A FORWARD ALGORITHM FOR SOLVING OPTIMAL STOPPING PROBLEMS

ALBRECHT IRLE,* *University of Kiel*

Abstract

We consider the optimal stopping problem for $g(Z_n)$, where Z_n, $n = 1, 2, \ldots$, is a homogeneous Markov sequence. An algorithm, called forward improvement iteration, is presented by which an optimal stopping time can be computed. Using an iterative step, this algorithm computes a sequence $B^0 \supseteq B^1 \supseteq B^2 \supseteq \cdots$ of subsets of the state space such that the first entrance time into the intersection F of these sets is an optimal stopping time. Various applications are given.

Keywords: Optimal stopping; algorithm

2000 Mathematics Subject Classification: Primary 60G40; 60J05

1. Introduction

Let us consider the following general situation for a problem of optimal stopping. Starting with some probability space $(\Omega, \mathcal{A}, P)$, we have a filtration $(\mathcal{A}_n)_n$, where $n = 0, 1, 2, \ldots$, and an adapted real-valued stochastic process $(X_n)_n$ such that each X_n is integrable and $X_\infty = \limsup X_n$.

A stopping rule is a mapping $\tau \colon \Omega \to \{0, 1, 2, \ldots, \infty\}$ satisfying $\{\tau = n\} \in \mathcal{A}_n$ for all n, and we let $\mathcal{T}$ denote the set of all stopping rules. We assume that $\operatorname{E} X_\tau$ exists (possibly infinite) for all $\tau \in \mathcal{T}$. It is our aim to find a stopping rule τ^* satisfying

$$\operatorname{E} X_{\tau^*} = \sup_{\tau \in \mathcal{T}} \operatorname{E} X_\tau.$$

The following algorithm, whose conception goes back to Howard's policy improvement of dynamic programming, was treated in Irle (1980). We shall refer to this algorithm as *forward improvement iteration* (FII). To describe FII and state the result of Irle (1980), we use the following notation. Let

$$\mathcal{M} = \{\mathcal{C} = (C_n)_n \colon C_n \in \mathcal{A}_n \text{ for all } n, C_\infty = \Omega\}$$

and, for $\mathcal{C} \in \mathcal{M}$,

$$\mathcal{T}(\mathcal{C}) = \{\tau \in \mathcal{T} \colon \{\tau = n\} \subseteq C_n \text{ for all } n\},$$
$$\tau_n(\mathcal{C}) = \inf\{k \geq n \colon I(C_k) = 1\}, \quad \text{with } \tau(\mathcal{C}) = \tau_0(\mathcal{C}).$$

Here, $I(A)$ denotes the indicator of a set A; hence, $\tau_n(\mathcal{C})(\omega) = \inf\{k \geq n \colon \omega \in C_k\}$. We shall use the convention $\inf \varnothing = \infty$.

For $\mathcal{C} \in \mathcal{M}$, we define $\mathcal{C}^* \in \mathcal{M}$ by

$$C_n^* = \{\mathrm{E}(X_{\tau_{n+1}(\mathcal{C})} \mid \mathcal{A}_n) \leq X_n\} \cap C_n, \qquad C_\infty^* = \Omega.$$

Clearly, $\tau_n(\mathcal{C}), \tau_n(\mathcal{C}^*) \in \mathcal{T}(\mathcal{C})$. FII proceeds in the following way. Let

$$\mathcal{C}^0 = (\Omega)_n \quad \text{and, by induction,} \quad \mathcal{C}^k = (\mathcal{C}^{k-1})^*.$$

Define $\mathcal{D}$ by

$$D_n = \bigcap_k C_n^k.$$

Then it is obvious that $\tau(\mathcal{C}^k) \leq \tau(\mathcal{C}^{k+1})$ for any k and $\lim \tau(\mathcal{C}^k) = \tau(\mathcal{D})$. The following result on the validity of FII was given by Irle (1980).

