WORKING PAPERS
McMASTER U

FACULTY OF BUSINESS

Optimal Control of an Ornstein-Uhlenbeck Diffusion Process, with Applications

McMaster University Library
JAN 15 1985

NORMAN P. ARCHER

INNIS LIBRARY
NON-CIRCULATING

FACULTY OF BUSINESS

McMASTER UNIVERSITY

HAMILTON, ONTARIO, CANADA L8S 4M4

Innis
[redacted]
HB
74.5
.R47
no.231

Research and Working Paper Series No. 231
November, 1984

OPTIMAL CONTROL OF AN ORNSTEIN-UHLENBECK DIFFUSION PROCESS, WITH APPLICATIONS*

Norman P. Archer

Faculty of Business

McMaster University

Hamilton, Ontario, Canada L8S 4M4

Telephone (416) 525-9140, Ext. 3944

*This work was supported by a grant from the Natural Sciences and Engineering Research Council of Canada.

McMASTER UNIVERSITY LIBRARY

Optimal Control Of An Ornstein-Uhlenbeck Diffusion Process, With Applications

Summary

The Ornstein-Uhlenbeck diffusion process presents an opportunity for the development of approximation models of many real processes in business and industry because it is the continuous analog of the first order autoregressive process. The Ornstein-Uhlenbeck process has two decision variables: one relates to the target level of the process trajectory, and the other relates to the dispersion of the process sample paths. An optimization model of the process is developed, which includes holding or carrying costs, control costs and penalty costs. The penalty costs are related to the nearness of the process trajectory to reflecting barriers which are included in the model.

Key words: Ornstein-Uhlenbeck diffusion

Diffusion model

Buffer stock model

First order autoregressive model

1. Introduction

In recent years a considerable amount of interest has been aroused in developing diffusion approximations for models of stochastic processes in management science. The interest in diffusion approximations is due to the relative ease with which solutions may be obtained to complex stochastic problems which are often intractable if modeled exactly. Diffusion process models of inventory systems and (storage systems) were first developed by Bather[2,3]. A sample of the many references for diffusion approximations of storage system models are references 11 and 14. Related examples of diffusion models for financial operations appear in references 8, 9, 10, and 12.

Bhat, Shalaby and Fischer[4] have also published a survey of approximation techniques for queuing systems, including a large number of references for diffusion approximations.

In the Brownian diffusion processes used to model storage systems of various types, control is usually applied as an impulse at certain instants when the process trajectory approaches some pre-defined position(s). For example, in an (s,S) inventory control diffusion model of inventory position with negative drift (the rate of drift being used to model the rate of use of inventory), more stock is ordered when the process crosses the order boundary S. This results in an instantaneous jump away from the lower boundary in terms of stock position. In the Brownian diffusion model of a dam, an instantaneous control is applied by modeling a reduction in the water level by releasing flow when the level is too high.

In many real processes, control is not applied instantaneously as the stochastic process trajectory approaches a boundary. Rather, control is applied more or less continuously so as to maintain a limiting distribution of the process sample paths within the normal operating region. Many continuous processes in the chemical industry, such as mixing, heating, reacting, etc. are of this type. Some discrete processes are also subject to continuous control activities, including buffer stock level control in multi-stage production lines.

One feature of processes in which continuous control may be applied is that discrete time series measurements of the process trajectories reveal that these trajectories can often be analyzed by means of autoregressive models. This paper will discuss the Ornstein-Uhlenbeck[17] (O.U.) diffusion process which exhibits continuous control, and demonstrate its applicability to modeling certain processes. Some of these processes have a discrete time series behavior which may be fitted by first order autoregressive models, and the stochastic differential equation describing the O.U. model turns out to be the continuous analog for first order autoregressive models.

2. The Diffusion Model

Most diffusion approximations make use of the Wiener process (Brownian motion) in one dimension. Brownian motion is a diffusion with a generator which is the linear second-order differential operator

$$\frac{\sigma^2}{2}\frac{d^2}{dx^2} + \mu\frac{d}{dx}$$

and domain equal to the twice continuously differentiable functions on the real line. Here, $\sigma^2/2$ is the diffusion coefficient and μ is the drift coefficient. On the other hand, the Ornstein-Uhlenbeck (O.U.) process is a diffusion with generator

