

Convolution Rather Than Monte Carlo Simulation to Price a Barrier Option

Daniel T. Cassidy

Department of Engineering Physics, McMaster University, Hamilton, ON, Canada

Email: cassidy@mcmaster.ca

How to cite this paper: Cassidy, D.T. (2025) Convolution Rather Than Monte Carlo Simulation to Price a Barrier Option. *Journal of Mathematical Finance*, 15, 550-578.
<https://doi.org/10.4236/jmf.2025.153023>

Received: June 7, 2025

Accepted: August 15, 2025

Published: August 18, 2025

Copyright © 2025 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Repeated convolution and truncation of a truncated fat-tailed distribution, instead of Monte Carlo simulation, for pricing a discrete, simple barrier option is presented. The parameters for the truncated fat-tailed distribution are obtained by fitting the sum of a Student's t distribution plus a normal distribution to the one-day returns obtained from the adjusted closing values for the S&P 500. It is argued that truncation of the fat-tailed, one-day returns distribution is a physically reasonable action and evidence to support this truncation is provided. The steps to price a discrete up-and-out barrier on an European call option through repeated convolution and truncation are given.

Keywords

Option Pricing, Convolution, Truncation, Student's t Distribution, Fat Tails, S&P 500 Returns

1. Introduction

In this paper, the use of repeated convolution and truncation of a truncated one-day probability density function (pdf) to price a discrete, simple barrier option is discussed.

There are multiple methods to price barrier options, and Kabby [1] categorised methods to value barrier options into groups and provided references for early papers in the groups. The groups of Kabby are partial differential equation methods [2], binomial and trinomial tree methods [3] [4], Monte Carlo simulations methods [3]-[7], transform methods [3] [8] [9], and backward stochastic differential equations. The citations in the listing of groups are to some recent publications that fall into the groups.

One potential shortcoming of many of the works is the use of Gaussian (normal) statistics given the fact that it is well known that market prices are not nor-

mally distributed over short time frames [10]–[12]. Lee *et al.* [8] and Shevchenko and Moral [6] explicitly state that their approaches can be used with jump-diffusion processes or other underlying statistical processes.

Termens [4] used a basic Monte Carlo approach of simulating many paths and averaging to obtain the expected value to price the option, and used a path-integral approach from quantum mechanics to value barrier options, both using geometric Brownian motion.

Shevchenko and Moral [6] used a sequential Monte Carlo approach, wherein paths that hit the barrier are restarted with appropriate weighting, to improve the efficiency of the simulation and mentioned methods, such as Brownian bridge simulation, importance sampling, and control variates, to minimise the impacts of finite sampling errors.

Nouri and Abbasi [5] developed a Monte Carlo method of using a uniform random variable and an exit probability to improve efficiency and to mitigate sampling error.

Li and Yan [7] investigated the use of Monte Carlo simulation and machine learning to predict the value of an option.

Blanda [3] evaluated the use of Monte Carlo simulation and techniques such as importance sampling and antithetic variables to increase simulation efficiency, binomial and trinomial trees, and finite difference method solutions to a partial differential equation to price foreign exchange barrier options.

Interestingly, there is no mention of numerical convolution to replace Monte Carlo techniques. Convolution is an approach that considers all paths to a given value. There are no statistical sampling errors with convolution, numerical convolution works with finite length sequences from any (reasonably well-behaved) function, and truncation of insignificant values in the tails can reduce considerably the computation time for convolution.

Basnarkov *et al.* [13] used convolution to price a European option. These authors used the n^{th} power of the characteristic function for a Student's t probability density function [14] with $\nu = 3$ degrees of freedom to obtain the n -day characteristic function for the n -day return. A fast Fourier transform (FFT) algorithm [15] was used to compute numerically the pdf for the n -day return. From this sequence of numbers representing the n -day pdf, the authors were able, after truncation of the sequence, to compute prices for European call options.

In this manuscript, numerical convolution is suggested to obtain the development in time of the distributions that are required to price an up-and-out European call option. The approach taken here is somewhat different than the approach of Basnarkov *et al.* [13].

First, Basnarkov *et al.* [13] considered a European call option. A European call option depends only on the beginning and end points; the path is not of interest. This allowed Basnarkov *et al.* [13] to compute the n -day pdf in one step as the n^{th} power of a characteristic function. A barrier option is path dependent, so it is not possible to compute an end point in one step, although the characteristic function

approach could easily provide the required steps.

Second, Basnarkov *et al.* [13] used the convolution theorem to work in the Fourier domain [16] and used an inverse FFT algorithm to compute the required pdf. Basnarkov *et al.* [13] chose to truncate the n -day return and used a storage array of length $2^{18} = 262144$. In this work, a brute-force calculation of the convolution along with truncation of the one-day return pdf and of n -fold convolutions (separate from the truncations required to implement the barrier) was used with three arrays of lengths $< 2^{12} = 4096$. Truncations of the n -fold convolutions reduce the sizes of the storage arrays and the time required to compute the convolutions.

Third, in this work, it is argued that truncation of the one-day return distribution is more than just a convenience in computing the convolution but is required to match the time development of the n -day pdf by n -fold convolution of one-day return distributions. In contrast, Basnarkov *et al.* [13] argued that options prices were only weakly dependent on the support and thus truncation is not a major impediment to find a fair price and that the point of truncation can be chosen based on convenience, as information on tail events is sparse.

1.1. Convolution

The convolution of two functions $f(x)$ and $g(x)$ is given by

$$h(x) = (f * g)(x) = f(x) * g(x) = \int_{-\infty}^{\infty} f(\xi) g(x - \xi) d\xi \quad (1)$$

where $f(x) * g(x)$ is a common engineering/science notation for convolution.

Convolution has some interesting properties [16]. Convolution is commutative, associative, and distributive over addition. The area under a convolution $h = f * g$, defined as $\int h(\xi) d\xi$, equals the area under $f(x)$ times the area under $g(x)$. If the variances for $f(x)$ and $g(x)$ exist, then the variances add under convolution. Similarly, if the means for $f(x)$ and $g(x)$ exist, then the means add under convolution. The Fourier transform of a convolution $f(x) * g(x)$ is the product of the Fourier transforms of the two functions $f(x)$ and $g(x)$.

If $f_1(x)$ is the probability density function (pdf) for a random variable (r.v.) x_1 and $f_2(x)$ is the pdf for a r.v. x_2 that is independent of x_1 , then the pdf for $x_T = x_1 + x_2$ is $f_T(x) = f_1(x) * f_2(x)$ [16] [17].

If, e.g., daily returns are independent, then the pdf for a two-day return is given by $f_T = f_1 * f_2$ where f_1 is the pdf for the one-day return for the first of the two days and f_2 is the pdf for the one-day return of the second day. If the returns are assumed to be stationary (*i.e.*, $f = f_1 = f_2$, the one-day return is assumed to be same for each day), then the analysis simplifies and $f_T = f * f$ and the n -day return is the n -fold self convolution of the pdf for the one-day returns.

Clearly, the time frame can be any interval provided that the time intervals are non-overlapping and that the returns over these time intervals are independent.

Given the underlying pdf, the distribution of returns for n days (or n time in-

tervals) can be estimated by convolution without the need for Monte Carlo simulations. It is the replacement of Monte Carlo simulation with convolution that is investigated here. The convolution approach provides a uniformly-sampled estimate of the n -day pdf. The steps to price an up-and-out European call option are outlined.

The intent is to calculate the convolution numerically. Returns over “short” time frames are fat tailed and there is likely no analytic expression for the convolution integral. The convolution of two Gaussian distributions is a Gaussian distribution, so the case of an underlying Gaussian pdf is trivial, but Gaussian pdfs do not fit short time-frame returns, such as one-day returns, well [10]–[12] [18]–[23]. Fits of Student’s t and Gaussian distributions to one-day returns obtained from adjusted closing values for the S&P 500 are given and are used to obtain best-fit parameters for the one-day pdfs. One could use the normalised histogram of the returns (*i.e.*, the raw data) as the pdf to convolve, but these histograms have random fluctuations (noise) so for now the smoother best-fit pdfs are used.

The one-day pdf for returns must be truncated to represent reality (returns of $\pm\infty$ are not physically possible) and truncation of n -day pdfs can be employed for fast calculation of the convolutions. Truncation of the one-day pdf is discussed, and evidence to support the truncation is given. n -fold convolutions of truncated, Student’s t distributions rapidly approach normal distributions as n increases from unity [24]. The transformation of the distribution of n -day returns from fat-tailed distributions for small n to Gaussian (or normal) distributions for larger n , $n \gg 16$, is a known feature or “stylized” fact of returns [12].

1.2. Outline

In this paper, pricing of a simple barrier option using repeated convolution and truncation of a truncated fat-tailed distribution is considered. Parameters for the fat-tailed distribution are obtained by non-linear, minimum mean square error (mmse, or least squares [25] [26]) fits of Student’s t and normal distributions to frequency of occurrence data for adjusted closing values of the S&P 500. **Appendix A** provides information on fits to the data. The quality of fit for a fit of a sum of a Student’s t distribution plus a normal distribution (a $t+$ distribution for short) to the one-day returns for the S&P 500 is remarkable. This sum fits better than a fit of just a Student’s t or a normal distribution. For n -day returns with $2 \leq n \leq 8$, it was found that only a Student’s t distribution fit the data well; adding a Gaussian component did not improve the fit.

Section 2 provides background information and definitions that are needed throughout. This section also presents an argument for truncation of the fat-tailed distribution for the one-day returns. **Appendix A** provides additional evidence for the need for truncation of the fat-tailed $t+$ distribution, based on the match of n -fold convolutions of the best-fit one-day distribution with n -day returns calculated from historical data.

Section 3 presents pricing based on repeated convolution.

