N'} \to \overline{\mathbb{R}}$. Let L be an $N \times N'$ matrix. Define $(Lu) : \mathbb{R}^N \to \overline{\mathbb{R}}$ by*

$$(Lu)(x) = \sup_{c \in \mathbb{R}^{N'}} \{u(c) : Lc = x\}, \quad \forall\, x \in \mathbb{R}^N. \tag{7}$$

Then

$$(Lu)_*(p) = u_*(L^\mathsf{T} p), \quad \forall\, p \in \mathbb{R}^N. \tag{8}$$

Lemma 2.2 (Hiriart-Urruty [10], p. 484): *Let $f, g \in F$ be such that* $\mathrm{dom} f = \mathrm{dom} g = \mathbb{R}^N_+$. *Suppose that both f and g are bounded on $\mathbb{R}^N_+$. Then* $\mathrm{dom} f^* = \mathrm{dom} g^* = \mathbb{R}^N_+$, *and*

$$\sup_{x \in \mathbb{R}^N_+} |f(x) - g(x)| = \sup_{p \in \mathbb{R}^N_+} |f_*(p) - g_*(p)|. \tag{9}$$

The following result is proved in the Appendix.

Lemma 2.3: *Let A be an invertible $N \times N$ matrix. Let $\beta \in \mathbb{R}_{++}$ and $S = A^{-1}/\beta$. Let $f, v : \mathbb{R}^N \to \overline{\mathbb{R}}$ be such that $f(x) = \beta v(Ax)$ for all $x \in \mathbb{R}^N$. Then*

$$f_*(p) = \beta v_*(S^\mathsf{T} p). \tag{10}$$

3. Preliminaries II: dynamic programming

In this section we present the general framework for dynamic programming used in our analysis, and show a standard result based on the contraction mapping theorem. Our exposition here is based on Stokey and Lucas [24] and Kamihigashi [13,14].

Let $N \in \mathbb{N}$ and $X \subset \mathbb{R}^N$. Let $x \in X$. We need the following definitions from Stokey and Lucas [24, pp. 56–57]. A correspondence $\Gamma : X \to 2^X$ is called *lower hemi-continuous* at x if $\Gamma(x) \neq \emptyset$ and if for any $y \in \Gamma(x)$ and sequence $\{x_n\}_{n=0}^\infty$ with $\lim_{n\uparrow\infty} x_n = x$, there exist $N \in \mathbb{N}$ and a sequence $\{y_n\}_{n=N}^\infty$ such that $\lim_{n\uparrow\infty} y_n = y$ and $y_n \in \Gamma(x_n)$ for all $n \geq N$. A correspondence $\Gamma : X \to 2^X$ is called *upper hemi-continuous* at x if $\Gamma(x) \neq \emptyset$ and if for any sequence $\{x_n\}_{n=0}^\infty$ with $\lim_{n\uparrow\infty} x_n = x$ and sequence $\{y_n\}_{n=0}^\infty$ with $y_n \in \Gamma(x_n)$ for all $n \in \mathbb{Z}_+$, there exists a subsequence $\{y_{n_i}\}_{i=0}^\infty$ of $\{y_n\}$ with $\lim_{i\uparrow\infty} y_{n_i} \in \Gamma(x)$.

A correspondence $\Gamma : X \to 2^X$ is called *continuous* at x if it is both lower hemi-continuous and upper hemi-continuous at x. It is called *continuous* if it is continuous at each $x \in X$.

Consider the following problem:

$$\sup_{\{x_{t+1}\}_{t=0}^\infty} \sum_{t=0}^\infty \beta^t r(x_t, x_{t+1}) \tag{11}$$

$$\text{s.t.} \quad x_0 \in X \text{ given,} \tag{12}$$

$$\forall\, t \in \mathbb{Z}_+, \quad x_{t+1} \in \Gamma(x_t). \tag{13}$$

In this section we maintain the following assumption.

Assumption 3.1: (i) $\beta \in (0,1)$. (ii) $\Gamma : X \to 2^X$ is a nonempty, compact-valued, continuous correspondence. (iii) X and gphΓ are convex sets, where gphΓ is the graph of Γ:

$$\text{gph}\Gamma = \{(x,y) \in \mathbb{R}^N \times \mathbb{R}^N : y \in \Gamma(x)\}. \tag{14}$$

(iv) $r : \text{gph}\Gamma \to \mathbb{R}$ is continuous, bounded, and concave.