$$\frac{\sigma^2}{2}\frac{d^2}{dx^2} - \beta x\frac{d}{dx}$$

with domain the same as for Brownian motion. Here, β is the constant of proportionality for a controlling or restoring force which will be referred to later. It is restricted to $\beta > 0$. The density of the O.U. process which satisfies the related forward Kolmogorov equation

$$\frac{\sigma^2}{2}\frac{\partial^2 f(x,x_0;t)}{\partial x^2} + \frac{\beta\partial x f(x,x_0;t)}{\partial t} = \frac{\partial f(x,x_0;t)}{\partial t}$$

in the absence of barriers for $t \geq 0$, is a Gaussian diffusion with mean

$$E(X) = x_0\exp(-\beta t) \tag{1a}$$

and variance

$$\mathrm{Var}(X) = \sigma^2\{1-\exp(-2\beta t)\}/2\beta \tag{1b}$$

Here, x_0 is the starting point of the process. If $\beta = 0$, then the O.U. process can be shown to be equivalent to a Wiener process with zero drift. Also note that, while both the O.U. and Wiener processes are Gaussian, the asymptotic mean and variance for the O.U. process are 0 and $\sigma^2/2\beta$ respectively, but the mean of the Wiener process is $x_0 + \mu t$ and its variance grows without limit. Although both processes are Markovian, the Wiener process has independent increments, unlike the O.U. process.

If the trajectory of the O.U. process is given by x_t at time t, then there is an associated cost arising from the sample path which is of the magnitude $c(x_t)dt$ in the time interval $[t, t + dt]$. Hence the total cost of the process over some time interval $[t_1, t_2]$ is

$$\int_{t_1}^{t_2} c(x_t)\, dt$$

In the limit, as $t \to \infty$, we are interested in the expected cost C per unit time from an equilibrium process, which can be determined from the asymptotic O.U. distribution. This is given by

$$C = \int_x f(x) \lim_{t \to \infty} \frac{1}{t} \int_0^t c(x_s) ds\, dx + B(\beta)$$

$$= \int_x f(x)\, c(x)\, dx + B(\beta)$$

where $$f(x) = \frac{1}{\sqrt{2\pi\gamma^2}} \exp(-x^2/2\gamma^2)$$

and $\gamma^2 = \sigma^2/2\beta$, the variance of the limiting distribution of the O.U. process. $B(\beta)$ is the cost contribution from control activities which are a function of β alone.

Consider the first order stochastic differential equation

$$\frac{dX(t)}{dt} + \beta X(t) = \frac{d\sigma z(t)}{dt} \tag{2}$$

where z(t) is a scalar white noise. This equation is known as the Langevin equation in fluid dynamics, and the equilibrium solution is a Gaussian process with mean and variance given by the asymptotic forms of (1a) and (1b). The autocovariance of the O.U. process is[1]

$$\mathrm{Cov}(X,X+s) = \frac{\sigma^2}{2\beta}\exp(-\beta s).$$

Thus, the autocorrelation function for the O.U. process is

$$\rho(s) = \exp(-\beta s).$$

Let us consider the measurement of time series, which are important in many aspects of business and industry. Time series are discrete samples of the levels of either continuous or discrete systems at uniform time intervals of Δ. We may obtain a continuous approximation to the process described by a discrete time series by equating the autocorrelation function of the time series to that of the approximating diffusion process. Equation 2 is the continuous analog of the discrete AR(1) (first order autoregressive) model described below.[6] For this model, the time series is given by

$$x_t = \phi_1 x_{t-1} + \sigma_a z_t.$$

Here, ϕ_1, is a constant to be estimated from the time series, with $-1 < \phi_1 < 1$, and the observations x_t are taken at time increments of Δ. However, the restriction $\beta > 0$ for a stationary process also restricts $0 < \phi < 1$ for the continuous approximation to be valid. z_t is a pure noise process with mean zero and variance one, and σ_a^2 is the variance of the inherent noise. Clearly, disturbances in the AR(1) process decay exponentially with the passage of time, when $0 < \phi < 1$ as in the equivalent continous process above when $\beta > 0$. The autocorrelation function for the AR(1) process is

$\rho_1 = \phi_1$.

