Partially exact and bounded approximations for arithmetic Asian options

Roger Lord[1]

First version: August 27th, 2003
This version: October 12th, 2006

ABSTRACT

This paper considers the pricing of European Asian options in the Black-Scholes framework. All approaches we consider are readily extendable to the case of an Asian basket option. We consider three methods for evaluating the price of an Asian option, and contribute to all three. Firstly, we show the link between the approaches of Rogers and Shi [1995], Andreasen [1999], Hoogland and Neumann [2000] and Večeř [2001]. For the latter formulation we propose two reductions, which increase the numerical stability and reduce the calculation time. Secondly, we show how a closed-form expression can be derived for Curran's and Rogers and Shi's lower bound for the general case of multiple underlyings. Thirdly, we considerably sharpen Thompson's [1999a,b] upper bound such that it is tighter than all known upper bounds. Finally, we consider analytical approximations and combine the traditional moment matching approximations with Curran's conditioning approach. The resulting class of partially exact and bounded approximations can be proven to lie between a sharp lower and upper bound. In numerical examples we demonstrate that they outperform all current state-of-the-art bounds and approximations.

Keywords: Asian option, average price option, basket option, lower bound, upper bound, analytical approximation.

This manuscript was written while the author was partially employed by the Erasmus University Rotterdam. The author would like to express his gratitude to Antoon Pelsser for providing valuable discussions and suggestions. Thanks also to Jeroen van der Hoek for pointing out Curran's approximation, Michèle Vanmaele and Griselda Deelstra for carefully reading and correcting the manuscript, Reimer Beneder, Walter Foppen, Olivier Laseroms, Joeri Potters, Erik van Raaij, Matthew Steiner, Sacha van Weeren, Mike Curran, Jiri Hoogland and Mike Staunton for comments and suggestions. Finally, we are grateful to seminar participants at Rabobank International, the Tinbergen PhD lunch seminar, the ECMI 2004 conference and the 3rd Actuarial and Financial Mathematics Day in Brussels for fruitful comments.

[1] Modelling and Research (UC-R-355), Rabobank International, P.O. Box 17100, 3500 HG Utrecht, The Netherlands (e-mail: roger.lord@rabobank.com, Tel. +31-(0)30-2166566).

1. Introduction

This paper deals with the pricing of arithmetic European Asian options. Asian options[2], also referred to as average rate or average price options, are financial derivatives depending on the average of a certain underlying asset over a prespecified time interval. Interest rates, exchange rates, bonds, commodities, stocks or indices act as the underlying asset.

Asian options come in numerous flavours. If the strike price depends on a fixed quantity, the option is referred to as a fixed strike Asian option or an average price or rate option. If instead the strike price is proportional to the asset price itself, the contract is called a floating strike Asian option or average strike option. A further distinction can be made on basis of the nature of the average. This can be either arithmetic or geometric, both with possibly varying weights for past observations. The average itself can be discretely sampled, i.e. based on a finite number of past realisations, or continuously sampled. In practice all contracts are based on the discretely sampled arithmetic average, although a vast amount of papers deal with the continuously sampled variety.

There are many reasons for the existence of Asian options. Whereas contracts depending only on one snapshot of an asset price are vulnerable to sudden large shocks or price manipulation, Asian options are much more robust against such phenomena. End-users may prefer Asian options as hedging instruments as they may be exposed to the average performance of the underlying over time. In addition, Asian options are cheaper than their plain vanilla counterparts and are easier to hedge. The latter can easily be seen if we consider the volatility of the average: it will typically be lower than that of the underlying asset. In addition, the closer we are to the maturity date, the smaller the uncertainty in the average will be. This implies a lesser dependence of the option on the spot price than a plain vanilla option with the same maturity.

Asian options can be embedded in more complex financial structures. To counter the price manipulation issue, many exotic options contain so-called “Asian tails”. This entails nothing else than that the final payoff is based on the average price of the underlying over a certain interval prior to expiry. Another example of such a structure is unit-linked insurance, as mentioned by e.g. Nielsen and Sandmann [2003] and Schrager and Pelsser [2004]. In its most basic form this can be described as an investment plan with a long maturity where periodic payments are invested in risky investment funds. An Asian option on the average return can be used as a rate of return guarantee. Schrager and Pelsser mention that many insurance companies have supplied these forms of insurance without realising the risk attached to the embedded options. With the fair value calculations of insurance contracts currently at the center of attention, quantifying the risk attached to these embedded options is of the utmost importance.

