Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
Journal of Economic Dynamics & Control 28 (2004) 1437–1460
JOURNAL OF Economic Dynamics & Control
www.elsevier.com/locate/econbase

Investment under uncertainty: calculating the value function when the Bellman equation cannot be solved analytically

Thomas Dangl[a], Franz Wirl[b,*]

[a] *Vienna University of Technology, Theresianumgasse 27, A-1040 Vienna, Austria*
[b] *University of Vienna, Brűnnerstr. 72, A-1210 Vienna, Austria*

Abstract

The treatment of real investments parallel to financial options if an investment is irreversible and facing uncertainty is a major insight of modern investment theory well expounded in the book *Investment under Uncertainty* by Dixit and Pindyck. The purpose of this paper is to draw attention to the fact that many problems in managerial decision making imply Bellman equations that cannot be solved analytically and to the difficulties that arise when approximating the required value function numerically. This is so because the value function is the saddlepoint path from the entire family of solution curves satisfying the differential equation. The paper uses a simple machine replacement and maintenance framework to highlight these difficulties (for shooting as well as for finite difference methods). On the constructive side, the paper suggests and tests a properly modified projection algorithm extending the proposal in Judd (J. Econ. Theory 58 (1992) 410). This fast algorithm proves amendable to economic and stochastic control problems (which one cannot say of the finite difference method).

JEL classification: D81; C61; E22

Keywords: Stochastic optimal control; Investment under uncertainty; Projection method; Maintenance

1. Introduction

Without doubt, the book of Dixit and Pindyck (1994) is one of the most stimulating books concerning management sciences. The basic message of this excellent book is the following: Considering investments characterized by uncertainty (either with

* Corresponding author. Tel.: +43-1 4277 38101; fax: +43-1 4277 38104.
E-mail address: franz.wirl@univie.ac.at (F. Wirl).

0165-1889/$ - see front matter
doi:10.1016/S0165-1889(03)00110-6

respect to the investment costs or the value of the project) and irreversibility, standard calculations of the (expected) net present value (NPV) of profits can lead to highly misleading suggestions, because waiting has a potentially positive value. For example, a project may be profitable on a NPV basis, but waiting until the investment costs fall sufficiently avoids the potential loss (included with a corresponding probability in the calculation of expected present value) and thus may increase profits. The basic fault of the standard (or orthodox as Dixit and Pindyck prefer to say) NPV method is the implicit, tacit but implausible assumption that a project (e.g., an additional plant, capacity expansion) must be either carried out today or cannot be carried out at all.

The purpose of this paper is not to question this important insight but to draw attention to the following fact: Although all what a stochastic dynamic optimization problem requires is to solve "a non-linear differential equation, which needs numerical solutions in most cases", Dixit (1993, p. 39 for a particular example), (i) standard shooting algorithms using either an Euler or higher order algorithms such as the Runge–Kutta algorithm are likely to fail to approximate the required solution, (ii) finite difference methods which are applied in valuing financial options suffer from very slow convergence and (iii) collocation methods (such as the Chebyshev projection method proposed in Judd, 1992) offer a fast and robust alternative.

Real option analysis reduces to the valuation of American put or call options. However, managerial decision-making is in many cases more complex than finding the optimal exercise strategy of an American style options contract. While financial option valuation usually assumes a complete and arbitrage free market in which the individual investor is a price taker, managerial decision-making often influences the dynamics of the state variable. This generally results in non-linear differential equations that have to be satisfied by the value function of the corresponding real option. The non-linearity usually prevents analytical, closed form solutions and, furthermore, leads to poor convergence behavior of shooting algorithms. Tree methods and finite difference methods can be adapted,[1] but require extensive computations if the real option's time to maturity is large (or infinite as in many applications), because these methods have to evaluate huge trees.

This paper is organized as follows. Section 2 introduces a simple example, the optimal shutdown of a machine, which is in fact taken from Dixit and Pindyck (1994) and which allows for an explicit analytical solution. This example is used to highlight the problems of standard numerical procedures. Almost all genuine optimization problems (except for the stopping problems to which Dixit and Pindyck (1994) and most follow ups confine by and large their analysis) lead to non-linear differential equations that do not allow for an analytical solution. This fact excludes analytical techniques and restricts, according to the results of Section 2, the application of standard numerical routines. Section 3 extends the example of optimal shutdown by allowing for continuous maintenance, where the Bellman equation leads to a non-linear differential equation. Section 4 sketches an algorithm using projection methods and Chebyshev polynomials

[1] For an introduction to the binomial tree method in financial options valuation see Cox et al. (1979), the trinomial tree method is studied, e.g., in Kamrad and Ritchken (1991), and finite difference methods in this context are presented in Brennan and Schwartz (1977, 1978) and Hull and White (1990).

based on Judd (1992) that is applied in Section 5 to the optimization problem introduced in Section 3.

2. Optimal shutdown – no maintenance

The following framework – the optimal shutdown of a machine (or any kind of a profit generating equipment) – serves as *pars pro toto* to highlight the difficulties in calculating the value function using standard differential equation algorithms in the absence of the analytical solution of the Bellman differential equation.

Let $\pi(t)$ denote the profit flow in period t from using a particular equipment and it is assumed that this profit follows a Brownian motion (with drift $a < 0$)

$$\mathrm{d}\pi(t) = a\,\mathrm{d}t + \sigma\,\mathrm{d}z, \;\; \pi(0) = \pi_0. \tag{1}$$

The deterministic part, $a\,\mathrm{d}t$, states that the expected change follows a linear depreciation. The second term, $\sigma\,\mathrm{d}z$, adds a stochastic element to the actual evolution of the profits in which $\mathrm{d}z$ is the increment of a standardized Wiener-process z and the parameter σ denotes the instantaneous standard deviation of the profit flow. Fig. 1 shows a few realizations of this process for the base case parameters in Table 1, all starting with a profit of \$1, an average depreciation of 10 cents/a (thus expected lifetime equals 10 years) and a standard deviation of 20 cents/a. One of these realizations follows by and large the expected evolution, one shows even increasing profits, while two others

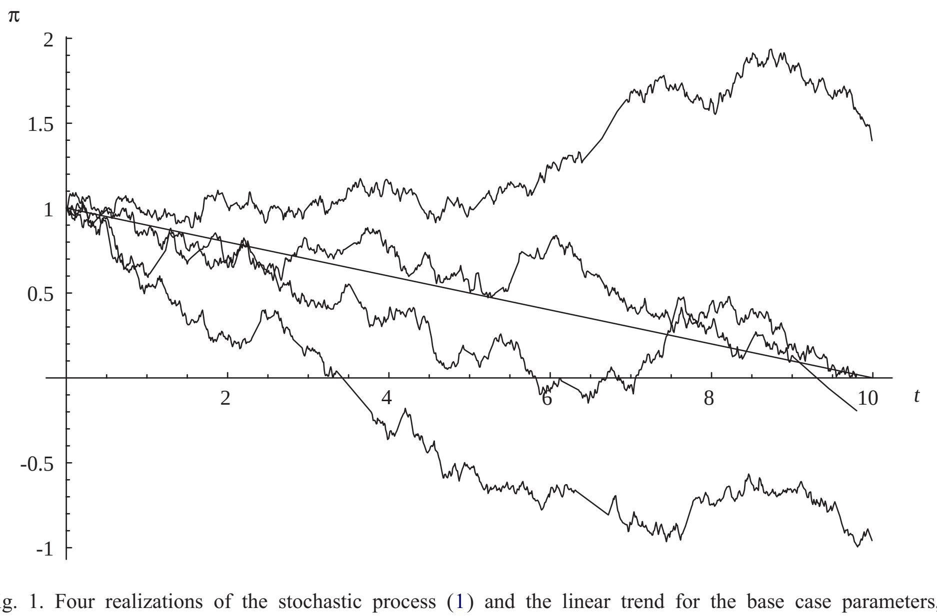

Fig. 1. Four realizations of the stochastic process (1) and the linear trend for the base case parameters, $a = -0.1$, $\sigma = 0.2$, and $\pi_0 = 1$.

Table 1
Base case parameters

a	−0.1
σ	0.2
r	0.1
π_0	1

lead to deficits already after 3 and 6 years, respectively, yet one of them immediately returns to be profitable which highlights that deficits are not irreversible.

It is assumed that equipment that is once removed cannot be reinstalled and this irreversibility requires a forward looking behavior when deciding to eliminate the equipment. Let F and r denote the present value of the claim on the profit flow π (i.e., the value function of the equipment) and the discount rate respectively.[2] Then F is determined by

$$F(\pi,t) = \max_{\{T\}} E \int_0^T e^{-rt}\pi(t)\,dt, \tag{2}$$

where T is a random stopping time at which the machine is eliminated (at a salvage value of 0).