Equating the autocorrelation functions for the discrete and approximating continuous processes, with $s = \Delta$, gives

$$\phi_1 = \exp(-\beta\Delta), \text{ or}$$

$$\beta = -\frac{\ln\phi_1}{\Delta}.$$

In the following we will be concerned with the properties of an O.U. process not only in open one-dimensional space as discussed above, but also with its properties when confined between reflecting barriers $x_w \leq x \leq x_m$. Sweet and Hardin[16] have shown that the asymptotic density for such an O.U. process is given by

$$f(s) = \exp(-s^2/2) \Big/ \int_{r(x_w)}^{r(x_m)} \exp(-r^2/2)dr \qquad s_w \leq s \leq s_m \tag{3}$$

where $s(x) = r(x) = (x - \alpha)/\gamma$.

In the O.U. model, with $x_w \leq \alpha \leq x_m$, α corresponds to the mode of the distribution $f(s)$ which is a truncated gaussian. If α is outside this range, then the mode is x_w or x_m if $\alpha < x_w$ or $\alpha > x_m$ respectively. β is restricted to positive values and is the proportionality constant of the restoring force which maintains the trajectory around the target level α.

The denominator in (3) will be denoted as $U(\alpha,\beta)$, and is easily shown to be

$$U(\alpha,\beta) = \sqrt{\frac{\pi}{2}}\left[\mathrm{erf}\{(\beta/\sigma^2)^{1/2}|\alpha-x_w|\}\mathrm{sgn}(\alpha-x_w) + \mathrm{erf}\{(\beta/\sigma^2)^{1/2}|x_m-\alpha|\}\mathrm{sgn}(x_m-\alpha)\right]$$

where $\mathrm{erf}(v) = \frac{2}{\sqrt{\pi}} \int_0^v e^{-t^2} dt$

3. Cost Model

A cost model can be developed for the O.U. process which will allow optimization of the process with respect to the decision variables α and β.

For generality we will assume that the process is confined between reflecting barriers, either or both of which may be removed if desired.

The objective function will be assumed to consist of cost contributions from three sources: carrying or holding costs H, control costs B, and penalty costs G_w and G_m related to the distance of the trajectory from lower and upper boundaries respectively. Then the optimization problem is

$$\min_{\alpha,\beta} C = H + B + G_w + G_m \tag{4}$$

subject to $\beta > 0$.

Note that the holding cost contribution is a function of the absolute level of the process at any time t, while the control cost is a function of the decision variable β which controls the dispersion of the process $|x_t - \alpha|$.

To demonstrate that optimal solutions may be obtained for the model, certain functional forms will be assumed for these cost functions in the following discussion.

The holding cost per unit time may be given by

$$\begin{aligned} H &= \int_{s(x_w)}^{s(x_m)} \gamma c_h s f(s) ds + c_h \alpha \\ &= c_h\{\gamma E(S) + \alpha\} \\ &= c_h \gamma\{\exp[-(\beta/\sigma^2)(x_w-\alpha)^2] - \exp[-(\beta/\sigma^2)(x_m-\alpha)^2]\}/U(\alpha,\beta) + c_h\alpha. \end{aligned}$$

Here, c_h is the cost per unit time to carry each unit of X.

The value of β is inversely proportional to the variance of the process. Thus, to hold the process trajectory within a narrower range about the target level α, the value of β must be increased. Increased values of β imply tighter control which tends to be more costly in a manufacturing or physical process. In many contexts, this cost may be expected to increase more rapidly than as a linear function of β. In general, any suitable increasing function of β may be used. In this example we will assume that control cost B is a function of β^2, or

$$B = c_b\beta^2.$$

Here, c_b is the cost per unit time per unit increase in the control function.

The general nature of the penalty functions g_w and g_m which contribute to the penalty cost functions G_w and G_m is that they will tend to increase sharply near the boundaries x_w and x_m respectively, and to tail off quickly as the process trajectory moves away from the boundaries. A variety of functions may be used in this case, depending upon how well the operational situation is approximated. In this example, we will assume that two separate penalty functions are used for the upper and lower boundaries respectively, and that the penalty functions decrease exponentially with distance from the boundaries. Then for the lower boundary penalty function we have

$$g_w = a_w \exp\{-d_w(x-x_w)\},$$

leading to

$$G_w = \frac{\int_{x_w}^{x_m} a_w \exp\{-d_w(x-x_w)\}\exp-\{\frac{(x-\alpha^2)}{2\gamma^2}\}dx}{\gamma\, U(\alpha,\beta)}$$

where a_w is the maximum penalty per unit time incurred at the lower boundary of the trajectory, and d_w is the fall-off rate of the penalty per unit distance above the lower boundary.

