Optimal exercise of executive stock options [1]

L.C.G. Rogers & José Scheinkman

University of Cambridge & Princeton University

1 Introduction.

In the absence of frictions if a portfolio strategy replicates the payoff of one unit of a claim, an appropriately scaled strategy replicates any amount of the claim. If assets are priced by arbitrage, the value per-unit is invariant to the amount of the asset considered. In particular, in the case of American claims, the optimal exercise time is independent of the amount of the claim that is considered. In this paper we show that this result does not necessarily hold in the presence of portfolio constraints or other frictions. We produce an example in which the absence of short sales leads the holder of an American option on a (possibly non-dividend paying) stock to exercise parts of his option over time.

There has been a lot of interest in the valuation of American type securities with portfolio constraints (*e.g.* Cvitanic and Karatzas [1993], Detemple and Sundaresan [1999]). However, in this literature, it is assumed that there is a single unit of the derivative securities, and one studies the optimal exercise time for that unit. It is implicitly assumed that the optimal strategy is independent of the amount held.

To fix ideas, we will consider an executive who holds an American-style call option on the stock of his firm; we make the (realistic) assumption that he is forbidden to short sell the underlying stock. For simplicity we will actually assume that the executive also cannot hold the stock, although it is obvious that this constraint is not binding. The holder of the option can however exercise *parts* of the option at any time, and deposit the proceeds from the exercise in a risk-free account. Risk-aversion leads naturally to early exercise, and because the executive's wealth fluctuates with changes in the price of the underlying, the optimal amount of options in the executive's portfolio changes over time. If this optimal amount increases there is nothing the executive can do. However if it decreases the executive will exercise some of his remaining options. This leads to an exercise boundary relating the time remaining to expiration, the price of the stock, and the number of unexercised options. The optimal policy consists of exercising

[1]We thank Karim Belkheir for assistance with the computations of this paper. José Scheinkman thanks the National Science Foundation for support through grant 0001647. The bulk of the research reported here was done in 2002 while Scheinkman held a Chaire Blaise Pascal de l'État et de la Région Ile de France.

enough options to stay below the boundary.

Others have studied the problem faced by executives holding stock options that they are not allowed to hedge. Lambert, Larcker and Verechia (1991), Carpenter (1998) and Hall and Murphy (2002) consider an executive that must fully exercise a block of options. In contrast, Ingersoll (2006) derives a model for the marginal value of an option. None of these authors consider the possibility of partial exercise. An exception is Jain and Subramanian (2004) who allow for partial exercise in a discrete time framework.

In the next section we shall formulate the problem in a general setting, and present our approach, which is to discretise the problem immediately. Since we expect that most (if not all) interesting examples will admit no closed-form solution, we discretise the problem, and prove that the optimal solutions to the discretisations do indeed converge to the optimal solution of the original problem. Section 3 presents a few examples which can be solved numerically. Special scaling properties that result from assuming CARA or CRRA utility functions are exploited to reduce the dimension of the problems, since in general the value depends on four variables; the underlying state, the current value of realized options, the number of options still to be exercised, and the time until expiry. Computation is possible for the reduced three-dimensional problem, though care is needed. The examples are used to illustrate some comparative statics results. As the number of remaining options get larger, exercise will occur at lower prices. Increases in risk-aversion have the same effect. The effect of the time-to-expiry is more subtle. When time-to-expiry is very small, the agent is keen to seize whatever value he can, so he will exercise for a small premium; when time-to-expiry is large, he will be willing to exercise for less than he would require when time-to-expiry is moderate, because of the interest that will accrue on the exercised option value. In the examples these two forces combined yield a non-monotonic pattern - agents are most conservative at exercising for intermediate values of the time-to-go. Our emphasis on numerical methods is because our aim is not to show that some behaviour is impossible - for that, only analytics will do - but rather to show that certain types of behaviour *can* happen, and for that it is sufficient to show a numerical example where the behaviour *does* happen.

Section 4 studies the limiting form of the optimal rule as the expiry gets ever bigger. In this situation, the solution depends on just three variables, and is thus quite a lot simpler. An example shows however that, for reasonable parameter values, the finite-horizon solution differs substantially from the infinite-horizon solution even when there are decades to go before expiration.

2 The problem.

The holder of the executive stock options faces an optimal control problem of the following form. There is some given adapted[2] non-negative[3] process φ^0, and the executive has to choose a right-continuous increasing adapted process m, started at zero, so as to maximise

$$EU(x_T), \tag{1}$$

where U is a concave increasing function and

$$x_T = x_0 + \int_0^T \beta_s \, \varphi_s^0 \, dm_s \equiv x_0 + \int_0^T \varphi_s \, dm_s \tag{2}$$

is the time-0 value of the terminal wealth from exercise of the options, and x_0 is the initial wealth. The discount factor $\beta_s \equiv \exp(-\int_0^s r_u \, du)$ discounts all gains from exercise back to time-0 values; for brevity, we have absorbed this into the payoff by writing $\beta\varphi^0 \equiv \varphi$. The process m records the cumulative total of options exercised as time evolves, and so must satisfy the constraint

$$m_T \leq A, \tag{3}$$

where A is the total number of options initially held. Since $\varphi \geq 0$, we may (and shall) without loss of generality suppose that the bound (3) holds with equality, with any remaining options at time T being optimally instantaneously exercised at that time. We use the notation

$$V^* = V^*(T, x_0, A) = \sup EU(x_0 + \int_0^T \varphi_s \, dm_s) \tag{4}$$

subject to the constraint

$$m_T = A. \tag{5}$$

In any particular application, V^* will depend on variables other than T, x_0 and A, of course. To incorporate this, it is conventional at this stage to introduce some underlying diffusion or Markov process, and show that V^* satisfies some more or less complicated Hamilton-Jacobi-Bellman (HJB) equation, which in the case of a diffusion will be a second-order non-linear PDE, with suitable boundary conditions. This leads to a number of questions of a technical nature:

(i) Does this PDE have a unique solution?

[2].. to the filtration $(\mathcal{F}_t)_{t\geq 0}$ of some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, P)$ satisfying the usual conditions; see, for example, Rogers & Williams (2000)

[3]If φ^0 ever took negative values, the holder would certainly not exercise at any such time, so it is clear that we could replace φ^0 by its positive part without altering the value to the holder.

(ii) If so, does the unique solution provide the value function of the problem? This is usually a straightforward if tedious verification result.

(iii) How can we solve the equation? If we can find a closed-form solution, then (i) and (ii) were to a large extent unnecessary; if not, the only recourse is to numerical methods.

(iv) If we have to use numerical methods to solve the problem, can we show that the solutions to the approximating problems converge to the solution of the original problem?

Any or all of these questions may be challenging. For our problem, however, only in the most exceptional situations will there be any closed-form solution. Answering questions (i) and (ii) will therefore lead no further than answers to questions (i) and (ii), and if we are to attempt to solve the problems by numerical means then we will be left with (iv).