Appendix B gives information on run times for draws of random numbers, as would be used in Monte Carlo simulations, and on time to perform convolutions. Unlike Basnarkov *et al.* [13], discrete Fourier transforms (DFT's or the FFT implementation) were not used to calculate the convolutions and thereby possibly reduce the time taken to compute the convolution. With sufficient padding of the sequences, the circular convolution inherent to a DFT will give the desired linear convolution [15] [16]. However, neglect of small terms in the brute-force method to calculate the convolution reduces the run time and memory requirements for the brute-force method and the DFT approach might not be advantageous.

A conclusion is given in Section 4.

2. Background and Definitions

Let the probability density function (pdf) for a random variable (r.v.) x be $f_x(x)$. The probability that a measurement of the r.v. x takes a value between x and $x+dx$, $P\{x < x \leq x+dx\}$, is given by $f_x(x)dx$. The cumulative density function (CDF) $F_x(a) = P\{x \leq a\} = \int_{-\infty}^a f_x(\xi)d\xi$ [17].

Let S_T be the adjusted close value of an asset for day (or time) T . Let the one-day return for day (or time) T be $R_{T,1} = \ln(S_T/S_{T-1})$ and the n -day return for day (or time) T be $R_{T,n} = \ln(S_T/S_{T-n})$. Note that

$$\begin{aligned} R_{T,n} &= \ln(S_T/S_{T-n}) = \ln(S_T/S_{T-1} \times S_{T-1}/S_{T-2} \times \cdots \times S_{T-n+1}/S_{T-n}) \\ &= R_{T,1} + R_{T-1,1} + R_{T-2,1} + \cdots + R_{T-n+1,1} \end{aligned} \quad (2)$$

and, for simplicity in notation, let $R_T = R_{T,1}$.

To avoid uninteresting leading zeros, per mille returns are used (with symbol ‰), wherein the return is multiplied by 1000 to obtain a per mille return.

When returns over non-lapping time periods are independent, the pdf for an n -day return is the n -fold convolution of the pdf of the one-day return since the n -day return is, as shown in Equation (2), the sum of n one-day returns [16] [17]. It is known that, for the most part, daily returns are independent [12].

Variances add under convolution [16], and hence the variance of n -day returns is n times the variance that describes the distribution of the one-day returns. Since the normal distribution is stable under self-convolution, if the one-day return distribution is a normal distribution with a mean μ and a variance σ^2 , then the pdf for an n -day return is a normal distribution with a mean $= n\mu$ and with a variance $= n\sigma^2$.

A Student's t distribution is not necessarily stable under self-convolution. The convolution of two Student's t distributions retains the fat tails of the original distributions, but the distribution is not necessarily the same form as the original distributions, which makes pricing with Student's t distributions difficult [13] [23] [27]–[29]. However, a Student's t distribution describes well the returns of many stocks and indices, and thus Student's t distributions cannot be ignored [18] [20]–[23].

A Student's t pdf, $f_t(x; \nu, \beta, \mu)$, with ν degrees of freedom, scale factor β ,

and mean μ , is evaluated as

$$f_t(x; \nu, \beta, \mu) dx = \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\Gamma\left(\frac{\nu}{2}\right) \sqrt{\pi \nu \beta^2}} \left(1 + \frac{(x-\mu)^2}{\nu \beta^2}\right)^{-\frac{\nu+1}{2}} dx \quad (3)$$

whereas a Gaussian (or normal) pdf, $f_G(x; \sigma, \mu)$, with scale factor σ and mean μ , is evaluated as

$$f_G(x; \sigma, \mu) dx = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) dx. \quad (4)$$

Over the range $-\infty \leq x \leq \infty$ (*i.e.*, support $[-\infty, +\infty]$) the variance of the normal distribution $f_G(x; \sigma, \mu)$ is σ^2 ; over the range $-\infty \leq x \leq \infty$, the variance of the Student's t distribution $f_t(x; \nu, \beta, \mu)$ is $\nu\beta^2/(\nu-2)$, provided that $\nu > 2$. A Cauchy or Lorentz distribution is $f_t(x; \nu=1, \beta, \mu)$, and is stable under self-convolution. In the limit as $\nu \rightarrow \infty$, $f_t(x; \nu, \beta, \mu) \rightarrow f_G(x; \sigma, \mu)$, which for many purposes holds for $\nu > 30$ [30].

2.1. Distribution of Closing Values S Given a Distribution of Returns

If one assumes that $S_{(T-n)} = S_0$, then the pdf for $R_{T,n}$ can be transformed to a pdf for S_T [31] [32]. Under the assumption that S_0 is a constant (*i.e.*, a sharp value), $R_{T,n} = \ln(S_T) - \ln(S_0)$ and $S_T = S_0 \exp(R_{T,n})$ ([31], pp 187, 281).

By substitution of variables in Equation (3), the pdf for S_T , $f_{S_T}(S_T)$, is found to be

$$f_{S_T}(S_T) dS_T = \Lambda(\nu, \beta) \left(1 + \frac{(\ln(S_T/S_0) - \mu)^2}{\nu \beta^2}\right)^{-(\nu+1)/2} \frac{dS_T}{S_T}, \quad (5)$$

and is a log Student's t distribution with normalisation $\Lambda(\nu, \beta)$ given by

$$\Lambda(\nu, \beta) = \frac{\Gamma((\nu+1)/2)}{\Gamma(\nu/2) \sqrt{\pi \nu \beta^2}} \times \frac{1}{p_b - p_a}. \quad (6)$$

In the normalisation for the Student's t distribution, support over the interval $(x_a, x_b]$ has been assumed, *i.e.*, the pdf is assumed to be truncated and is only non-zero in the range $(x_a, x_b]$, where x_a solves $p_a = F_{S_T}(x_a)$, x_b solves $p_b = F_{S_T}(x_b)$, with $p_b > p_a$.

The probability that a value of the asset lies in the interval $(S_l, S_u]$ is thus given by

$$F_{S_T}(S_u) - F_{S_T}(S_l) = \int_{S_l}^{S_u} f_{S_T}(\xi) d\xi = \int_{\ln(S_l/S_0)}^{\ln(S_u/S_0)} \Lambda(\nu, \beta) \left(1 + \frac{(\xi - \mu)^2}{\nu \beta^2}\right)^{-(\nu+1)/2} d\xi. \quad (7)$$

A similar approach can be used to find the distribution of closing values given

normally distributed returns. The distributions for closing values of an asset are required to price barrier and European call options, as presented in Section 3.

2.2. Truncation

Truncation of the support of a pdf from $[-\infty, +\infty]$ to $[a, b]$ where $-\infty < a < b < \infty$ and the pdf equals zero outside of the interval $[a, b]$ can impart some useful properties to the pdf. Truncation is generally not needed for a Gaussian (normal) pdf $f_G(x; \sigma, \mu)$, as $f_G(x; \sigma, \mu)$ exhibits a sharp roll off with increasing x .

The situation is different for Student's t distributions $f_t(x; \nu, \beta, \mu)$. If $f_t(x; \nu, \beta, \mu)$ is truncated, then the variance is finite (*i.e.*, exists) for all ν and one can think of, e.g., Student's t random walks [33] and pricing options based on returns that are distributed as $f_t(x; \nu, \beta, \mu)$ [13] [23] [34].

Effective truncation can be obtained by multiplication by a normal pdf with a large variance [35] [36]. This approach is not pursued here, although it is interesting to note that an effectively truncated Student's t distribution arises from truncation of the distribution of the variance in a mixing integral that yields a Student's t distribution [37].

Truncation of $f_t(x; \nu, \beta, \mu)$ leads to results that are physically realisable. For example, the amount of capital is limited, thus an infinite positive return makes no practical sense and the distribution that describes the returns should be truncated or capped [23]. Similarly, the price of an asset is not likely to be negative. Thus returns should be greater than -1000% and $\lll +\infty$.

Probably the most apparent and useful property imparted by truncation of a Student's t distribution is as a means to explain the development of the pdf for one-day returns to a normal distribution for n -day returns, where $n \gg 16$ [24]. Returns over short time periods of less than several days are known to be described well by fat-tailed distributions such as Student's t distributions [18] [20]–[24]. Returns for longer time-intervals are known to be described by normal distributions [12]. For one-day returns described by a Student's t distribution, truncation and n -fold convolution of the one-day t distribution of returns to create n -day returns quickly becomes Gaussian-like (*i.e.*, normally distributed), as was demonstrated [24]. Note that without truncation, n -fold convolutions of a Student's t retain the fat tails of the underlying t -distribution, as mentioned earlier, and one is left trying to explain the observed transition from a fat-tailed distribution from short time frames to a normal distribution for longer time frames.

The development of an n -day return from n -fold convolution of a truncated, 1-day Student's t distribution is salient and worth repeating here.

Figure 1 shows 128-fold convolutions of the best fits of a t -distribution, of a normal distribution, and of a t distribution to the one-day returns of the S&P 500 closing values; see **Appendix A** for an overview of the data and fits. For comparison, the one-day best fit values are also plotted. The figure shows that the 128-fold convolutions of the fat-tailed distributions are Gaussian-like at the peak, in the

vicinity of “a”; that repeated convolution of a Student’s t distribution retains the fat tails (see curves in the vicinity of “b”); and, that truncation causes the n -fold convolution to appear Gaussian like (see curves in the vicinity of “c”) in the wings. For generation of this figure, the support for the one-day pdf was taken as $8\times$ the width of the observed support for the S&P 500 one-day returns and the n -fold convolutions were truncated at 10^{-9} .