Within the Black-Scholes framework, in which we will be based, already no closed-form solutions are available for arithmetic Asian options. In this framework the underlying is assumed to follow a geometric Brownian motion, which amounts to a lognormal distribution for each asset price. Unlike the geometric average, which as a product of lognormal random variables is itself lognormally distributed, the arithmetic average is a sum of correlated lognormal random variables. Unfortunately no closed-form expression is available for the probability law of this sum. The exact same problem occurs when pricing a basket option, whose price depends on the arithmetic average of several assets.

Consequently, the literature has explored a large variety of ways to find a price for the value of an arithmetic Asian option, and research is still ongoing. One can broadly distinguish between methods based on the solution of a partial differential equation (PDE), analytical approximations, lower and upper bounds, tree methods, Monte Carlo methods and transform methods. In the

---

[2] The only reason that they are referred to as “Asian” options is that the first known transaction occurred in Tokyo, as is noted in Falloon and Turner [1999].

following we by no means aim to give a complete overview. In particular, as this paper contributes to the first three areas of research, the subsequent sections will deal with these methods in much greater detail than we do here.

For now it suffices to say that PDE methods are most probably the most flexible and efficient way of valuing Asian options, especially if more exotic features are included in the financial structure. Over the past years several authors, starting with Rogers and Shi [1995], have demonstrated that the price of an Asian option can be found by solving a PDE in one space dimension, we mention Andreasen [1998], Hoogland and Neumann [2000a,b] and Večeř [2001]. Though the PDE methods can be extended to the case where the option depends on a basket of underlyings, this is less advisable if the number of factors is large. Our contribution is to relate all mentioned PDE approaches to each other by a change of variable. For Večeř's formulation, which seems to be the most stable in practice, we propose two reductions that increase the numerical stability and reduce the calculation time.

Analytical approximations can be very useful to quickly generate a hopefully accurate estimate of the option value and its sensitivities. In contrast with the PDE methods, once we have an analytical approximation for an Asian option it is usually relatively straightforward to come up with an approximation for a basket option that has the same computational cost. The approximations we will consider in this paper combine the moment matching approaches, which date back to Levy [1992], and the conditioning approaches of Curran [1994]. We introduce the class of partially exact and bounded (PEB) approximations, which apply conditional moment matching. The class is a logical extension of Curran's approximation. The advantage of conditional moment matching is that, in contrast with the traditional moment matching approaches, the size of the error made lies between a sharp lower and upper bound. We show that the error tends to zero when the strike price tends to zero or to infinity. Though the latter is a very natural criterion, we show that it is not satisfied by Curran's approximation; there the call price diverges as the strike price tends to infinity.

The final line of research that we contribute to in this paper is concerned with deriving lower and upper bounds for the value of an Asian option. Bounds can themselves serve as an approximation if they are sufficiently tight, with the advantage that the sign of the error is known a priori. The sharpest lower bound that is currently known is due to Curran [1992, 1994] and Rogers and Shi [1995]. Though a closed-form expression is already available for the lower bound in case of an Asian option, Deelstra, Liinev and Vanmaele [2004] already noted that this expression does not always carry over when there are multiple underlyings. In general one therefore has to resort to a numerical integral over a discontinuous integrand, something which is undesirable. We derive a closed-form expression for the general case, which requires at most three numerical searches. As the previous lower bound was already very sharp, subsequent research efforts have mainly focused on deriving a sharp upper bound. We revisit an upper bound proposed by Thompson [1999a], and show how his upper bound incorporates many of the upper bounds considered later in the literature. More importantly, we enhance Thompson's upper bound and produce a new bound which is much tighter than all currently reported upper bounds. The new upper bound still performs well when the total volatility is high, i.e. for high volatility environments and/or for long maturities. As the previously mentioned unit-linked insurance typically has long maturities, this is an important finding.