We will pass directly to (iv) by introducing discrete approximation immediately. If we let $\mathcal{A}$ denote the class of all right-continuous increasing adapted processes m such that $m_0 = 0$ and $m_T = A$, then $\mathcal{A}$ is the set of admissible controls. If $m \in \mathcal{A}$ is a generic control, we define the stopping times τ_a for $0 \leq a \leq A$ by

$$\begin{aligned} \tau_a &\equiv \inf\{t : m_t > a\} \qquad (0 \leq a < A) \\ &\equiv T \qquad\qquad\qquad\quad (a = A). \end{aligned}$$

Now we introduce the subsets of $\mathcal{A}$ defined by

$$\mathcal{A}_n \equiv \{m \in \mathcal{A} : m_t \in 2^{-n}A\mathbb{Z}^+ \text{ for all } t \in [0,T], \tau_a \in 2^{-n}T\mathbb{Z}^+ \text{ for all } a \in [0,A]\}. \tag{6}$$

This apparently clumsy definition masks a simple reality; the processes in $\mathcal{A}_n$ are just *staircase processes* of the kind illustrated in Figure 1. They increase only by jumps, which must be multiples of $2^{-n}A$, and which must occur at multiples of the time $2^{-n}T$. With this notation, we have the following simple but central result.

Proposition 1 *Assume that the process φ has right-continuous paths and satisfies the condition*

$$EU(x_0 + A\bar{\varphi}_T) < \infty, \tag{7}$$

where $\bar{\varphi}_t \equiv \sup\{\varphi_s : 0 \leq s \leq t\}$. Then

$$\sup_{m \in \mathcal{A}_n} EU(x_T) \uparrow \sup_{m \in \mathcal{A}} EU(x_T) \qquad (n \uparrow \infty). \tag{8}$$

PROOF. Since evidently $\mathcal{A}_n \subseteq \mathcal{A}_{n+1} \subseteq \mathcal{A}$, it is immediate that the approximate values $\sup_{m\in\mathcal{A}_n} EU(x_T)$ increase with n to a limit which is no bigger than the true value $\sup_{m\in\mathcal{A}} EU(x_T)$, and the only issue is to prove that the limit is equal to the true value.

To do this, we show that given $m \in \mathcal{A}$ we can find approximating $m^{(n)} \in \mathcal{A}_n$ converging to m in such a way that the values of the strategies $m^{(n)}$ converge to the value of m. Define the two staircase functions[4]

$$\theta_n(t) \equiv T2^{-n}[T^{-1}2^n t], \qquad f_n(x) \equiv A2^{-n}[A^{-1}2^n x],$$

and then the approximations to m

$$m_t^{(n)} \equiv f_n(m(\theta_n(t))),$$

with corresponding right-continuous inverses

$$\begin{aligned} \tau_a^{(n)} &\equiv \inf\{t : m_t^{(n)} > a\} \qquad (0 \le a < A). \\ &\equiv T \qquad\qquad\qquad\qquad (a = A). \end{aligned}$$

Because of the monotonicity of m and the fact that $f_n \le f_{n+1}$, $\theta_n \le \theta_{n+1}$, it is clear that $m^{(n)}$ increase, and that $\tau^{(n)}$ decrease. Although there may be isolated points t where the limit of $m_t^{(n)}$ is not m_t, it is true that for *every* a, $\tau_a^{(n)} \downarrow \tau_a$, because for any $a < A$

$$\begin{aligned} \tau_a < t &\Rightarrow \text{for some } \varepsilon > 0,\ m_s > a \text{ for all } s \in (t-\varepsilon, T) \\ &\Rightarrow \text{for all large enough } n,\ m_s^{(n)} > a \text{ for all } s \in (t-\varepsilon/2, T) \\ &\Rightarrow \text{for all large enough } n,\ \tau_a^{(n)} \le t - \varepsilon/2 \\ &\Rightarrow \downarrow \lim_n \tau_a^{(n)} < t. \end{aligned}$$

Now by simple change-of-variables,

$$\begin{aligned} x_0 + \int_0^T \varphi_s\, dm_s &= x_0 + \int_0^A \varphi(\tau_a)\, da \\ &= x_0 + \lim_n \int_0^A \varphi(\tau_a^{(n)})\, da \\ &= x_0 + \lim_n \int_0^T \varphi_s\, dm_s^{(n)} \end{aligned}$$

using the right-continuity of φ and dominated convergence (since $\varphi(\tau_a^{(n)}) \le \bar{\varphi}$.) Finally,

$$EU(x_0 + \int_0^T \varphi_s\, dm_s) = \lim_n EU(x_0 + \int_0^T \varphi_s\, dm_s^{(n)})$$

[4]For any real x, we use $[x]$ to denote the largest integer not greater than x.

once again using dominated convergence and the integrability condition (7).

□

REMARKS. The importance of Proposition 1 is that it reduces the original problem to discrete time and discrete quantities of the option. This way we can solve a 'stack' of optimal stopping problems, the stopping value for each being defined by the value on the level below. This is illustrated in Section 3. It is worth noting that the equal spacing of the time and m grids is not really necessary; all that is required is a sequence of refinements with mesh tending to zero. We shall make use of this extension later.

3 Some examples.

We consider here the motivating example of an executive who holds A American call options on the stock of his firm. The options all have strike K and expiry T. He is forbidden to trade the stock, but is free to exercise the options at will through the interval $[0, T]$. We shall take the model for the log stock price Y^0 to be the standard log-Brownian model

$$dY_t^0 = \sigma dW_t + \mu_0 dt$$

where the volatility σ and the drift μ_0 are assumed constant. It turns out to be better[5] for the numerical work to use instead the log *discounted* price process

$$dY_t = dY_t^0 - r dt = \sigma dW_t + \mu dt,$$

with $\mu = \mu_0 - r$. Write a for the amount of options still available for exercise. We expect that the value function $V_T(t, y, x, a)$, for this problem[6], defined by

$$V(t,y,x,a) \equiv \sup E\Big[\, U(x + \int_t^T (e^{Y_u} - e^{-ru}K)^+ \, dm_u) \;\Big|\; m_t = A - a, Y_t = y \,\Big], \quad (9)$$

should satisfy the differential inequality

$$\max\Big\{ \frac{\partial V}{\partial t} + \mathcal{G}V, -\frac{\partial V}{\partial a} + (e^y - e^{-rt}K)^+ \frac{\partial V}{\partial x} \Big\} = 0, \quad (10)$$

with boundary conditions

$$V(T,y,x,a) = U(x + a(e^y - e^{-rT}K)^+), \quad (11)$$

$$V(t,y,x,0) = U(x), \quad (12)$$

[5]The integrand $(e^{Y_u} - e^{-ru}K)^+ = e^{-ru}(e^{Y_u^0} - K)^+$ in (9) remains $O(1)$ if Y stays in some fixed interval, but if Y^0 stays in some fixed interval the integrand gets extremely small for large u.