Let $\hat{v} : X \to \mathbb{R}$ be the *value function* of the problem (11)–(13); i.e. for $x_0 \in X$ we define

$$\hat{v}(x_0) = \sup_{\{x_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t r(x_t, x_{t+1}) \quad \text{s.t. (13)}. \tag{15}$$

It is well-known that $\hat{v}$ satisfies the optimality equation (see [12,14]):

$$\hat{v}(x) = \sup_{y \in \Gamma(x)} \{r(x,y) + \beta\hat{v}(y)\}, \quad \forall\, x \in X. \tag{16}$$

Thus $\hat{v}$ is a fixed point of the *Bellman operator* B defined by

$$(Bv)(x) = \sup_{y \in \Gamma(x)} \{r(x,y) + \beta v(y)\}, \quad \forall\, x \in X. \tag{17}$$

Let $\mathscr{C}(X)$ be the space of continuous, bounded, concave functions from X to $\mathbb{R}$ equipped with the sup norm $\|\cdot\|$. The following result is proved using standard arguments in the Appendix. We claim no originality and state it here as a theorem only for reference purposes.

Theorem 3.1: *Under Assumption* 3.1, *the following statements hold*:

(a) *The Bellman operator B is a contraction on $\mathscr{C}(X)$ with modulus β; i.e. B maps $\mathscr{C}(X)$ into itself, and for any $v, w \in \mathscr{C}(X)$ we have*

$$\|Bv - Bw\| \le \beta\|v - w\|. \tag{18}$$

(b) *B has a unique fixed point $\tilde{v}$ in $\mathscr{C}(X)$. Furthermore, for any $v \in \mathscr{C}(X)$,*

$$\forall\, i \in \mathbb{N}, \quad \|B^i v - \tilde{v}\| \le \beta^i \|v - \tilde{v}\|. \tag{19}$$

(c) *We have $\tilde{v} = \hat{v}$, where $\hat{v}$ is defined by* (15).

Recent results on convergence and uniqueness that require neither continuity nor concavity can be found in Kamihigashi [14] and Kamihigashi et al. [15].

4. The dual Bellman operator

In this section we introduce the 'dual Bellman operator', which traces the iterates of the Bellman operator in a dual space for a special case of (11)–(13). In particular we consider the following problem:

$$\max_{\{c_t,x_{t+1}\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{20}$$

$$\text{s.t.} \quad x_0 \in \mathbb{R}^N_+ \text{ given,} \tag{21}$$

$$\forall\, t \in \mathbb{Z}_+, \quad x_{t+1} = Ax_t - Dc_t, \tag{22}$$

$$c_t \in \mathbb{R}^{N'}_+,\ x_{t+1} \in \mathbb{R}^N_+, \tag{23}$$

where A is an $N \times N$ matrix, D is an $N \times N'$ matrix with $N' \in \mathbb{N}$, and c_t is a $N' \times 1$ vector. Throughout this section we maintain the following assumption.

Assumption 4.1: (i) $\beta \in (0, 1)$. (ii) $u : \mathbb{R}^{N'}_+ \to \mathbb{R}$ is continuous, bounded, and concave. (iii) A is a nonnegative monotone matrix (i.e. $Ax \in \mathbb{R}^N_+ \Rightarrow x \in \mathbb{R}^N_+$). (iv) D is a nonzero nonnegative matrix.

It is well-known (e.g. [7, p. 137]) that a square matrix is monotone if and only if it is invertible and its inverse is nonnegative. Furthermore, a nonnegative square matrix is monotone if and only if it has exactly one nonzero element in each row and in each column [16]. Thus the latter property is equivalent to part (iii) above. See Subsection 6.1 for a simple economic model satisfying this and other properties.

Under Assumption 4.1, the optimization problem (20)–(23) is a special case of (11)–(13) with

$$X = \mathbb{R}^N_+, \tag{24}$$

$$\Gamma(x) = \{y \in \mathbb{R}^N_+ : \exists c \in \mathbb{R}^{N'}_+, y = Ax - Dc\}, \quad \forall\, x \in X, \tag{25}$$

$$r(x, y) = \max_{c \in \mathbb{R}^{N'}_+} \{u(c) : y = Ax - Dc\}, \quad \forall\, (x, y) \in \operatorname{gph}\Gamma. \tag{26}$$

It is easy to see that Assumption 3.1 holds under Assumption 4.1 and (24)–(26).

Before proceeding, we introduce a standard convention for extending a function defined on a subset of $\mathbb{R}^n$ to the entire $\mathbb{R}^n$. Given any function $g : E \to \overline{\mathbb{R}}$ with $E \subset \mathbb{R}^n$, we extend g to $\mathbb{R}^n$ by setting

$$g(x) = -\infty, \quad \forall\, x \in \mathbb{R}^n \setminus E. \tag{27}$$

Note that in general, for any extended real-valued function f defined on a subset of $\mathbb{R}^n$, we have

$$\sup_{x \in \mathbb{R}^n} f(x) = \sup_{x \in \operatorname{dom} f} f(x), \tag{28}$$

where the function f on the left-hand side is the version of f extended according to (27).