Clearly, the above methods are not all one can use. Tree methods (see Hull and White [1993] and more recently Klassen [2001]) make it possible to value certain path-dependent contracts in a tree. Monte Carlo methods can be very useful, as one can combine control variates (see e.g. Kemna and Vorst [1990], who demonstrate that the geometric average option is highly correlated with the arithmetic average option), importance sampling and stratification (see Glasserman, Heidelberger and Shahabuddin [1990] for an application to Asian options) to significantly reduce the variance. Finally, we should also mention transform methods. Whereas Geman and Yor [1993], using Laplace inversion, derive a closed-form expression for the value of a continuously sampled Asian option, Carverhill and Clewlow [1990] consider the discretely sampled case. Their work, revisited more recently by Benhamou [2002] and Den Iseger and Oldenkamp [2006], shows how to write the arithmetic average as a product of independent random variables, so that its distribution can be found by convolution techniques. Unfortunately these approaches heavily hinge on the assumption that the stock increments are independent, which clearly is only true in the Black-Scholes model when there is only one underlying asset. As such the algorithm cannot easily be adapted to the more general case of a basket option.

The remainder of this paper is organised as follows. In section 2 we briefly describe the model of the financial market that will be used throughout this document. Section 3 shows the link between Rogers and Shi's, Andreasen's and Večeř's PDE, and proposes two reductions for the latter PDE. In section 4 we consider lower bounds, and show how to arrive at a closed-form expression for Curran's and Rogers and Shi's lower bound, which remains valid when multiple underlyings are present. Section 5 deals with Thompson's upper bound, and sharpens it considerably. In section 6 the class of PEB approximations are introduced. Section 7 concludes with the numerical comparison of the new upper bound and some elements of the class of PEB approximations to state-of-the-art bounds and approximations. Večeř's PDE is used to generate the "exact" prices for the numerical examples. Deltas, gammas and vegas are calculated to demonstrate the accuracy of the PEB approximations.

Finally, it cannot be stressed enough that the approximations and bounds considered in this paper extend to more general situations where the underlying can be expressed as a sum of lognormal random variables. Basket options and the pricing of interest rate swaptions in a Gaussian term structure model are hence two other examples of financial products for which the approximations and bounds will be valid.

2. The model

Throughout this document we will constrain ourselves to the Black-Scholes framework. As mentioned earlier, this simple framework already does not yield closed-form expressions for the value of an Asian option. For ease of exposure we will work with a constant and deterministic interest rate, volatility and growth rate of the asset. All results remain valid when the interest rate, volatility and growth rate are deterministic functions of time. The results could even be extended to the case where the term structure of interest rates is Gaussian.

In the Black-Scholes framework the underlying asset and the money market account evolve according to the following stochastic differential equation:

$$\begin{aligned} dS(t) &= \mu S(t)dt + \sigma S(t)dW(t) \\ dB(t) &= rB(t)\,dt \end{aligned} \tag{2.1}$$

where μ is the growth rate, r the interest rate, σ the volatility of the stock, and W(t) is a Brownian motion under the risk-neutral probability measure. For dividend protected assets no arbitrage restrictions enforce the growth rate to be equal to the risk-free rate, whereas when the asset under consideration has a constant dividend yield equal to q, the growth rate μ must be equal to r-q. Throughout the document we will assume, without loss of generality, that the current date is 0. The arithmetic average at the maturity date T will be defined as:

$$A(T) = \int_0^T S(t)\,\rho(t)\,dt \tag{2.2}$$

where ρ is a non-negative weighting function, which integrates to 1 over the interval [0,T]. As an example, the continuously sampled arithmetic average with equal weights is represented by $\rho(t) = \frac{1}{T}$, whereas the discretely sampled arithmetic average with fixing dates $0 < t_1 \leq ... \leq t_N = T$ and equal weights is obtained when $\rho(t) = \frac{1}{N}\sum_{i=1}^{N}\delta(t_i - t)$. Here δ is Dirac's delta function.