Theorem 1. *Assume that*

$$\mathrm{E}\, X_{\lim \sigma_n} = \lim \mathrm{E}\, X_{\sigma_n},$$

for any increasing sequence of stopping rules $(\sigma_n)_n$. Then we obtain

$$\mathrm{E}\, X_{\tau(C^0)} \leq \mathrm{E}\, X_{\tau(C^1)} \leq \cdots \uparrow \mathrm{E}\, X_{\tau(D)}$$

and

$$\mathrm{E}\, X_{\tau(\mathcal{D})} = \sup_{\tau \in \mathcal{T}} \mathrm{E}\, X_\tau, \quad \textit{i.e. } \tau \textit{ is optimal.}$$

If $\mathcal{C}^{k+1} = \mathcal{C}^k$ for some k then $\mathcal{C}^{k+n} = \mathcal{C}^k$ for all n; hence, $\mathcal{D} = \mathcal{C}^k$, and the algorithm has found the optimal stopping rule in k steps.

We note that $\mathcal{C}^1 = (\{\mathrm{E}(X_{n+1} \mid \mathcal{A}_n) \leq X_n\})_n$; hence, $\tau(\mathcal{C}^1)$ is the 1-stage look-ahead rule, also known as the myopic rule. Theorem 1 may thus be used to provide a proof of optimality of this rule for monotone stopping problems in the sense of Chow *et al.* (1971). We come back to this in Section 3.2.

It is the aim of this paper to adapt the FII algorithm to the Markovian situation where $X_n = g(Z_n)$ for a homogeneous Markovian sequence. In the general situation, FII works in the space Ω, which may be viewed as the space of infinite sequences. Here, we will show in Section 2 that, in the Markovian setting, we can adapt FII so that it works in the state space of the Markovian sequence. This makes the algorithm easily applicable, and we shall provide various examples in Section 3. An attractive feature of the algorithm is its property that it can be viewed as a Monte Carlo algorithm which only uses simulations but no numerical calculations. Its performance in some specific situations will be described in Section 4.

Let us remark that the proofs of Theorems 1 and 2 are self-contained in the sense that they do not use any results from the theory of optimal stopping.

2. Markovian stopping problems

Let $(Z_n)_n$ be a homogeneous Markov process with respect to the underlying filtration. The measurable state space will be denoted by $(S, \mathcal{S})$. So we have

$$\mathrm{P}(Z_{n+1} \in B \mid \mathcal{A}_n) = Q(B, Z_n), \quad \text{for all } B \in \mathcal{S},$$

with respect to a kernel $Q \colon \mathcal{S} \times S \to [0, 1]$.

Let $g \colon S \to \mathbb{R}$ be measurable. We consider the problem of optimal stopping for

$$X_n = g(Z_n), \qquad n = 0, 1, \ldots.$$

To simplify the notation, we set

$$g(Z_\infty) = \limsup g(Z_n); \quad \text{hence,} \quad X_\infty = g(Z_\infty),$$

and we have $X_\tau = g(Z_\tau)$ for any stopping rule τ.

We write P_z for $P(\cdot \mid Z_0 = z)$ and E_z for $E(\cdot \mid Z_0 = z)$, where $z \in S$. As before, we assume that $E_z\, g(Z_\tau)$ exists for all stopping rules τ and that $z \in S$. We are looking for a stopping rule τ^* with

$$E_z\, g(Z_{\tau^*}) = \sup_{\tau \in \mathcal{T}} E_z\, g(Z_\tau), \quad \text{for all } z \in S.$$

We set

$$v(z) = \sup_\tau E_z\, g(Z_\tau), \qquad z \in S.$$

In the Markovian case, FII only has to work in the systems of subsets of the state space S. Let $B \subseteq S$, where B is measurable, and set

$$\mathcal{T}(B) = \{\tau \in \mathcal{T} : Z_\tau \in B \text{ on } \{\tau < \infty\}\},$$
$$\tau_n(B) = \inf\{k \geq n \colon Z_k \in B\}, \quad \text{with } \tau(B) = \tau_0(B).$$

Define $B^* \subseteq S$ by

$$B^* = \{z \colon g(z) \geq E_z\, g(Z_{\tau_1(B)})\} \cap B.$$

Clearly, $\tau_n(B), \tau_n(B^*) \in \mathcal{T}(B)$. We set

$$B^0 = S, \qquad B^k = (B^{k-1})^*, \qquad F = \bigcap_k B^k.$$

Obviously, $\tau(B^0) \leq \tau(B^1) \leq \cdots$ and $\lim \tau(B^k) = \tau(F)$.