Letting $L = A^{-1}D$, we can express (22) as

$$Lc_t = x_t - A^{-1}x_{t+1}. \tag{29}$$

In view of this and (26), we have

$$r(x,y) = \max_{c\in\mathbb{R}_+^{N'}} \{u(c) : Lc = x - A^{-1}y\}, \quad \forall\,(x,y) \in \mathrm{gph}\Gamma. \tag{30}$$

By Assumption 4.1 and (24)–(27), the Bellman operator B defined by (17) can be written as

$$(Bv)(x) = \sup_{z\in\mathbb{R}^N} \{(Lu)(x-z) + \beta v(Az)\}, \quad \forall\, x \in X, \tag{31}$$

where Lu is defined by (7) and (30) with $z = A^{-1}y$. The constraint $y \in \Gamma(x)$ in (17) is implicitly imposed in (31) by the effective domains of u and v, which require, respectively, that there exist $c \in \mathbb{R}_+^{N'}$ with $Lc = x{-}z$ and that $y = Az \in X$. Following the convention (27), we set

$$(Bv)(x) = -\infty, \quad \forall\, x \in \mathbb{R}^N \setminus X. \tag{32}$$

For any $f : \mathbb{R}^N \to \overline{\mathbb{R}}$ with $\mathrm{dom} f = \mathbb{R}_+^N$, we write $f \in \mathscr{C}(X)$ if f is continuous, bounded, and concave on X. Since u is bounded, $(Bv)(x)$ is well-defined for any $x \in \mathbb{R}^N$ and $v : X \to \overline{\mathbb{R}}$. In particular, for any $v \in \mathscr{C}(X)$, we have $Bv \in \mathscr{C}(X)$ by Theorem 3.1. The following result shows that the Bellman operator B becomes a simple algebraic rule in the 'dual' space of conjugates.

Lemma 4.1: *Let $S = A^{-1}/\beta$. For any $v \in \mathscr{C}(X)$ we have*

$$(Bv)_*(p) = u_*(L^{\mathsf{T}}p) + \beta v_*(S^{\mathsf{T}}p), \quad \forall\, p \in \mathbb{R}^N. \tag{33}$$

Proof: Let $f(z) = \beta v(Az)$ for all $z \in \mathbb{R}^N$. Let $g = Lu$. We claim that

$$Bv = f\#g. \tag{34}$$

It follows from (31) that $(Bv)(x) = (f\#g)(x)$ for all $x \in X$. It remains to show that $(Bv)(x) = (f\#g)(x)$ for all $x \in \mathbb{R}^N \setminus X$ or, equivalently, $(f\#g)(x) > -\infty \Rightarrow x \in X$. Let $x \in \mathbb{R}^N$ with $(f\#g)(x) > -\infty$. Then there exists $z \in \mathbb{R}^N$ with $g(x-z) + f(z) = (Lu)(x-z) + \beta v(Az) > -\infty$; i.e.

$$x - z \in \mathrm{dom}(Lu), \quad Az \in \mathrm{dom} v. \tag{35}$$

Since A^{-1} and D are nonnegative by Assumption 4.1, L is nonnegative; thus $\mathrm{dom}(Lu) \subset \mathbb{R}_+^N$. Hence $x - z \geq 0$ and $Az \geq 0$ by (35). Since the latter inequality implies that $z \geq 0$ by monotonicity of A, it follows that $x \geq z \geq 0$; i.e. $x \in X$. We have verified (34).

By Lemma 2.1 we have $(Bv)_* = (f\#g)_* = f_* + g_*$. Thus for any $p \in \mathbb{R}_+^N$,

$$(Bv)_*(p) = f_*(p) + g_*(p) = \beta v_*(S^{\mathsf{T}}p) + u_*(L^{\mathsf{T}}p), \tag{36}$$

where the second equality uses Lemmas 2.3 and 2.1. Now (33) follows. ■

We call the mapping from v_* to $(Bv)_*$ defined by (33) the *dual Bellman operator* B_*; more precisely, for any $f : \mathbb{R}^N \to \overline{\mathbb{R}}$, we define $B_* f$ by

$$(B_* f)(p) = u_*(L^\mathsf{T} p) + \beta f(S^\mathsf{T} p), \quad \forall\, p \in \mathbb{R}^N. \tag{37}$$

Using the dual Bellman operator B_*, (33) can be written simply as

$$(Bv)_* = B_* v_*, \quad \forall\, v \in \mathscr{C}(X). \tag{38}$$

According to the convention (27), the domain (rather than the effective domain) of a function f defined on $\mathbb{R}^N_+$ can always be taken to be the entire $\mathbb{R}^N$. Although this results in no ambiguity when we evaluate $\sup_x f(x)$ (recall (28)), it causes ambiguity when we evaluate $\|f\|$. For this reason we specify the definition of $\|\cdot\|$ as follows:

$$\|f\| = \sup_{x \in \mathbb{R}^N_+} \|f(x)\|. \tag{39}$$

We use this definition of $\|\cdot\|$ for the rest of the paper. The following result establishes the basic properties of the dual Bellman operator B_*.

