

NAG Library Chapter Introduction

S – Approximations of Special Functions

Contents

1	Scope of the Chapter	2
2	Background to the Problems	2
2.1	Functions of a Single Real Argument	2
2.2	Approximations to Elliptic Integrals	4
2.3	Bessel and Airy Functions of a Complex Argument	5
2.4	Option Pricing Routines	5
2.5	Hypergeometric Functions	6
3	Recommendations on Choice and Use of Available Routines	6
3.1	Vectorized Routine Variants	6
3.2	Elliptic Integrals	6
3.3	Bessel and Airy Functions	7
3.4	Confluent Hypergeometric Function ${}_1F_1$	7
4	Functionality Index	7
5	Auxiliary Routines Associated with Library Routine Parameters	11
6	Routines Withdrawn or Scheduled for Withdrawal	11
7	References	11

1 Scope of the Chapter

This chapter is concerned with the provision of some commonly occurring physical and mathematical functions.

2 Background to the Problems

The majority of the routines in this chapter approximate real-valued functions of a single real argument, and the techniques involved are described in Section 2.1. In addition the chapter contains routines for elliptic integrals (see Section 2.2), Bessel and Airy functions of a complex argument (see Section 2.3), complementary error function of a complex argument, hypergeometric functions and various option pricing routines for use in financial applications.

2.1 Functions of a Single Real Argument

Most of the routines provided for functions of a single real argument have been based on truncated Chebyshev expansions. This method of approximation was adopted as a compromise between the conflicting requirements of efficiency and ease of implementation on many different machine ranges. For details of the reasons behind this choice and the production and testing procedures followed in constructing this chapter see Schonfelder (1976).

Basically, if the function to be approximated is $f(x)$, then for $x \in [a, b]$ an approximation of the form

$$f(x) = g(x) \sum_{r=0}^n C_r T_r(t)$$

is used (\sum' denotes, according to the usual convention, a summation in which the first term is halved), where $g(x)$ is some suitable auxiliary function which extracts any singularities, asymptotes and, if possible, zeros of the function in the range in question and $t = t(x)$ is a mapping of the general range $[a, b]$ to the specific range $[-1, +1]$ required by the Chebyshev polynomials, $T_r(t)$. For a detailed description of the properties of the Chebyshev polynomials see Clenshaw (1962) and Fox and Parker (1968).

The essential property of these polynomials for the purposes of function approximation is that $T_n(t)$ oscillates between ± 1 and it takes its extreme values $n + 1$ times in the interval $[-1, +1]$. Therefore, provided the coefficients C_r decrease in magnitude sufficiently rapidly the error made by truncating the Chebyshev expansion after n terms is approximately given by

$$E(t) \simeq C_n T_n(t).$$

That is, the error oscillates between $\pm C_n$ and takes its extreme value $n + 1$ times in the interval in question. Now this is just the condition that the approximation be a minimax representation, one which minimizes the maximum error. By suitable choice of the interval, $[a, b]$, the auxiliary function, $g(x)$, and the mapping of the independent variable, $t(x)$, it is almost always possible to obtain a Chebyshev expansion with rapid convergence and hence truncations that provide near minimax polynomial approximations to the required function. The difference between the true minimax polynomial and the truncated Chebyshev expansion is seldom sufficiently great enough to be of significance.

The evaluation of the Chebyshev expansions follows one of two methods. The first and most efficient, and hence the most commonly used, works with the equivalent simple polynomial. The second method, which is used on the few occasions when the first method proves to be unstable, is based directly on the truncated Chebyshev series, and uses backward recursion to evaluate the sum. For the first method, a suitably truncated Chebyshev expansion (truncation is chosen so that the error is less than the *machine precision*) is converted to the equivalent simple polynomial. That is, we evaluate the set of coefficients b_r such that