In our analysis we will only consider newly issued, non-forward-starting, fixed strike arithmetic Eurasian calls. This is no loss of generality. Put options can be priced via the Asian put-call parity, as we will show. In Hull [2005, pp. 538-540] it is shown how to treat running average Eurasian options as newly issued ones. Similarly, when interest rates are deterministic, forward-starting options can be dealt with easily as well. As for floating strike options, symmetry results between floating and fixed strike Asian options were first shown to exist in Hoogland and Neumann [2000a], and later in Henderson and Wojakowski [2002] and Vanmaele et al. [2006].

For ease of exposure we will mostly deal with forward prices in our analysis. The forward price of the arithmetic Eurasian fixed strike option is equal to its expected value under the risk-neutral probability measure $\mathbb{Q}$, conditional upon all information known at time 0:

$$c_A(T,K) = \mathbb{E}_0^{\mathbb{Q}}[(A(T) - K)^+] \tag{2.3}$$

If no confusion can arise, we will leave out the superscript indicating the measure and the subscript indicating at which time the expectation is evaluated. Current prices can easily be obtained by discounting with the risk-free interest rate. The Asian put-call parity (in terms of forward prices) states that:

$$p_A(T,K) = c_A(T,K) + K - \mathbb{E}[A(T)] \tag{2.4}$$

From (2.4) it is evident that lower and upper bounds for calls also translate to lower and upper bounds for puts. A final quantity of interest is the geometric average. This is defined as:

$$G(T) = \exp\left(\int_0^T \ln S(t)\rho(t)\,dt\right) \tag{2.5}$$

An application of the weighted Jensen's inequality shows that $A(T) \geq G(T)$, with equality attained if and only if all components of the average (where $\rho(t) \neq 0$) are equal.

3. The partial differential equation approach

In general the price of an Asian option can be found by solving a PDE in two space dimensions, see Wilmott [2006]. In the seminal paper of Rogers and Shi [1995], a variable reduction was used to find a PDE in one space dimension for the value of an Asian claim, for both fixed and floating strikes. A problem associated with this PDE is that for both discretely sampled Asians and floating strike options the Dirac delta function appears as a coefficient in the PDE. Zvan, Forsyth and Vetzal [1997/98] applied techniques from the field of computational fluid dynamics to this PDE to improve the numerical accuracy. Recently, Hoogland and Neumann [2000a,b] and Večeř [2001] arrive at a different one-dimensional PDE for the value of an Asian claim, which does not have the problems associated with the Rogers and Shi PDE. Hoogland and Neumann use the notion of local scale invariance, whereas Večeř considers options on a traded account, and demonstrates that Asian options are a special case hereof.

Basing ourselves on Večeř's derivation of the PDE we will demonstrate, as Hoogland and

Neumann [2000a] briefly mention, that it is related to the Rogers and Shi's PDE by a simple change of variable, hereby eliminating the Dirac delta function from the PDE. Similarly it can be shown that Andreasen's [1998] PDE, which was applied only for discretely sampled Asians, is also related to both PDE's. In the second paragraph we propose two reductions for this PDE formulation, which increase the numerical stability and reduce the calculation time required.

3.1. The equivalence of various PDE approaches

We will start by introducing some concepts related to Večeř's PDE, as this will be insightful for the reductions to follow in the next paragraph. This PDE is formulated for options on a traded account. A traded account can be viewed as a bank account in which we are allowed to invest in stocks during the life of the option, within certain restrictions. The remaining cash position is invested against a constant interest rate. An option on this so-called traded account is a contract that promises to pay the value of the traded account at maturity if this is positive. If the trading strategy has not been prosperous, and the value of the traded account is negative, the holder receives nothing. Options on a traded account generalise the concept of many options. European, American, passport and vacation options can be shown to be special cases of options on a traded account. For the remainder we will assume that the trading strategy is known a priori and that the interest rate earned on the remaining cash position is zero. This is sufficient for our purposes. If we denote the trading strategy at time t by q(t) and the value of the traded account by X(t), then the terminal value of the traded account equals:

$$X(T) = X(0) + \int_0^T q(t)dS(t) \tag{3.1}$$

The payoff of the option on the traded account is equal to $X(T)^+$ at maturity. Note that q(t) = 1 and initial wealth equal to S(0) – K will yield a final payoff equal to that of a European call option, a result we will use later. The link between Asian options and options on a traded account is found by relating the trading strategy q(t) to the weighting function ρ(t) in (2.2):

$$q(s) = \int_s^T \rho(t)\,dt \qquad 0 \le s \le T \tag{3.2}$$

As ρ is non-negative and integrates to 1, q(0) = 1 and q(T) = 0. Partial integration then yields:

$$\begin{aligned} X(T) &= X(0) + q(T)S(T) - q(0)S(0) + \int_0^T S(t)\,\rho(t)\,dt \\ &= X(0) - S(0) + \int_0^T S(t)\,\rho(t)\,dt \end{aligned} \tag{3.3}$$

Setting the initial value of the traded account to be equal to X(0) = S(0) – K, we see that the payoff of the option on the traded account is equal to that of an arithmetic Eurasian call.

Let us introduce the value of the traded account relative to the stock price, Z(t) = X(t)/S(t). The value of the option at time t is now given by:

$$\begin{aligned} V\big(t,S(t),X(t)\big) &= e^{-r(T-t)}\,\mathbb{E}_t^{Q}[X(T)^+] \\ &= S(t)\,e^{-(r-\mu)(T-t)}\,\mathbb{E}_t^{\$}[\tfrac{X(T)}{S(T)}^+] = S(t)\,e^{-(r-\mu)(T-t)}\,\mathbb{E}_t^{\$}[Z(T)^+] \end{aligned} \tag{3.4}$$

where $\mathbb{S}$ is the probability measure associated with taking the stock price (including all accumulated dividends) as the numeraire asset. By deriving the dynamics of Z(t) and applying the Feynman-Kač theorem one finally obtains the following PDE:

$$\frac{\partial u}{\partial t}+\mu\big(q(t)-Z(t)\big)\frac{\partial u}{\partial Z}+\tfrac{1}{2}\sigma^2\big(q(t)-Z(t)\big)^2\frac{\partial^2 u}{\partial Z^2}=0 \tag{3.5}$$

which has to be solved subject to the terminal condition $u\big(T,Z(T)\big)=Z(T)^+$. The option price follows from (3.4) after solving the PDE. The link with Rogers and Shi's PDE will now be obtained by a change of variable. To this end we introduce:

$$Y(t)=q(t)-Z(t) \tag{3.6}$$

The reparametrised option price is denoted as $f\big(t,Y(t)\big)=u\big(t,Z(t)\big)$. The PDE then becomes:

$$\frac{\partial f}{\partial t}-\big(\mu Y(t)+\rho(t)\big)\frac{\partial f}{\partial Y}+\tfrac{1}{2}\sigma^2 Y(t)^2\frac{\partial^2 f}{\partial Y^2}=0 \tag{3.7}$$

with the boundary condition being $f\big(T,Y(T)\big)=Y(T)^-$ for the fixed strike option. Rogers and Shi only considered the case $\mu = r$; in this case the PDE in (3.7) coincides completely with their PDE for the option price divided by the spot price.

Due to the appearance of $\rho(t)$, the Dirac delta function will enter the coefficients of the PDE for discretely sampled and floating strike options. In contrast, these disappear in the formulations of Hoogland and Neumann, and Večeř, simply by integrating over time. This does not leave Rogers and Shi's formulation without any merit. From a computational point of view it is clear that for continuously sampled arithmetic Asians with equal weights the PDE of Rogers and Shi may be preferred, as it has constant coefficients. This leads to a much quicker solution when using finite differences.

Finally, Andreasen's PDE is formulated for the option price divided by the stock price. Realising this, it is easy to see that we can obtain this PDE from Večeř's formulation by simply defining a new variable which is equal to –Z(t). We leave this for the interested reader.