Theorem 4.1: *For any $v \in \mathscr{C}(X)$, the following statements hold*:

(a) *For any $i \in \mathbb{N}$ we have*

$$(B^i v)_* = B^i_* v_*, \tag{40}$$

$$B^i v = (B^i_* v_*)_*, \tag{41}$$

$$B^i_* v_* \in \mathscr{C}(X), \tag{42}$$

where $B^i_* = (B_*)^i$.

(b) *The sequence $\{B^i_* v_*\}_{i \in \mathbb{N}}$ converges uniformly to $\hat{v}_*$ (the conjugate of the value function $\hat{v}$); in particular, for any $i \in \mathbb{N}$ we have*

$$\|B^i_* v_* - \hat{v}_*\| = \|B^i v - \hat{v}\| \le \beta^i \|v - \hat{v}\| = \beta^i \|v_* - \hat{v}_*\|. \tag{43}$$

(c) *We have $\hat{v} = (\hat{v}_*)_*$.*

Proof: (a) We first note that (40) implies (42) by Lemma 2.1 and Theorem 3.1(a). We show by induction that (40) holds for all $i \in \mathbb{N}$.

Note that for $i = 1$, (40) hods by (38). Suppose that (40) holds for $i = n - 1 \in \mathbb{N}$:

$$(B^{n-1} v)_* = B^{n-1}_* v_*. \tag{44}$$

Let us consider the case $i = n$. We have $B^{n-1} v \in \mathscr{C}(X)$ by Theorem 3.1. Thus using (38) and (44), we obtain

$$(B^n v)_* = (BB^{n-1} v)_* = B_* (B^{n-1} v)_* = B_* B^{n-1}_* v_* = B^n_* v_*. \tag{45}$$

Hence (40) holds for $i = n$. Now by induction, (40) holds for all $i \in \mathbb{N}$.

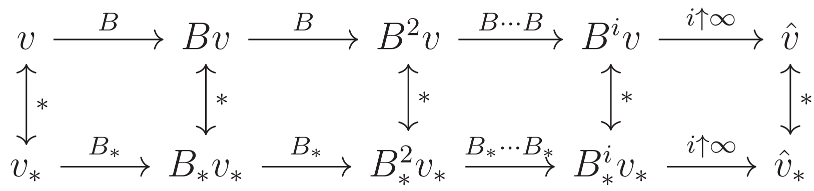

Figure 1. Bellman operator B and dual Bellman operator B^*.

To see (41), let $i \in \mathbb{N}$. Since $B^i v \in \mathscr{C}(X)$ by Theorem 3.1, we have $B^i v = (B^i v)_{**}$ by Lemma 2.1. Recalling (4), we have $B^i v = ((B^i v)_*)_* = (B_*^i v_*)_*$, where the second equality uses (40). We have verified (41).

(b) By Lemma 2.2 and (39), for any $i \in \mathbb{N}$ we have

$$\|B^i v - \hat{v}\| = \|(B^i v)_* - \hat{v}_*\| = \|B_*^i v_* - \hat{v}_*\|, \tag{46}$$

where the second equality uses (40). Thus the first equality in (43) follows; the second equality follows similiarly. By Theorem 3.1 the inequality in (43) holds. As a consequence, $\{B_*^i v_*\}_{i\in\mathbb{N}}$ converges uniformly to $\hat{v}_*$.

(c) Since $\hat{v} \in \mathscr{C}(X)$ by Theorem 3.1, we have $\hat{v} = \hat{v}_{**} = (\hat{v}_*)_*$ by Lemma 2.1. This completes the proof of Theorem 4.1. ■

Figure 1 summarizes the results of Theorem 4.1. The vertical bidirectional arrows between Bv and B_*v_*, B^2v and $B_*^2v_*$, etc, indicate that any intermediate result obtained by the Bellman operator B can be recovered through conjugacy from the corresponding result obtained by the dual Bellman operator B_*, and vice versa. This is formally expressed by statement (a) of Theorem 4.1. Statement (b) shows that both iterates $\{B^i v\}$ and $\{B_*^i v_*\}$ converge exactly the same way. In fact, as shown by Lemma 2.2, conjugacy preserves the sup norm between any pair of functions in F whose effective domains are $\mathbb{R}_+^N$. The right-most vertical arrow in Figure 1 indicates that the value function $\hat{v}$ can be obtained as the conjugate of the limit of $\{B_*^i v_*\}$, as shown in statement (c) of Theorem 4.1.