Removing the constant terms from the integral gives

$$G_w = \frac{a_w \exp(d_w x_w)}{\gamma U(\alpha,\beta)} \int_{x_w}^{x_m} \exp\{\frac{-(x-\alpha)^2}{2\gamma^2} - d_w x\}\,dx$$

Completing the square and simplifying leads to

$$G_w = R_w \int_{x_w}^{x_m} \exp[- \frac{1}{2x^2} \{x + \gamma^2 d_w - \alpha\}^2] dx$$

where

$$R_w = \frac{a_w \exp d_w\{x_w - \alpha + \frac{\gamma^2 d_w}{2}\}}{\gamma\ U(\alpha,\beta)}$$

and then

$$G_w = \sqrt{2}\ \gamma\ R_w \int_{\frac{1}{\sqrt{2}\gamma}\{x_w+(\gamma^2 d_w-\alpha)\}}^{\frac{1}{\sqrt{2}\gamma}\{x_m+(\gamma^2 d_w-\alpha)\}} \exp(-t^2)\ dt$$

$$= \gamma R_w \sqrt{\frac{\pi}{2}} [\mathrm{sgn}(y_m)\mathrm{erf}|y_m| + \mathrm{sgn}(y_w)\mathrm{erf}|y_w|]$$

where $y_m = \frac{1}{\sqrt{2}\gamma}(x_m + \gamma^2 d_w - \alpha)$

and $y_w = \frac{1}{\sqrt{2}\gamma}(\alpha - x_w - \gamma^2 d_w)$.

To determine G_m, where the penalty function g_m related to the upper boundary is given by $a_m \exp\{-d_m(x_m-x)\}$, a similar approach gives

$$G_m = \gamma R_m \sqrt{\frac{\pi}{2}} [\mathrm{sgn}(z_m)\mathrm{erf}|z_m| + \mathrm{sgn}(z_w)\mathrm{erf}|z_w|]$$

where $z_m = \frac{x_m - (\alpha + \gamma^2 d_m)}{\sqrt{2}\ \gamma}$,

$z_w = \frac{\alpha + \gamma^2 d_m - x_w}{\sqrt{2}\ \gamma}$,

and $R_m = \frac{a_m \exp\{d_m(\alpha - x_m + \frac{\gamma^2 d_m}{2})\}}{\gamma\ U\ (\alpha,\beta)}$.

4. <u>Optimization</u>

The objective function (4) is highly non-linear, and the GRG2 non-linear programming package of Lasdon, Waren et al[13] was used for solution. It is easy to show that the Hessian of the objective function is not generally positive semidefinite for the functions used in the foregoing example. It is therefore necessary to be cautious about the interpretation of results since, if there is a finite solution, more than one minimum may exist in certain cases. This was occasionally found to be the case, for example, when constant penalty functions were used over limited regions near the boundaries instead of the exponential penalty functions discussed in the foregoing section.

To demonstrate the sensitivity of decision variables to parameter values, two examples were set up for solution. The first example used symmetrical penalty functions and zero holding costs. Figure 1 shows the optimal values ($\hat{\beta}$) of β and the total cost obtained for different values of c_b and $d_w(=d_m)$. Note that in this symmetrical model with no holding cost, the optimal value ($\hat{\alpha}$) of α is always the midpoint between the two reflecting barriers at $x_w=0$ and $x_m=100$. Values used for the constant parameters were $\sigma=10$ and $a_m=a_w=100$. $\hat{\beta}$ first increases and then decreases as $d_w=d_m$ increases because, for higher values of d_w, the penalty regions are very narrow and close to the boundaries. $\hat{\beta}$ may be quite small here because the variance of the process can be relatively large with little penalty. As $d_w=d_m$ decreases, the penalty regions widen, forcing $\hat{\beta}$ up and hence the variance decreases. Eventually when d_w becomes quite small the penalty region is