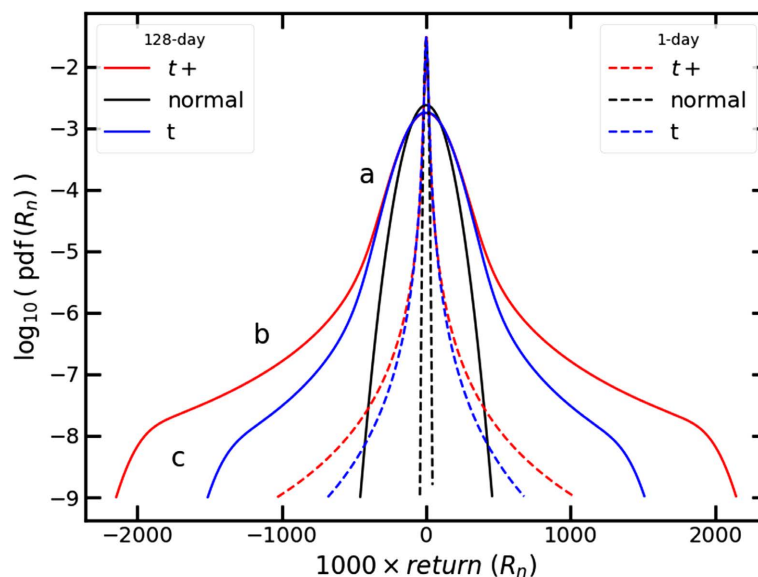


Figure 1. Plots of pdfs that represent one-day and 128-day returns. The 128-day pdfs were obtained by repeated convolution of the one-day pdf. Three features of interest are identified by “a”, “b”, and “c”.

Appendix A describes the data and data handling, shows fits to the returns, and the development of the distribution of the n -day return by n -fold convolution. The development of the n -day return distribution from the one-day return distribution and correspondence with observed n -day return distribution provides evidence for the physical necessity for truncation of the one-day, fat-tailed distribution of returns. Truncation of the one-day pdf imparts predictive power to the n -day distribution obtained by n -fold convolution.

3. Pricing

In this section, the use of convolution to price a European call option and an up-and-out European call option is addressed. The pricing of a European call option is straight forward with the convolution approach, as convolution provides an estimate, whether based on historical performance or intuition, of the pdf in the future when the call option is to be settled. In addition, convolution provides the pdf as a sequence of numbers at equally spaced returns, which facilitates numerical evaluation of the integrals required to price an option.

The price of a European call option is required for pricing an up-and-out European call option, since the European option is the default option if the price

does not cross the barrier. The pricing of a European call option is considered next, in Section 3.1.

3.1. European Option

Following the Gosset papers [23] [34], which were motivated by Chapter 10 in Ross [38] and belief in the physical reality of truncation, the price of a European call option at time T is $C_T = E\{(S_T - K_T)^+\}$ subject to the constraint that $E\{S_T\} = \exp(rT)S_0$. In the expression for C_T , $E\{(S_T - K_T)^+\}$ is the expectation of the maximum value of 0 or the difference $(S_T - K_T)$ of the value of the asset S_T and the strike price K_T at time T . At time $T=0$, when the option is purchased, the value of the option is $C_0 = \exp(-rT)C_T$ as the sale of the option is a cash transaction and therefore the value of the option at time T , C_T , is discounted by the time value of money, with r the risk-free rate.

The pdf $f(x)$ required to evaluate the expectations $C_T = E\{(S_T - K_T)^+\}$ and $E\{S_T\} = \exp(rT)S_0$ is known, from n -fold convolution of the truncated, one-day return pdf, as a sequence of finite length of uniformly spaced numbers. The expectations can be evaluated numerically over the finite sequence, which allows C_T to be evaluated.

Let $S_T = A_T \exp(\mathbf{R}_{T,n})$, and use the constraint on the expected value of S_T to find an expression for A_T :

$$E\{S_T\} = E\{A_T \exp(\mathbf{R}_{T,n})\} = \frac{\int_{x_a}^{x_b} A_T \exp(\xi) f(\xi) d\xi}{\int_{x_a}^{x_b} f(\xi) d\xi} = S_0 \exp(rT) \quad (8)$$

with x_a and x_b the points of truncation of the pdf $f(x)$.

If the pdf $f(x)$ is a Gaussian function, $f_G(x; \sigma\sqrt{T}, \alpha T)$, then support over the interval $[-\infty, \infty]$ can be taken with no harm incurred, the expectation for C_T can be expressed in terms of error functions, and the standard Black-Scholes formula for an European call option is obtained if one is careful with the drift rates [39].

One needs to set $\alpha = r$ in $f_G(x; \sigma\sqrt{T}, \alpha T)$ to obtain the standard Black-Scholes formula for a European call option. Forcing $\alpha = r$ changes the pdf and hence changes the value of $C_T = E\{(S_T - K_T)^+\}$. This risk neutral approach “works” for the most part because the magnitude of the risk premium $|(\alpha - r)T|$ is typically $\ll \sigma\sqrt{T}$. The drift owing to the risk premium is dwarfed by the random fluctuations (*i.e.*, the signal is buried in the noise) and hence the risk neutral approach makes only a small difference [39].

For a pdf equal to a Gaussian function, $f_G(x; \sigma\sqrt{T}, \alpha T)$, $E\{S_T\} = A_T \exp(\sigma_T^2/2 + \alpha T)$. $E\{S_T\}$ contains the infamous “noise rectification” term $\exp(\sigma_T^2/2)$ plus drift at the rate of α . The integrals in Equation (8)

can be understood to give similar information for a non-Gaussian pdf.

3.2. Barrier Option

In this section, the value of a barrier option, an up-and-out European call option, is evaluated from the distribution for the closing price of the asset, Equation (5), which is derived from the distribution of the one-day returns, Equation (2). The values for other barrier options are obtained by analogy. Restriction of the discussion to an up-and-out European call option allows the language of this section to be specific.

Let p_i be the probability that the closing price S on the i^{th} day (or interval) is $\leq B$ where B is the value of the barrier. Let $q_i = 1 - p_i$ be the probability that the closing price of the asset is $> B$. Thus

$$p_i = \int_{-\infty}^B f_{S_i}(S) dS \quad \text{and} \quad q_i = \int_B^{\infty} f_{S_i}(S) dS \quad (9)$$

where $f_{S_i}(S)$ is the pdf for the 1-day closing price S on the i^{th} day and $S_i = A_i \exp(R_{i,1})$ is the one-day closing price on the i^{th} day.

Let

$$\hat{f}_1(S) = f_{S_1}(S) \times (1 - H_e(S - B)) \quad (10)$$

be a truncated replica of $f_{S_1}(S)$ where $\hat{f}_1(S) = f_{S_1}(S)$ for $S \leq B$ and $\hat{f}_1(S) = 0$ for $S > B$. In the definition of $\hat{f}_1(S)$, $H_e(S - B)$ is the Heaviside step function [16].

The probability that an option has not crossed a given price, called the barrier, is a useful number.

It is a property of convolution that the area of the convolution of two functions equals the product of the areas of the two functions ([16], p. 118). This property is of use in calculating probabilities.

Assume q is an absorbing state. That is, once the closing value of the asset is $> B$, there is no further interest in following this branch of the development of the closing price of the asset. At closing on the first day, the probability that the option has not crossed the barrier is p_1 . At closing of the second day, the probability \mathcal{P}_2 , with a script \mathcal{P} , that the closing value of the asset has not crossed the barrier on either the first or second days is

$$\begin{aligned} \mathcal{P}_2 &= \int_{-\infty}^B (f_{S_1}(S)(1 - H_e(S - B))) * f_{S_2}(S) dS \\ &= \int_{-\infty}^B \hat{f}_1(S) * f_{S_2}(S) dS \\ &= \int_{-\infty}^{+\infty} (\hat{f}_1(S) * f_{S_2}(S)) \times (1 - H_e(S - B)) dS \\ &= p_1 - p_1 q_2 = p_1(1 - q_2) = p_1 p_2. \end{aligned} \quad (11)$$

$\hat{f}_1(S) * f_{S_2}(S)$ is the convolution of the pdf's $\hat{f}_1(S)$ and $f_{S_2}(S)$, and gives

the pdf for the closing price of the second day for which the closing price of the first day did not exceed the barrier B . The truncation of the pdf of the closing price for the previous day selects the “paths” for the time development of the closing price of the asset that do not cross the barrier. The approach presented here thus calculates the value of a path-dependent option without the need to simulate paths using a Monte Carlo approach.

Note in Equation (11) that $\mathcal{P}_2 = p_1 p_2$ as expected since p_1 and p_2 are independent one-day probabilities that the price remained $\leq B$. The fact that

$$\begin{aligned} p_1 q_2 &= \int_{-\infty}^{+\infty} \left(\hat{f}_1(S) * f_{S_2}(S) \right) \times (H_e(S-B)) dS \\ &= \int_B^{+\infty} \left(\hat{f}_1(S) * f_{S_2}(S) \right) dS \\ &= \int_B^{+\infty} \int_{-\infty}^B \left(\hat{f}_1(u) f_{S_2}(S-u) \right) du dS \end{aligned} \quad (12)$$

in Equations (11)-(12) is because $\hat{f}_1(S) = 0, S > B$, the integral over the convolution is for $S > B$, and the support for $f_{S_2}(S)$ includes the support for $\hat{f}_1(S)$ since the width of a convolution is greater than the widths of the functions.

One method to obtain the result is to use the serial product method described by Bracewell ([16], Ch. 3). Break the functions into small rectangles and follow one rectangle of $f_{S_2}(S')$ as it is swept across $\hat{f}_1(u)$ by the convolution. The rectangle of $f_{S_2}(S')$ is essentially a constant and can be taken outside of the integration over the dummy variable of integration u . The integral over $\hat{f}_1(u)$ gives a contribution to p_1 and the integral over S gives a contribution to q_2 .