5. Fast value iteration

We exploit the relations expressed in Figure 1 to construct a numerical algorithm. The upper horizontal arrows in Figure 1 illustrate the standard value iteration algorithm, which approximates the value function $\hat{v}$ by successively computing $Bv, B^2v, B^3v, \ldots$ until convergence. The same result can be obtained by successively computing $B_*v_*, B_*^2v_*, B_*^3v_*, \ldots$ until convergence and by computing the conjugate of the last iterate. Theorem 4.1(b) suggests that this alternative method can achieve convergence in the same number of steps as value iteration, but it is considerably faster since each step is a simple algebraic rule without optimization; recall (37).

Algorithm 1, which we call 'Fast Value Iteration', implements this procedure with a finite number of grid points, using nearest-grid-point interpolation to approximate points not on the grid. To be precise, we take n grid points $p_1, \ldots, p_n$ in $\mathbb{R}_+^N$ as given, and index them by $j \in J \equiv \{1, \ldots, n\}$. Recall from (42) that it suffices to consider the behaviour of $B_*^i v_*$ on $X = \mathbb{R}_+^N$. We also take as given a function $\rho : \mathbb{R}_+^N \to \{p_1, \ldots, p_n\}$ that maps each point

Algorithm 1: Fast Value Iteration

let n grid points in $\mathbb{R}^N_+$ be given by $p_1, \ldots, p_n \in \mathbb{R}^N_+$
initialize $a, b, w : J \to \mathbb{R}$ (i.e. $\forall\, j \in J, a(j), b(j), w(j) \in \mathbb{R}$)
initialize $g : J \to J$ (i.e. $\forall\, j \in J, g(j) \in J$)
compute u_* on $L^\mathsf{T} p_1, \ldots, L^\mathsf{T} p_n$
for $j = 1, \ldots, n$
  $b(j) = 0$
  $w(j) = u_*(L^\mathsf{T} p_j)$
  $g(j) = \lambda(S^\mathsf{T} p_j)$
fix $\epsilon > 0$
$d = 2\epsilon$
while $d > \epsilon$
  $a = b$
  for $j = 1, \ldots, n$
    $b(j) = w(j) + \beta a(g(j))$
  $d = \max_{j \in J}\{|a(j) - b(j)|\}$
compute b_*
return b_*

$p \in \mathbb{R}^N_+$ to a nearest grid point. We define $\lambda : \mathbb{R}^N_+ \to J$ by $\rho(p) = p_{\lambda(p)}$; i.e. $\lambda(p)$ is the index of the grid point corresponding to p.

Algorithm 1 requires us to compute the conjugate of the return function u at the beginning as well as the conjugate of the final iterate at the end. To compute these conjugates, we employ the linear-time algorithm (linear in the number of grid points) presented in Lucet [19], which computes the conjugate of a concave function on a box grid. Since the rate of convergence for $\{B^i_* v_*\}$ is determined by β (as shown in Theorem 4.1(b)) and the number of algebraic operations required for each grid point in each iteration of the 'while' loop in Algorithm 1 is independent of the number of grid points, it follows that FVI is a linear-time algorithm.

6. Numerical examples

6.1. The AK model

As a simple example satisfying all our assumptions, consider the 'AK model' of economic growth used in macroeconomics (e.g. [1]):

$$\max_{\{c_t\}_{t=0}^{\infty}} \quad \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{47}$$

$$\text{s.t.} \quad x_0 \geq 0 \text{ given,} \tag{48}$$

$$\forall\, t \in \mathbb{Z}_+, \quad c_t + x_{t+1} = Ax_t, \tag{49}$$

$$c_t, x_t \geq 0, \tag{50}$$

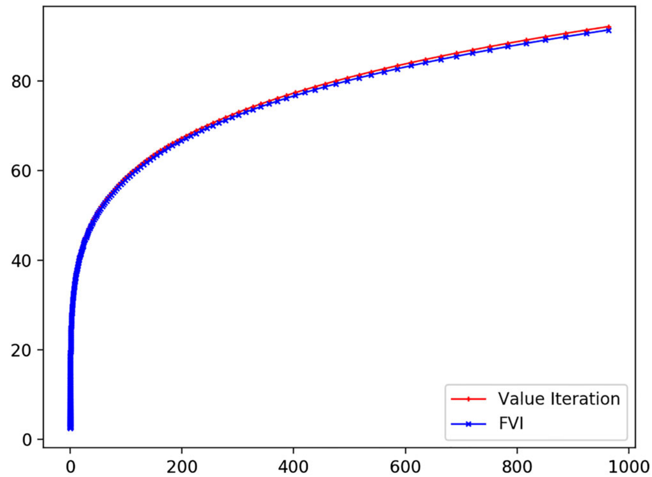

Figure 2. FVI vs. value iteration for the AK model.