Theorem 2. *Assume that*

$$E_z\, g(Z_{\lim \sigma_n}) = \lim_n E_z\, g(Z_{\sigma_n}),$$

for all z and all increasing sequences of stopping rules $(\sigma_n)_n$. Then, for any z, we obtain

(i) $E_z\, g(Z_{\tau(B^0)}) \leq E_z\, g(Z_{\tau(B^1)}) \leq \cdots \uparrow E_z\, g(Z_{\tau(F)})$,

(ii) $E_z\, g(Z_{\tau(F)}) = v(z)$, *i.e. $\tau(F)$ is optimal.*

Proof. The proof adapts the ideas of the proof of Theorem 1 of Irle (1980) to the Markovian situation and the state space setting. We outline the basic steps. Details may be added by comparison with the above reference.

(a) Consider sets B and B^* as above. For $\sigma \in \mathcal{T}(B)$, set

$$\sigma^* = \inf\{n \geq \sigma \colon Z_n \in B^*\}.$$

Then we can show that

$$E_z\, g(Z_{\sigma^*}) \geq E_z\, g(Z_\sigma),$$

from which (i) immediately follows.

(b) Using (a) we show in the next step that, for any stopping time σ, there exists $\tau \in \mathcal{T}(F)$ such that

$$\mathrm{E}_z\, g(Z_\tau) \geq \mathrm{E}_z\, g(Z_\sigma).$$

Such a τ is given by

$$\tau = \inf\{n \geq \sigma : Z_n \in F\}.$$

(c) In this step, we consider $\rho, \tau \in \mathcal{T}(F)$, such that $\rho \leq \tau$. We can prove that

$$\mathrm{E}_z\, g(Z_\rho) \geq \mathrm{E}_z\, g(Z_\tau).$$

(d) For the final step, let σ be any stopping time. Then by (b) there exists $\tau \in \mathcal{T}(F)$ such that

$$\mathrm{E}_z\, g(Z_\sigma) \leq \mathrm{E}_z\, g(Z_\tau).$$

By definition of $\tau_0(F)$ we have $\tau_0(F) \leq \tau$; hence, by (c), we obtain

$$\mathrm{E}_z\, g(Z_{\tau_0(F)}) \geq \mathrm{E}_z\, g(Z_\tau) \geq \mathrm{E}_z\, g(Z_\sigma).$$

Remark 1. If $B^{k+1} = B^k$ for some k, then

$$B^k = F.$$

So, if the state space S is finite, i.e. $|S| = n$, then, necessarily,

$$F = B^n;$$

hence, the algorithm terminates in at most n steps.

Remark 2. Let

$$F^* = \{z : g(z) = v(z)\}.$$

It follows, by induction, that

$$F^* \subseteq B^k \quad \text{for all } k; \quad \text{hence,} \quad F^* \subseteq F.$$

Since $\tau(F)$ is optimal in the situation of Theorem 2, $g(z) = v(z)$ for $z \in F$, showing that $F \subseteq F^*$ and

$$F = F^*.$$

Theorem 2 thus provides a proof of the optimality of $\tau(F^*)$ which is independent of the general results of optimal stopping theory.

Remark 3. Suppose that we have already found a set B^0 for which we know that

$$B^0 \supseteq \{z : g(z) = v(z)\} = F.$$

So we only have to look for an optimal stopping time within $\mathcal{T}(B^0)$. We may then start FII with this B^0 and obtain $F' = \bigcap_k B^k$. The proof of Theorem 2 applies and we obtain $\tau(F')$ as an optimal stopping time. As in Remark 2, it follows that $F' = \{z : g(z) = v(z)\} = F$.

Remark 4. To apply FII, we have to compute a sequence of sets

$$B^0 \supseteq B^1 \supseteq B^2 \supseteq \cdots .$$

At each step, we go from a certain subset B of the state space to the set B^* given by

$$B^* = \{z \colon g(z) \geq \mathrm{E}_z\, g(Z_{\tau_1(B)})\} \cap B.$$

For practical purposes, this can be done by simulating paths of the process $(Z_n)_n$, which provide simulated values of $g(Z_{\tau_1(B)})$ and Monte Carlo estimates for $\mathrm{E}_z\, g(Z_{\tau_1(B)})$. These estimates can then be used in the step to obtain B^*.