$$y(t) = \sum_{r=0}^{n-1} b_r t^r = \sum_{r=0}^{n-1} C_r T_r(t).$$

The polynomial can then be evaluated by the efficient Horner's method of nested multiplications,

$$y(t) = (b_0 + t(b_1 + t(b_2 + \dots t(b_{n-2} + tb_{n-1})))) \dots).$$

This method of evaluation results in efficient routines but for some expansions there is considerable loss of accuracy due to cancellation effects. In these cases the second method is used. It is well known that if

$$\begin{aligned} b_{n-1} &= C_{n-1} \\ b_{n-2} &= 2tb_{n-1} + C_{n-2} \\ b_j &= 2tb_{j+1} - b_{j+2} + C_j, \quad j = n-3, n-4, \dots, 0 \end{aligned}$$

then

$$\sum_{r=0}^n C_r T_r(t) = \frac{1}{2}(b_0 - b_2)$$

and this is always stable. This method is most efficiently implemented by using three variables cyclically and explicitly constructing the recursion.

That is,

$$\begin{aligned} \alpha &= C_{n-1} \\ \beta &= 2t\alpha + C_{n-2} \\ \gamma &= 2t\beta - \alpha + C_{n-3} \\ \alpha &= 2t\gamma - \beta + C_{n-4} \\ \beta &= 2t\alpha - \gamma + C_{n-5} \\ &\vdots \\ \text{say } \alpha &= 2t\gamma - \beta + C_2 \\ \beta &= 2t\alpha - \gamma + C_1 \\ y(t) &= t\beta - \alpha + \frac{1}{2}C_0 \end{aligned}$$

The auxiliary functions used are normally functions compounded of simple polynomial (usually linear) factors extracting zeros, and the primary compiler-provided functions, sin, cos, ln, exp, sqrt, which extract singularities and/or asymptotes or in some cases basic oscillatory behaviour, leaving a smooth well-behaved function to be approximated by the Chebyshev expansion which can therefore be rapidly convergent.

The mappings of $[a, b]$ to $[-1, +1]$ used range from simple linear mappings to the case when b is infinite, and considerable improvement in convergence can be obtained by use of a bilinear form of mapping. Another common form of mapping is used when the function is even; that is, it involves only even powers in its expansion. In this case an approximation over the whole interval $[-a, a]$ can be provided using a mapping $t = 2(x/a)^2 - 1$. This embodies the evenness property but the expansion in t involves all powers and hence removes the necessity of working with an expansion with half its coefficients zero.

For many of the routines an analysis of the error in principle is given, namely, if E and ∇ are the absolute errors in function and argument and ϵ and δ are the corresponding relative errors, then

$$E \simeq |f'(x)|\nabla$$

$$E \simeq |xf'(x)|\delta$$

$$\epsilon \simeq \left| \frac{xf'(x)}{f(x)} \right| \delta.$$

If we ignore errors that arise in the argument of the function by propagation of data errors, etc., and consider only those errors that result from the fact that a real number is being represented in the computer in floating point form with finite precision, then δ is bounded and this bound is independent of the magnitude of x . For example, on an 11-digit machine

$$|\delta| \leq 10^{-11}.$$

(This of course implies that the absolute error $\nabla = x\delta$ is also bounded but the bound is now dependent on x .) However, because of this the last two relations above are probably of more interest. If possible the relative error propagation is discussed; that is, the behaviour of the error amplification factor $|xf'(x)/f(x)|$ is described, but in some cases, such as near zeros of the function which cannot be extracted explicitly, absolute error in the result is the quantity of significance and here the factor $|xf'(x)|$ is described. In general, testing of the functions has shown that their error behaviour follows fairly well these theoretical

error behaviours. In regions where the error amplification factors are less than or of the order of one, the errors are slightly larger than the above predictions. The errors are here limited largely by the finite precision of arithmetic in the machine, but ϵ is normally no more than a few times greater than the bound on δ . In regions where the amplification factors are large, of order ten or greater, the theoretical analysis gives a good measure of the accuracy obtainable.