where c_t is the consumption in period t, $u(c_t)$ is the utility derived from consumption, x_t is the capital stock at the beginning of period t, and $A > 0$ is a constant. We assume that the utility function is given by

$$u(c) = \frac{c^q}{q}, \quad q \in (0, 1). \tag{51}$$

In Figure 2 we compare the solutions computed by FVI and value iteration using a log-spaced grid of 500 points from 10^{-6} to 1000, and a convergence tolerance of 10^{-6}. The parameter values used are $\beta = 0.9, q = 0.2$, and $A = 0.9$. For both algorithms, the initial guess for the value function is $u(x)$. As expected from our theoretical results, the solutions produced by the two methods appear to be reasonably close to each other. Our experience suggests that accuracy can be improved with a larger number of grid points and a larger domain for $u(c)$. We make a more serious comparison in the next subsection.

6.2. Numerical comparison

To illustrate the efficiency of FVI, we compare the performance of FVI with that of modified policy iteration (MPI), which is a standard method to accelerate value iteration [22, Ch. 6.5]. In what follows, we assume the following in (20)–(23):

$$u(c_1, c_2) = -(c_1 - 10)^2 - (c_2 - 10)^2, \quad \beta = 0.9, \quad D = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}. \tag{52}$$

Although u above is not bounded, it is bounded on any bounded region that contains $L^\intercal p_1, \ldots, L^\intercal p_n$; thus we can treat u as a bounded function for our purposes. Concerning the matrix A, we consider two cases:

$$\text{(a) } A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \text{(b) } A = \begin{bmatrix} 0 & 1.1 \\ 1 & 0 \end{bmatrix}. \tag{53}$$

The grid points for MPI are evenly spread over $[0, 20] \times [0, 20]$. For FVI, the same number of grid points are evenly spread over a sufficiently large bounding box in the dual space.

We implement both FVI and MPI in Python, using the Scipy 0.13.3 package on a 2.40 GHz i7-3630QM Intel CPU. For MPI, we utilize C++ to find a policy that achieves the maximum of the right-hand side of the Bellman equation (31) by brute-force grid search. We use brute-force grid search because a discretized version of a concave function need not be concave (see [21]); we utilize C++ because brute-force grid search is unacceptably slow in Python. The resulting policy is used to update the approximate value function 100 times, and the resulting approximate value function is used to find a new policy.

Table 1. Number of iterations to convergence, time to convergence in seconds, maximum relative difference, and average relative difference for FVI and MPI algorithms.

		grid size	40 × 40	80 × 80	120 × 120	160 × 160	200 × 200	240 × 240	280 × 280
(a)	FVI	iterations	120	141	178	191	209	203	219
		CPU time	0.045	0.139	0.299	0.587	0.954	1.398	2.086
	MPI	iterations	7	7	8	8	8	8	8
		CPU time	1.636	7.964	34.901	93.305	214.673	434.626	764.269
	diff	max	3.72E−03	3.36E−03	2.31E−03	2.29E−03	1.84E−03	1.79E−03	1.98E−03
		mean	1.20E−03	1.27E−03	8.77E−04	6.91E−04	5.44E−04	7.20E−04	6.34E−04
(b)	FVI	iterations	131	193	192	190	205	214	207
		CPU time	0.048	0.184	0.311	0.536	0.943	1.446	1.977
	MPI	iterations	8	8	8	8	8	8	8
		CPU time	1.888	9.690	35.887	96.848	220.004	448.523	796.127
	diff	max	7.95E−03	5.62E−03	5.01E−03	4.74E−03	4.60E−03	4.52E−03	4.46E−03
		mean	2.63E−03	1.19E−03	1.02E−03	1.40E−03	1.19E−03	1.07E−03	1.19E−03

Notes: Case (a) assumes (52) and (53)(a), while case (b) assumes (52) and (53)(b).