Let

$$f_{S_2}(S-u) = \sum_j \bar{f}_2(S'_j) \quad (13)$$

where $\bar{f}_2(S'_j)$ is composed of contiguous rectangles of small width such that $\bar{f}_2(S'_j)$ is essentially constant with respect to S and u over the width of the rectangle.

In the integral definition of the convolution, such as Equation (1), x is a parameter chosen by the user; it is the point for which the user wishes to know the convolution. In the last line of Equation (12), the parameter for the convolution is S . Chose S such that $f_{S_2}(S-u)$ is constant over the integral $\int \hat{f}_1(u) du$. If this choice is made, then

$$\begin{aligned} p_1 q_2 &= \int_B^{+\infty} \left(\hat{f}_1(S) * f_{S_2}(S) \right) dS \\ &= \sum_j \int_B^{+\infty} \int_{-\infty}^B \left(\hat{f}_1(u) \bar{f}_2(S'_j) \right) du dS \\ &= \sum_j p_1 q_{2j} = p_1 q_2. \end{aligned} \quad (14)$$

$\hat{f}_2(S)$ is defined in a similar manner to $\hat{f}_1(S)$, with $\hat{f}_2(S)$ the truncation

of a convolution

$$\hat{f}_2(S) = (\hat{f}_1(S) * f_{S_2}(S)) \times (1 - H_e(S - B)), \quad (15)$$

which leads to general expressions for \hat{f}_i , p_i , and $\mathcal{P}_i = p_1 p_2 \cdots p_i$, where i , $i > 1$, is the i^{th} day of an n -day option, and thus to the ability to price a barrier option. The required pdf's are built by sequentially convolving and truncating:

$$\hat{f}_i(S) = (\hat{f}_{i-1}(S) * f_{S_i}(S)) \times (1 - H_e(S - B)) \quad (16)$$

$$\mathcal{P}_i = \int_{-\infty}^B \hat{f}_{i-1}(S) * f_{S_i}(S) dS = \int_{-\infty}^B \hat{f}_i(S) dS. \quad (17)$$

If it is assumed that the one-day returns are stationary in that the distribution of returns is the same for each day that the option is alive, then the notation can be simplified as $S = S_i$ for all i . In this case $\mathcal{P}_i = p^i, i = 1, \dots, n$ where

$p = \int_{-\infty}^B f_S(S) dS$ is the probability of crossing the barrier on any single day. \mathcal{P}_i gives the probability that the option has not crossed the barrier at the beginning of the $(i+1)^{\text{th}}$ interval.

There is no need for Monte Carlo simulation of S to find the price of a barrier option; there is no need to analyse an ensemble of calculated paths to find probabilities. n -fold truncation and convolution of the one-day distribution gives the n -day return pdf as a sequence of uniformly sampled numbers. This pdf as a sequence of numbers can be integrated to find the probabilities and expectations required to price options. The pdf also does not contain random fluctuations and non-uniform coverage owing to draws from a random number generator.

Drift is handled in a straightforward manner with the method of repeated convolution and truncation of the pdfs for the adjusted closing value. $f_{i+1}(S)$ contains the memory of all drift from the previous i days. Note that α_i , the rate of drift for each time interval, need not be same for each day. Similarly, different pdf's could be used for each interval of time, which is a day in this work.

Figure 2 is composed of semi- \log_{10} plots of the pdfs for the closing value of the asset S based on best fits of a normal distribution (black) and the best-fit $t+$ distribution (*i.e.*, the sum of a Student's t distribution plus a Gaussian distribution) (red) to the one-day returns for the S&P 500. The convolution of a one-day barrier-truncated $t+$ distribution (shown in red to the left of the green line) with the one-day closing value (the full red curve) is shown in blue. The barrier truncation was for values $> B = \$53.00$. The truncation of the blue curve at the barrier level is shown in green. The area under the blue curve to the left of the green line gives the probability that the closing value after two days will be $\leq B$.

The convolutions were calculated numerically. In the calculation of the convolution, any final results that were $< 10^{-7}$ were set to zero. This reduces the amount of storage and the time that are required to calculate the convolutions while keeping the uncertainty $\approx 10^{-5}$ for calculation of the closing value after 100 days. The steep roll-off in the wings of the blue curve (which is visible on the right-hand

side) shows that there is negligible area in the truncation for values of S that occur with the pdf having a value $< 10^{-7}$. Nevertheless, the pdf after the i^{th} convolution and before truncation at the barrier value was renormalised to the value of \mathcal{P}_{i-1} to avoid accumulation of the loss of probability through truncation of final results that were $< 10^{-7}$.

The pdfs for the one-day (red) and two-day (blue) closing values of the asset are truncated (as discussed in the preceding paragraph), although not as harshly as the barrier truncation shown in **Figure 2**. The truncations for the red and blue curves are shown only on the left-hand side as vertical lines. The other truncations are not shown to avoid visual distraction. Nevertheless, repeated convolution of the truncated one-day pdf leads rapidly to a normal-like distribution [24]. This truncation of the pdf is physical and necessary to price options with fat tails [23] [34]. This is argued in Section 2.2 and also shown in **Appendix A**.

The best-fit $t+$ distribution to the one-day returns was truncated to the interval $[S_0 - 14.13, S_0 + 7.05]$ such that the minimum return of -229% and the maximum return of 109.6% were included. The interval $[S_0 - 14.13, S_0 + 7.05]$ corresponds to per mille returns of $[-332.1, 13.19]$. Option prices for assets with returns that follow a Student's t distribution or other fat-tailed distribution can be obtained if the distributions are truncated or the return is capped [23] [28] [29] [34]. Truncation of the pdf was chosen.




Figure 2. Plots of probability density functions (pdfs) for values of S as a function of distance from $S_0 = 50$. Best-fit pdfs to one-day returns, $f_{s_1}(S)$, for normal (black) and for $t+$ (red) distributions to the S&P 500. The curves in grey (normal distribution) and blue ($t+$ distribution) display the predicted pdfs for two-day returns, $f_{s_2}(S)$, based on convolutions $\hat{f}_1(S) * f_{s_2}(S)$. Areas to the left of the green line and under the pdfs give the probability for the value of S to remain below the barrier of 53. $\hat{f}_1(S)$ is the red curve to the left of the vertical, green line. $\hat{f}_2(S)$ is the blue curve to the left of the vertical, green line.

Table 1 lists the probabilities $\mathcal{P}_T = p_1 p_2 \dots p_T$ that the closing prices of the asset are all below a barrier B for T closing days for three values of the barrier. The initial price of the asset was assumed to be $S_0 = \$50.00$ with an average drift of $\alpha = 0.5 \times 10^{-3}$ per day (*i.e.*, 0.5‰ per day). Also listed are the values for an up-and-out European call option with strike price $K_T = \$53.00$. C_T was calculated as $C_T = E\{(S_T - K_T)^+\}$. The T -fold convolution and truncation was used to determine numerically the pdf for the T -day return, and this pdf was used to calculate the expectation of the maximum value of $S_T - K_T$ and zero. Calculations assuming a t + distribution and a normal distribution are presented in the table. Each column for \mathcal{P}_T and C_T has two entries. The value on the left of a column is the value obtained using a T -fold convolution and truncation of a best-fit t + distribution for the one-day returns. The value on the right of a column is the value obtained using T -fold convolution and truncation of a best-fit normal distribution for the one-day returns. The data and fits are discussed in **Appendix A**.

The calculated probabilities of unity for $B = \infty$ demonstrate the accuracy of the calculations and approach of neglecting the results of the convolutions that are $<10^{-7}$. The value of a 1-day call is zero to a good approximation as there is a small probability of $<7 \times 10^{-4}$ that the closing value of the asset will increase from $\$50.00$ to $>\$53.00$ in one day, as shown by the red curve in **Figure 2**. The $B = \infty$ results show the difference in prices of a European call option based on repeated self-convolution of a t + (red curve, **Figure 2**) or based on repeated convolution of a normal distribution (black curve, **Figure 2**). The fat tails of the log Student's t distribution are responsible for the larger price of the option as compared to the predictions based on a normal pdf.

It is somewhat surprising that the probabilities \mathcal{P}_T calculated for the two distributions are not more different given the large differences between the two distributions, which can be observed in **Figure 2**.

The options prices for $B = K_T$ are zero since the up-and-out nature of the barrier sets the pdf zero for $S > K_T$ and thus $C_T = E\{(S_T - K_T)^+\} = 0$ for $B \leq K_T$.

Table 1. Table of probabilities \mathcal{P}_T and price C_T for $K_T = \$52.00$, $S_0 = \$50.00$, and $\alpha = 0.5\%$.

T days	barrier = ∞				barrier = $S_0 + \$3.00$				barrier = $K_T = S_0 + \$2.00$			
	\mathcal{P}_T		C_T		\mathcal{P}_T		C_T		\mathcal{P}_T		C_T	
30	1.000	1.000	0.003	0.000	0.999	1.000	0.001	0.000	0.997	1.000	0.000	0.000
	1.000	1.000	0.552	0.338	0.722	0.802	0.030	0.036	0.539	0.600	0.000	0.000
	1.000	1.000	1.276	0.919	0.495	0.561	0.015	0.022	0.344	0.378	0.000	0.000
	1.000	1.000	1.989	1.538	0.371	0.413	0.009	0.014	0.251	0.267	0.000	0.000

For comparison, **Table 2** has $\alpha = 0$. In general, for $\alpha = 0$ as compared to $\alpha = 0.5 \times 10^{-3}$, the probability of survival until time T , \mathcal{P}_T , is increased (as the

barrier is further from the value of the asset) and the price of the call option is decreased (as there is a smaller probability that $S_T > K_T$ since S is not drifting towards K_T).

Table 2. Table of probabilities \mathcal{P}_T and price C_T for $K_T = \$52.00$, $S_0 = \$50.00$, and $\alpha = 0.0$.