The Bellman equation for this optimal stopping problem is

$$F(\pi,t) = \max\{0, \pi\,dt + (1 + r\,dt)^{-1}E[F(\pi + d\pi, t + dt)]\}. \tag{3}$$

That is, the value function $F(\pi,t)$ describes the maximum expected present value of profits given the current flow is π. The first argument, 0, applies if the machine is eliminated, otherwise the second argument describes the value from continuation. Thus, the value F of the machine is a stochastic process that is determined endogenously contingent on the realization of the process π, so that $F(\pi,t) = F(\pi)$, because there is no explicit dependence on (calendar) time t. The firm's optimal decision can be characterized by a threshold π^*, such that it is optimal to use the equipment (or plant) as long as $\pi > \pi^*$, and to shut it down when the profit flow hits his threshold for the first time.[3]

In the domain of continuation the value function F has to satisfy the Bellman equation[4]

$$rF = \pi + aF_\pi + 1/2\sigma^2 F_{\pi\pi}. \tag{4}$$

[2] The uncertainty in the evolution of the profitability of a single machine has an important idiosyncratic component that cannot be replicated by a portfolio of traded financial contracts. This prevents the application of preference free valuation techniques, and therefore, we assume that the decision maker is risk neutral.

[3] i.e., the random stopping time T is determined by $T = \inf\{\tau > 0 | \pi(\tau) \leqslant \pi^*\}$.

[4] Eq. (4) is a special case of the general Bellman equation which will be derived in Section 3. We use subscripts to indicate differentiation instead of primes in order to refer to the argument and because the Bellman equation is in general a partial differential equation.

Therefore, the so far unknown value function $F(\pi)$ is the solution of a second order differential equation that requires corresponding boundary conditions: The salvage value is zero, therefore

$$F(\pi^*(t)) = 0. \tag{5}$$

Of course, a highly profitable plant will not be shut down[5] and moreover is unlikely to be abandoned in the near future due to properties of the diffusion process (1). Therefore, the value of the opportunity to shut down vanishes for $\pi \to \infty$. Consequently, the value function must converge to the value of a unit that cannot be mustered out, even if it produces losses. This implies a boundary condition on the upper side (i.e., for very large profitability)[6]

$$\lim_{\pi \to \infty} [F(\pi) - (a/r^2 + \pi/r)] = 0, \tag{6}$$

where the term between parentheses is the value of a perpetually operating unit (see Appendix A).

Condition (5) is applied at a free boundary π^*, because the firm is free to choose the threshold level at which to stop. The first order condition of optimal choice of the exit threshold π^* is

$$F_\pi(\pi^*(t)) = 0. \tag{7}$$

For a derivation of this 'smooth pasting' condition see, e.g., Dixit (1993).

Given these three conditions (5)–(7), this simple stopping problem can be solved analytically which is not possible for the extended model in Section 3. The analytical solution of the pure stopping problem serves as a benchmark to demonstrate the problems of standard numerical procedures to approximate the value function and to compute the optimal policy (i.e., the optimal stopping threshold π^*).

The general solution of the inhomogeneous, linear differential equation (4) is:

$$F(\pi) = [(a/r^2) + (\pi/r)] + c_1 \exp(\lambda_1 \pi) + c_2 \exp(\lambda_2 \pi). \tag{8}$$

The term between the squared brackets in (8) is the value if the flow cannot be stopped, which is a particular solution of (4) (differentiation proves this). $\lambda_1 > 0$ and $\lambda_2 < 0$ denote the roots of the characteristic polynomial of the corresponding homogeneous differential equation:

$$\lambda_{12} = (-a \pm \sqrt{a^2 + 2\sigma^2 r})/\sigma^2. \tag{9}$$

The value of the machine given in (8) can therefore be interpreted as the sum of the value of the perpetual profit flow (the value between squared brackets) and the value of the option to terminate this flow when the profit flow hits π^* (characterized by the

[5] Generally, the firm has also the opportunity to abandon production at some upper threshold. However, the fact that a productive machine is never shut down pushes this upper threshold towards infinity; see boundary condition (6).

[6] See Appendix A for a derivation of the value of the perpetual flow, $a/r^2 + \pi/r$.

two exponential terms). Since $\lambda_1 > 0$, boundary condition (6) can only be satisfied for

$$c_1 = 0. \tag{10}$$

This means that the option to shut down is worthless for a very high profitable unit, which is an alternative way to state that it is not optimal to close it. $c_1 \neq 0$ would add exponentially growing gains (or losses, depending on the sign of c_1) that are inconsistent with the fundamentals. Economists call this the exclusion of 'Ponzi-games', or of speculative bubbles. Hence, the closed form solution of the value function is

$$F(\pi) = \begin{cases} \dfrac{a}{r^2} + \dfrac{\pi}{r} + c_2 \exp(\lambda_2 \pi) & \text{for } \pi \geqslant \pi^*, \\ 0 & \text{for } \pi < \pi^*. \end{cases} \tag{11}$$

Application of the boundary condition (5), the optimality condition (7), and the parameters of Table 1 to (11) reproduces the result quoted in Dixit and Pindyck (1994), $\pi^* = -0.17082039\ldots$ and $c_2 = 10.1188\ldots$. This states that an equipment producing up to a 17% loss (compared with the initial profit of \$1) should be kept operating. The reason is that a shutdown is irreversible, but the machine might still yield profits in the future (see the trajectory in Fig. 1 which after losses around $t \approx 6$ reverses to profits).

A first numerical approach is a modified shooting method employing a Runge–Kutta algorithm.[7] Starting at a guess of π^* we apply conditions (5) and (7) at this boundary and compute the corresponding value function using a Runge–Kutta algorithm. The asymptotic behavior required by condition (6) serves as a criterion for the quality of the guess. Since the guess certainly deviates from the actually optimal threshold, the implicitly determined constant c_1 deviates from zero. Hence, this numerical solution contains an exponentially growing function and will thus not show the asymptotic behavior, no matter how good the guess is. That is, the value function is the (unique) saddlepoint path from the entire family of solutions of this differential equation. While an analytical solution allows picking the saddlepoint path (this is essentially the technique applied in Dixit–Pindyck (1994) that amounts here to set $c_1 = 0$), it is impossible to determine a saddlepoint path by guessing initial conditions and this is shown in Fig. 2 for two 'good' guesses.

The sign of the deviation from the asymptotic solution indicates whether the guess is above or below the actual optimal threshold and this allows applying search algorithms. Due to the linearity of the differential equation, this shooting approach converges to the analytically determined optimum π^*. However, the quality of the approximation of the value function is restricted to a small interval (say $\pi < 3$) even if the computed threshold deviates from the optimum by less then 10^{-15}. Fortunately, this poor approximation of the value function in Fig. 2 is of minor relevance for this stopping problem,

[7] The standard shooting method is formulated for boundary value problems where the value of the function is given at the boundaries but the initial slope has to be determined, see e.g., Deuflhard and Bornemann (2002). Our problem is a free boundary problem. Both, the value and the slope of the value function are given (see conditions (5) and (7)) at free boundary, π^*, which has to be determined.

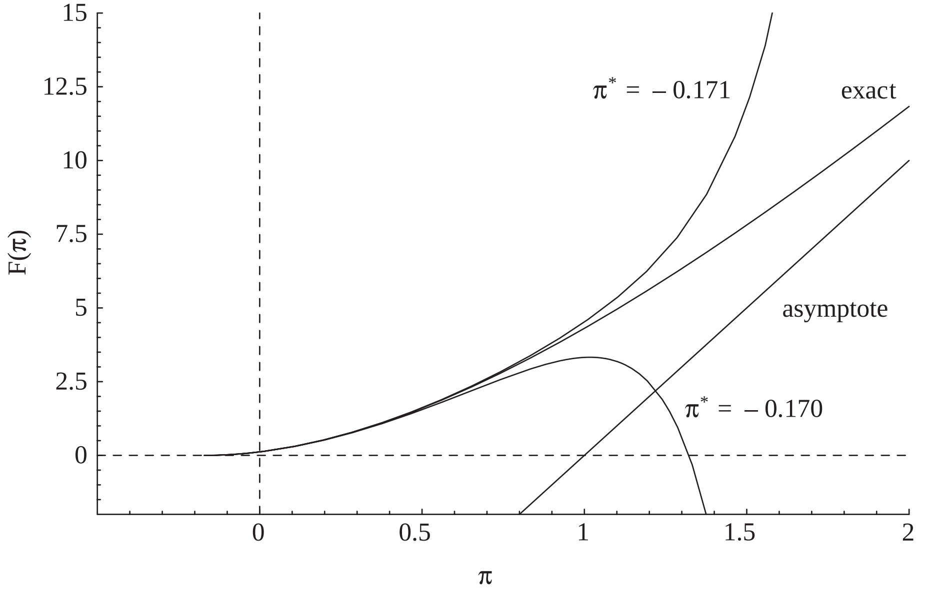

Fig. 2. Value function $F(\pi)$ and numerical solutions of the differential equation (4) with the boundary conditions (5) and (7) and guesses for π^* ($a=-0.1$, $\sigma=0.2$, $r=0.1$).

because the value function itself is of no interest for decision making, but only the threshold π^*.