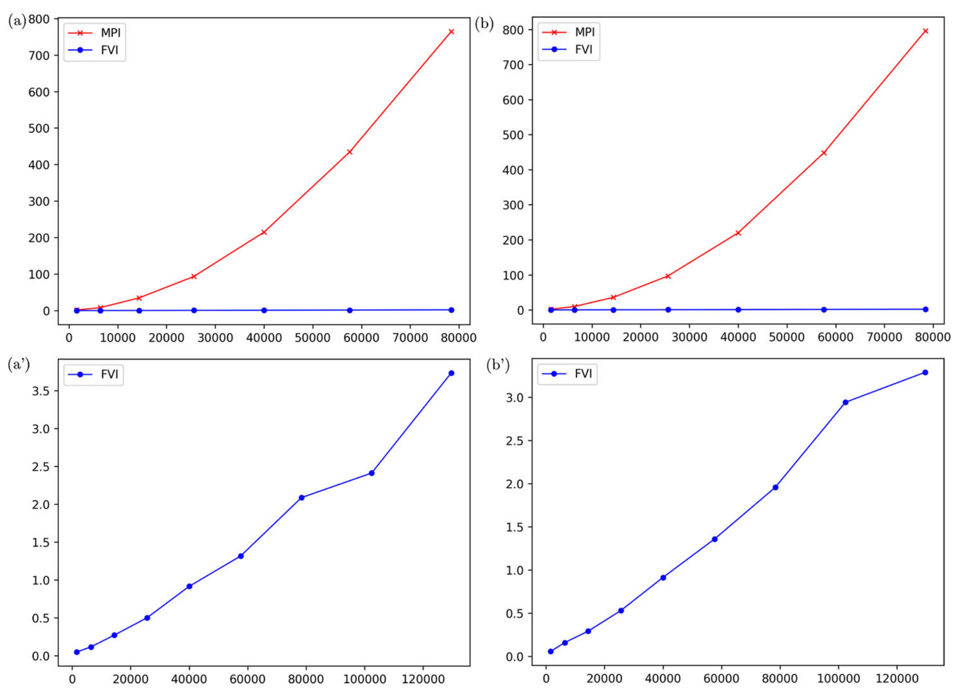

Figure 3. Time to convergence in seconds vs. number of grid points. Panels (a) & (a') assume (52) & (53)(a), while panels (b) & (b') assume (52) & (53)(b).

Table 1 shows the number of iterations and total CPU time for FVI and MPI to converge to a tolerance of 10^{-5}. For each grid size, the final approximate value functions from FVI and MPI are compared by computing, at each grid point, the absolute difference divided by the largest absolute value of the MPI value function; we report the maximum and average values of this difference over all grid points.

Panels (a) and (b) in Figure 3 plot the time to convergence of FVI and MPI against the number of grid points using the data in Table 1. Panels (a') and (b') show the performance of FVI for an extended range of grid point sizes. These plots indicate that FVI is a linear-time algorithm, as discussed above. In terms of CPU time, FVI clearly has a dramatic advantage.

7. Concluding comments

In this paper we proposed an algorithm called 'Fast Value Iteration' (FVI) to compute the value function of a deterministic infinite-horizon dynamic programming problem in discrete time. FVI is an efficient algorithm that offers a dramatic computational advantage for a class of problems with concave return (or convex cost) functions and linear constraints.

The algorithm we presented is based on the theoretical results shown for continuous state problems, but in practice, numerical errors are introduced through discretization and computation of conjugates. Although precise error estimates are yet to be established, our numerical experiments suggest that the difference between the approximate value functions computed using FVI and MPI, respectively, is rather insignificant.

In practice, one can combine FVI with other numerical methods to achieve a desired combination of speed and accuracy. For example, to obtain essentially the same MPI value function while economizing on time, one can apply FVI until convergence first and then switch to MPI. As in this algorithm, FVI can be used to quickly compute a good approximation of the value function.

In conclusion, we should point out that the theoretical results shown in Section 4 can be extended to problems with more general and nonlinear constraints using a general formula for the conjugate of a composite function [11]. New algorithms based on such an extension are left for future research.

Notes

1. This is a different concept from proper maps in topological spaces. For example, the function $f : \mathbb{R} \to \overline{\mathbb{R}}$ defined by $f(x) = -x^{-1}$ for $x > 0$ and $f(x) = -\infty$ for $x \leq 0$ is proper as a concave function but not proper in the topological sense.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

Financial support from the Japan Society for the Promotion of Science (JSPS KAKENHI Grant Number 15H05729) is gratefully acknowledged.

References

[1] D. Acemoglu, *Introduction to Modern Economic Growth*, Princeton University Press, Princeton, 2009.

[2] M. Akian and S. Gaubert, *Spectral theorem for convex monotone homogeneous maps, and ergodic control*, Nonlinear Anal. 52 (2003), pp. 637–679.

[3] R. Bellman and W. Karush, *Mathematical programming and the maximum transform*, SIAM J. Appl. Math. 10 (1962), pp. 550–567.

[4] R. Bellman and W. Karush, *On the maximum transform and semigroups of transformations*, Bull. Amer. Math. Soc. 68 (1962), pp. 516–518.

[5] R. Bellman and W. Karush, *Functional equations in the theory of dynamic programming XII: An application of the maximum transform*, J. Math. Anal. Appl. 6 (1963), pp. 155–157.