[6]We use the abbreviation $V \equiv V_T$ whenever the time horizon T does not need to be emphasised in the notation.

where

$$\mathcal{G} \equiv \frac{1}{2}\sigma^2 \frac{\partial^2}{\partial y^2} + \mu \frac{\partial}{\partial y}$$

is the generator of Y. We make no attempt to prove that the value function (9) solves (10), (11), (12); we know of no interesting examples where the value function can be found in closed form, so our philosophy is to investigate the problem through its discrete approximations which can at least be evaluated numerically. Proposition 1 establishes that the numerical values found will be 'close' to the true values. However, even for this simple example, it can be hard to grasp the features of the numerical solution, being as it is a function of four variables; so we shall content ourselves with some interesting but special cases where the problem simplifies. Nevertheless, one feature is obvious; for fixed (t, x, a), there will be a critical $\eta(t, x, a)$ such that we will choose to exercise if $y > \eta(t, x, a)$ and not to exercise if $y < \eta(t, x, a)$, since φ is increasing in y. We translate this exercise boundary into log price (rather than log discounted price) terms, by setting

$$\eta^0(t, x, a) = rt + \eta(t, x, a).$$

This is easier to interpret, in that $\eta^0(t, x, a) > \log K$ always. With no real loss of generality, we shall in the numerical examples always take

$$K = 1;$$

this means that values of η^0 are greater than 0, and often (for large a) very close to 0. In the plots, we shall display $\log \eta^0$ to make the shape of the surface more evident when the number of remaining options is large.

Example 1: exponential utility. If we assume that the utility is

$$U(x) \equiv -\exp(-\gamma x)$$

for some constant γ, then a familiar argument shows that the value function simplifies to

$$\begin{aligned} V(t, y, x, a) &= \exp(-\gamma x) V(t, y, 0, a) \\ &\equiv \exp(-\gamma x) v(t, y, a) \end{aligned}$$

so computing the value function requires us to find a function of only three variables. Since the current wealth x factors out of the problem, we see that the critical value $\eta(t, x, a)$ of y at which exercise happens will depend only on (t, a); but how does this critical boundary η look?

In order to study this, we make a discrete approximation. In more detail, we divide the interval $[0, T]$ into N_t equal intervals of length $\Delta t = T/N_t$ and take a spatial discretisation of Y onto a grid

$$\{l\Delta y : L_- \leq l \leq L_+\}$$

with spacing Δy. Finally, we take a finite sequence $a_0 = 0 < a_1 < \ldots < a_J = A$ of discretisation points in the variable a. There is no need for these to be equally spaced (and indeed it is not necessary to set the grid points in time at equal spacings), but because of our numerical approach we *do* need the spacings of Y to be equal. We then compute the array

$$\{v(n,l,j) : 0 \leq n \leq N_t,\ L_- \leq l \leq L_+,\ 0 \leq j \leq J\}$$

where $v(n,l,j)$ approximates $V(n\Delta t, l\Delta y, 0, a_j)$. The boundary condition (12) gives us that

$$v(n,l,0) = -1$$

for all n and l, and we have for each $j \geq 1$ the American-style optimal stopping problem, where stopping at grid point (n,l) produces reward

$$\exp(\ -\gamma(a_j - a_{j-1})(e^{l\Delta y} - e^{-rn\Delta t}K)^+\ )\, v(n,l,j-1).$$

The dynamics of the approximating process can be modelled in various different ways, but our method was to approximate the steps of Y by their *exact* $N(\mu\Delta t, \sigma^2\Delta t)$ distribution; this was computed by convoluting the discretised values of v with the discretised values of the transition density. The computations (carried out in Scilab) use the (fast and accurate) FFT algorithm. Not only can the numerical values be expected to be quite accurate, but it would also be possible with minimal changes to the code to replace the log-Brownian dynamics with *any* log-Lévy dynamics.

We present some numerical results here. The plots show the level of $\log \eta^0$ at which the agent should exercise, as a function of time, and the number of remaining options (on log scale). For each set of parameter values we present five sections through the surface, each representing a distinct number of remaining options. Notice that as the number of remaining options gets larger, the critical threshold gets lower, as we would expect. The first set of figures, that we will use as a baseline, Figures 2 and 3, also show that the critical threshold is not always monotone in time-to-go. As time-to-go gets small, we find the threshold gets small quite fast, because the agent is keen to get some value, however small, from his options before they expire. On the other hand, when time-to-go is quite large, the agent is willing to cash the options in for less than he would get by waiting, in order to gain from the interest on the earlier-realised cash.

Figure 4 show the effect of an increase in risk-aversion. As should be expected, the critical values go down for every value of a, the number of remaining options, and time to expiry. Figure 5 exhibits the consequences of an increase in volatility. Notice that, in this example, critical prices go down, indicating that the effect of the increase in risk exceeds that of the increase in optionality. In Figure **??** we choose a negative value of μ, but hold all other parameters as in Figure **??**. As

expected critical prices go down. The next Figure **??**, differing from Figure **??** only by increasing σ from 0.3 to 2.0, is interesting in that it shows that in some cases an increase in volatility can lead to higher critical prices, indicating that the increase in optionality dominates the increase in risk. However, the effect is not uniform across different values of a; comparing the initial values of the η^0 curves in Figures 6 and 7, we see that the top five curves start higher in Figure 7, whereas the bottom one starts lower, by 0.0194.

Example 2: CRRA utility. Taking the utility to be

$$U(x) = x^{1-R}/(1-R),$$

where $R \neq 1$ is a positive constant, we obtain a situation rather like that of the previous example. Indeed, a moment's thought shows that the value function is homogeneous of degree $1-R$:

$$V(t, y, \lambda x, \lambda a) = \lambda^{1-R} V(t, y, x, a) \tag{13}$$

for any $\lambda > 0$. Negative values of x are not ruled out, but it is clear that if at any time we get $x + a\varphi < 0$ then the value to the agent will be $-\infty$, since there is a positive probability that at all times between now and T the total $x_t + a_t\varphi_t$ will be strictly negative. To simplify the analysis, we shall assume henceforth that we are in the interesting case, where $x > 0$, and introduce the reduced value function v defined via

$$V(t, y, x, a) = x^{1-R} V(t, y, 1, a/x) \equiv x^{1-R} v(t, y, a/x) \equiv x^{1-R} v(t, y, s), \tag{14}$$

with the notation $s \equiv a/x$. To compute a numerical approximation to v, we shall assume that s is only allowed to take finitely many values $0 = s_0 < s_1 < \ldots < s_N$, and when an exercise takes place, the current value s_n of a/x drops to s_{n-1}. The number Δa of options that are exercised to achieve this is easily shown to be

$$\Delta a = x \frac{s_n - s_{n-1}}{1 + \varphi s_{n-1}}.$$

Hence the value of x gets scaled up by a factor

$$\frac{x + \varphi \Delta a}{x} = \frac{1 + \varphi s_n}{1 + \varphi s_{n-1}}$$

The computed values of η display similar qualitative features to the CARA example, so we omit these.

4 Limiting behaviour as $T \to \infty$.

We expect that as the horizon T recedes into the indefinite future, the form of the optimal strategy and the value should both settle down to some limit. Indeed, it is clear that as T increases, the value

$$V_T(t,y,x,a) \equiv \sup E\Big[\, U(x + \int_t^T e^{-ru}\varphi_u^0\, dm_u) \;\Big|\; m_t = A-a, Y_t = y \,\Big], \qquad (15)$$

will increase[7]. Let us remark immediately that the limiting form of the problem (15) will be ill posed if $\mu_0 \geq r$; if this condition applies, then the discounted asset process Y will reach arbitrarily high values with probability 1, so a good (though not optimal) policy would be to wait until Y rises to 10^{90} and then exercise all options.