In the proof of Theorem 2, the condition $\lim \mathrm{E}_z\, g(Z_{\sigma_k}) = \mathrm{E}_z\, g(Z_\sigma)$ for increasing sequences $(\sigma_k)_k$ with $\lim \sigma_k = \sigma$ is used. We shall now show that this condition is not needed in the case of finite S.

Corollary 1. *Let S be finite. Then $\tau(F)$ is optimal, i.e.*

$$\mathrm{E}_z\, g(Z_{\tau(F)}) = v(z), \quad \text{for all } z.$$

Proof. As in Remark 1, we look at $F^* \subseteq F$ and $\tau(F^*)$. According to a well-known result of the theory of optimal stopping for Markov chains (see Shiryayev (1978, Chapter 2, Theorem 4)), we know that $\tau(F^*)$ is optimal and almost surely finite for all z. Looking at the proof of Theorem 2, it is immediate that we only have to obtain part (c) of the proof for $\tau = \tau(F^*)$; hence, for P_z-almost surely finite τ. Furthermore (see Irle (1980)), this follows if it can be shown that

$$\mathrm{E}_z\, g(Z_\sigma) \geq \liminf \mathrm{E}_z\, g(Z_{\sigma_n}),$$

for any increasing sequence σ_n with $\sigma = \lim \sigma_n \leq \tau$. Noting that $\sigma < \infty$ P_z-almost surely, we have

$$g(Z_\sigma) = \lim g(Z_{\sigma_n}), \quad \mathrm{P}_z\text{-almost surely,}$$

and

$$\mathrm{E}_z\, g(Z_\sigma) = \lim \mathrm{E}_z\, g(Z_{\sigma_n})$$

by dominated convergence.

3. Applications

3.1. Time-dependent pay-offs

3.1.1. Generalities. We assume that we have a pay-off of the form

$$g(Z_n, n), \qquad n = 0, 1, \ldots,$$

for a mapping $g \colon S \times \{0, 1, \ldots\} \to \mathbb{R}$. Then we may apply our result to the space–time Markov process (Z_n, n), $n \geq 0$. The iterative step is going from $B \subseteq S \times \{0, 1, \ldots\}$ to

$$B^* = \{(z, n) \in B \colon g(z, n) \geq \mathrm{E}_z\, g(Z_{\tau_1(B)}, n + \tau_1(B))\}.$$

Note that, by going over to the space–time process, we can also treat inhomogeneous Markov sequences in this manner.

3.1.2. *Finite horizon.* In the case of a stopping problem with finite time horizon N and pay-off

$$g(Z_n, n), \qquad n = 0, 1, \ldots, N,$$

we may set

$$g(z, k) = -\infty, \quad \text{for } k = N + 1, N + 2, \ldots.$$

Then, of course,

$$B^* \subseteq S \times \{N, N + 1, \ldots\},$$

and the FII only considers stopping times which are less than or equal to N.

3.1.3. *Linear costs.* Consider, for some $c > 0$,

$$g(Z_n, n) = f(Z_n) - cn, \qquad n = 0, 1, \ldots.$$

Then the condition defining B^* becomes

$$\begin{aligned}
f(z) - cn &\geq \mathrm{E}_z \, g(Z_{\tau_1(B)}, n + \tau_1(B)) \\
&= \mathrm{E}_z \, f(Z_{\tau_1(B)}) - cn - c \, \mathrm{E}_z \, \tau_1(B);
\end{aligned}$$

hence,

$$f(z) \geq \mathrm{E}_z \, f(Z_{\tau_1(B)}) - c \, \mathrm{E}_z \, \tau_1(B).$$

This reflects the well-known fact that the optimal stopping rule does not have explicit time dependence for linear costs.

3.1.4. *Discounting.* Let, for some $\alpha > 0$,

$$g(Z_n, n) = \alpha^n f(Z_n), \qquad n = 0, 1, \ldots.$$

Then B^* is defined by the condition

$$\begin{aligned}
\alpha^n f(Z) &\geq \mathrm{E}_z \, g(Z_{\tau_1(B)}, n + \tau_1(B)) \\
&= \alpha^n \, \mathrm{E}_z \, \alpha^{\tau_1(B)} f(Z_{\tau_1(B)}),
\end{aligned}$$

which becomes

$$f(z) \geq \mathrm{E}_z \, \alpha^{\tau_1(B)} f(Z_{\tau_1(B)}),$$

again without explicit time dependence.