T days	barrier = ∞				barrier = $S_0 + \$3.00$				barrier = $K_T = S_0 + \$2.00$			
	\mathcal{P}_T		C_T		\mathcal{P}_T		C_T		\mathcal{P}_T		C_T	
30	1.000	1.000	0.003	0.000	0.999	1.000	0.001	0.000	0.997	1.000	0.000	0.000
	1.000	1.000	0.353	0.177	0.802	0.884	0.022	0.024	0.641	0.721	0.000	0.000
	1.000	1.000	0.720	0.417	0.637	0.734	0.012	0.016	0.481	0.558	0.000	0.000
	1.000	1.000	1.028	0.630	0.541	0.637	0.008	0.011	0.400	0.470	0.000	0.000

The time required to calculate the pdf every \$0.01 for $T = 100$ days, which is 99 convolutions, was ≈ 1.8 s for the best fit t^+ distribution and required 3 arrays with <4096 8-byte storage elements. The time required to calculate the pdf every 0.01 for $T = 100$ days for the best-fit normal distribution was ≈ 0.2 s. The normal distribution is fast to work with as the severe roll-off in the tails means fewer points to keep in the convolutions. Note that the convolution approach gives the pdf sampled on a regular basis. Integrals required in the work were obtained by summing over all data points. The sampling pitch was set arbitrarily at \$0.01.

4. Conclusions

The pricing of a discrete, simple barrier option by repeated numerical convolution and truncation of a truncated one-day return distribution is presented. n -fold convolution of a truncated one-day return distribution gives the distribution of prices at the end of day (or time interval) $n+1$ as a sequence of uniformly-spaced numbers. Functions of these numbers can be integrated numerically to price the discrete barrier option.

Pricing of a discrete, up-and-out European option based on fits to the one-day returns of the S&P 500 is used as an example to demonstrate the process. It is reported that a superposition of a Student's t distribution and a Gaussian (normal) distribution fit the S&P 500 one-day returns remarkably well. Data handling, fits, and results of n -fold convolutions are described in **Appendix A**.

Truncation of the one-day return distribution and of the convolutions reduces significantly the storage and time requirements. The truncation of the distributions is in addition to the truncation of the distribution that is required by the up-and-out nature of the barrier option.

It is argued that truncation of the one-day return distribution is not a convenience but is required to impart realistic properties to the n -fold convolution. The Student's t component of the best-fit distribution to the one-day return data is fat-tailed with degrees of freedom $\nu = 2.52 \pm 0.2$ (see **Table A1**). It is known that in the longer term, say $n \gg 16$ for the S&P data used here, the n -day return is

Gaussian-like. Convolution of a truncated version of the best-fit distribution to the one-day returns to have a similar region of support as the S&P 500 returns data takes the fat tails of the one-day return data to a normal-like distribution for the n -day return as n increases. Data, in the form of n -fold convolutions of the truncated one-day returns, are given to support the argument.

Convolution of probability density functions gives the probability of all paths possible. Thus, convolution can replace Monte Carlo simulations of the possible paths. With truncation of the convolutions for small values (values $< 10^{-7}$) of the pdf's, calculation of the convolution is fast and storage efficient. In addition, the results of the convolution do not suffer from sampling noise and are sequences of uniformly spaced numbers ready for numerical quadrature.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

- [1] Kabby, W. (2022) The Valuations of Barrier Options: A Method Review. *Revue Maghrébine Management des Organisations*, **7**, 1-14.
- [2] Wang, H., Zhang, J. and Zhou, K. (2022) On Pricing of Vulnerable Barrier Options and Vulnerable Double Barrier Options. *Finance Research Letters*, **44**, Article 102100. <https://doi.org/10.1016/j.frl.2021.102100>
- [3] Blanda, V. (2023) FX Barrier Option Pricing. MA Thesis, Imperial College.
- [4] Guarch Termens, M.E. (2022) Pricing Methods for Barrier Options. MA Thesis, Universitat de Barcelona.
- [5] Nouri, K. and Abbasi, B. (2017) Implementation of the Modified Monte Carlo Simulation for Evaluate the Barrier Option Prices. *Journal of Taibah University for Science*, **11**, 233-240. <https://doi.org/10.1016/j.jtusci.2015.02.010>
- [6] Shevchenko, P.V. and Del Moral, P. (2015) Valuation of Barrier Options Using Sequential Monte Carlo. *Journal of Computational Finance*. <https://doi.org/10.2139/ssrn.2529539>
- [7] Li, Y. and Yan, K. (2023) Prediction of Barrier Option Price Based on Antithetic Monte Carlo and Machine Learning Methods. *Cloud Computing and Data Science*, **4**, 77-86. <https://doi.org/10.37256/ccds.4120232110>
- [8] Lee, H., Lee, G. and Song, S. (2022) Multi-Step Reflection Principle and Barrier Options. *Journal of Futures Markets*, **42**, 692-721. <https://doi.org/10.1002/fut.22306>
- [9] Guardasoni, C., Rodrigo, M.R. and Sanfelici, S. (2018) A Mellin Transform Approach to Barrier Option Pricing. *IMA Journal of Management Mathematics*, **31**, 49-67. <https://doi.org/10.1093/imaman/dpy016>
- [10] Mandelbrot, B. (1963) The Variation of Certain Speculative Prices. *The Journal of Business*, **36**, 394-419. <https://doi.org/10.1086/294632>
- [11] Fama, E.F. (1965) The Behavior of Stock-Market Prices. *The Journal of Business*, **38**, 34-105. <https://doi.org/10.1086/294743>
- [12] Cont, R. (2001) Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, **1**, 223-236. <https://doi.org/10.1080/713665670>
- [13] Basnarkov, L., Stojkoski, V., Utkovski, Z. and Kocarev, L. (2019) Option Pricing with

- Heavy-Tailed Distributions of Logarithmic Returns. *International Journal of Theoretical and Applied Finance*, **22**, Article 1950041. <https://doi.org/10.1142/s0219024919500419>
- [14] Dreier, I. and Kotz, S. (2002) A Note on the Characteristic Function of the t -Distribution. *Statistics & Probability Letters*, **57**, 221-224. [https://doi.org/10.1016/s0167-7152\(02\)00032-9](https://doi.org/10.1016/s0167-7152(02)00032-9)
 - [15] Oppenheim, V. and Schaffer, R.V. (1975) Digital Signal Processing. Prentice Hall, Inc.
 - [16] Bracewell, R.N. (2000) The Fourier Transform and Its Applications. 3rd Edition, McGraw-Hill Companies.
 - [17] Papoulis (1965) Probability, Random Variables, and Stochastic Processes. McGraw Hill, Inc.
 - [18] Blattberg, R.C. and Gonedes, N.J. (1974) A Comparison of the Stable and Student Distributions as Statistical Models for Stock Prices. *The Journal of Business*, **47**, 244-280. <https://doi.org/10.1086/295634>
 - [19] Amaral, L.A.N., Plerou, V., Gopikrishnan, P., Meyer, M. and Stanley, H.E. (2000) The Distribution of Returns of Stock Prices. *International Journal of Theoretical and Applied Finance*, **3**, 365-369. <https://doi.org/10.1142/s0219024900000218>
 - [20] de Jong, C. and Huisman, R. (2000) From Skews to a Skewed-T: Modelling Option-Implied Returns by a Skewed Student-T. *Proceedings of the IEEE/IAFE/INFORMS 2000 Conference on Computational Intelligence for Financial Engineering (CIFER)*, New York, 28-28 March 2000, 132-142. <https://doi.org/10.1109/cifer.2000.844611>
 - [21] Platen, E. and Rendek, R. (2008) Empirical Evidence on Student- t Log-Returns of Diversified World Stock Indices. *Journal of Statistical Theory and Practice*, **2**, 233-251. <https://doi.org/10.1080/15598608.2008.10411873>
 - [22] Galbraith, J.W. and Zhu, D. (2009) A Generalized Asymmetric Student- t Distribution with Application to Financial Econometrics. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.1499914>
 - [23] Cassidy, D.T., Hamp, M.J. and Ouyed, R. (2010) Pricing European Options with a Log Student's t -Distribution: A Gosset Formula. *Physica A: Statistical Mechanics and its Applications*, **389**, 5736-5748. <https://doi.org/10.1016/j.physa.2010.08.037>
 - [24] Cassidy, D.T. (2011) Describing n -Day Returns with Student's t -Distributions. *Physica A: Statistical Mechanics and Its Applications*, **390**, 2794-2802. <https://doi.org/10.1016/j.physa.2011.03.019>
 - [25] Bevington, P.R. and Robinson, D.K. (1992) Data Reduction and Error Analysis for the Physical Sciences. 2nd Edition, McGraw Hill, Inc.
 - [26] Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P. (1992) Numerical Recipes in Fortran: The Art of Scientific Computing. 2nd Edition, Cambridge University Press.
 - [27] Bouchaud, J.-P. and Sornette, D. (1994) The Black-Scholes Option Pricing Problem in Mathematical Finance: Generalization and Extensions for a Large Class of Stochastic Processes. *Journal de Physique I*, **4**, 863-881. <https://doi.org/10.1051/jp1:1994233>
 - [28] McCauley, J.L., Gunaratne, G.H. and Bassler, K.E. (2007) Martingale Option Pricing. *Physica A: Statistical Mechanics and its Applications*, **380**, 351-356. <https://doi.org/10.1016/j.physa.2007.02.038>
 - [29] Bouchaud, J. and Potters, M. (2003) Theory of Financial Risk and Derivative Pricing. 2nd Edition, Cambridge University Press. <https://doi.org/10.1017/cbo9780511753893>