We shall therefore assume that $\mu_0 < r$ for this section. For this problem, it is better to be working with Y_t^0, the log of the asset price, rather than Y_t, the log of the discounted asset price. Introducing the notation $d\tilde{m}_u = e^{-rt}dm_u$, we have

$$\begin{aligned}
V_T^0(t,y,x,a) &\equiv \sup E\Big[\, U(x + e^{-rt}\int_t^T e^{-r(u-t)}\varphi_u^0\, dm_u) \;\Big|\; m_t = A-a, Y_t^0 = y \,\Big] \\
&= \sup E\Big[\, U(x + \int_t^T e^{-r(u-t)}\varphi_u^0\, d\tilde{m}_u) \;\Big|\; \int_t^T d\tilde{m}_u = e^{-rt}a, Y_t^0 = y \,\Big] \\
&= \sup E\Big[\, U(x + \int_0^{T-t} e^{-ru}\varphi_u^0\, d\tilde{m}_u) \;\Big|\; \int_0^{T-t} d\tilde{m}_u = e^{-rt}a, Y_0^0 = y \,\Big] \\
&\equiv f_{T-t}(y,x,e^{-rt}a) \\
&\uparrow \sup E\Big[\, U(x + \int_0^{\infty} e^{-ru}\varphi_u^0\, d\tilde{m}_u) \;\Big|\; \tilde{m}_\infty = e^{-rt}a, Y_0^0 = y \,\Big] \\
&\equiv f(y,x,e^{-rt}a)
\end{aligned}$$

as $T \uparrow \infty$. Notice that

$$V_T(t,y,x,a) = f_{T-t}(y + rt, x, ae^{-rt}).$$

The optimality characterisation, that for T fixed the process

$$V_T(t, Y_t, x_0 + \int_0^t e^{-ru}\varphi_u^0 dm_u, (A - m_t)) \quad \text{is a supermartingale}$$

and a martingale under optimal control, translates in the limit as $T \uparrow \infty$ to the statement

$$f(Y_t^0, x_0 + \int_0^t e^{-ru}\varphi_u^0 dm_u, \alpha_t) \quad \text{is a supermartingale} \qquad (16)$$

[7]We are assuming a constant interest rate throughout this section.

and a martingale under optimal control, where $\alpha_t \equiv (A - m_t)e^{-rt}$. The dynamics of the system are easy to describe. The process Y^0 diffuses as a Brownian motion with constant and volatility, and the residual number of options α_t decays at constant rate r. When $\Delta\alpha$ of the residual options are exercised at time t, the increment in the time-0 value of the cash held is $\Delta\alpha(\exp(Y_t^0) - K)^+$. As in Section 2, we make a discretisation where time moves in steps of Δt, and changes in α can only be made at multiples of Δt. When such changes are made, the change in $z \equiv \log(\alpha)$ must be a multiple of $\Delta z \equiv r\Delta t$.

We propose to solve the infinite-horizon problem numerically for the two examples we have dealt with earlier, exploiting the scaling relationships to reduce the dimension of the problem by 1. In the CARA example, we have the scaling relationship

$$F(y, x, z) = \exp(-\gamma x)F(y, 0, z) \equiv \exp(-\gamma x)F(y, z) \tag{17}$$

and in the CRRA example we have

$$F(y, x, z) = x^{1-R}F(y, 1, z - \log x) \equiv x^{1-R}F(y, z - \log x) \tag{18}$$

at the expense of a slight notational abuse.

In more detail, we discretise the problem by setting a grid $z_0 < z_1 < \ldots < z_J$ in the variable z with equal spacing Δz, and solve an American-style problem; when z reaches z_0, all remaining options are immediately exercised, and when $z = z_k$, $k > 0$, the decision is taken (based on the current value of Y^0) *either* to allow the process (Y^0, z) to diffuse for further time $\Delta t = \Delta z/r$ *or* to exercise enough options to jump z down immediately to z_{k-1}, with the appropriate scaling of the value of F.

The results of the calculations for the second example considered above (where the infinite-horizon problem is well posed) are now displayed in Figures 8 (for the CARA example) and 9 (for the CRRA example). In order to display everything on the same picture, we have replaced the time parameter by $\tau/(5 + \tau)$ in the plot, where τ denotes time-to-go. Thus the infinite-horizon limit corresponds to $\tau/(5 + \tau) = 1$ in the plots. Notice how these infinite-horizon results (computed using the method of this Section) match up well with the finite-horizon results computed using the (quite different) method of the previous Section. Notice also that for the numerical examples chosen here (with fairly typical parameter values) the finite-horizon result is noticeably different from the infinite-horizon result, even when time-to-go is of the order of 10 years. This suggests that the infinite-horizon result may not be appropriate for a given finite-horizon problem. For a finite horizon, the optimal policy is *not* approximately to wait until the price rises to a level which depends on the number of remaining options but not on time; even the *form* of the optimal rule in the finite-horizon problem is very different from what would be expected from the infinite-horizon analysis.

5 Conclusion

In this paper we showed that optimal exercise policies of American-call options when the option holder cannot trade on the underlying may involve partial exercise. There is an optimal exercise boundary that relates the time to expiration, the price of the underlying and the number of options held. The optimal policy consists of exercising enough options to stay below the boundary. We also computed the optimal exercise boundary for some examples and discussed the effect of changes in parameters on this boundary.

The insights in these paper should apply to a much larger set of problems. Risk averse agents facing incomplete markets will typically partially exercise a derivative security as the price of the underlying changes.

6 References

Carpenter, J. [1998] "The exercise and valuation of executive stock options" *Journal of Financial Economics* **48**, 12758.

Cvitanic, J. and I. Karatzas [1993] "Hedging contingent claims with constrained portfolios" *Annals of Applied Probability* **3**, 652-680.

Detemple, J. and S. Sundaresan [1999] "Nontraded asset valuation with portfolio constraints: A binomial approach" *Review of Financial Studies* **12**, 835-872.

Hall, B. J., and K. J. Murphy [2002] "Stock options for undiversified executives" *Journal of Accounting and Economics* **33**, 342

Ingersoll, J. [2006] "The subjective and objective evaluation of incentive stock options" *Journal of Business* **79**, 453-487

Jain, A. and A. Subramanian [2004] "The Intertemporal Exercise and Valuation of Employee Options" *The Accounting Review* **79**, 705-744

Lambert, R., D. Larcker and R. Verrecchia [1991] "Portfolio considerations in valuing executive compensation" *Journal of Accounting Research*, **29**, 129-49.

Rogers, L. C. G. and Williams, D. [2000] *Diffusions, Markov Processes, and Martingales, Volume 2*, Cambridge University Press, Cambridge.

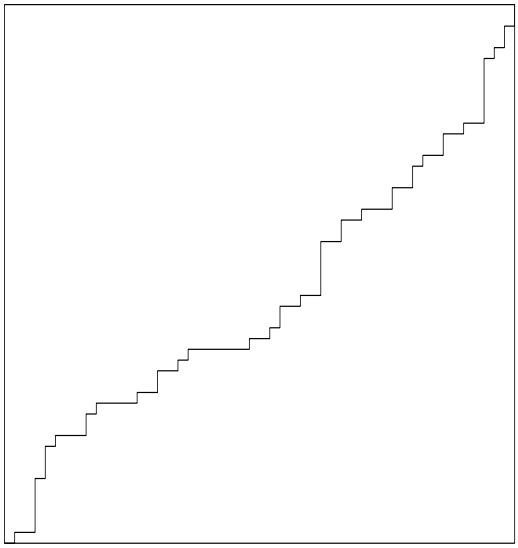

Figure 1: A typical path from $\mathcal{A}_n$.