It should be noted that the definitions and notations used for the functions in this chapter are all taken from Abramowitz and Stegun (1972). You are strongly recommended to consult this book for details before using the routines in this chapter.

2.2 Approximations to Elliptic Integrals

Four functions provided here are symmetrised variants of the classical (Legendre) elliptic integrals. These alternative definitions have been suggested by Carlson (1965), Carlson (1977b) and Carlson (1977a) and he also developed the basic algorithms used in this chapter.

The symmetrised elliptic integral of the first kind is represented by

$$R_F(x, y, z) = \frac{1}{2} \int_0^\infty \frac{dt}{\sqrt{(t+x)(t+y)(t+z)}}$$

where $x, y, z \geq 0$ and at most one may be equal to zero.

The normalization factor, $\frac{1}{2}$, is chosen so as to make

$$R_F(x, x, x) = 1/\sqrt{x}.$$

If any two of the variables are equal, R_F degenerates into the second function

$$R_C(x, y) = R_F(x, y, y) = \frac{1}{2} \int_0^\infty \frac{dt}{(t+y)\sqrt{t+x}},$$

where the argument restrictions are now $x \geq 0$ and $y \neq 0$.

This function is related to the logarithm or inverse hyperbolic functions if $0 < y < x$, and to the inverse circular functions if $0 \leq x \leq y$.

The symmetrised elliptic integral of the second kind is defined by

$$R_D(x, y, z) = \frac{3}{2} \int_0^\infty \frac{dt}{\sqrt{(t+x)(t+y)(t+z)^3}}$$

with $z > 0$, $x \geq 0$ and $y \geq 0$, but only one of x or y may be zero.

The function is a degenerate special case of the symmetrised elliptic integral of the third kind

$$R_J(x, y, z, \rho) = \frac{3}{2} \int_0^\infty \frac{dt}{\sqrt{(t+x)(t+y)(t+z)(t+\rho)}}$$

with $\rho \neq 0$ and $x, y, z \geq 0$ with at most one equality holding. Thus $R_D(x, y, z) = R_J(x, y, z, z)$. The normalization of both these functions is chosen so that

$$R_D(x, x, x) = R_J(x, x, x, x) = 1/(x\sqrt{x}).$$

The algorithms used for all these functions are based on duplication theorems. These allow a recursion system to be established which constructs a new set of arguments from the old using a combination of arithmetic and geometric means. The value of the function at the original arguments can then be simply related to the value at the new arguments. These recursive reductions are used until the arguments differ from the mean by an amount small enough for a Taylor series about the mean to give sufficient accuracy when retaining terms of order less than six. Each step of the recurrences reduces the difference from the mean by a factor of four, and as the truncation error is of order six, the truncation error goes like $(4096)^{-n}$, where n is the number of iterations.

The above forms can be related to the more traditional canonical forms (see Section 17.2 of Abramowitz and Stegun (1972)), as follows.

If we write $q = \cos^2 \phi$, $r = 1 - m \sin^2 \phi$, $s = 1 - n \sin^2 \phi$, where $0 \leq \phi \leq \frac{1}{2}\pi$, we have the classical elliptic integral of the first kind:

$$F(\phi | m) = \int_0^\phi (1 - m \sin^2 \theta)^{-\frac{1}{2}} d\theta = \sin \phi R_F(q, r, 1);$$

the classical elliptic integral of the second kind:

$$\begin{aligned} E(\phi | m) &= \int_0^\phi (1 - m \sin^2 \theta)^{\frac{1}{2}} d\theta \\ &= \sin \phi R_F(q, r, 1) - \frac{1}{3} m \sin^3 \phi R_D(q, r, 1) \end{aligned}$$

the classical elliptic integral of the third kind:

$$\begin{aligned} \Pi(n; \phi | m) &= \int_0^\phi (1 - n \sin^2 \theta)^{-1} (1 - m \sin^2 \theta)^{-\frac{1}{2}} d\theta \\ &= \sin \phi R_F(q, r, 1) + \frac{1}{3} n \sin^3 \phi R_J(q, r, 1, s). \end{aligned}$$