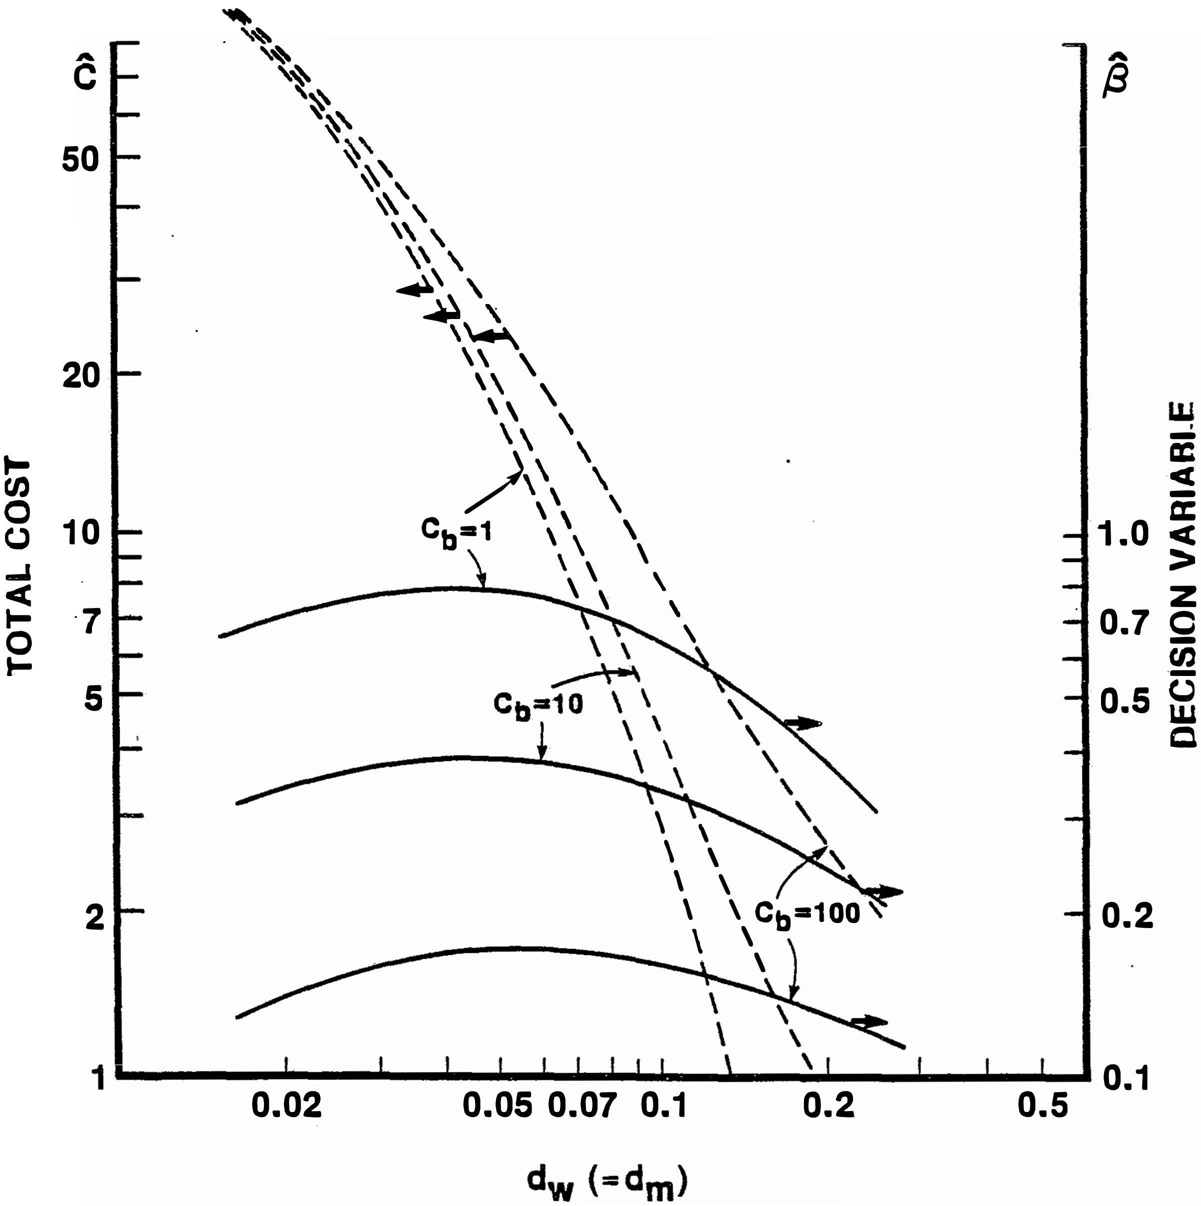

Figure 1. Variation of Optimal $\hat{\beta}$ and total cost C with penalty function parameter d_w (= d_m for penalty function symmetric about x = 50). $a_m = a_w = 100$, $x_w = 0$, $x_m = 100$, $\sigma = 10$, and $\hat{\alpha} = 50$ in all cases. No holding cost.