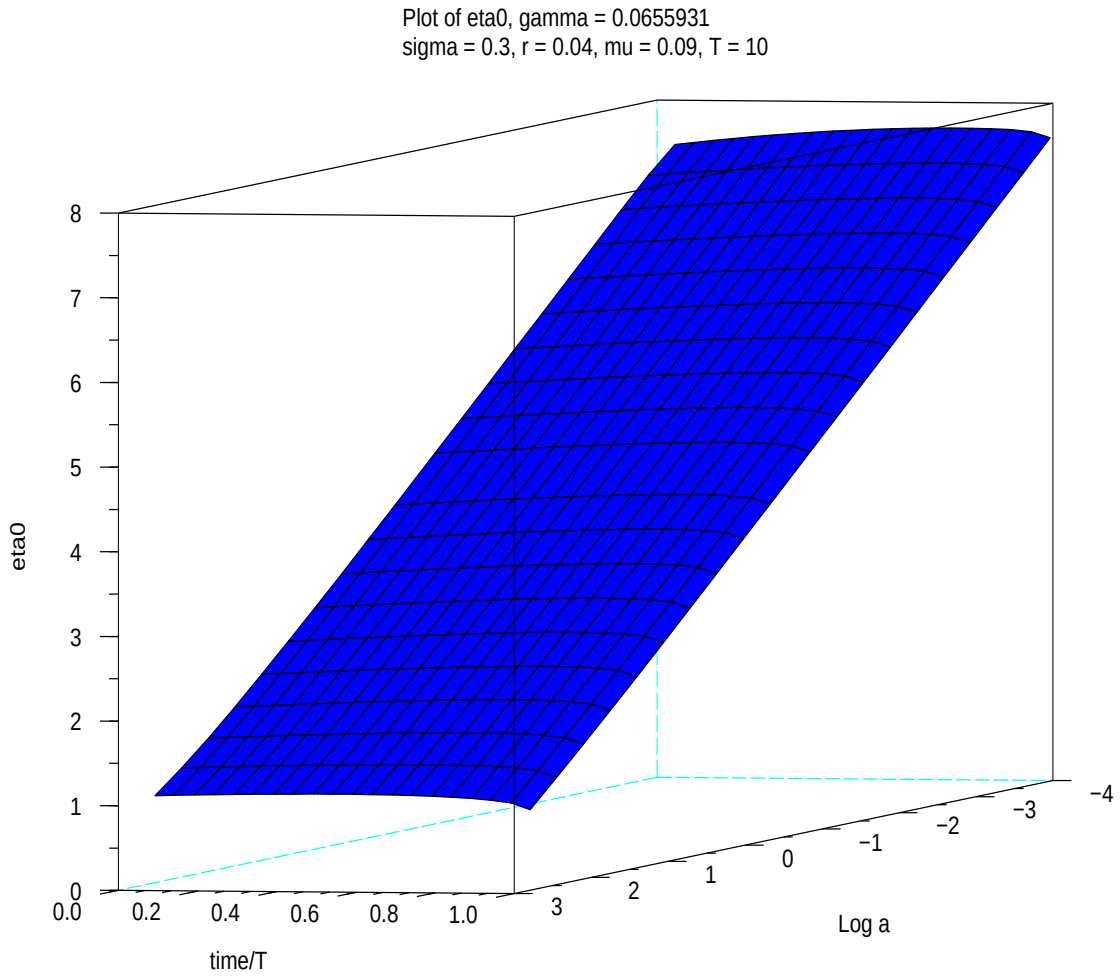

Figure 2: The critical level as a surface, base case.

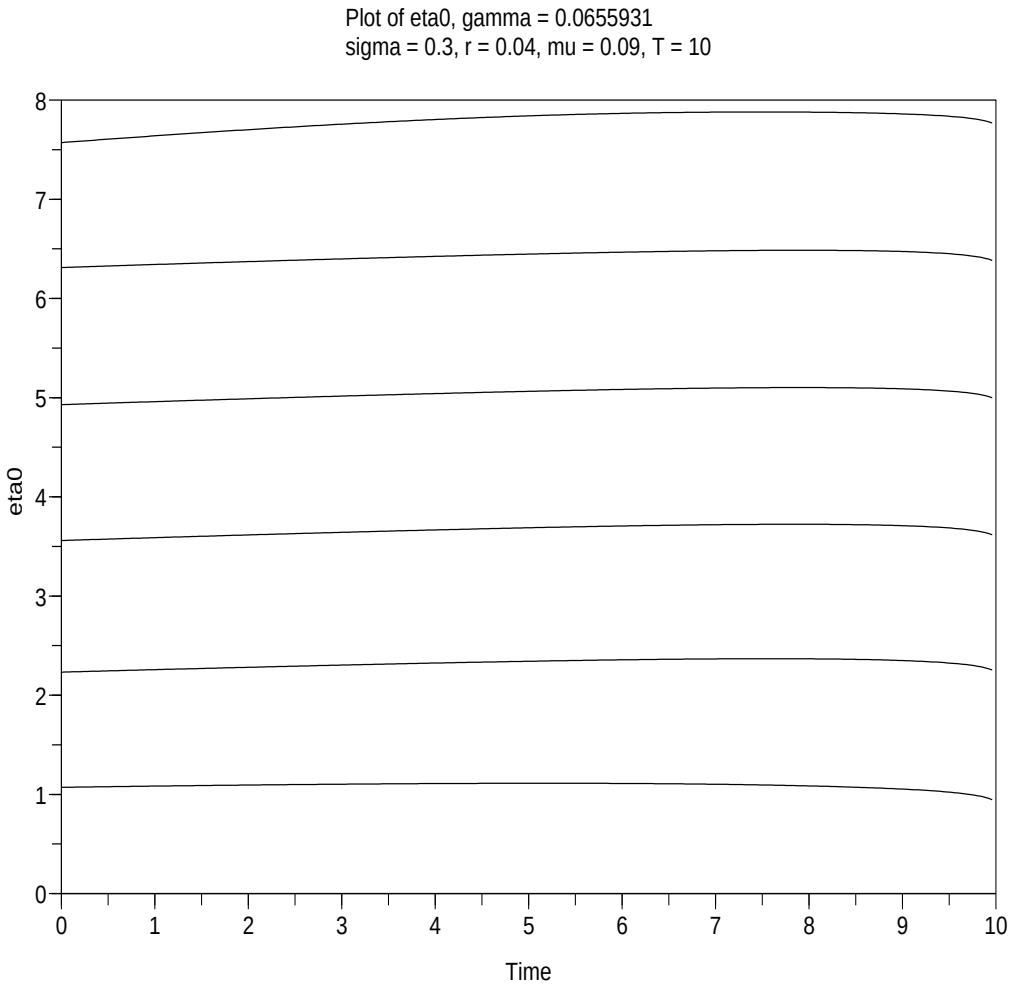

Figure 3: *Base case:* Critical level of η^0, for $a = 16, 4, 1, 1/4, 1/16, 1/64$ from bottom to top.