The second numerical approach is the explicit finite difference method proposed in Brennan and Schwartz (1977) for the valuation of American style financial options. This method computes the value function at discrete points in the state/time space (analogous to tree methods). Starting from a terminal condition at maturity of the contract it works backwards in time applying finite difference approximations for the derivatives. To deal with the perpetual real option to shut down, we have first to reformulate the problem (for details see Appendix A) by adding an artificial expiration date at some $\bar{T}$. That is, it is assumed that the firm is allowed to shut down the machine only before $\bar{T}$, if the machine is not eliminated before $\bar{T}$ then it has to run forever such that the corresponding terminal condition is $F(\pi,\bar{T})=\max\{0,a/r^2+\pi/r\}$. Sufficient accuracy requires small increments in time and state, and approximating the stationary solution calls for a large $\bar{T}$ (such that the obtained solution is 'invariant' with respect to the choice of $\bar{T}$). These requirements make the application of finite difference methods computationally expensive. Fig. 3 shows the convergence behavior of this method by plotting the absolute error of the determined exit threshold as a function of the computing time for several choices of increments (the calculations are performed using GNU gcc on an AMD 1200 MHz PC running SuSE Linux 7.3). The time to maturity is set to $\bar{T}=100a$. The number of time steps ranges from 1000 to 512000. The step size in the state space is set to $\Delta\pi=\sigma\sqrt{\Delta t}$, in which case the finite difference method is equivalent to a trinomial tree approach. We see that calculation times of up

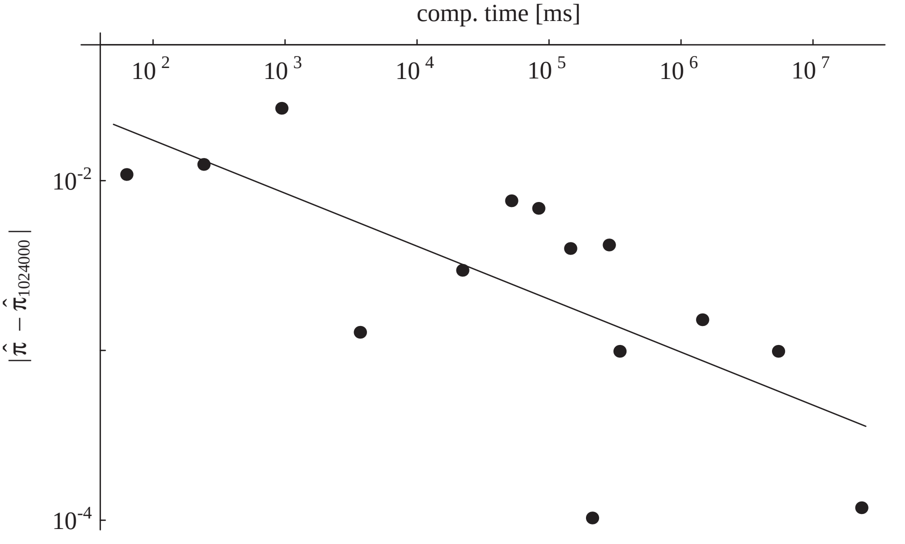

Fig. 3. Absolute error of the exit threshold determined by the finite difference method versus computing time.

to 6.5 h have to be accepted in order to approximate the exit threshold by an error less than 10^{-4} (measured not against the true π^* but against the approximation obtained for discretizing time up to $\bar{T}$ into $2^{10}10^3$ steps to account for the limit imposed by the choice of $\bar{T}$, denoted $\pi_{1024000}$ in Fig. 3; the 'true' error is even larger). Of course, one can speed up the convergence of finite difference methods (e.g., using the Richardson extrapolation), but the convergence behavior of the projection approach presented in Section 4 is superior even if the finite difference method could be accelerated by several orders of magnitude.

3. Optimal maintenance and shutdown – model

The profit obtained from the operation of a machinery depreciates linearly at the rate a, but this depreciation can now be reduced by care, maintenance, repair, etc., denoted u. However this maintenance is costly, $C(u)$, $C' \geqslant 0$, $C'' > 0$. Although it is plausible to assume that the costs become very large if maintenance tries to stop the natural decay completely, i.e., $C \to \infty$ for $u \to -a$, we assume for reasons of simplicity quadratic costs, $C(u) = 1/2cu^2$. In addition, the same stochastic term as in (1) affects the evolution of profit flow π, see (13). Summarizing, allowing for maintenance yields the following dynamic, stochastic optimization problem:

$$\max_{\{u(t),t\in[0,T],T\}} E\int_0^T \mathrm{e}^{-rt}\left(\pi(t)-\frac{1}{2}\,cu^2(t)\right)\,\mathrm{d}t, \tag{12}$$

$$\mathrm{d}\pi = (a+u)\,\mathrm{d}t + \sigma\,\mathrm{d}z. \tag{13}$$

The Bellman equation for this problem (12) and (13) is

$$rF = \max_u \{\pi - 1/2cu^2 + (a+u)F_\pi + 1/2\sigma^2 F_{\pi\pi}\}. \tag{14}$$

The optimization on the right hand side of (14) yields the necessary and sufficient condition $C'(u) = F_\pi$, which is economically plausible: the marginal costs of maintenance must equal the expected present value of marginal profits. This condition can be explicitly solved for the optimal maintenance strategy due to the quadratic cost function [8]

$$u^* = F_\pi / c. \tag{15}$$

Hence, optimal maintenance depends on the derivative of the so far unknown value function F. As a consequence, the value function must be calculated very accurately in order to determine optimal maintenance. Substitution of (15) into (14) yields now a non-linear, inhomogeneous, second order differential equation

$$\begin{aligned} rF &= \pi - 1/2c(F_\pi/c)^2 + (a + F_\pi/c)F_\pi + 1/2\sigma^2 F_{\pi\pi} \\ &= \pi + aF_\pi + 1/2F_\pi^2/c + 1/2\sigma^2 F_{\pi\pi}, \end{aligned} \tag{16}$$

which has to be solved subject to the two boundary conditions (5) and (7). The transversality condition,

$$\lim_{\pi\to\infty} (F(\pi) - [a/r^2 + 1/(2cr^3) + \pi/r]) = 0, \tag{6'}$$

is a generalization of (6). The term between the squared brackets is a particular solution of (16), which corresponds to the case that the machine cannot be eliminated irrespective of the losses it might generate (see Appendix A). At high levels of profitability, the value of the machine must converge to the value of this perpetual flow. Despite this simple linear solution if a shutdown is infeasible, it seems to be impossible (at least for us) to obtain an explicit, analytical solution of (16) that meets the boundary conditions (5), (6′), and (7). The fact that the Bellman equation does not allow for an analytical solution is not just a peculiar feature of our model but applies to numerous genuine stochastic, continuous time optimizations and can be found in several models, e.g., see Raman and Chatterjee (1995), Eq. (10) and footnote 5. As a consequence, we cannot eliminate the ‘wrong’ exponential term since we lack the corresponding analytical solution.

Due to the non-linearity of (16) the shooting method described in Section 2 fails to converge to an optimal exit threshold π^*, no matter how good the initial guess is, and consequently the computed value function diverges almost immediately to $+\infty$ or $-\infty$. In contrast to the pure stopping problem in Section 2, the value function (actually its derivative) is required to determine the maintenance efforts ($u^* = F_\pi/c$, see (15)).

[8] Optimal maintenance $u(t)$ is an endogenous stochastic process depending on the realization of π, $u(t) = u(\pi(t))$.