Also the classical complete elliptic integral of the first kind:

$$K(m) = \int_0^{\frac{\pi}{2}} (1 - m \sin^2 \theta)^{-\frac{1}{2}} d\theta = R_F(0, 1 - m, 1);$$

the classical complete elliptic integral of the second kind:

$$E(m) = \int_0^{\frac{\pi}{2}} (1 - m \sin^2 \theta)^{\frac{1}{2}} d\theta = R_F(0, 1 - m, 1) - \frac{1}{3} m R_D(0, 1 - m, 1).$$

For convenience, Chapter S contains routines to evaluate classical **and** symmetrised elliptic integrals.

2.3 Bessel and Airy Functions of a Complex Argument

The routines for Bessel and Airy functions of a real argument are based on Chebyshev expansions, as described in Section 2.1. The routines provided for functions of a complex argument, however, use different methods. These routines relate all functions to the modified Bessel functions $I_\nu(z)$ and $K_\nu(z)$ computed in the right-half complex plane, including their analytic continuations. I_ν and K_ν are computed by different methods according to the values of z and ν . The methods include power series, asymptotic expansions and Wronskian evaluations. The relations between functions are based on well known formulae (see Abramowitz and Stegun (1972)).

2.4 Option Pricing Routines

The option pricing routines evaluate the closed form solutions or approximations to the equations that define mathematical models for the prices of selected financial option contracts. These solutions can be viewed as special functions determined by the underlying equations. The terminology associated with these routines arises from their setting in financial markets and is briefly outlined below. See Joshi (2003) for a comprehensive introduction to this subject. An option is a contract which gives the holder the right, but not the obligation, to buy (if it is a call) or sell (if it is a put) a particular asset, S . A European option can be exercised only at the specified expiry time, T , while an American option can be exercised at any time up to T . For Asian options the average underlying price over a pre-set time period determines the payoff.

The asset is bought (if a call) or sold (if a put) at a pre-specified strike price X . Thus, an option contract has a payoff to the holder of $\max\{(S_T - X), 0\}$ for a call or $\max\{(X - S_T), 0\}$, for a put, which depends on whether the asset price at the time of exercise is above (call) or below (put) the strike, X . If at any moment in time a contract is currently showing a theoretical profit then it is deemed ‘in-the-money’; otherwise it is deemed ‘out-of-the-money’.

The option contract itself therefore has a value and, in many cases, can be traded in markets. Mathematical models, such as the Black–Scholes model, give theoretical prices for particular option contracts using a number of assumptions about the behaviour of financial markets. Typically, the price, S_t , of the underlying asset at time t , is modelled as the solution of a stochastic differential equation for the return, dS_t/S_t , on the asset price over a time interval, dt ,

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t,$$

where dW_t is a Brownian motion. The drift, μ , defines the trend in the movements of S , while the volatility, σ , measures the risk and may be taken to be the standard deviation of the returns on the asset price. In addition the model requires a riskless money market account or bond with value, B_t , at time t and risk-free rate, r , such that

$$dB_t = rB_t dt.$$

This leads to the determination of the Black–Scholes option price, P , by a martingale method or via the derivation of the Black–Scholes partial differential equation,

$$\frac{\partial P}{\partial t} + \frac{\partial P}{\partial S} rS + \frac{1}{2} \frac{\partial^2 P}{\partial S^2} \sigma^2 S^2 - rP = 0.$$

For this case a closed form solution exists which is evaluated by S30AAF.

A number of different option types where the solution exists in closed form or as a closed form approximation are presented in this chapter. See Haug (2007) for an extensive listing of option pricing formulae.