3.2. Reducing calculations

Having shown the relation between the various PDE approaches, we will now turn to its numerical implementation. We will again work from Večeř's PDE. Two reductions are proposed which will increase the numerical stability and the calculation time required. The reductions are derived with a finite difference solution method in mind, as demonstrated in Večeř's article. The first reduction is the equivalent of the fact that the value of an Asian with the average taken over only one observation is equal to that of its plain vanilla counterpart. The second reduction, which can be found in a slightly different form for discretely sampled options in Andreasen [1998], is based on the fact that a large number of values are already known beforehand.

Reducing calculations: elimination of final fixing date

It can be shown that in the case of discretely sampled Eurasians, when N fixing dates are present we only have to build a grid for N-1 fixing dates. Assume we have fixing dates t_i, where as before $0 < t_1 \leq ... \leq t_N = T$. Suppose that q_j is the trading strategy[3] on the interval $[t_{j-1}, t_j)$. If we focus on the final interval $[t_{N-1},t_N)$ and use the change of variable $Y(t) = \frac{1}{q_N} Z(t)$, we have to solve the following PDE to calculate the value of the option here:

$$\frac{\partial u}{\partial t} + \mu(1 - Y(t))\frac{\partial u}{\partial Y} + \tfrac{1}{2}\sigma^2(1 - Y(t))^2 \frac{\partial^2 u}{\partial Y^2} = 0 \qquad (3.8)$$

whereas the constraint changes into: $u(T, Y(T)) = q_N \cdot Y(T)^+$. As mentioned earlier, when the trading strategy is equal to 1, we are dealing with a European call. Here we thus have q_N European calls. The risk-free rate and volatility are specified, so all that we have to figure out are the correct spot and strike price. As Z(0) = X(0)/S(0) and X(0) = S(0) – K, it follows that the strike price must equal K = S(0) · (1-Z(0)). Fixing the initial spot price at 1, an initial wealth of Z(t) corresponds to a strike price of 1-Z(t). We end up with:

$$u(t_{N-1}, Z(t_{N-1})) = q_N e^{-r(t_N - t_{N-1})} \mathbb{E}^Q_{t_{N-1}}\left[\left(\frac{S(t_N)}{S(t_{N-1})} - \left(1 - \tfrac{1}{q_N} Z(t_{N-1})\right)\right)^+\right] \qquad (3.9)$$

This formula can be calculated straightforwardly using the Black-Scholes formula. Concluding, it is evident that we only need to build a grid for the first N-1 fixing dates. Formula (3.9) takes the role of the boundary condition. Although this may not speed up calculations, the boundary condition now no longer has any discontinuities, which increases the numerical stability. Unfortunately, using this approach for the remaining fixing dates will not yield a closed-form formula for discretely sampled Eurasians. The result in this paragraph is therefore no different from the fact that an Asian option with one fixing is equal to its plain vanilla counterpart. In the next subparagraph we show how to reduce the calculations significantly by using known values on the grid.

Reducing calculations: known values on the grid

The reductions in this subparagraph apply to all types of arithmetic Eurasians. Consider that we have arrived at time s in the grid. Assuming the trading-strategy is non-zero, we propose the following change of variable:

$$Y(t) = \tfrac{1}{q(s)} Z(t) \qquad s \leq t \leq T \qquad (3.10)$$

As in the previous subparagraph it is easy to see that only two things change:

- The boundary condition becomes $u(T, Y(T)) = q(s) \cdot Y(T)^+$;
- The trading strategy at time t becomes $q(t)/q(s)$.

The boundary condition is no problem, as this only implies we are pricing q(s) options on a traded account. Via the change of variable the weighting function on [s,T] turns into $\rho(t)/q(s)$, which by definition is non-negative and integrates to 1. We can thus consider the value of the option at

---

[3] Note that the trading strategy is always positive for a discrete arithmetic Eurasian fixed strike call.

time s as that of a newly issued option with time to maturity equal to T-s, in conjunction with the adjusted boundary condition and trading strategy.