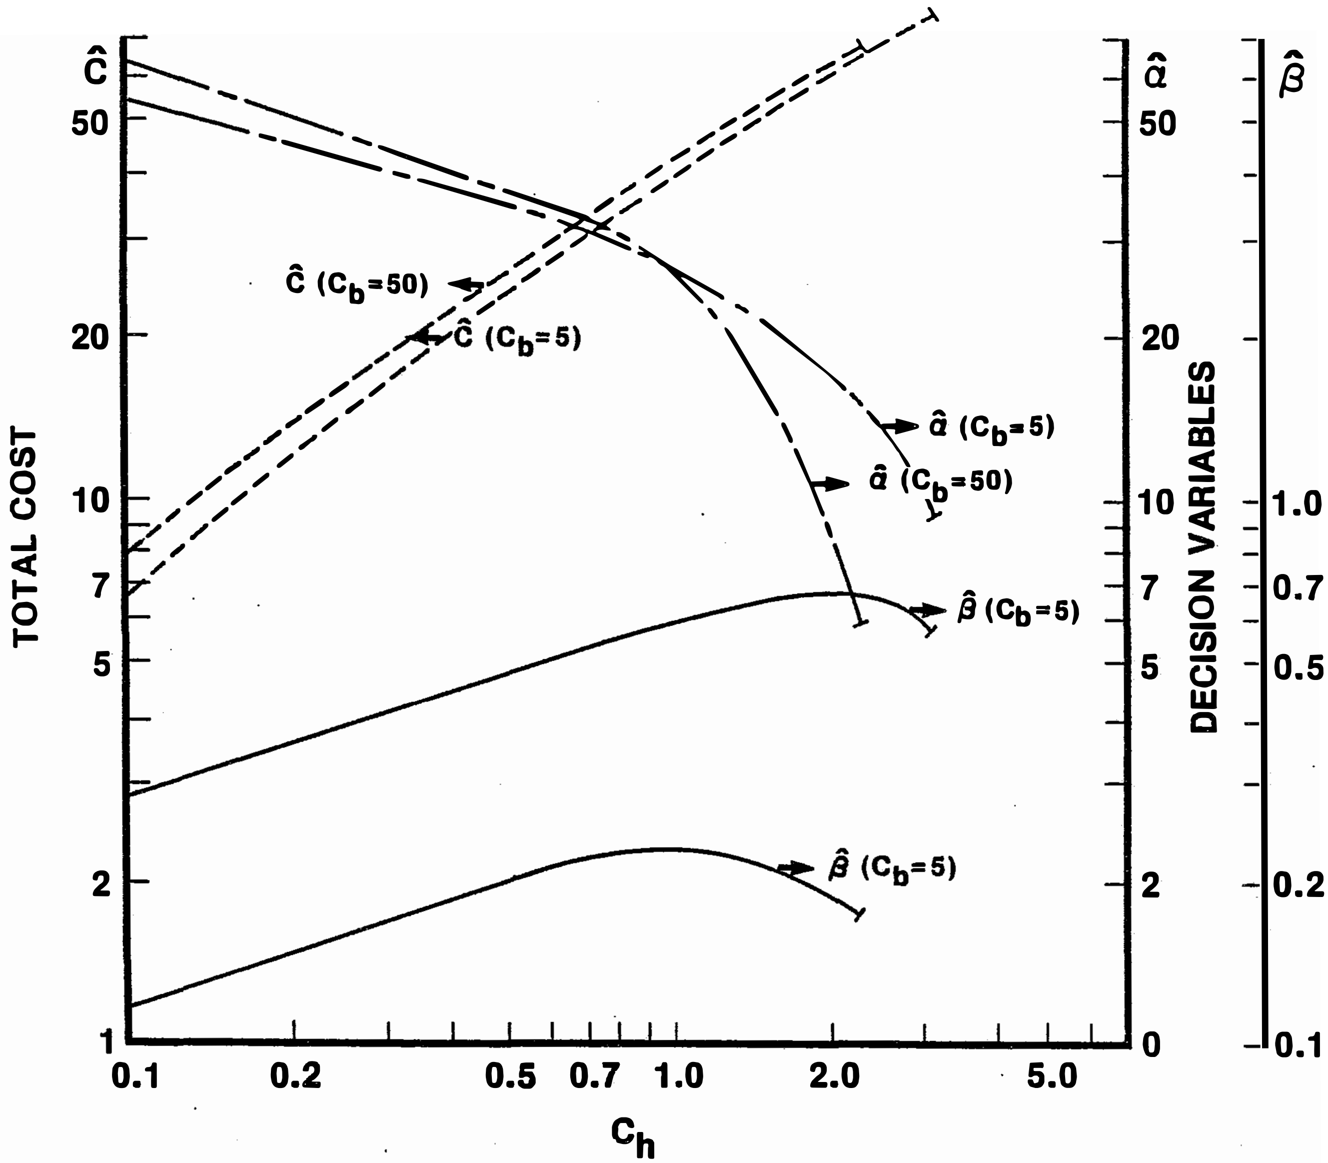

Figure 2. Variation of optimal α, β and total cost C with holding cost parameter c_h. No upper penalty function, and $x_w = 0$, $a_w = 100$, $d_w = 0.1$ and $\sigma = 10$.

quite flat across the entire range and β again falls because there is little cost incurred by increasing the variance in this region.

The second example included a non-zero holding cost, but with only a lower boundary. Figure 2 shows β, $\hat{\alpha}$ and the total cost as a function of c_h and c_b with fixed d_w=0.1 and a_w=100. In this example, σ=10. Note that the algorithm diverged above certain values of c_h, depending upon the value of c_b, indicating that no minimum exists for finite values of $\hat{\alpha}$ in this region.

5. Conclusions

It is clear that there is a wide scope of application for models using the O.U. diffusion process, since it can be used to approximately model systems which have been found by time series analysis to be first order autoregressive. These represent a wide class of real systems and may be considered to be under continuous control. Assuming that penalty functions used in the O.U. model are sufficiently large to prevent process trajectories from approaching the boundaries more than a small fraction of the time, (this was the case in the majority of the examples demonstrated here) then processes confined within reflecting barriers such as the models developed in section 3, are good approximations to AR(1) type processes when $0 < \phi_1 < 1$. For example, Steudel[15] has shown that AR(1) models are good representations of buffer stocks in high-rise storage for a multi-stage production line, and this situation certainly has upper and lower boundaries with associated penalty costs. Diffusion approximations are well suited to buffer stock level models. In terms of the functional measure for