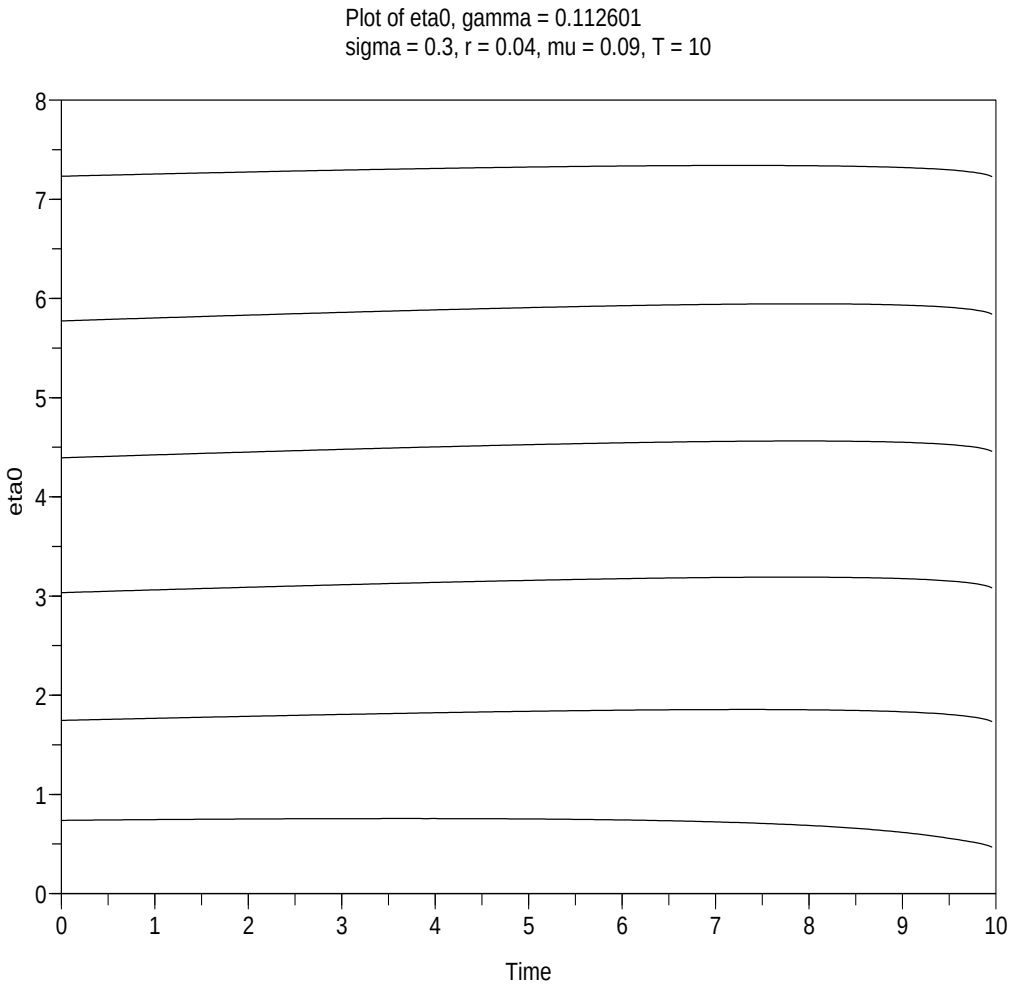

Figure 4: *Higher absolute risk aversion*: Critical level of η^0, for $a = 16, 4, 1, 1/4, 1/16, 1/64$ from bottom to top.

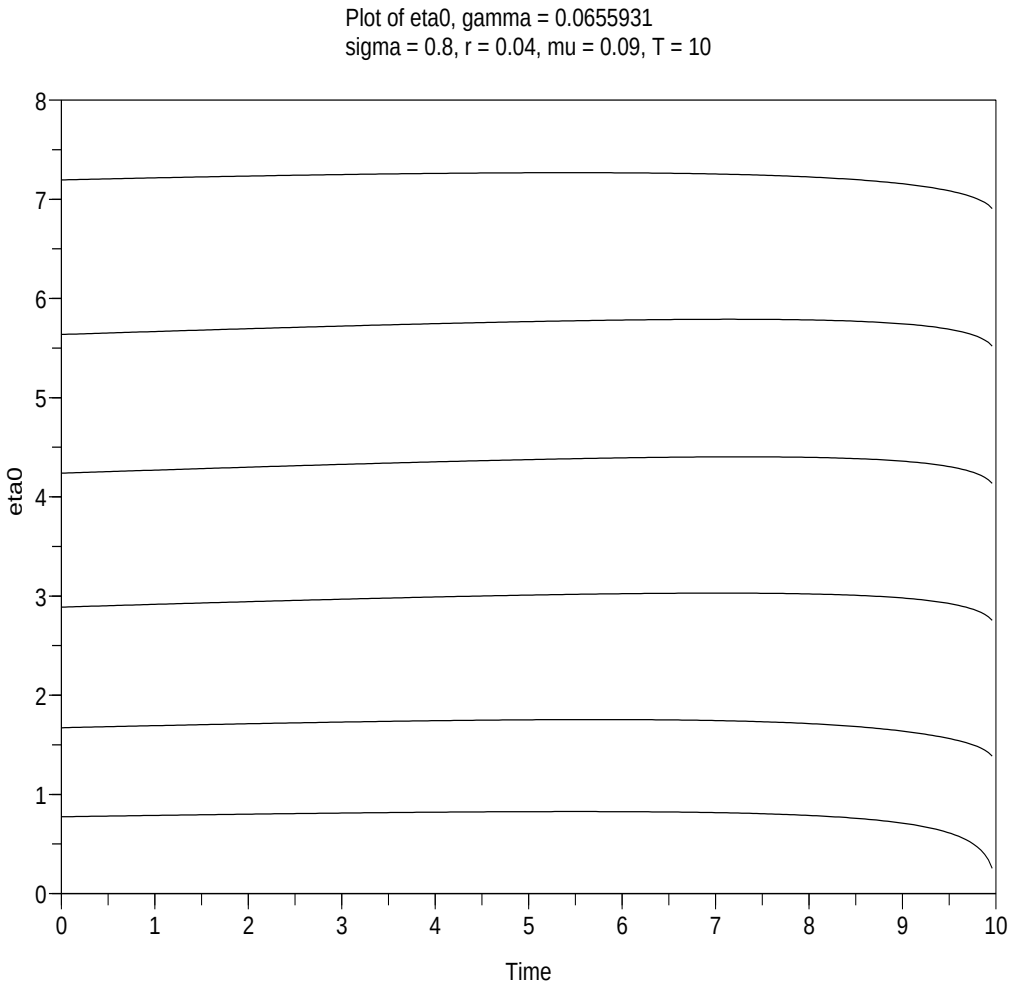

Figure 5: *Higher volatility*: Critical level of η^0, for $a = 16, 4, 1, 1/4, 1/16, 1/64$ from bottom to top.