3.2. The monotone case

There is a class of simple optimal stopping problems for which the 1-stage look-ahead stopping rule, which compares the current gain with the gain to be expected in the next step, is optimal under certain conditions of regularity; see Chow *et al.* (1971). These so-called monotone stopping problems have the following form in the Markovian setting. For our sequence $B^0, B^1, \ldots$, we have

$$B^1 = \{z \colon g(z) \geq \mathrm{E}_z \, g(Z_1)\},$$

and the 1-stage look-ahead rule is $\tau(B^1)$. We speak of a monotone-case stopping problem if

$$\mathrm{P}_z(Z_1 \in B^1) = 1, \quad \text{for all } z \in B^1.$$

This implies that

$$P_z(\tau_1(B^1) = 1) = 1, \quad \text{for all } z \in B^1;$$

hence,

$$B^1 = B^2 = F \quad \text{and} \quad \tau(B^1) = \tau(F).$$

So for monotone-case stopping problems, Theorem 2 and Corollary 1 show optimality of the 1-stage look-ahead rule.

The FII algorithm may also be described in terms of 1-stage look-ahead rules. We assume that FII has arrived at some set $B^k = B$, which is assumed to be a recurrent set for this discussion. We consider the restricted chain Z^B having state space B and transition probabilities $P_z(Z_1^B \in \cdot) = P_z(Z_{\tau_1(B)} \in \cdot)$. Then $B^* = \{z \in B : g(z) \geq E_z \, g(Z_1^B)\}$ and $\tau(B^*)$ is the 1-stage look-ahead rule for Z^B.

In some cases, FII may be performed analytically to find the optimal stopping time in explicit form. The following example is of this type.

3.3. Example

We look at a stopping problem which arises in mathematical finance. It amounts to finding the price of the American put option in a Cox–Ross–Rubinstein model. For some constant $a > 1$, the stock price is modelled by a multiplicative random walk $(Z_n)_n$ on the grid $\mathbb{Z} = \{a^l : l \in \mathbb{Z}\}$. We have, for some p, $0 < p < 1$,

$$P_z(Z_1 = az) = p = 1 - P_z\left(Z_1 = \frac{z}{a}\right).$$

For a discounting factor α, $0 < \alpha < 1$, and some $\gamma > 0$, we consider the stopping problem for

$$g(Z_n, n) = \alpha^n (\gamma - Z_n)^+, \qquad n = 0, 1, \ldots,$$

where $(\cdot)^+$ denotes the positive part of its argument. As discussed in Section 3.1.4, the explicit time dependence vanishes. The optimal stopping set F is a subset of the state space $\{a^l : l \in \mathbb{Z}\}$ and the FII operates on this set.

We assume that $E_1 Z_1 \geq 1$ and $E_1 \log Z_1 \leq 0$; see Beibel and Lerche (1997) for the corresponding assumptions in the continuous-time analogue for geometric Brownian motion. In particular, we have $\liminf Z_n = 0$ P_z-almost surely for any z. It is obvious that

$$F \subseteq \{a^l : a^l < \gamma\},$$

so we start FII with

$$B^0 = \{a^l : a^l < \gamma\} = \{a^l : l \leq k\},$$

say; see Remark 3. The condition defining $B^1 = (B^0)^*$ is

$$\gamma - a^j \geq \alpha[p(\gamma - a^{j+1}) + (1 - p)(\gamma - a^{j-1})], \quad \text{for } j < k,$$

or, equivalently,

$$\gamma(1 - \alpha) \geq a^j(1 - \alpha(pa + (1 - p)a^{-1})) = a^j(1 - \alpha E_1 Z_1).$$