2.5 Hypergeometric Functions

The confluent hypergeometric function $M(a, b, x)$ (or ${}_1F_1(a; b; x)$) requires a number of techniques to approximate it over the whole parameter (a, b) space and for all argument (x) values. For x well within the unit circle $|x| \leq \rho < 1$ (where $\rho = 0.8$ say), and for relatively small parameter values, the function can be well approximated by Taylor expansions, continued functions or through the solution of the related ordinary differential equation by an explicit, adaptive integrator. For values of $|x| > \rho$, one of several transformations can be performed (depending on the value of x) to reformulate the problem in terms of a new argument x' such that $|x'| \leq \rho$. If one or more of the parameters is relatively large (e.g., $|a| > 30$) then recurrence relations can be used in combination to reformulate the problem in terms of parameter values of small size (e.g., $|a| < 1$).

Approximations to the hypergeometric functions can therefore require all of the above techniques in sequence: a transformation to get an argument well inside the unit circle, a combination of recurrence relations to reduce the parameter sizes, and the approximation of the resulting hypergeometric function by one of a set of approximation techniques.

All the techniques described above are based on those described in Pearson (2009).

3 Recommendations on Choice and Use of Available Routines

3.1 Vectorized Routine Variants

Many routines in Chapter S which compute functions of a single real argument have variants which operate on vectors of arguments. For example, S18AEF computes the value of the I_0 Bessel function for a single argument, and S18ASF computes the same function for multiple arguments. In general it should be more efficient to use vectorized routines where possible, though to some extent this will depend on the environment from which you call the routines. See Section 4 for a complete list of vectorized routines.

3.2 Elliptic Integrals

IMPORTANT ADVICE: users who encounter elliptic integrals in the course of their work are strongly recommended to look at transforming their analysis directly to one of the Carlson forms, rather than to the traditional canonical Legendre forms. In general, the extra symmetry of the Carlson forms is likely to simplify the analysis, and these symmetric forms are much more stable to calculate.

The routine S21BAF for R_C is largely included as an auxiliary to the other routines for elliptic integrals. This integral essentially calculates elementary functions, e.g.,

$$\ln x = (x - 1) R_C\left(\left(\frac{1+x}{2}\right)^2, x\right), \quad x > 0;$$

$$\arcsin x = x R_C(1 - x^2, 1), |x| \leq 1;$$

$$\operatorname{arcsinh} x = x R_C(1 + x^2, 1), \text{ etc.}$$

In general this method of calculating these elementary functions is not recommended as there are usually much more efficient specific routines available in the Library. However, S21BAF may be used, for example, to compute $\ln x/(x - 1)$ when x is close to 1, without the loss of significant figures that occurs when $\ln x$ and $x - 1$ are computed separately.

3.3 Bessel and Airy Functions

For computing the Bessel functions $J_\nu(x)$, $Y_\nu(x)$, $I_\nu(x)$ and $K_\nu(x)$ where x is real and $\nu = 0$ or 1, special routines are provided, which are much faster than the more general routines that allow a complex argument and arbitrary real $\nu \geq 0$. Similarly, special routines are provided for computing the Airy functions and their derivatives $\operatorname{Ai}(x)$, $\operatorname{Bi}(x)$, $\operatorname{Ai}'(x)$, $\operatorname{Bi}'(x)$ for a real argument which are much faster than the routines for complex arguments.

3.4 Confluent Hypergeometric Function

Two routines are provided for the confluent hypergeometric function ${}_1F_1$. Both return values for ${}_1F_1(a; b; x)$ where parameters a and b , and argument x , are all real, but one variant works in a scaled form designed to avoid unnecessary loss of precision. The unscaled routine S22BAF is easier to use and should be chosen in the first instance, changing to the scaled routine S22BBF only if problems are encountered.