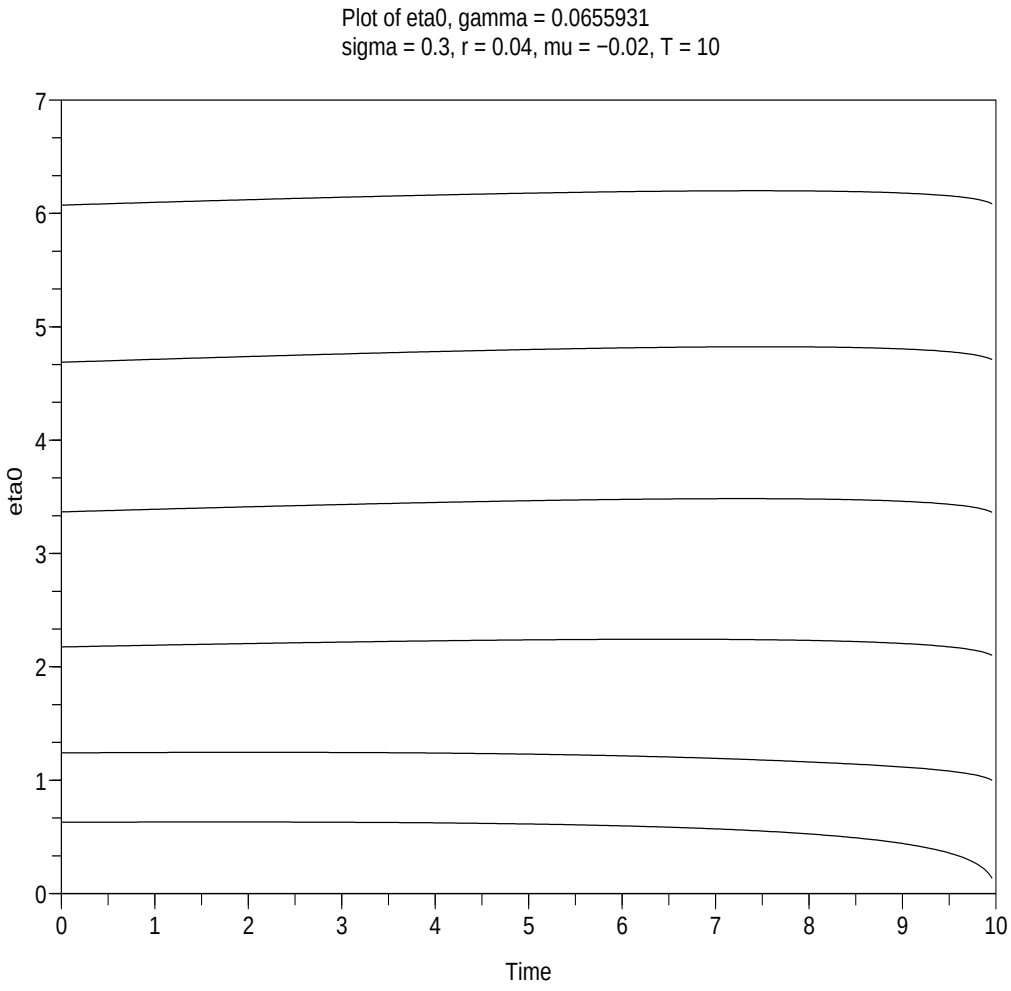

Figure 6: *Lower growth rate*: Critical level of η^0, for $a = 16, 4, 1, 1/4, 1/16, 1/64$ from bottom to top.

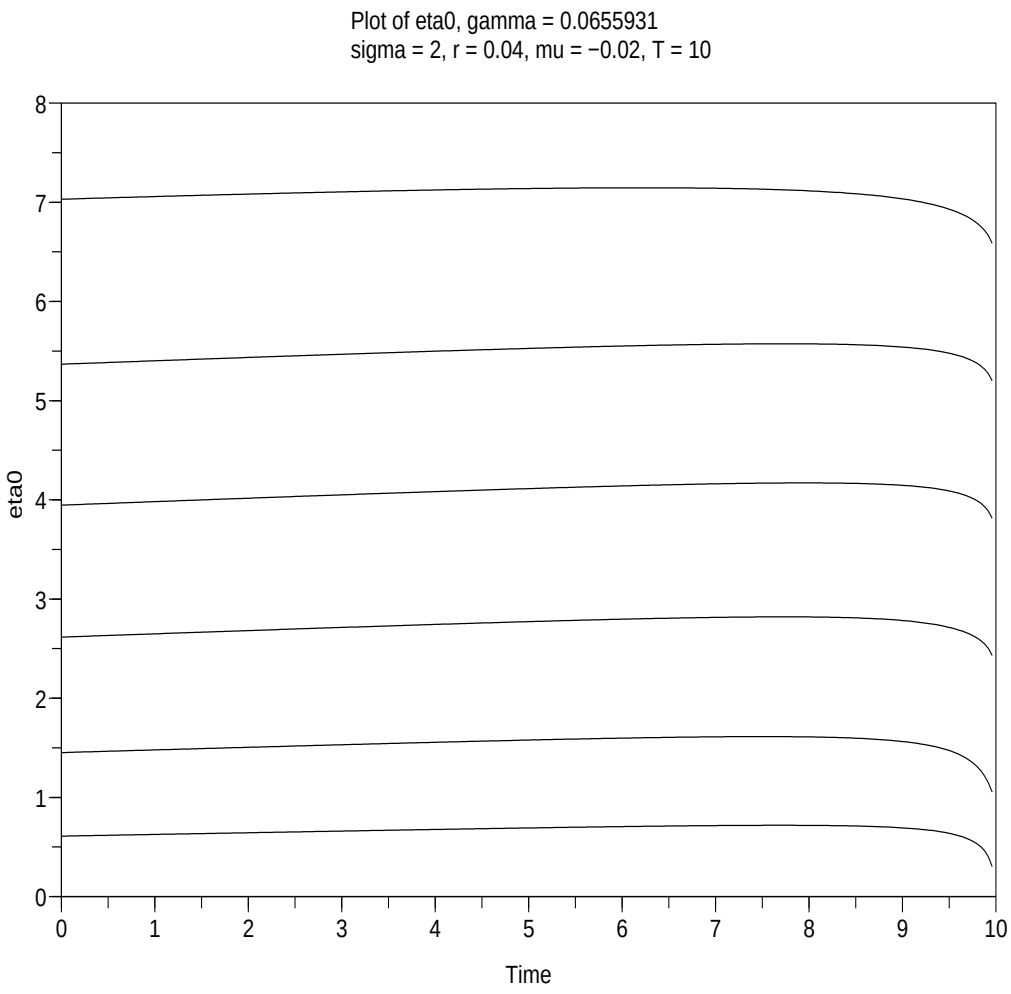

Figure 7: *Lower growth rate, higher volatility*: Critical level of η^0, for $a = 16, 4, 1, 1/4, 1/16, 1/64$ from bottom to top.