control costs, Davis and Taylor[7] have discussed the balancing of in-process inventories or buffer stocks as an on-going feature of production line control. Resource shifting and balancing is a necessary management activity which maintains buffer stocks at appropriate levels, and costs may be determined for these control activities.

References

1. Arnold, Ludwig, Stochastic Differential Equations: Theory And Applications, Wiley, New York (1974).

2. Bather, J., "A Continuous Time Inventory Model", Journal of Applied Probability 3, 538-549 (1966).

3. Bather, J., "A Diffusion Model For The Control of A Dam", Journal of Applied Probability 5, 55-71 (1968).

4. Bhat, M.N., Shalaby, M. and Fischer, M.J., "Approximation Techniques in the Solution of Queueing Problems", Naval Research Logistics Quarterly, Vol. 26 (1979), pp. 311-326.

5. Box, George E.P. and Jenkins, Gwilym M., Time Series Analysis Forecasting And Control, Holden-Day, San Francisco (1970), Ch. 3.

6. Cox, D.R. and Miller, H.D., The Theory of Stochastic Processes, Chapman and Hall, London (Paperback Edition, 1977), Ch. 5.

7. Davis, K. Roscoe and Taylor, Bernard W. III, "A Heuristic Procedure for Determining In-Process Inventories", Decision Sciences, Vol. 9 (1978), pp. 452-466.

8. Doshi, Bharat T., "Controlled One Dimensional Diffusions With Switching Costs-Average Cost Criterion", Stochastic Processes And Their Application, Vol. 8 (1978), pp. 211-227.

9. Foster, W.F., "Optimal Reserves Growth Under a Wiener Process", J. Appl. Prob., Vol. 12 (1975), pp. 457-465.

10. Harrison, J. Michael, "Ruin Problems With Compounding Assets", Stochastic Processes And Their Applications, Vol. 5 (1977), pp. 67-79.

11. Harrison, J. Michael and Taylor, Allison, A.J., "Optimal Control of a Brownian Storage System", Stochastic Processes And Their Applications", Vol. 6, 1978, pp. 179-194.