4 Functionality Index

Airy function,

Ai, real argument,	
scalar	S17AGF
vectorized	S17AUF
Ai or Ai', complex argument, optionally scaled	S17DGF
Ai', real argument,	
scalar	S17AJF
vectorized	S17AWF
Bi, real argument,	
scalar	S17AHF
vectorized	S17AVF
Bi or Bi', complex argument, optionally scaled.....	S17DHF
Bi', real argument,	
scalar	S17AKF
vectorized	S17AXF

Arccos,

inverse circular cosine.....	S09ABF
------------------------------	--------

Arccosh,

inverse hyperbolic cosine.....	S11ACF
--------------------------------	--------

Arcsin,

inverse circular sine.....	S09AAF
----------------------------	--------

Arcsinh,

inverse hyperbolic sine.....	S11ABF
------------------------------	--------

Arctanh, inverse hyperbolic tangent	S11AAF
Bessel function, I_0 , real argument, scalar	S18AEF
vectorized	S18ASF
I_1 , real argument, scalar	S18AFF
vectorized	S18ATF
I_ν , complex argument, optionally scaled	S18DEF
J_0 , real argument, scalar	S17AEF
vectorized	S17ASF
J_1 , real argument, scalar	S17AFF
vectorized	S17ATF
$J_{\alpha \pm n}(z)$, complex argument.....	S18GKF
J_ν , complex argument, optionally scaled.....	S17DEF
K_0 , real argument, scalar	S18ACF
vectorized	S18AQF
K_1 , real argument, scalar	S18ADF
vectorized	S18ARF
K_ν , complex argument, optionally scaled	S18DCF
Y_0 , real argument, scalar	S17ACF
vectorized	S17AQF
Y_1 , real argument, scalar	S17ADF
vectorized	S17ARF
Y_ν , complex argument, optionally scaled.....	S17DCF
beta function, incomplete	S14CCF
Complement of the Cumulative Normal distribution	S15ACF
Complement of the Error function, complex argument, scaled.....	S15DDF
real argument	S15ADF
real argument, scaled.....	S15AGF
Cosine, hyperbolic.....	S10ACF
Cosine Integral.....	S13ACF
Cumulative Normal distribution function.....	S15ABF
Dawson's Integral.....	S15AFF
Digamma function, scaled.....	S14ADF
Elliptic functions, Jacobian, sn, cn, dn, complex argument	S21CBF
real argument	S21CAF

Elliptic integral, general, of 2nd kind, $F(z, k', a, b)$	S21DAF
Legendre form, complete of 1st kind, $K(m)$	S21BHF
complete of 2nd kind, $E(m)$	S21BJF
of 1st kind, $F(\phi m)$	S21BEF
of 2nd kind, $E(\phi m)$	S21BFF
of 3rd kind, $\Pi(n; \phi m)$	S21BGF
symmetrised, degenerate of 1st kind, R_C	S21BAF
of 1st kind, R_F	S21BBF
of 2nd kind, R_D	S21BCF
of 3rd kind, R_J	S21BDF
Erf, real argument	S15AEF
Erfc, complex argument, scaled	S15DDF
real argument	S15ADF
erfcx, real argument	S15AGF
Exponential, complex	S01EAF
Exponential Integral	S13AAF
Fresnel integral, <i>C</i> , scalar	S20ADF
vectorized	S20ARF
<i>S</i> , scalar	S20ACF
vectorized	S20AQF
Gamma function	S14AAF
Gamma function, incomplete	S14BAF
Generalized factorial function	S14AAF
Hankel function $H_\nu^{(1)}$ or $H_\nu^{(2)}$, complex argument, optionally scaled	S17DLF
Hypergeometric functions, ${}_1F_1(a; b; x)$, confluent, real argument	S22BAF
${}_1F_1(a; b; x)$, confluent, real argument, scaled form	S22BBF
Jacobian theta functions $\theta_k(x, q)$, real argument	S21CCF
Kelvin function, bei x , scalar	S19ABF