How is this advantageous? We are going to use the fact that the price of an arithmetic Eurasian call is known analytically for strike prices equal to zero. When the strike price is zero, the contract is simply a forward on the arithmetic average, for which closed-form expressions are readily available. For strike prices smaller than zero, we only need to add a zero-coupon bond with a notional equal to the absolute value of the strike price. In formulas:

$$e^{-r(T-t)}\,\mathbb{E}_t^Q[(A(T)-K)^+] = e^{-r(T-t)}\cdot\left(\mathbb{E}_t^Q[A(T)]-K\right) \qquad K \le 0 \tag{3.11}$$

As we showed in the previous subparagraph, an initial wealth of Y(t) corresponds to a strike price of 1-Y(t), if we keep the spot price fixed at 1. This is smaller than or equal to zero when:

$$1-Y(t) = 1-\tfrac{1}{q(s)}Z(t) \le 0 \Leftrightarrow Z(t) \ge q(s) \tag{3.12}$$

For the time interval on which the trading strategy is equal to q(s) we only need to build a grid for those points not satisfying (3.12). All grid points that do satisfy this inequality can simply be assigned the known value. The latter only requires the knowledge of the forward price of the arithmetic average.

To quantify the reduction, suppose we have a uniform grid in space and time for the PDE, i.e. $z_i = z_0 + i\,dz$ and $t_j = j\,dt$, for $0 \le i \le M$ and $0 \le j \le N$. The space step and timestep are represented by dz and dt, respectively. For the final timepoint we have $t_N = T$. Note that the choice $z_M = 1$ corresponds to an option with a strike price equal to zero. For the continuously sampled call we have $q(s) = \tfrac{T-s}{T}$ for $0 \le s \le T$, so that the reduction for the continuously sampled Eurasian call with equal weights is approximately equal to:

$$\tfrac{1}{T}\int_0^T \frac{z_M - q(t)}{z_M(z_M - z_0)}dt = \frac{z_M - \tfrac{1}{2}}{z_M(z_M - z_0)} \tag{3.13}$$

Assuming that z_M is larger than or equal to 1, and z_0 is smaller than zero, $(z_M - q(t))/z_M$ is approximately the fraction of positive grid points that do not have to be calculated. We divide by $z_M - z_0$ to obtain the fraction of all grid points that are known. Note that e.g. the Crank-Nicolson scheme requires O(M) calculations at each timestep, where M is the amount of space points, so that this is an appropriate measure of the reduction. Subsequently we average over all timepoints by integrating over time, and divide by the maturity. Večeř chooses[4] $z_0 = -1$ and $z_M = 1$, so that the reduction in this case is approximately 25% and thus very significant.

In addition to speeding up calculations, this reduction also increases the numerical stability. Typically, one would assume that for large values of z, i.e. for high positions of the normalised traded account, the option value will be linear in z. Estimating the point from where onwards this will be approximately valid is already very difficult, as this requires prior knowledge of the function we are trying to solve. Using a known value as the boundary condition bypasses this problem, making the solution more numerically stable.

More on the implementation of Večeř's PDE will follow later in the section on numerical results, where we numerically solve the PDE in order to compare the exact prices of Asian options to upper bounds and approximations. We now continue with lower bounds for the value of an Asian option.

---

[4] Note that this choice will not be suitable for an arbitrary choice of parameter values, but is sufficient for the parameters chosen by Večeř.

4. Lower bounds via conditioning

As mentioned a line of research on Asian options has dealt with deriving lower and upper bounds for the value of these options. In this section we will focus on lower bounds. The first article to our knowledge to derive a lower bound for the value of an Asian option was that of Vorst [1992]. Vorst uses the knowledge that the geometric average is always smaller than or equal to the corresponding average. This is all that is required to see that the value of an arithmetic Eurasian fixed strike call is bounded below by that of its geometric counterpart. Subsequently Curran [1992, 1994] and Rogers and Shi [1995] managed to derive a very tight lower bound by conditioning and applying Jensen's inequality. The resulting lower bound is very tight and in most cases is in fact a better estimate of the option value than a large number of analytical approximations. As the underlying ideas will feature prominently in the rest of this paper, we provide its derivation in section 4.2. Prior to this we give some preliminary results on conditioning in a Gaussian setting. Finally, section 4.3 gives a closed-form expression for the lower bound when applied to an arbitrary sum of lognormal random variables. Though a closed-form expression was already available for the case of arithmetic Asian options, this was not the case for basket options.