12. Iglehart, D.L. "Diffusion Approximations in Collective Risk Theory", J. Appl. Prob., Vol. 6 (1969), pp. 285-292.

13. Lasdon, L.S., A.D. Waren, A. Jain and M. Ratner, "Design And Testing of A Generalized Reduced Gradient Code For Nonlinear Programming", A&M Transactions on Mathematical Software, March 1978, pp. 34-50.

14. Puterman, Martin L., "A Diffusion Process Model For A Storage System", in Logistics, Murray A. Geisler (Ed.), American Elsevier, New York, (1975), pp. 143-159.

15. Steudel, Harold J., "Monitoring and Controlling In-Process Inventories Within High-Rise Storage Via Time Series Analysis", Int. J. Prod. Res., Vol. 15 (1977), pp. 383-390.

16. Sweet, A.L. and Hardin, J.C., "Solutions for Some Diffusion Processes With Two Barriers", J. Appl. Prob., Vol. 7 (1970), pp. 423-431.

17. Uhlenbeck, G.E. and Ornstein, L.S., "On The Theory of Brownian Motion", Physical Review, Vol. 36 (1930), pp. 823-841.

Faculty of Business

McMaster University

WORKING PAPERS - RECENT RELEASES

207. Min Basadur and Carl T. Finkbeiner, "Identifying Attitudinal Factors Related to Ideation in Creative Problem Solving", June 1983.

208. Min Basadur and Carl T. Finkbeiner, "Measuring Preference for Ideation in Creative Problem Solving", June 1983.

209. George Steiner, "Sequencing on Single Machine with General Precedence Constraints - The Job Module Algorithm", June 1983.

210. Varouj A. Aivazian, Jeffrey L. Callen, Itzhak Krinsky and Clarence C.Y. Kwan, "The Demand for Risky Financial Assets by the U.S. Household Sector", July 1983.

211. Clarence C.Y. Kwan and Patrick C.Y. Yip, "Optimal Portfolio Selection with Upper Bounds for Individual Securities", July 1983.

212. Min Basadur and Ron Thompson, "Usefulness of the Ideation Principle of Extended Effort in Real World Professional and Managerial Creative Problem Solving", October 1983.

213. George Steiner, "On a Decomposition Algorithm for Sequencing Problems with Precedence Constraints", November 1983.

214. Robert G. Cooper and Ulrike De Brentani, "Criteria for Screening Industrial New Product Ventures", November 1983.

215. Harish C. Jain, "Union, Management and Government Response to Technological Change in Canada", December 1983.

216. Z. Drezner, G. Steiner, G.O. Wesolowsky, "Facility Location with Rectilinear Tour Distances", March 1984.

217. Latha Shanker and Jack S.K. Chang, "The Hedging Performance of Foreign Currency Options and Foreign Currency Futures: A Comparison", March 1984.

218. Ulrike de Brentani and Robert G. Cooper, "New Product Screening Decisions in High Technology Firms: An Empirical Study", March 1984.

219. Paul D. Dowling and Robert F. Love, "Bounding Methods for Facilities Location Algorithms", March 1984.

220. George Steiner, "A Decomposition Approach for Finding the Setup Number of a Partial Order", April 1984.

221. Joseph B. Rose and Gary N. Chaison, "The State of the Unions: United States and Canada", April 1984.

222. Naresh C. Agarwal, "Economic Costs of Employment Discrimination", May, 1984.

223. R.J. Adams, "Estimating the Extent of Collective Bargaining in Canada", May, 1984.

224. Robert G. Cooper, "New Product Strategies: What Distinguishes The Top Performers", May 1984.

225. Norman P. Archer, "Selection of a Simulation Language for an MBA Modeling and Simulation Course", June 1984.

226. Robert F. Love, William G. Truscott, John H. Walker, "A Facility Location Decision: Thibodeau-Finch Transport Ltd.", June 1984.

227. R.J. Adams, "Two Policy Approaches to Labour-Management Decision Making at the Level of the Enterprise: A comparison of the Wagner Model and Statutory Works Councils", July 1984.

228. George Saunders, "Employment and the Productivity Slowdown 1958 to 1980", September 1984.

229. George Saunders, "Industrial Relations in Canada", September 1984.

230. George Saunders, "Wages and Interest Arbitration: A Case Study of the Canadian Federal Public Service", September 1984.

Innis
REF
HB
74.5
.R47
no. 231

1235801