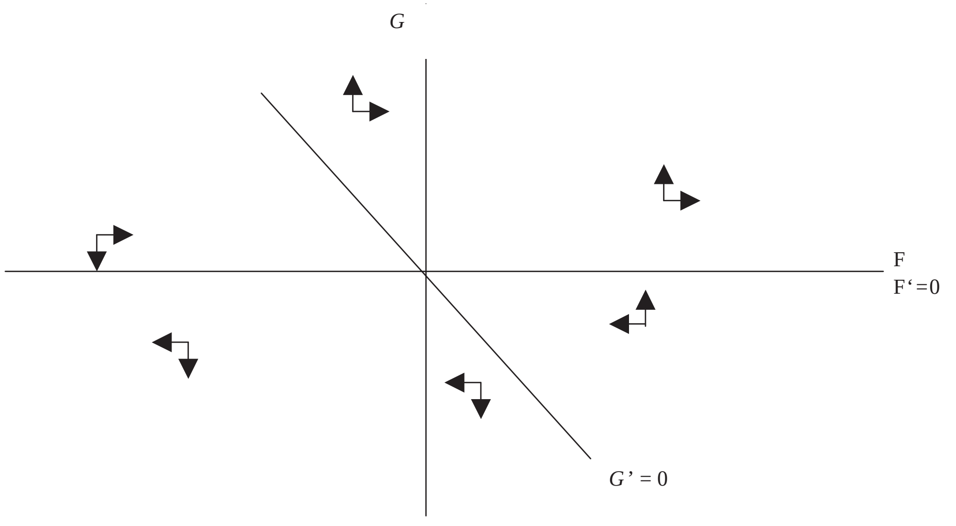

Fig. 4. Phase diagram of the autonomous system (17).

One possibility to overcome this failure of shooting approaches applied to non-linear differential equations is to apply multiple shooting. These methods divide the state space into several subintervals. On each of these intervals a simple shooting method is performed, however, the derivation of the associated equation system that connects the subintervals to ensure joint convergence is highly complex (see, e.g., Deuflhard and Bornemann (2002) for a recent discussion of multiple shooting methods). The finite difference method introduced in the previous section converges, but again, a sufficiently accurate approximation requires substantial computing time. However, we show in Sections 4 and 5 that the algorithm developed by Judd (1992) is suitable to solve stochastic, dynamic optimization problems numerically.

Finally we give a geometric explanation why direct (shooting) attempts to solve the differential equation (16) subject to the boundary conditions (5) and (7) fail. To highlight the point that the value function is a saddlepoint, consider the autonomous and linearized part of the second order differential equation (16) and define $G = F'$, thus $G' = F''$:

$$\begin{pmatrix} F' \\ G' \end{pmatrix} = \begin{bmatrix} 0 & 1 \\ \dfrac{2r}{\sigma^2} & -\dfrac{1}{\sigma^2}(2a + G/c) \end{bmatrix} \begin{pmatrix} F \\ G \end{pmatrix}. \tag{17}$$

The determinant of the Jacobian, $-2r/\sigma^2$, is negative so that the corresponding equilibrium (the origin in this autonomous system) is a saddlepoint. Therefore, only a single solution curve will converge to the origin (see Fig. 4) and no matter how close we start to the saddlepoint with a Runge–Kutta algorithm, we will not recover the saddlepoint solution.

4. A numerical solution based on Judd (1992)

Since shooting methods as well as finite difference approaches to solve (16) subject to the corresponding boundary conditions (5), (6′), and (7) either fail to determine the value function sufficiently accurately or are extremely slow, other means to approximate the value function deem necessary. We follow the suggestion of Judd (1992) to use Chebyshev polynomials and projection methods (sometimes referred to as Chebyshev collocation approach) to approximate value functions (Judd's paper on endogenous growth models sketches already briefly the application to a stochastic, but discrete time, model). The reason why the projection method succeeds is the fact that, in contrast to shooting methods, projection methods do not follow the flow by locally solving the equation. These approaches parameterize the entire problem and approximate the solution using non-linear equation solvers (like a simple Newton method in our implementation).

A brief description of the algorithm follows; for further details concerning the projection method see Judd (1992, 1998). We choose the first M Chebyshev polynomials

$$T_n(x) = \cos(n \arccos(x)) \tag{18}$$

and inspect the space $[T_i]_{i=0}^{i<M}$ spanned by these polynomials. Using the following orthogonal relation on $[-1, 1]$

$$\sum_{k=0}^{M-1} T_i(z_k^M) T_j(z_k^M) = \begin{cases} 0 & i \neq j \\ 1/2M & i = j \neq 0 \\ M & i = j = 0 \end{cases} \qquad i, j < M. \tag{19}$$

where z_k^M $(k = 0, \ldots, M-1)$ are the M roots of the Chebyshev polynomial T_M, we can define the projection $I_M^g(x)$ of a function g on $[T_i]_{i=0}^{i<M}$

$$I_M^g(x) = -1/2c_0 + \sum_{j=1}^{M-1} c_j T_j(x). \tag{20}$$

The coefficients c_j are determined by

$$c_j = \langle g | T_j \rangle = \frac{2}{M} \sum_{k=0}^{M-1} g(z_k^M) T_j(z_k^M). \tag{21}$$

The Chebyshev Interpolation Theorem (see Rivlin, 1990) states that $I_M^g(x)$ is the optimal approximation (with respect to $\|.\|_\infty$) of g by a linear combination of polynomials of a degree $< M$. Therefore the projection $I_M^g(x)$ is often called the Chebyshev interpolation of g. Two functions are identical with respect to their projection on $[T_i]_{i=0}^{i<M}$ if all their coefficients c_j are identical. The restriction to the interval $[-1, 1]$ is of no loss in generality, because each function defined on a bounded interval can be transformed to a function defined over $[-1, 1]$.

Referring to the differential equation (16), which we have to solve, we define the functional N

$$N(g)(\pi) = 1/2\sigma^2 g''(\pi) + ag'(\pi) + \frac{g'(\pi)^2}{2c} - rg(\pi) + \pi. \tag{22}$$

If g is a solution of (16), then:

$$N(g) \equiv 0. \tag{23}$$

The projection method simplifies the original problem (23) by replacing the required identity with the identity with respect to the projection on $[T_i]_{i=0}^{i<M}$. That is, the objective of the numerical method is to compute a function $\hat{f}_M$ (which is chosen from $[T_i]_{i=0}^{i<M}$ with coefficients c_i) such that the Chebyshev interpolation $N(\hat{f}_M)$ is identical to the Chebyshev interpolation of the function that is constant zero. The corresponding interpolation is, of course, the constant zero itself (see (21)) and so the coefficients of the Chebyshev polynomials vanish. This can be expressed as

$$p_i = \langle N(\hat{f}_M)|T_n\rangle = 0, \quad n = 0, \dots, M-1. \tag{24}$$

Let $C \in \mathfrak{R}^{\mathrm{M}}$ denote the vector of the projection coefficients c_i of $\hat{f}_M$ and $P \in \mathfrak{R}^{\mathrm{M}}$ the vector of the coefficients p_i of $N(\hat{f}_M)$. P is a non-linear operator, $P(C) : \mathfrak{R}^{\mathrm{M}} \to \mathfrak{R}^{\mathrm{M}}$, which reduces the problem of solving the functional equation (23) to solving M non-linear equations:

$$P(C) = 0. \tag{25}$$

If the coefficients of $\hat{f}_M$ are stable with respect to M for M exceeding a certain threshold $\underline{M}$ (implying that high order coefficients are of a negligible magnitude), then $\hat{f}_{M \geqslant \underline{M}}$ is a candidate for a suitable numerical approximation of g. This candidate has to pass further tests, see below.

The non-linear equations system (25) can be solved iteratively starting with a guess $C^0 = (c_j^0)$, e.g., using Newton's method: $C^{k+1} = C^k - (A_{C^k})^{-1} P(C^k)$, where $A_C = \mathrm{d}p_i/\mathrm{d}c_j = \partial p_i/\partial c_j$ is the Jacobian of the operator P evaluated at the respective point C^k. In order to speed up convergence, we add a relaxation parameter h:

$$C^{k+1} = C^k - h[(A_{C^k})^{-1} P(C^k)], \quad 0 < h \leqslant 1. \tag{26}$$

So far, the two boundary conditions (5) and (7) as well as the transversality condition (6′), which determine the critical level π^* and identify the economically stable solution, were neglected. A boundary condition like (5) or (7) can be included through a simple modification. We pick one of the c_i's (say c_k) and interpret it as function of the remaining set of coefficients[9]

$$c_k = c_k(c_1, c_2, \dots, c_{k-1}, c_{k+1}, \dots, c_{M-1}), \tag{27}$$

drop one of the projection conditions $p_i = 0$ (say p_l), and adapt the Jacobian

$$A_C = \left(\frac{\mathrm{d}p_i}{\mathrm{d}c_j}\right) = \left(\frac{\partial p_i}{\partial c_j} + \frac{\partial p_i}{\partial c_k}\frac{\mathrm{d}c_k}{\mathrm{d}c_j}\right), \quad i \neq l, \; j \neq k. \tag{28}$$

[9] The level π^* at which this condition has to be satisfied is a free boundary, which is not determined a priori.