vectorized	S19APF
ber x ,	
scalar	S19AAF
vectorized	S19ANF
kei x ,	
scalar	S19ADF
vectorized	S19ARF
ker x ,	
scalar	S19ACF
vectorized	S19AQF
Legendre functions of 1st kind $F_n^m(x)$, $\overline{F}_n^m(x)$	S22AAF
Logarithm of $1 + x$	S01BAF
Logarithm of beta function,	
real	S14CBF
Logarithm of gamma function,	
complex	S14AGF
real	S14ABF
real, scaled	S14AHF
Option Pricing,	
American option, Bjerksund and Stensland option price	S30QCF
Asian option, geometric continuous average rate price	S30SAF
Asian option, geometric continuous average rate price with Greeks	S30SBF
binary asset-or-nothing option price	S30CCF
binary asset-or-nothing option price with Greeks	S30CDF
binary cash-or-nothing option price	S30CAF
binary cash-or-nothing option price with Greeks	S30CBF
Black–Scholes–Merton option price	S30AAF
Black–Scholes–Merton option price with Greeks	S30ABF
European option, option prices, using Merton jump-diffusion model	S30JAF
European option, option price with Greeks, using Merton jump-diffusion model	S30JBF
floating-strike lookback option price	S30BAF
floating-strike lookback option price with Greeks	S30BBF
Heston’s model option price	S30NAF
Heston’s model option price with Greeks	S30NBF
standard barrier option price	S30FAF
Polygamma function,	
$\psi^{(n)}(x)$, real x	S14AEF
$\psi^{(n)}(z)$, complex z	S14AFF
Psi function	S14ACF
Psi function derivatives, scaled	S14ADF
Scaled modified Bessel function(s),	
$e^{- x }I_0(x)$, real argument,	
scalar	S18CEF
vectorized	S18CSF
$e^{- x }I_1(x)$, real argument,	
scalar	S18CFF
vectorized	S18CTF
$e^x K_0(x)$, real argument,	
scalar	S18CCF

vectorized.....	S18CQF
$e^x K_1(x)$, real argument,	
scalar.....	S18CDF
vectorized.....	S18CRF
Sine,	
hyperbolic.....	S10ABF
Sine Integral.....	S13ADF
Tangent,	
circular.....	S07AAF
hyperbolic.....	S10AAF
Trigamma function, scaled.....	S14ADF
Zeros of Bessel functions $J_\alpha(x)$, $J'_\alpha(x)$, $Y_\alpha(x)$, $Y'_\alpha(x)$,	
scalar.....	S17ALF

5 Auxiliary Routines Associated with Library Routine Parameters

None.

6 Routines Withdrawn or Scheduled for Withdrawal

None.

7 References

- Abramowitz M and Stegun I A (1972) *Handbook of Mathematical Functions* (3rd Edition) Dover Publications
- Carlson B C (1965) On computing elliptic integrals and functions *J. Math. Phys.* **44** 36–51
- Carlson B C (1977a) *Special Functions of Applied Mathematics* Academic Press
- Carlson B C (1977b) Elliptic integrals of the first kind *SIAM J. Math. Anal.* **8** 231–242
- Clenshaw C W (1962) Chebyshev Series for Mathematical Functions *Mathematical tables* HMSO
- Fox L and Parker I B (1968) *Chebyshev Polynomials in Numerical Analysis* Oxford University Press
- Haug E G (2007) *The Complete Guide to Option Pricing Formulas* (2nd Edition) McGraw-Hill
- Joshi M S (2003) *The Concepts and Practice of Mathematical Finance* Cambridge University Press
- Pearson J (2009) Computation of hypergeometric functions *MSc Dissertation, Mathematical Institute, University of Oxford*
- Schonfelder J L (1976) The production of special function routines for a multi-machine library *Softw. Pract. Exper.* **6(1)**
-