Because $E_1 Z_1 \geq 1$, this condition is fulfilled for any $j < k$. For $j = k$, the condition becomes

$$\gamma - a^k \geq E_{a^k} \alpha^{\tau_1(B^0)} (\gamma - Z_{\tau_1(B^0)})^+$$
$$= \alpha[(1 - p)(\gamma - a^{k-1}) + p(\gamma - a^k) E_{a^{k+1}} \alpha^\tau],$$

where $\tau = \inf\{n\colon Z_n = a^k\}$. Obviously the distribution of τ with respect to $P_{a^{k+1}}$ is independent of k and we write

$$e(\alpha) = E_{a^{k+1}} \alpha^\tau.$$

Note that

$$e(\alpha) = (1 - p)\alpha + p\alpha e(\alpha)^2 = \frac{1}{2p\alpha} - \sqrt{\frac{1}{4p^2\alpha^2} - \frac{1 - p}{p}};$$

see, for example, Feller (1957, Chapter XIV.4). Then the condition defining $B^1 = (B^0)^*$ takes the form

$$\gamma(1 - [(1 - p)\alpha + p\alpha e(\alpha)]) \geq a^k(1 - [(1 - p)\alpha a^{-1} + p\alpha e(\alpha)]).$$

Thus,

$$B^1 = B^0 = \{a^l : l \leq k\} \quad \text{if } a^k \leq \gamma \frac{1 - ((1 - p)\alpha + p\alpha e(\alpha))}{1 - ((1 - p)\alpha a^{-1} + p\alpha e(\alpha))}$$

and

$$B^1 = \{a^l : l \leq k - 1\} \quad \text{if } a^k > \gamma \frac{1 - ((1 - p)\alpha + p\alpha e(\alpha))}{1 - ((1 - p)\alpha a^{-1} + p\alpha e(\alpha))}.$$

Repeating this argument, we see that the optimal stopping region is given by

$$F = \left\{ a^l : a^l \leq \gamma \frac{1 - ((1 - p)\alpha + p\alpha e(\alpha))}{1 - ((1 - p)\alpha a^{-1} + p\alpha e(\alpha))} \right\}.$$

4. Simulation studies

As the following examples will show, FII provides a viable algorithm for solving stopping problems with infinite time horizon without approximating using finite time horizon problems.

This may also be seen as a possibility for the reverse approximation. Consider a problem with a very large time horizon where backwards induction methods are no longer computationally feasible. We may solve the corresponding infinite-horizon problem by using FII and use the solution as an approximate solution for the large finite-horizon problem.

In the following examples we have an infinite time horizon with discounting. For each iteration, 10 000 simulated paths in each state are generated. A PC with a 1900MHz processor was used.

4.1. Examples – part 1

The state space of the chain $(Z_n)_n$ is

$$\{0, 1, \ldots, 20\} \times \{0, 1, \ldots, 20\}.$$

The pay-off function is

$$g(Z_n, n) = \alpha^n f(Z_n),$$

with

$$f(5, 5) = 10,$$
$$f(5, 15) = f(15, 15) = 0,$$
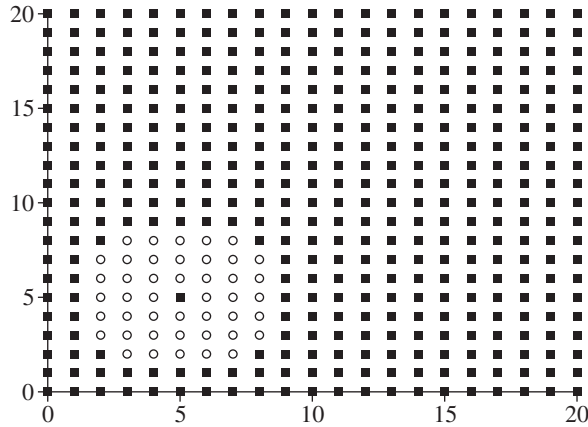$$f(x, y) = 5 \quad \text{otherwise.}$$

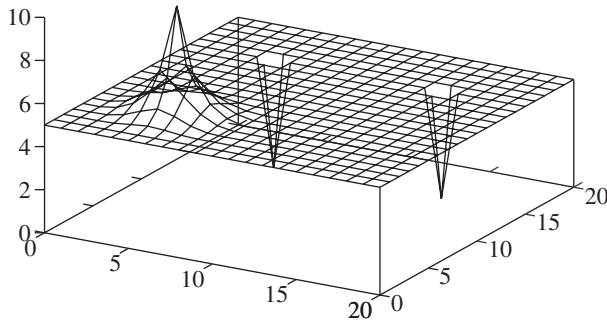FIGURE 1: Continuation points (circles) and stopping points (squares).

FIGURE 2: Optimal values.