A consequence of this modification is that only $M-1$ projections $\langle N(\hat{f}_M)|T_i\rangle$ are treated by the algorithm so that convergence of Newton's method does not guarantee that the projection of $N(\hat{f}_M)$ is identical 0.

In order to compute an approximation of the economically stable solution (i.e., $\hat{f}_M$ which satisfies (6′)) we start with an initial guess that implies a stable evolution of the system. This procedure – solving systems of non-linear equations – will converge to the stable solution, at least if we start sufficiently close. We know that unstable solutions diverge to $\pm\infty$. Hence, an initial guess that satisfies (6′), i.e., which has the corresponding linear asymptote, should be sufficiently close to the stable solution to guarantee convergence to the stable manifold. This was confirmed in all our computations: e.g., choosing the deterministic (thus stable) solution as our initial guess ensured convergence; however, convergence is possible for poor guesses too. Although convergence theorems that guarantee the quality and stability of the approximation do not exist, this does not invalidate the projection approach. Together with a proper test, this pragmatic 'compute and verify' approach is fast and robust (see Judd, 1998); after all, the numerically value function satisfies indeed the differential equation and the boundary condition (approximately up to the required degree).

The projection method relaxes the condition (23) by introducing a number of properties that have to be met. The accuracy of the approximation is measured by the maximum error, $R=\max|N(g)(\pi)|$, over the interval. The approximation is calculated over an interval far larger than the actual domain of interest in order to check that the approximation does not diverge but converges to the asymptote as required in (6′).

To determine the optimal stopping threshold π^* we proceed as follows. Since the boundary condition (7) is a first-order-condition for the optimality of π^* (see Dixit, 1993), we first drop it and impose only (5) at an arbitrarily chosen level $\hat{\pi}$. There exists one and only one economically stable solution of (16) for each $\hat{\pi}$ satisfying condition (5), i.e., $F(\hat{\pi})=0$, and the algorithm converges reliably.[10] Hence, we search for the optimal level π^* (e.g. applying a simple search algorithm) where the optimality condition (7) is also satisfied.

In order to demonstrate the suitability and accuracy of this algorithm, we draw on the model presented in Section 2, because we can compare the computed approximation with the explicit and analytical solution (11). This requires a re-definition of N (see (22)) in order to account for the lack of maintenance:

$$N(g)(\pi)=1/2\sigma^2 g''(\pi)+ag'(\pi)-rg(\pi)+\pi. \tag{22′}$$

Since (22′) is linear, convergence of the projection method is independent of the initial guess, thus, we start with the guess $g(\pi)=0$.

Table 2 lists the Chebyshev coefficients of the analytical solution and of the approximations for M=25 and 35. The critical level and the value of the free projection together with the residual R are listed in the header of Table 2. This table shows that the residuals are of diminishing order of magnitude and the critical level computed by the algorithm is identical to the analytical value up to at least fifteen digits.

[10] This solution corresponds to the decision to operate the machine until $\hat{\pi}$ is reached for the first time, irrespective of the optimality of this strategy.

Table 2
No maintenance

	Analytical solution	$M = 25$	$M = 35$
$\hat{\pi}$	0.170820393249937	−0.170820393249937	−0.170820393249937
$\lvert\hat{\pi} - \hat{\pi}^*\rvert$	0.0	$< 1 \times 10^{-15}$	$< 1 \times 10^{-15}$
p_{M-1}		$-8.078261 \times 10^{-17}$	$-4.070708 \times 10^{-20}$
R		$< 1 \times 10^{-16}$	$< 2 \times 10^{-18}$
c_0	79.022113607059	79.022113607059	79.022113607058
c_1	47.007651771983	47.007651771983	47.007651771982
c_2	5.619348408761	5.619348408761	5.619348408761
c_3	−3.207440799833	−3.207440799833	−3.207440799833
c_4	1.522617028486	1.522617028486	1.522617028485
c_5	−0.614406519018	−0.614406519018	−0.614406519018
c_6	0.214690333554	0.214690333554	0.214690333554
c_7	−0.065976364635	−0.065976364635	−0.065976364635
c_8	0.018062947260	0.018062947260	0.018062947260
c_9	−0.004453551505	−0.004453551504	−0.004453551505
c_{10}	0.000997936236	0.000997936236	0.000997936236
c_{11}	−0.000204809337	−0.000204809337	−0.000204809337
c_{12}	0.000038756442	0.000038756442	0.000038756442
c_{13}	−0.000006801339	−0.000006801339	−0.000006801339
c_{14}	0.000001112480	0.000001112480	0.000001112480
c_{15}	−0.000000170359	−0.000000170359	−0.000000170359
c_{16}	0.000000024520	0.000000024520	0.000000024520
c_{17}	−0.000000003329	−0.000000003329	−0.000000003329
c_{18}	0.000000000428	0.000000000428	0.000000000428
c_{19}	−0.000000000052	−0.000000000052	−0.000000000052
c_{20}	0.000000000006	0.000000000006	0.000000000006
c_{21}	−0.000000000001	−0.000000000001	−0.000000000001
c_{22}	0.000000000000	0.000000000000	0.000000000000
c_{23}	0.000000000000	0.000000000000	0.000000000000
c_{24}	0.000000000000	0.000000000000	0.000000000000
c_{25}	0.000000000000		0.000000000000
c_{26}	0.000000000000		0.000000000000
c_{27}	0.000000000000		0.000000000000
c_{28}	0.000000000000		0.000000000000
c_{29}	0.000000000000		0.000000000000
c_{30}	0.000000000000		0.000000000000
c_{31}	0.000000000000		0.000000000000
c_{32}	0.000000000000		0.000000000000
c_{33}	0.000000000000		0.000000000000
c_{34}	0.000000000000		0.000000000000

Chebyshev coefficients (over the interval $[-1, 10]$) of the analytical solution (11) compared with the approximation obtained by the projection method for $M = 25$ and $M = 35$. The computed critical levels $\hat{\pi}$, their absolute errors, the free projection coefficient and the maximum error R are listed in the column-headers.

Fig. 5 plots the absolute error versus computing time (again using GNU gcc on an AMD 1200 MHz PC running SuSE Linux 7.3). It takes only 1.5 s (and a basis of 25 Chebychev polynomials) to compute π^* up to 14 significant digits accurately. Compared

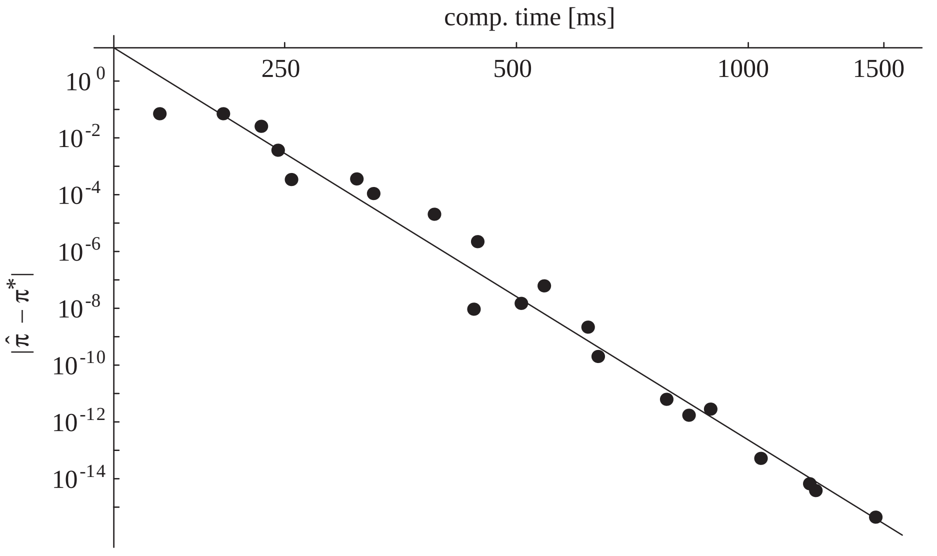

Fig. 5. Absolute error of the threshold $\hat{\pi}$ for the stopping problem in Section 2 computed by the Chebyshev projection method for $M = 5$ to $M = 25$ versus computing time.

with the finite difference method, it is 10 digits more accurate and the computational effort is of several orders of magnitude (around 10^3) less. The convergence of the approximated value function is of the same order of magnitude.