- [30] Evans, M., Hastings, N. and Peacock, B. (1993) *Statistical Distributions*. 2nd Edition, John Wiley and Sons.
- [31] Lax, M., Cai, W. and Xu, M. (2006) *Random Processes in Physics and Finance*. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198567769.001.0001>
- [32] Coffey, W.T., Kalmykov, Y.P. and Waldron, J.T. (2004) *The Langevin Equation: With Applications to Stochastic Problems in Physics, Chemistry and Electrical Engineering*. 2nd Edition, World Scientific Publishing Co. Pte. Ltd. <https://doi.org/10.1142/9789812795090>
- [33] Cassidy, D.T. (2016) Student's t Increments. *Open Journal of Statistics*, **6**, 156-171. <https://doi.org/10.4236/ojs.2016.61014>
- [34] Cassidy, D.T., Hamp, M.J. and Ouyed, R. (2013) Log Student's t -Distribution-Based Option Sensitivities: Greeks for the Gosset Formulae. *Quantitative Finance*, **13**, 1289-1302. <https://doi.org/10.1080/14697688.2012.744087>
- [35] Moriconi, L. (2007) Delta Hedged Option Valuation with Underlying Non-Gaussian Returns. *Physica A: Statistical Mechanics and its Applications*, **380**, 343-350. <https://doi.org/10.1016/j.physa.2007.01.018>
- [36] Lim, G.C., Martin, G.M. and Martin, V.L. (2006) Pricing Currency Options in the Presence of Time-Varying Volatility and Non-Normalities. *Journal of Multinational Financial Management*, **16**, 291-314. <https://doi.org/10.1016/j.mulfin.2005.08.004>
- [37] Cassidy, D.T. (2012) Effective Truncation of a Student's t -Distribution by Truncation of the Chi Distribution in a Chi-Normal Mixture. *Open Journal of Statistics*, **2**, 519-525. <https://doi.org/10.4236/ojs.2012.25067>
- [38] Ross, S.M. (2007) *Introduction to Probability Models*. 9th Edition, Elsevier.
- [39] Cassidy, D.T. (2018) Risk-Neutral Pricing of European Call Options: A Specious Concept. *Journal of Mathematical Finance*, **8**, 335-348. <https://doi.org/10.4236/jmf.2018.82022>
- [40] Greenwood, P.E. and Nikulin, M.S. (1996) *A Guide to Chi-Squared Testing*. John Wiley and Sons.
- [41] Ross, S. (2006) *A First Course in Probability*. 7th Edition, Pearson Prentice Hall.

Appendices

A. S&P 500 Data to 2 February 2025

Adjusted daily close prices were obtained from the web for the S&P 500 index for the dates of 1950-01-03 to 2025-02-02. Natural logarithm returns were calculated from the data for $2^{(n-1)}$, $n = 1, \dots, 8$ -day returns, using Equation (2). The data were binned, using an adaptive binning procedure, to ensure at least five returns in each bin [40]. A non-linear minimum mean square error (or least squares) fitting routine, based on the Levenberg-Marquardt method [25] [26], was implemented in Maple to fit Student's t distributions, Gaussian distributions, and a sum of these two distributions to the histogram of returns. The derivatives required to construct the matrices in the fitting routine were calculated symbolically.

The best-fit parameters for the distributions, $a_1 \cdots a_{N_{fit}}$, were obtained by minimising the value of reduced chi-square, χ_{df}^2 , through an iterative Newton root-finding procedure.

A.1. Define χ_{df}^2

Reduced chi-square is defined here as

$$\chi_{df}^2 = \sum_{i=1}^N \frac{\left(y_i - N f_{fit} \left(x_i : a_1 \cdots a_{N_{fit}} \right) dx_i \right)^2}{df \sigma_i^2} \quad (\text{A-1})$$

where y_i is the number of data points (*i.e.*, returns) in the i th bin; N is the total number of data points distributed across the bins; $dx_i = x_{i+1} - x_i$ is the width of the i th bin; df is the number of degrees of freedom, which equals $N - N_{fit}$ with N_{fit} the number of parameters derived from the data;

$$f_{fit} \left(x_i : a_1 \cdots a_{N_{fit}} \right) = \frac{f \left(x_i : a_1 \cdots a_{N_{fit}} \right) + 4 f \left(\frac{x_i + x_{i+1}}{2} : a_1 \cdots a_{N_{fit}} \right) + f \left(x_{i+1} : a_1 \cdots a_{N_{fit}} \right)}{6}$$

where $f \left(x : a_1 \cdots a_{N_{fit}} \right)$ is a pdf that approximates the distribution of the returns amongst the bins; x_i is the independent variable; and, $a_1 \cdots a_{N_{fit}}$ are parameters, such as scale and location parameters, that define the pdf.

f_{fit} uses Simpson's 1 4 1 rule to approximate an integral of the pdf $f \left(x : a_1 \cdots a_{N_{fit}} \right)$ over the width of a bin.

Poisson or counting statistics are assumed for the number of counts in each bin. This means that $\sigma_i = \sqrt{y_i}$ [25] [30] [41].

A measure of the uncertainty in each calculation of χ_{df}^2 was required to decide if the changes were meaningful or not. The uncertainty in reduced chi-square for the fit, χ_{df}^2 , was taken as the uncertainty as the 95% confidence interval for a zero-mean chi square variable with df degrees of freedom. The contribution to χ_{df}^2 from the quality of the fit was ignored in assessing if the

changes in calculations of chi-square were likely to be statistically significant or not.

A.2. Fits to One-Day Returns

Figure A1 plots, on a semi \log_{10} scale, the number of returns per bin, the best fit Student's t distribution, $f_t(x: \nu = 3.16 \pm 0.1, \beta = (117.6 \pm 7)\%, \mu = (0.548 \pm 0.075)\%$), (blue line), and the best fit Gaussian distribution (black line) for one-day returns for the S&P 500. The solid dots are plotted at the mid point of each bin and show the number of points in that bin. The grey line outlines the bins. The fit of a Student's t distribution to the one-day returns is remarkable good, and significantly better than the best-fit Gaussian distribution.

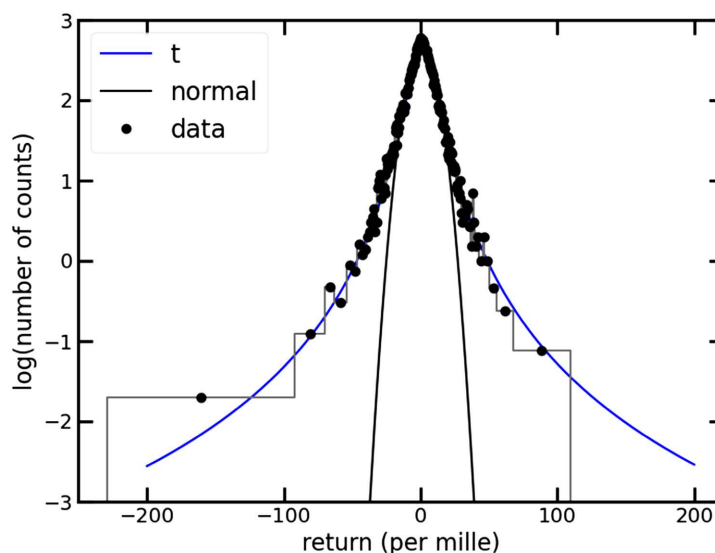


Figure A1. Best fit Student's t distribution to the one-day returns. The residues squared for the fit are shown in **Figure A2**. The bins for the counts are shown in grey with the number of counts in each bin displayed with a solid circle in the middle of each bin.

Figure A2 plots the residues squared for the fit of the Student's t distribution that is shown in **Figure A1**, where the residue is defined as

$$\left(y_i - N f_{fit} \left(x_i : a_1 \cdots a_{N_{fit}} \right) dx_i \right) / \sigma_i.$$
 The residue squared, summed over all data points, and divided by df equals χ^2_{df} , Equation (A-3). This plot of the residue squared has little intrinsic value on its own. The value of this plot is in comparison to results obtained for fits to a sum of Student's t plus normal distributions (*i.e.*, to a $t+$ distribution), and to a normal distribution.

Figure A3 plots, using a base ten logarithmic scale, the best fit pdf $t+$ distribution (red line) composed of a sum of a Student's t distribution plus a Gaussian distribution, and for comparison, a best fit Gaussian distribution (black line). As before, the solid dots show the number of returns in the bin and at the midpoint of the bin. The grey line outlines the bins. The sum of a Student's t distribution plus a Gaussian distribution fits the data better than just a Student's t distribution.

For this fit of a superposition of pdfs, $\nu = 2.52 \pm 0.2$ for the Student's- t component, which is less than the value found for fitting just a Student's t distribution. Adding a Gaussian component to the fit has improved the quality of the fit and given the best-fit pdf fatter tails.

The improvement in the fit can be observed from a plot of the residue squared for this fit.




Figure A2. Residue squared for a best fit of a Student's t distribution to 18890 natural logarithm one-day returns based on the adjusted close values for the S&P 500 from 1950-01-03 to 2025-01-31, inclusive. The returns were separated into 163 bins, with the bin widths in the tails adjusted for a minimum of five returns in a bin. The residues, summed over all 163 values and divided by the number of degrees of freedom give reduced chi-square $\chi^2_{df} = 1.43$. The parameters of the best-fit distribution are those parameters that minimise the sum of the residues squared over all 163 bins. The bins are outlined with a thin, grey line.




Figure A3. Best fit of the one-day returns to a Student's t distribution plus a normal distribution (*i.e.*, to a $t+$ distribution). The residues squared for the fit are shown in **Figure A4**.

Figure A4 plots the residue squared for the fit shown in **Figure A3**. The residues for a fit of a superposition of a Student's t plus a Gaussian distribution are smaller than for a fit of only a Student's t distribution.