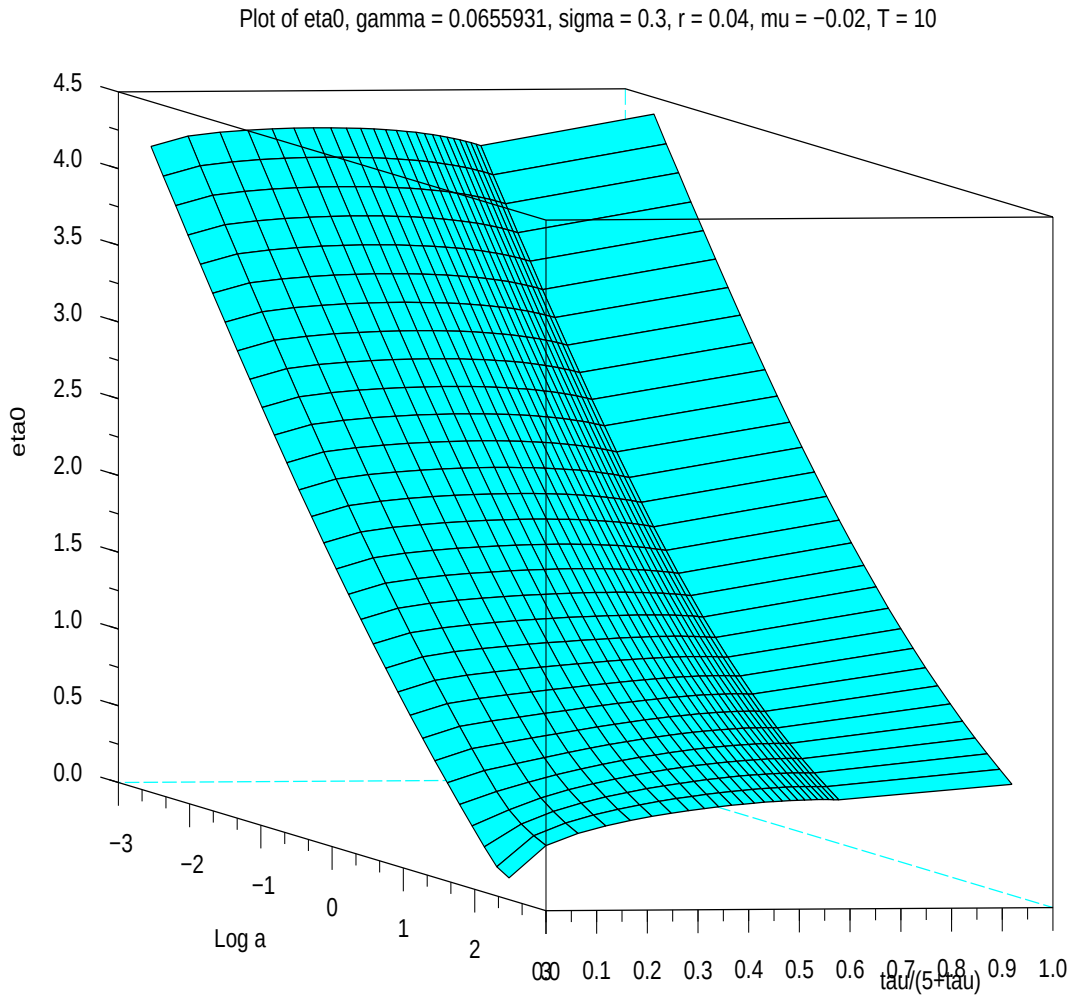

Figure 8: Critical surface as a function of $\tau/(5+\tau)$, where τ is time-to-go. values of a, parameters as given, CARA example.

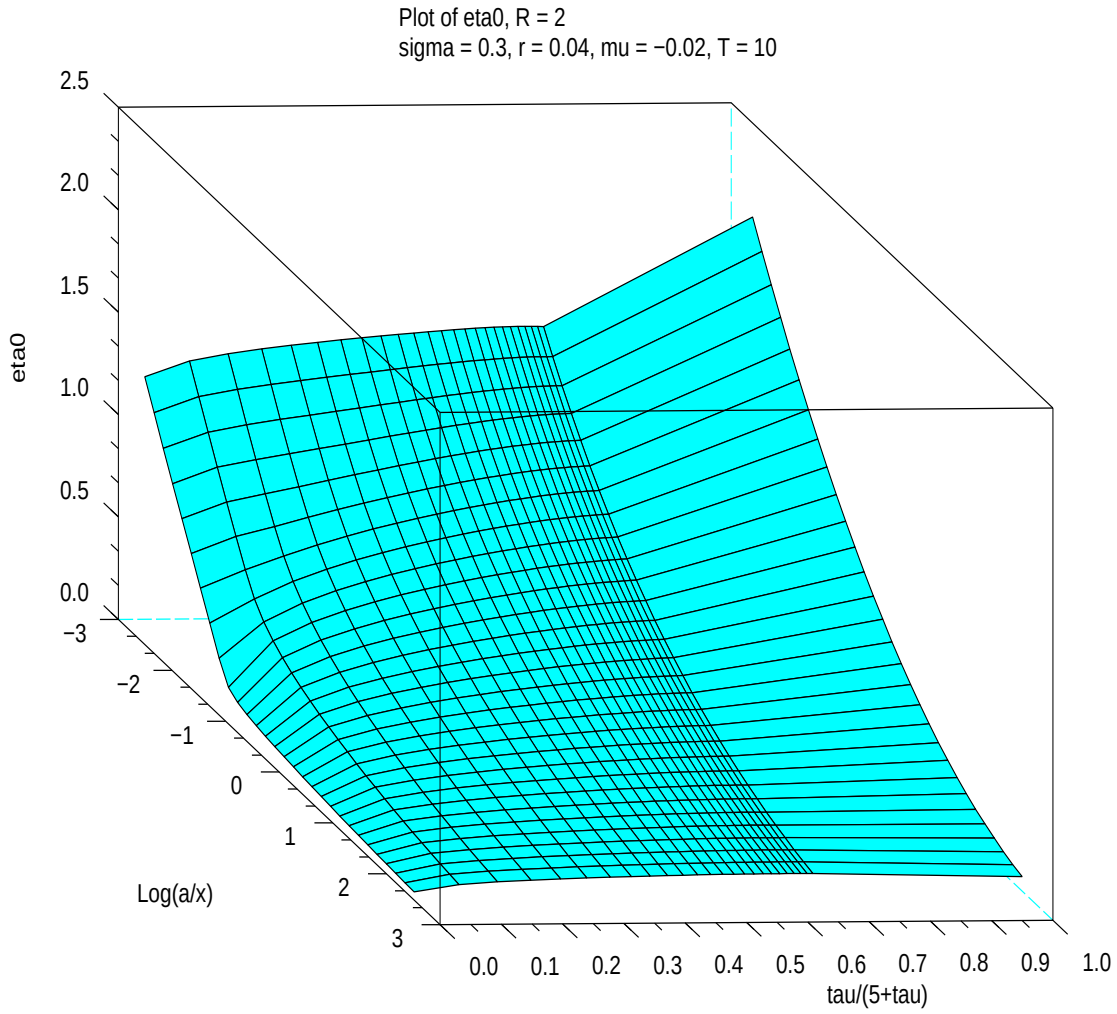

Figure 9: Critical surface as a function of $\tau/(5+\tau)$, where τ is time-to-go. values of a, parameters as given, CRRA example.