Finally we test the algorithm with the maintenance problem (12)–(13) for $c = 200$, see Table 3. The non-linearity of (16) implies that the projection algorithm has to solve a system of non-linear equations (25). As a consequence, the convergence is slower and the number of iterations required depends crucially on the initial guesses for the value function and π^*. The approximations of the respective value functions are stable and have sufficiently small residuals. For the maintenance problem the Newton algorithm inside the collocation method converges up to M=54; the Jacobi matrix is badly conditioned for higher order polynomials. Since there is no analytical solution for π^* the exit threshold $\hat{\pi}_{54}$ serves as a benchmark. Fig. 6 plots the absolute deviation of the exit threshold from $\hat{\pi}_{54}$ versus computing time for M=5 to M=53. This illustrates that the projection method converges reliably, yet due to the non-linearity the computing time is now one order of magnitude higher than for the pure stopping model: the approximation for M=53 takes less than 24 s.

5. Optimal maintenance and shutdown – results

The management problem introduced in Section 3 is now investigated – how the solution is affected by variations of parameters – in order to illustrate the workability of the above sketched algorithm. Fig. 7 shows the value function – setting the cost

Table 3
Maintenance

	$M = 25$	$M = 35$
$\hat{\pi}$	-0.1794460350381411	-0.1794460360744784
$\lvert\hat{\pi} - \hat{\pi}_{54}\rvert$	$< 1 \times 10^{-8}$	$< 1 \times 10^{-10}$
p_{M-1}	5.357081×10^{-7}	1.736056×10^{-10}
R	$< 7 \times 10^{-7}$	$< 2 \times 10^{-9}$
c_0	82.164379821594	82.173273023540
c_1	48.395313627255	48.403993418998
c_2	5.011191370217	5.019260557768
c_3	−3.232601206749	−3.225453482366
c_4	1.666221173918	1.672255588013
c_5	−0.712455059206	−0.707596902328
c_6	0.189410838442	0.193142798480
c_7	−0.021666476949	−0.018928884326
c_8	−0.032434148767	−0.030514755392
c_9	0.015786230239	0.017073848005
c_{10}	−0.009684306937	−0.008856771607
c_{11}	−0.000910953079	−0.000400665345
c_{12}	0.000349112815	0.000651563181
c_{13}	−0.001489533145	−0.001316837406
c_{14}	0.000123706407	0.000218963246
c_{15}	−0.000186048468	−0.000135100595
c_{16}	−0.000144022760	−0.000117489761
c_{17}	0.000023993479	0.000037541507
c_{18}	−0.000044955212	−0.000038133095
c_{19}	−0.000004846017	−0.000001435776
c_{20}	−0.000000048893	0.000001669102
c_{21}	−0.000006194736	−0.000005302678
c_{22}	0.000001279959	0.000001462959
c_{23}	−0.000000757151	−0.000000656083
c_{24}	−0.000001007460	−0.000000367655
c_{25}		0.000000230031
c_{26}		−0.000000189382
c_{27}		0.000000023688
c_{28}		0.000000004019
c_{29}		−0.000000025319
c_{30}		0.000000010235
c_{31}		−0.000000005085
c_{32}		−0.000000001117
c_{33}		0.000000001211
c_{34}		−0.000000001103

Chebyshev coefficients (over the interval $[-1, 10]$) computed with the projection method for $M = 25$ and 35. The computed critical levels $\hat{\pi}$, its absolute deviation from the exit threshold $\hat{\pi}_{54}$ (computed with $M = 54$ and thus our best guess of π^*), the free projection coefficient and the maximum error R are listed in the column-headers.

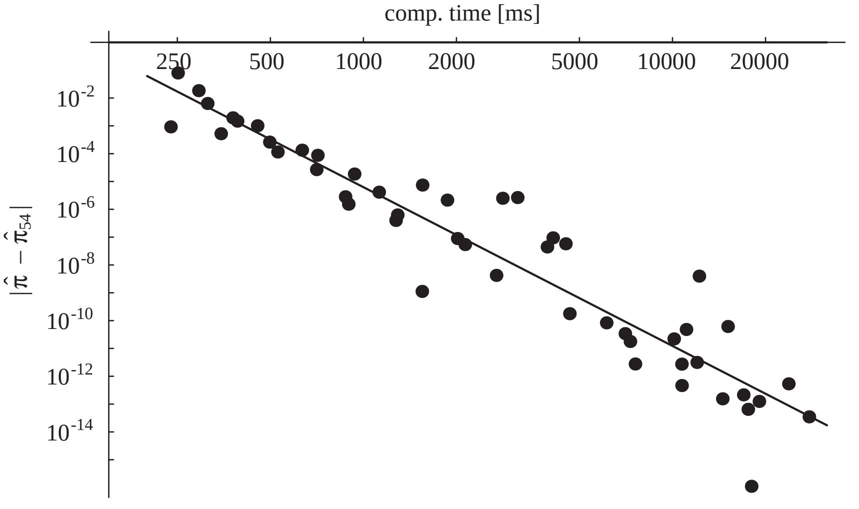

Fig. 6. Absolute error of the threshold $\hat{\pi}$ resulting from the application of Chebyshev projection method to the maintenance problem for $M=5$ to $M=53$; $\hat{\pi}_{54}$ (i.e. for $M=54$) serves as benchmark for the optimal π^*.

parameter[11] $c = 200$ (this applies to all subsequent computations) – and compares this result with the case of no maintenance displayed in Fig. 2 (thus the remaining parameters are as in Section 2). While value functions and optimal maintenance are sensitive to model parameters, the critical levels where to abandon the unit are all very similar and these values of π^* are tabulated below (Fig. 7). Furthermore, the possibility of maintenance has little impact on the shutdown decision but increases the value of an equipment significantly.

Fig. 7 illustrates that the value function shows the asymptotic behavior demanded by the transversality condition (6′), i.e., very high profitability implies that it is unlikely that this equipment will be eliminated in the near future and its value is therefore insensitive with respect to the 'option' of stopping. However, at lower profits, the value of a machine that can be eliminated exceeds considerably the value of a machine that must be kept forever. The difference between the two functions has the characteristic shape of the value of a put option (see Figs. 8 and 9) and can be interpreted as the value of the real option to eliminate the machine. In contrast to financial options theory this is not a simple stopping problem where one observes the underlying as an exogenous random process, instead, the stochastic process (13) is controlled and thus endogenous. At high profit flows, the value of the real option converges to zero, because nobody would eliminate the machine, i.e., exercise the put option before long. At low profits, the real option increases in value. The less the machine is maintained (due to a higher cost parameter c) the higher is the value of the option to liquidate. This is so because

[11] Note that for $c < 100$, maintenance is so cheap that it is 'economical' to reverse the natural depreciation, i.e., to improve the (expected) profitability of the machinery, which seems implausible, so that only $c \geqslant 100$ are sensible.

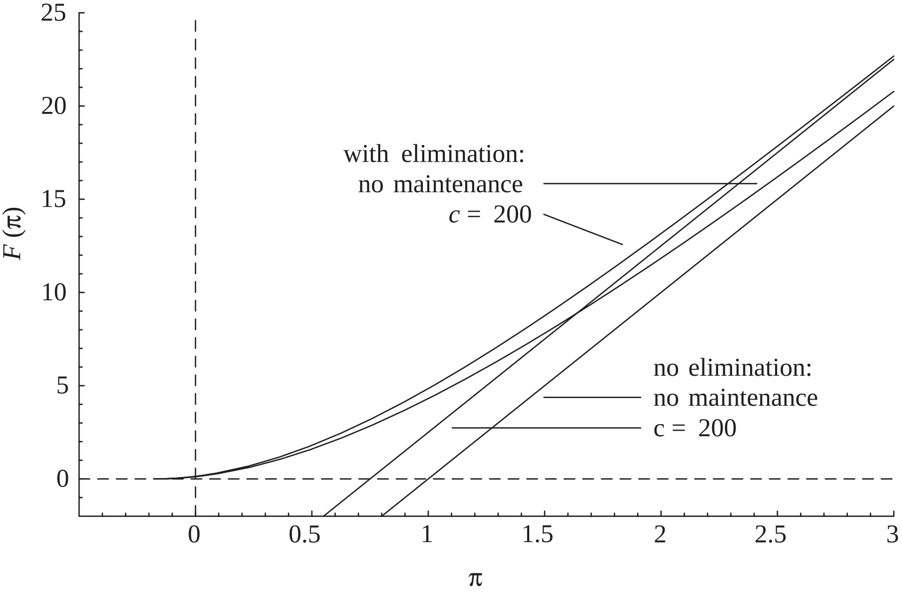

Fig. 7. Value functions of a maintained unit (c = 200) compared with one that cannot be maintained, for $a = -0.1$, $\sigma = 0.2$ and $r = 0.1$; the optimal level $\pi^* = -0.1708$ without and -0.1794 for maintenance. In addition, the value functions of equipments that cannot be eliminated (with and without maintenance) are also plotted.

a lower level of maintenance causes a more pronounced drift of the profit flow towards the liquidation threshold. For $\pi < \pi^*$, the put option will be exercised immediately. If both real options (of the maintained and of the not maintained machine) are in this region, their values differ by $1/(2cr^3)$.