Figure A4. Residue squared for a best fit of a Student's t distribution plus a normal distribution to the same histogram of returns as for **Figure A2**. The linear superposition of the pdf's yields $\chi^2_v = 1.02$, which is significantly smaller than for a fit of just a Student's t distribution or just a normal distribution. **Figure A2** and **Figure A4** are plotted on the same scale; the residue of **Figure A2** extends to ≈ 14 .

For comparison, a plot of the residue squared for fit to only a normal distribution is provided as **Figure A5**.

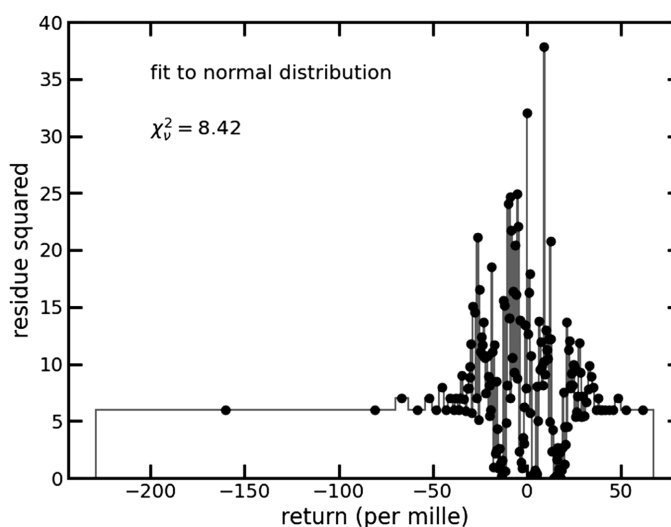


Figure A5. Residues for a best fit of a normal distribution to the same 1-day returns as used for **Figure A2** and **Figure A4**. The fit to a normal gives reduced chi-square of $\chi^2_v = 8.43$ and is not a good fit to the data. Note that the ordinate scale is $\approx 3\times$ greater for this figure than for **Figure A2** and **Figure A4**.

Figure A5 shows the residue squared for fits to a normal distribution only. A normal distribution does not fit the data well, particularly in the wings. Note that the ordinate scale has been changed relative to the other two residue squared plots, **Figure A2** and **Figure A4**.

Table A1 displays fit parameters for a fit of the superposition $f(x) = af_t(x; \nu, \beta, \mu) + (1-a)f_G(x; \sigma, \mu)$ to the n -day returns, where $f_t(x; \nu, \beta, \mu)$ is the pdf for a Student's t distribution and $f_G(x; \sigma, \mu)$ is the pdf for a Gaussian (or normal) distribution. Uncertainties are shown for two one-day fits. These uncertainties were calculated as $2\times$ the value of the corresponding diagonal element of the error matrix and should be satisfactory estimates of the 95% confidence intervals [16] [26].

It is interesting to note that addition of a Gaussian to a Student's t fitting function improved the fit by reducing χ^2_{df} from 1.4 to 1.02 and reduced ν from 3.2 to 2.5, making the tails fatter for a fit of the sum of the two functions.

For a one-day return, $1-a \approx 0.25$, and for two-day return, the expected contribution from $f_G(x)$ alone is $(1-a)^2 = 1/16$, since

$$f(x) * f(x) = a^2 f_t(x) * f_t(x) + 2a(1-a) f_t(x) * f_G(x) + (1-a)^2 f_G(x) * f_G(x). \quad (\text{A-2})$$

For an n -day return, the contribution from an n -fold convolution of $f_G(x)$ is $(1-a)^n$. **Table A1** shows $a \rightarrow 1$ as n increases, which is consistent with the expected $(1-a)^n$ behaviour and indicates that it is likely worthwhile to fit to only $f(x) = f_t(x; \nu, \beta, \mu)$ for $n > 1$. This conclusion is borne out by comparison of the reduced chi-squared values, χ^2_{df} , in **Table A1**.

Table A1. Results for fits of n -day returns to $f(x) = af_t(x) + (1-a)f_G(x)$ where $f_t(x)$ is a Student's t pdf and $f_G(x)$ is a Gaussian pdf.

n -day	ν	β	μ	a	σ_G	μ_G	χ^2_{df}
1-day	2.52 ± 0.2	5.04 ± 0.5	0.77 ± 0.1	0.753 ± 0.05	10.44 ± 0.5	-0.69 ± 0.5	1.020 ± 0.02
2-day	4.0324	9.6004	1.1029	1.0430	10.2148	1.8547	1.700
4-day	4.7674	14.2612	2.0422	1.0051	12.9264	-0.8312	1.918
8-day	5.1405	20.4503	3.7034	0.9547	14.8197	7.1134	2.607
1-day	3.16 ± 0.1	6.10 ± 0.2	0.548 ± 0.075	1			1.433 ± 0.02
2-day	4.0723	9.6494	1.0745	1			1.684
4-day	4.8116	14.2800	2.0408	1			1.869
8-day	5.2033	20.2040	3.8361	1			2.625
1-day				0	7.451 ± 0.08	0.5423	8.42 ± 0.02
2-day				0	11.27	0.9463	5.82
4-day				0	16.36	1.865	4.82
8-day				0	23.03	3.463	5.02

Only the one-day χ_{df}^2 shows a statistically significant difference between the fits for $a=1$ and $a \approx 0.75$. To draw this conclusion, write χ_{df}^2 as

$$\chi_{df}^2 = \sum_{i=1}^N \frac{(\Delta y_i^2 - 2\Delta y_i \delta_i + \delta_i^2)}{df \sigma_i^2} \quad (\text{A-3})$$

where: Δy_i is the difference, for the i th data point, between the true value and the best fit value obtained (in a thought experiment) by excluding the zero-mean random fluctuations δ_i ; N is the number of data points; σ_i^2 is the variance of the zero mean random fluctuations; and, the number of degrees of freedom, df , equals $N - N_{fit}$, where N_{fit} is the number of parameters determined from the fit.

To put Equation (A-3) into a form to estimate the uncertainty in χ_{df}^2 , make some simplifying assumptions such as uncorrelated terms and the contribution to χ_{df}^2 from the quality of the fit, *i.e.*, the contribution to χ_{df}^2 from $E\{\Delta y_i^2\}$, is independent of the random variable (r.v) δ_i . Take the expectation of χ_{df}^2 , where $E\{\cdot\}$ is the expectation operator, and simplify

$$\begin{aligned} E\{\chi_{df}^2\} &= \sum_{i=1}^N \frac{(E\{\Delta y_i^2\} - 2E\{\Delta y_i \delta_i\} + E\{\delta_i^2\})}{df \sigma_i^2} \\ &\approx \sum_{i=1}^N \frac{(E\{\Delta y_i^2\} - 0 + E\{\delta_i^2\})}{df \sigma_i^2} \\ &\approx (\chi_{df}^2(\Delta y_i) + 1) + E\{(\chi_{df}^2(\delta_i) - 1)\} \end{aligned} \quad (\text{A-4})$$

to see that, on average, χ_{df}^2 is comprised of a contribution from the quality of the fit, $\chi_{df}^2(\Delta y_i) + 1$, and a contribution from random fluctuations, $\chi_{df}^2(\delta_i) - 1$.

Equation (A-4) shows that the confidence interval for reduced chi-squared can, under certain assumptions, be used as a metric to decide if changes in the fit are statistically significant or not. Provided that $\chi_{df}^2(\Delta y_i)$ does not change greatly (*i.e.*, is independent of random fluctuations δ_i inherent in an ensemble of trials) with different δ_i in repeated trials, then changes in χ_{df}^2 can be ascribed to random fluctuations caused by sampling from the parent distribution for δ_i .

The 95% symmetric confidence interval for χ_{df}^2 is $0.980 \leq \chi_{df}^2 \leq 1.023$ for $df = 18887$.

Given the confidence interval for χ_{df}^2 and allowing for the assumptions listed above, it is clear that the differences in the six-term and three-term fits are only significant for the one-day data; the difference $1.433 - 1.020 = 0.413$ is significantly greater than 0.023, whereas, *e.g.*, the difference between the 2-day χ_{df}^2 's for three- and six-term fits, $1.700 - 1.684 = 0.016 < 0.023$, is not likely to be statistically significant.

Knowledge of the confidence interval for reduced chi-square χ_{df}^2 provides a metric to make a quick judgement on whether two estimates of χ_{df}^2 are statistically different or not.

Given the quality of the fit, the superposition

$f(x) = af_t(x; \nu, \beta, \mu) + (1-a)f_G(x; \sigma, \mu)$ with parameters as given in the first row of **Table A1** are used as the best-fit function to the one-day returns. For brevity, the distribution of $f(x)$ is called a $t+$ distribution rather than a sum or superposition of a Student's t distribution plus a Gaussian distribution.

A.3. Predictions and Comparisons of n -Day Returns

Figure A6 shows, on semi- \log_{10} plots, the development of the distribution of the returns for 2^n -day returns, $n = 1, \dots, 7$. Panel 1, which is located at the upper left of the figure, shows in red a histogram of the \log_{10} of the number of counts observed in each bin of the histogram. The left-most vertical edge of the histogram in panel 1 is located at $x = -229$ and the right-most vertical edge of the histogram is located at $x = 110$, where x is $1000 \times$ the one-day return (*i.e.*, $\%$ returns, see Equation (2) and nearby text), based on the adjusted closing values




Figure A6. Panel 1 shows semi- \log_{10} plots of the frequency of occurrence of 1-day returns as a histogram (red) and best fit curves (blue, green, and black) to the histogram. Panels $2^n, n = 1, \dots, 7$ show 2^n -fold convolutions of the one-day best fit pdfs (blue, green, black) and frequency of occurrence data for 2^n -day returns in red. The colour coding is: blue, Student's t distribution; green, superposition of a Student's t distribution plus a normal distribution, *i.e.*, a $t+$ distribution; and, black, normal distribution.

of the S&P 500. In addition the observed return distribution, the best fit Student's t distribution is shown in blue, the best fit superposition of a Student's t plus a Gaussian distribution is shown in green, and the best fit normal (Gaussian) distribution is shown in black. The Student's t distribution and the superposition fit the data well over a wide range but the normal distribution does not. Note that the distributions shown in panel 1 are truncated at $x = -249$ and at $x = 130$.