For the chain, we assume absorbtion at the points of minimal pay-off, i.e. at (5, 15) and (15, 15). Otherwise, the transition probabilities are given as

$$P(x + 1, y \mid x, y) = 0.25,$$
$$P(x - 1, y \mid x, y) = 0.25,$$
$$P(x, y + 1 \mid x, y) = 0.25,$$
$$P(x, y - 1 \mid x, y) = 0.25,$$

with obvious modifications at the boundaries; for example $P(1, y \mid 0, y) = 0.5$ and $P(19, y \mid 20, y) = 0.5$.

Example 1. We have a discount factor of $\alpha = 0.99$, and termination after eight iterations in 90.33 seconds; see Figures 1 and 2.

Example 2. We have a discount factor of $\alpha = (0.98)^{1/20}$, and termination after 21 iterations in 478.01 seconds; see Figures 3 and 4.
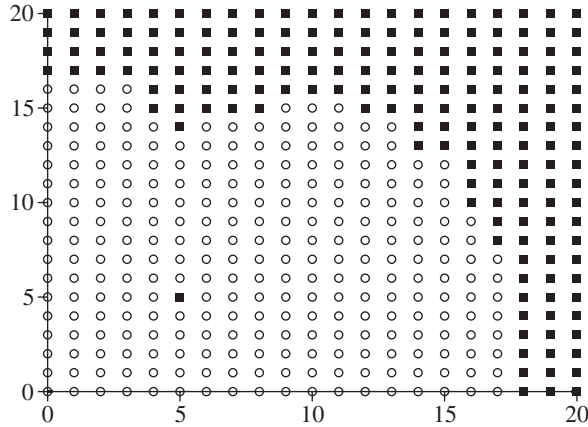
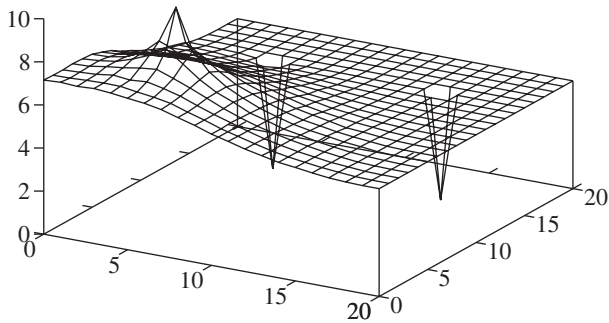FIGURE 3: Continuation points (circles) and stopping points (squares).

FIGURE 4: Optimal values.

4.2. Examples – part 2

The state space of the chain $(Z_n)_n$ is

$$\{0, 1, \ldots, 10\} \times \{0, 1, \ldots, 10\}.$$

We assume absorbtion at the boundaries. Otherwise, the transition probabilities are given as

$$P(x + 1, y \mid x, y) = 0.7 \times 0.6,$$
$$P(x - 1, y \mid x, y) = 0.7 \times 0.4,$$
$$P(x, y + 1 \mid x, y) = 0.3 \times 0.6,$$
$$P(x, y - 1 \mid x, y) = 0.3 \times 0.4.$$

Example 3. We have a pay-off of $f(x, y) = 0$ at the boundaries and $f(x, y) = (x + y - 5)^+$ otherwise. We have a discount factor of $\alpha = 0.9$, and termination after five iterations in 17.59 seconds; see Figures 5 and 6.
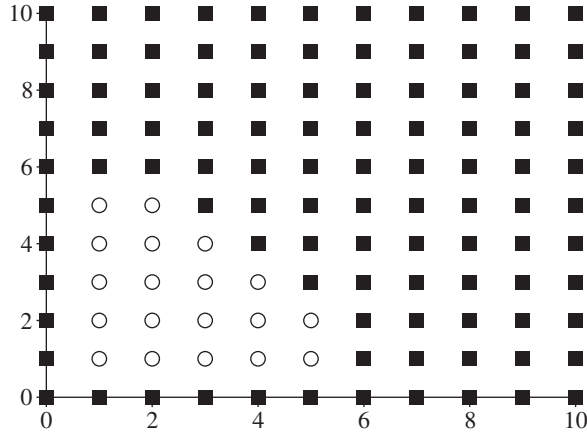
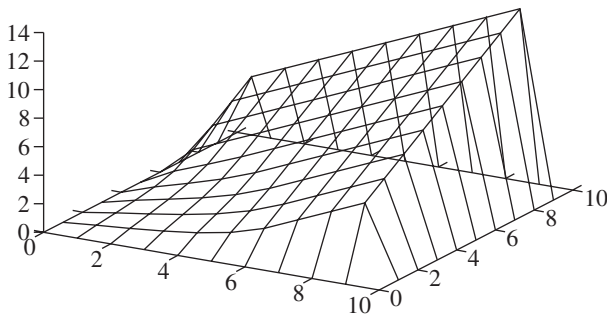FIGURE 5: Continuation points (circles) and stopping points (squares).