Fig. 9 shows the value of the real option for different levels of uncertainty together with the limiting deterministic case, $\sigma = 0$. A higher level of uncertainty results, as expected, in a higher value of the corresponding real option and in a lower value of π^*.

Finally, optimal maintenance u^* (see Fig. 10) must be determined from (15). Obviously, the required differentiation of the value function $F(\pi)$ calculated in Table 3 and shown in Fig. 7 is derived directly from the Chebyshev interpolation. For π large, the optimal maintenance converges to the constant strategy $u^* = 1/(rc)$, which is optimal if the machine has to run forever (see (6′)). Therefore, optimal maintenance of a highly profitable equipment is insensitive with respect to the degree of uncertainty σ and the depreciation rate a; see Fig. 10. Lower profitability reduces maintenance, which stops at the critical level π^*, where the machine is finally scrapped. Higher maintenance costs result in less maintenance and in an earlier deviation from the constant, asymptotic strategy.

Higher uncertainty, i.e., a larger σ, leads to a lower exit threshold and, thus, to a higher level of maintenance for low profit flows. However, at higher profitability,

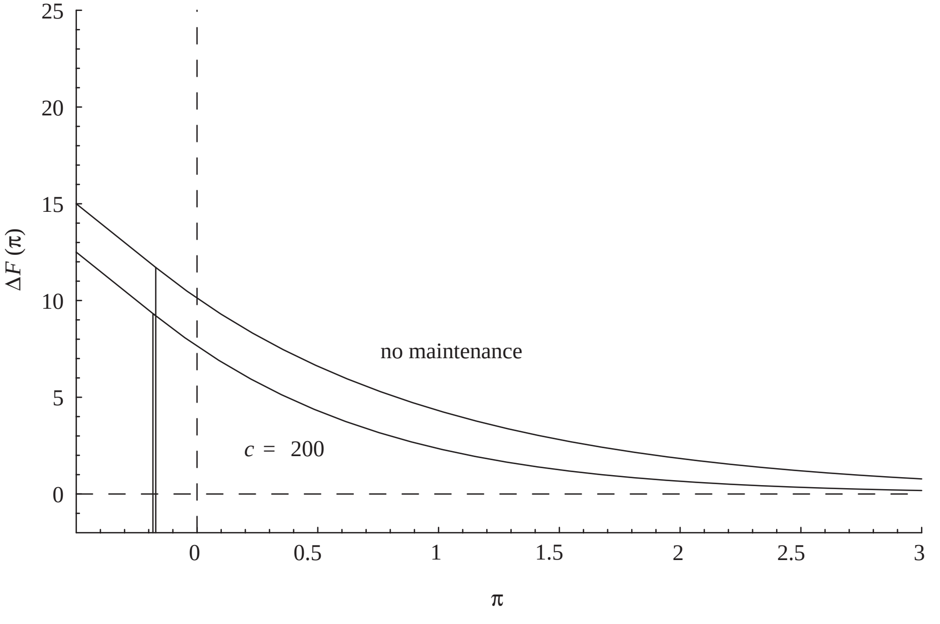

Fig. 8. Value of the real option to eliminate a machine with and without maintenance for $c = 200$ ($a = -0.1,\ \sigma = 0.2,\ r = 0.1$).

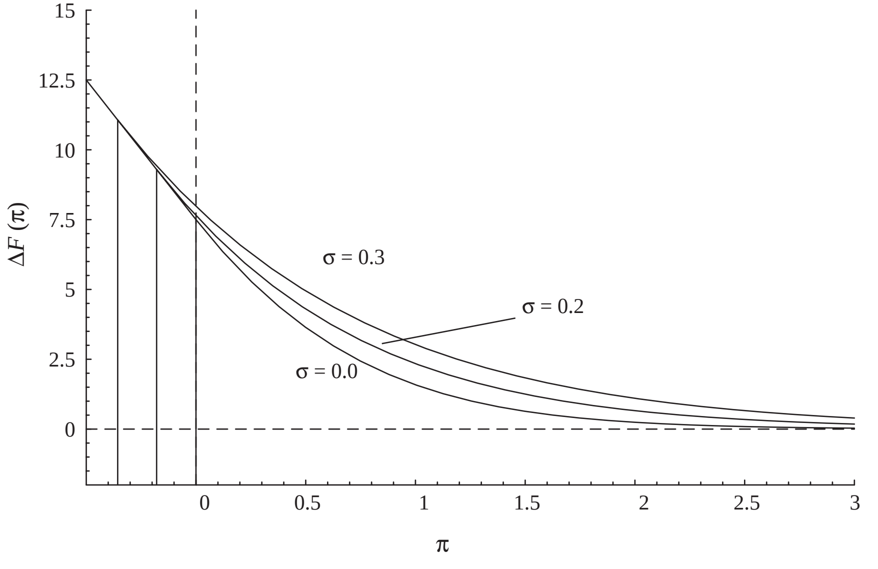

Fig. 9. Value of the real option to eliminate a machine for different values of σ, the critical levels π^* are 0.0, -0.1794 and -0.3567 for $\sigma = 0.0$, 0.2 and 0.3; and $c = 200$, $r = 0.1$, $a = -0.1$.

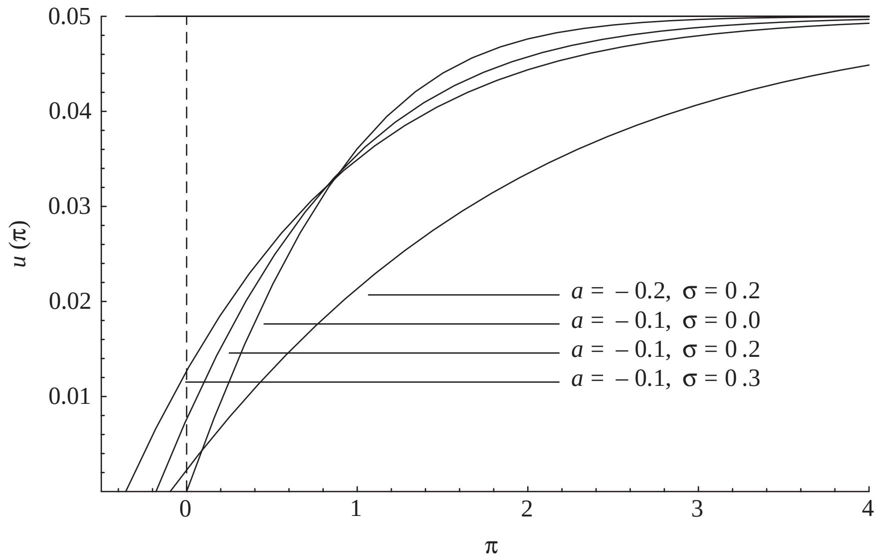

Fig. 10. Optimal maintenance for different levels of uncertainty and depreciation ($c = 200$).

uncertainty reduces maintenance. Therefore the impact of uncertainty on optimal maintenance can be positive or negative. The reason is that at low profitability a higher level of uncertainty increases the upside chance (the likelihood that the profit flow recovers) while the downside risk is limited (the worst situation is that π falls to π^* and the machine is eliminated). Therefore, uncertainty justifies higher maintenance at low profitability. At high profitability we face the opposite situation. Due to the common asymptote, uncertainty increases the downside risk that the profit flow decreases faster than expected whereas it cannot create an upside chance. The response is an earlier reduction of maintenance and thus larger depreciation at high profitability levels. An increase in average depreciation per period (i.e., an increase in a) leads to a global reduction in maintenance (yet approaching the same asymptote).

6. Summary

The purpose of this paper is to draw attention to the fact that standard numerical procedures for solving differential equations are not suitable for treating functional equations obtained from the Bellman equation of continuous time dynamic programming. The fact that solving these equations requires locating the ‘saddlepoint’ from the entire family of solution curves causes shooting methods to fail. Finite difference methods tend to converge but very slowly especially if the considered time horizon is large. Collocation methods, which are introduced to economics in Judd (1992) in the context of endogenous growth models, offer a fast and robust alternative. Both, the restrictions of conventional methods and the suitability of collocation methods (we use a projection method based on Chebyshev polynomials), were demonstrated for the management problem of maintaining and eventually eliminating a profit generating equipment. The presented sensitivity analysis and investigation of parameter variations, which are typical for many economic applications, underline the requirement for such fast and reliable algorithms (just recall that solving the simple stopping problem by the method of finite difference took above 6h computing time yielding an accuracy of just four digits).