The remaining panels show the observed 2^n -day returns, $n = 1, \dots, 7$, as histograms in red, similar to panel 1. The distributions shown in green, blue, and black were obtained by 2^n -fold convolution of the best-fit distributions shown in panel 1. The distributions shown in green, blue, and black in panels $2, 4, 8, \dots, 128$ are the predicted (by convolution) distributions of the returns given the best-fit distributions of the same colour shown in panel 1. Note that the panels do not share the same scaling on the abscissa.

Features similar to “a”, “b”, and “c” of **Figure 1** can be identified in **Figure A6**. The right-hand side of panel 2 is interesting in this regard. Note the sharp downturn of the blue and green curves near the abscissa. This downturn matches the data and arises because of the truncation of the one-day green and blue pdfs. The truncation can be seen in panel one as the blue and green curves ending in “free space”. The effect of the truncation on the left-hand side is not noticeable until panel 4, where a change in the slope of the green curve at the bottom left is just noticeable. On the left-hand side, the effect of the truncation and convolution is more noticeable in panels 8 and 16.

It is the combined effect of the truncation of the one-day return pdf and repeated convolution [24] that takes the fat-tailed distributions shown in panel 1 to Gaussian-like curves shown in panels 64 and 128. In panels 64 and 128, note that the green, blue, and black curves have similar shapes. The black curves in **Figure A6** are all Gaussian (or normal) distributions.

The conclusion from analysis of the trends shown in **Figure A6** is the transition from fat-tailed distributions to normal distributions provides evidence that truncation of the one-day pdf for returns is not an academic exercise, but is required to provide predictive value and to match the observed development of returns over time. The data in **Figure A6** extends the results of Ref. [24] by 14 years and corroborates what was reported 14 years earlier in Ref. [24].

Panels $2^n, n = 1, \dots, 7$ of **Figure A7** plot the development of returns for n -fold convolution of reconfigured returns data in panel 1. In panel 1, the returns as a histogram are shown in red on a semi- \log_{10} plot and the reconfigured data are shown in blue. The blue curve in panel 1 was reconfigured to have equal bin widths, which is required for the convolutions that are depicted in panels $2^n, n = 1, \dots, 7$. The white space between the vertical blue lines in panel 1 indicate regions where there are no returns.

The one-day distribution of returns is truncated in that there are no observed returns outside of the $[-229, 110]$ interval. With repeated self-convolutions, the

returns appear to approach a normal distribution as a comparison of panels 128 in **Figure A7** and **Figure A6** shows.

Figure A7 shows similar data as **Figure A6**. The n -fold convolutions of the interpolated one-day returns match the data well for $n \leq 32$. It should be possible to use numerical integration over the histograms to price options.

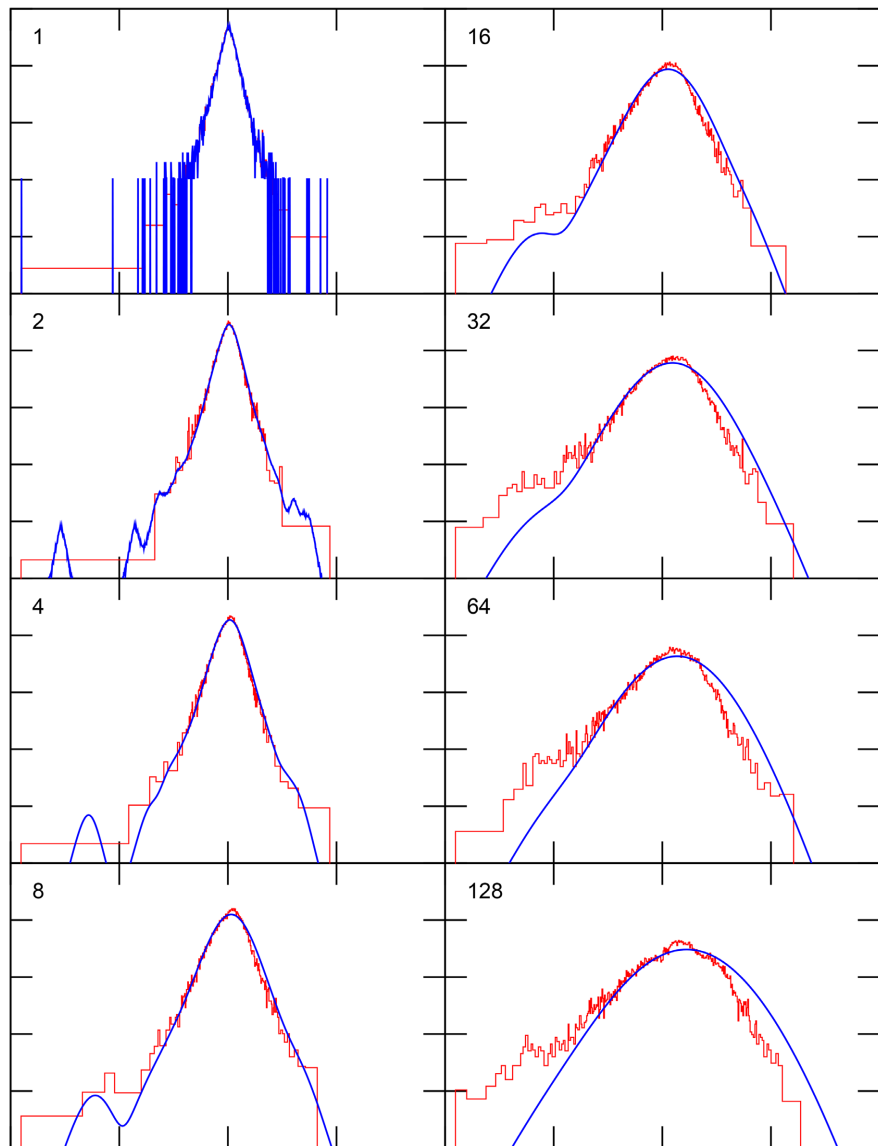


Figure A7. 2^n -fold convolutions of interpolated 1-day frequency of occurrence data (blue) and 2^n -day returns (red).

B. Run Times

100,000 random draws from a uniform PRNG (uniform deviates in the range $[0,1)$) takes $\lesssim 1$ ms (using a gfortran GNU compiler on an M3 Mac computer).

To create 100,000 normally distributed random numbers (normal deviates) using the Box-Muller method [25] [26] takes $\lesssim 3$ ms.

To create 100000 random numbers distributed as a Student's t distribution (Student's t deviates) takes $\lesssim 30$ ms. The inverse transformation method ([17], Ch. 5) ([41], p. 490) was used to create Student's t deviates from uniform deviates. To implement the method, a look-up table for the underlying cumulative density function for the desired Student's t distribution, Simpson's 1/4/1 numerical integration, and several iterations of a Newton-Raphson style root-finding technique were used. The execution time to create the look-up table was not included in the 30 ms time that was quoted.

Given a uniform deviate u from the interval $[0,1)$ and a pdf $f(x)$, then the solution x

$$u = \int_{-\infty}^x f(\xi) d\xi = P\{x \leq u\} \quad (\text{B-1})$$

is distributed as $f(x)$. A look-up table was used to provide an initial estimate for the solution x to the equation, Simpson's 1/4/1 quadrature was used to approximate the integral between the look-up point and u , and a Newton-Raphson style root-finding method was used to find a suitably-precise estimate of the solution x :

$$x_{n+1} = x_n - \frac{(P\{x \leq x_n\} - u)}{f(x_n)}. \quad (\text{B-2})$$

The time taken to find 100,000 solutions x for 100,000 uniform deviates u took twice as long without the denominator in Equation (B-2) as with inclusion of the denominator.

Equation (B-2) is obtained as follows. Let x_n be the n^{th} approximation to the solution to Equation (B-1) and δx be the difference between x_n and the true solution x . For small δx ,

$$\begin{aligned} u &= \int_{-\infty}^{x_n + \delta x} f(\xi) d\xi \\ &= \int_{-\infty}^{x_n} f(\xi) d\xi + \int_{x_n}^{x_n + \delta x} f(\xi) d\xi \\ &\approx P\{x \leq x_n\} + f(x_n) \delta x, \end{aligned} \quad (\text{B-3})$$

which can be solved for δx and used to form the $n+1$ approximation to the solution. Simpson's 1/4/1 rule was used to evaluate $P\{x \leq x_n\}$ over the interval from the nearest look-up point to the value x_n .

The execution time to produce the data for **Figure A6** was ≈ 60 ms. **Figure A6** required 21 convolutions to produce the returns distribution for 2^n , $n=1, \dots, 7$ -day returns. The code to calculate the convolutions was not optimised, so it is likely that the execution time could be reduced. The one-day return distribution was truncated (set to zero) for 20% points beyond the location of the minimum and maximum points, *i.e.*, the one-day distribution was taken as non-zero in the range (249%, 130%). The convolution was calculated for returns spaced every

0.5‰ points and calculation of the convolution in the wings was halted if the contribution ϵ was $<10^{-6}$. This $\epsilon = 10^{-6}$ was used for generation of the figure and could be reduced to decrease execution time. It should be noted that the convolution approach provides estimates of the return distribution that are evenly spaced (0.5‰ points for **Figure A6**) and does not include random spacing of returns and Poisson ‘noise’ as would be obtained with Monte Carlo simulation.