FIGURE 6: Optimal values.

Example 4. We have a pay-off of $f(x, y) = 0$ at the boundaries and $f(x, y) = ((x+y-5)^+)^2$ otherwise. We have a discount factor of $\alpha = 0.9$, and termination after five iterations in 23.93 seconds; see Figures 7 and 8.

4.3. Discussion

If we look at Section 4.1 without discounting, i.e. $\alpha = 1$, then the optimal stopping rule obviously takes the following form. Stop in the states (5, 5), (5, 15), and (15, 15), and in the eight neighbouring states of (5, 15) and (15, 15). Examples 1 and 2 show how the discount factor α substantially enlarges the optimal stopping region. Even a discount factor very close to 1, as in Example 2, leads to a stopping region which differs noticeably from that for $\alpha = 1$. We use $\alpha = 0.98^{1/20}$, the 20 stemming from the size of the state space, but of course any other α very close to 1 would lead to similar findings.

In Section 4.2, there is forced stopping at the boundaries and no stopping at interior points with $x + y \leq 5$. Examples 3 and 4 show how the square function pay-off enlarges the continuation set. We point out that, in terms of option pricing, the pay-off is that of a call option for two stocks with knock-out at the boundaries.
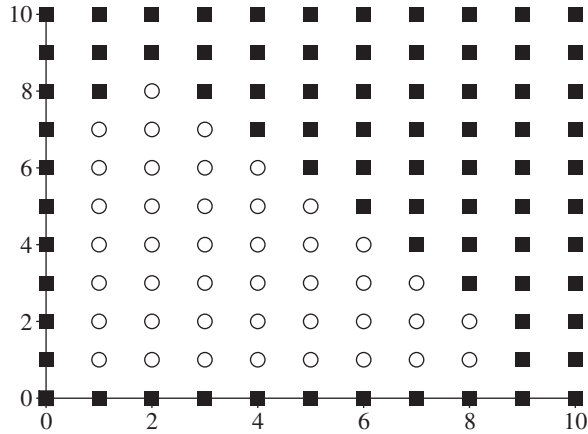
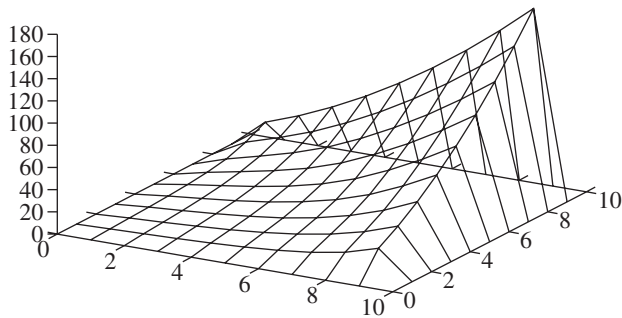FIGURE 7: Continuation points (circles) and stopping points (squares).

FIGURE 8: Optimal values.

Acknowledgement

References

BEIBEL, M. AND LERCHE, H. R. (1997). A new look at optimal stopping problems related to mathematical finance. *Statistica Sinica* **7,** 93–108.

CHOW, Y. S., ROBBINS, H. AND SIEGMUND, D. (1971). *Great Expectations: The Theory of Optimal Stopping.* Houghton Mifflin, Boston, MA.

FELLER, W. (1957). *An Introduction to Probability and Statistics,* Vol. I, 2nd edn. John Wiley, New York.

IRLE, A. (1980). On the best choice problem with random population size. *Z. Operat. Res.* **24,** 177–190.

SHIRYAYEV, A. N. (1978). *Optimal Stopping Rules.* Springer, New York.