Acknowledgements

Both authors acknowledge valuable suggestions from Winfried Auzinger, two referees and the editor, Professor Kenneth Judd.

Appendix A.

A.1. Deterministic

Setting $\sigma = 0$ the corresponding deterministic (and linear-quadratic) model results:

$$\max_{\{u(t),T\}} \int_0^T \mathrm{e}^{-rt} \left(\pi(t) - \frac{1}{2} c u^2(t) \right) \mathrm{d}t, \tag{12'}$$

$$\frac{\mathrm{d}\pi(t)}{\mathrm{d}t} = a + u(t), \qquad \pi(0) = \pi_0. \tag{13'}$$

The definition of the Hamiltonian (present value notation, λ denotes the costate),

$$H = \pi - 1/2cu^2 + \lambda(a + u), \tag{A.1}$$

gives the necessary and sufficient conditions for optimality (see e.g. Feichtinger and Hartl, 1986; or Léonard and Long, 1992):

$$H_u = -cu + \lambda = 0, \tag{A.2}$$

$$\dot{\lambda} = r\lambda - 1, \qquad \lambda(T) = 0, \tag{A.3}$$

$$H(T) = \pi(T) - 1/2cu^2(T) + \lambda(T)(a + u(T)) \Rightarrow \pi(T) = 0. \tag{A.4}$$

The implication in (A.4), which leads to the determination of the optimal stopping time T, follows from the transversality condition in (A.3) and from the maximum principle (A.1), because the optimal maintenance is given by $u = \lambda/c$. Substitution of this control into (13′) yields the canonical system of equations

$$\dot{\pi} = a + \lambda/c, \qquad \pi(0) = \pi_0,$$

$$\dot{\lambda} = r\lambda - 1, \qquad \lambda(T) = 0. \tag{A.5}$$

The system (A.5) can be solved explicitly:

$$\pi(t) = \pi_0 + at + \frac{1}{cr}\left[\frac{e^{-rT}}{r}(1 - e^{rt}) + t\right], \tag{A.6}$$

$$\lambda(t) = [1 - e^{-r(T-t)}]/r. \tag{A.7}$$

However, the stopping time T has to be calculated numerically by solving $\pi(T) = 0$. Eliminating the time t determines optimal maintenance as a function of the state, $u = u(\pi)$; which is displayed in Fig. 10.

If we assume an infinite horizon, i.e., the machine has to operate even at (growing) losses, the corresponding Bellman equation is

$$rF = \max_u \left(\pi - \frac{1}{2}cu^2 + F_\pi(a+u)\right) = \pi + aF_\pi + 1/2F_\pi^2/c, \tag{A.8}$$

and can be solved analytically. This linear quadratic control model with infinite planning horizon has a quadratic value function:

$$F(\pi) = k_0 + k_1\pi + 1/2k_2\pi^2. \tag{A.9}$$

Comparing coefficients (after substitution of the guess (A.9) into (A.8)) determines [12]

$$k_0 = a/r^2 + 1/(2cr^3),$$

$$k_1 = 1/r,$$

$$k_2 = 0, \tag{A.10}$$

so that $F(\pi) = (a/r^2) + 1/(2cr^3) + (\pi/r)$ results (which is equal to the right hand side of (6′)). This proves that the constant maintenance $u = 1/(rc)$ is optimal in the deterministic framework if it is impossible to eliminate the machine. For $c \to \infty$, $u^* = 0$, i.e., the maintenance model approaches the case without maintenance where $F(\pi) = (a/r^2) + (\pi/r)$.

A.2. Stochastic

Note that this deterministic result is also a particular solution of (16) even for $\sigma \neq 0$. And in fact, constant maintenance is the optimal strategy if it is impossible to get rid of the machine: Let π be a stochastic process defined by (13) and $\mu = \pi + h$, with constant h. Since the planning horizon is infinite, we can write (see (12)):

$$\begin{aligned} F(\mu) = F(\pi + h) &= \max_{\{u(t)\}} E\int_0^\infty e^{-rt}\left(\pi(t) + h - \frac{1}{2}cu^2(t)\right) dt \\ &= \int_0^\infty he^{-rt}\,dt + \max_{\{u(t)\}} E\int_0^\infty e^{-rt}\left(\pi(t) - \frac{1}{2}cu^2(t)\right) dt \\ &= F(\pi) + \frac{1}{r}h. \end{aligned} \tag{A.11}$$

[12] The positive root k_2 can be ruled out because it would cause, according to (15), unbounded maintenance for large π and thus yield an unstable solution.

Therefore $F(\pi)$ has the constant slope $1/r$, i.e. $F(\pi) = k + \pi/r$, and with (15) we get constant maintenance $u^* = 1/(rc)$. Substituting into (16) gives $k = a/r^2 + 1/(2cr^3)$ and $F(\pi) = (a/r^2) + 1/(2cr^3) + (\pi/r)$, which is the right hand side of (6′). Hence, without the possibility to eliminate the machine, the value function is independent of σ and the optimal maintenance is independent of both, σ and a.

A.3. Finite differences

To apply a finite difference method we have to reformulate the maintenance and stopping problem into a finite horizon problem. If $\bar{T}$ is the maturity of the shutdown option, then the corresponding Bellman equation is

$$rF = \pi + F_t + aF_\pi + 1/2F_\pi^2/c + 1/2\sigma^2 F_{\pi\pi}, \tag{16'}$$

which is now a partial differential equation. Consider a discrete grid in the time/state space with increments Δt and $\Delta\pi = \sigma\sqrt{\Delta t}$, and let $f_{i,j}$ denote the valuation at the grid points, where the indices i and j characterize the location of the point with respect to state and time, respectively. Approximation of the partial derivatives in (16′) by finite differences allows working backwards in time using the following iteration scheme

$$\begin{aligned} f_{i,j-1} = {} & (1 - r\Delta t) f_{i,j} + \pi\Delta t + \frac{f_{i+1,j} + f_{i-1,j} - 2f_{i,j}}{2} + \frac{a\sqrt{\Delta t}}{2\sigma}(f_{i+1,j} - f_{i-1,j}) \\ & + \frac{1}{8c\sigma^2}(f_{i+1,j} - f_{i-1,j})^2. \end{aligned} \tag{A.12}$$

The approximated value function is then given by $F(i\Delta\pi, j\Delta t) = f_{i,j}$. If maintenance is not allowed, the last term on the left hand side of (A.12) has to be dropped.

References

Brennan, M.J., Schwartz, E.S., 1977. The valuation of American put options. Journal of Finance 32 (2), 449–462.

Brennan, M.J., Schwartz, E.S., 1978. Finite difference methods and jump processes arising in the pricing of contingent claims: a synthesis. Journal of Financial and Quantitative Analysis 13, 461–474.

Cox, J., Rubinstein, M., Ross, S., 1979. Option pricing: a simplified approach. Journal of Financial Economics 7 (3), 229–263.

Deuflhard, P., Bornemann, F., 2002. Scientific Computing with Ordinary Differential Equations. Springer, Heidelberg.

Dixit, A.K., 1993. The Art of Smooth Pasting. Vol. 55 of series: Fundamentals of Pure and Applied Economics. J. Lesourne, H. Sonnenschein (Eds.), Harwood Academic, Reading, UK.

Dixit, A.K., Pindyck, R.S., 1994. Investment under Uncertainty. Princeton University, Princeton.

Feichtinger, G., Hartl, R.F., 1986. Optimale Kontrolle Ökonomischer Prozesse. de Gruyter, Berlin.

Hull, J., White, A., 1990. Valuing derivative securities using the explicit finite difference method. Journal of Financial and Quantitative Analysis 25 (1), 87–100.

Judd, K.L., 1992. Projection methods for solving aggregate growth models. Journal of Economic Theory 58, 410–452.

Judd, K.L., 1998. Numerical Methods in Economics. MIT, Cambridge, MA.

Kamrad, B., Ritchken, P., 1991. Multinomial approximating models for options with k state variables. Management Science 37 (12), 1640–1652.

Léonard, D., Long, N.V., 1992. Optimal Control Theory and Static Optimization in Economics. Cambridge University, Cambridge.

Raman, K., Chatterjee, R., 1995. Optimal monopolist pricing under demand uncertainty in dynamic markets. Management Science 41 (1), 144–162.

Rivlin, T.J., 1990. Chebyshev Polynomials: From Approximation Theory to Algebra and Number Theory. Wiley-Interscience, New York.