[6] R. Bellman and W. Karush, *On the maximum transform*, J. Math. Anal. Appl. 6 (1963), pp. 67–74.

[7] A. Berman and R.J. Plemmons, *Nonnegative Matrices in the Mathematical Sciences*, SIAM, Philadelphia, 1994.

[8] A.O. Esogbue and C.W. Ahn, *Computational experiments with a class of dynamic programming algorithms of higher dimensions*, Comput. Math. Appl. 19 (1990), pp. 3–23.

[9] W. Fenchel, *On conjugate convex functions*, Canad. J. Math. 1 (1949), pp. 73–77.

[10] J.B. Hiriart-Urruty, *A general formula on the conjugate of the difference of functions*, Canad. Math. Bull. 29 (1986), pp. 482–485.

[11] J.B. Hiriart-Urruty, A note on the Legendre-Fenchel transform of convex composite functions, Nonsmooth Mechanics and Analysis, Springer, New York, 2006, pp. 35–46.

[12] T. Kamihigashi, *On the principle of optimality for nonstationary deterministic dynamic programming*, Int. J. Econ. Theory 4 (2008), pp. 519–525.

[13] T. Kamihigashi, *An order-theoretic approach to dynamic programming: An exposition*, Econ. Theory Bull. 2 (2014), pp. 13–21.

[14] T. Kamihigashi, *Elementary results on solutions to the Bellman equation of dynamic programming: Existence, uniqueness, and convergence*. Economic Theory 56 (2014), pp. 1–23.

[15] T. Kamihigashi, K. Reffett, and M. Yao, *An application of Kleene's fixed point theorem to dynamic programming: A note*, Int. J. Econ. Theory 11 (2015), pp. 429–434.

[16] H. Kestelman, *Matrices with* $A \geq 0$ *and* $A^{-1} \geq 0$, Am. Math. Mon. E2379 (1973), pp. 1059–1060.

[17] C.M. Klein, *Conjugate duality and its implications in dynamic programming*, Math. Comput. Model.14 (1990), pp. 151–154.

[18] C.M. Klein and T.L. Morin, *Conjugate duality and the curse of dimensionality*, Eur. J. Oper. Res. 50 (1991), pp. 220–228.

[19] Y. Lucet, *Faster than the fast Legendre transform, the linear-time Legendre transform*, Numer. Algorithms 16 (1997), pp. 171–185.

[20] T.L. Morin and A.M.O. Esogbue, *The imbedded state space approach to reducing dimensionality in dynamic programs of higher dimensions*, J. Math. Anal. Appl. 48 (1974), pp. 801–810.

[21] K. Murota, *Discrete Convex Analysis*, SIAM, Philadelphia, 2003.

[22] M.L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, John Wiley & Sons, Hoboken, NJ, 2005.

[23] R.T. Rockafellar and R.J.B. Wets, *Variational Analysis*, Springer, Dordrecht, 2009.

[24] N. Stokey and R.E. Lucas Jr., *Recursive Methods in Economic Dynamics*, Harvard University Press, Cambridge, MA, 1989.

Appendix

A.1 Proof of Lemma 2.3

Let $p \in \mathbb{R}^N$. Note that $f_*(p) = \inf_{x \in \mathbb{R}^N}\{p^\intercal x - \beta v(Ax)\}$. Letting $y = Ax$ and noticing that $x = A^{-1}y$, we have

$$f_*(p) = \inf_{y \in \mathbb{R}^N}\{p^\intercal(A^{-1}y) - \beta v(y)\} \tag{A1}$$

$$= \beta \inf_{y \in \mathbb{R}^N}\{(p^\intercal A^{-1}/\beta)y - v(y)\} \tag{A2}$$

$$= \beta \inf_{y \in \mathbb{R}^N}\{((A^{-1}/\beta)^\intercal p)^\intercal y - v(y)\}. \tag{A3}$$

Now (10) follows.

A.2 Proof of Theorem 3.1

Let $C(X)$ be the space of continuous bounded functions from X to $\mathbb{R}$ equipped with the sup norm $\|\cdot\|$. Then statement (a) holds with $C(X)$ replacing $\mathscr{C}(X)$ by Stokey and Lucas [24, Theorem 4.6]. Thus if $v \in \mathscr{C}(X) \subset C(X)$, then $Bv \in C(X)$; furthermore, Bv is concave by a standard argument (e.g. [24, p. 81]). Thus B maps $\mathscr{C}(X)$ into itself. Hence statement (a) holds. It is easy to see that $\mathscr{C}(X)$ equipped with the sup norm $\|\cdot\|$ is a complete metric space; thus statement (b) follows by the contraction mapping theorem [24, p. 50]. Finally, statement (c) holds by Stokey and Lucas [24, Theorem 4.3].