4.1. Some preliminary results

Let Z be an arbitrary Gaussian conditioning variable. We introduce the following notation for ease of exposure:

$$\mu_Z = \mathbb{E}[Z] \quad \sigma_Z^2 = \mathrm{Var}(Z) \quad \sigma_Z(t) = \mathrm{Cov}\big(\ln S(t), Z\big) \tag{4.1}$$

All expectations are taken under $\mathbb{Q}$, conditional upon all information known at time 0. Using standard results about Gaussian random variables, the distribution of ln S(t) given Z equals:

$$\ln S(t) \mid Z \sim N\Big(\ln S(0) + \tfrac{\sigma_Z(t)}{\sigma_Z^2}(Z - \mu_S), \sigma^2 t - \tfrac{\sigma_Z^2(t)}{\sigma_Z^2}\Big) \tag{4.2}$$

The conditional expectation of S(t) given Z thus equals:

$$\mathbb{E}[S(t) \mid Z] = S(0)\exp\Big(\mu t + \tfrac{\sigma_Z(t)}{\sigma_Z^2}(Z - \mu_Z) + \tfrac{1}{2}(\sigma^2 t - \tfrac{\sigma_Z^2(t)}{\sigma_Z^2})\Big) \tag{4.3}$$

With this result in hand we can easily calculate the expectation of A(T) given Z. In the remainder of this section we will require the following expectation:

$$\mathbb{E}[S(t)1_{[Z\geq z]}] = \int_z^\infty \mathbb{E}[S(t) \mid Z = u]\ dF_Z(u) \tag{4.4}$$

where F_Z is the cdf of the Gaussian Z. To find a closed-form expression for this integral we will use the following result:

$$\mathbb{E}[\exp(tZ)1_{[Z\geq z]}] = \exp(\tfrac{1}{2}t^2)\cdot N(t - z) \tag{4.5}$$

where N is the normal cdf. Applying this to (4.4):

$$
\begin{aligned}
\frac{\mathbb{E}[S(t)1_{[Z\geq z]}]}{\mathbb{E}[S(t)|\,Z=0]} &= \int_{\lambda(K)}^{\infty} \exp\left(\tfrac{\sigma_Z(t)}{\sigma_Z^2}u\right) dF_Z(u) \\
&= \int_{\frac{z-\mu_Z}{\sigma_Z}}^{\infty} \exp\left(\tfrac{\sigma_Z(t)}{\sigma_Z^2}(\mu_Z+\sigma_Z u)\right) \varphi(u)\, du \\
&= \exp\left(\tfrac{\sigma_\Lambda(t)}{\sigma_\Lambda^2}\left(\mu_Z+\tfrac{1}{2}\sigma_Z(t)\right)\right)\cdot N\left(\tfrac{\sigma_Z(t)-(z-\mu_Z)}{\sigma_Z}\right)
\end{aligned}
\tag{4.6}
$$

where φ is the normal pdf. Using this closed-form solution it is trivial to find $\mathbb{E}[A(T)1_{[Z\geq z]}]$. In the following section it will become clear how the formulae derived here can be used to find a lower bound for arithmetic Asian options.

4.2. Derivation of the lower bound

Though Curran's and Rogers and Shi's lower bounds are derived in a different manner, they coincide when the same conditioning variable is used. Rogers and Shi's lower bound is more general in that it allows for an arbitrary conditioning variable. Curran's lower bound, as mentioned by Vanmaele et al., only works for Gaussian[5] variables Λ which have the convenient property that there is a threshold value λ(K) for which Λ ≥ λ(K) implies A(T) ≥ K. To distinguish general conditioning variables from such special conditioning variables, we introduce the following convention.

Notation:

Λ will denote a Gaussian random variable for which Λ ≥ λ(K) implies A(T) ≥ K
Z will denote a general Gaussian random variable

Clearly, if we have a random variable that is a lower bound for the arithmetic average, such a random variable Λ can be constructed. Two such examples follow.

Example:
As the geometric average G(T) in (2.5) is a lower bound for A(T), its logarithm obviously satisfies the previous criterion. We will denote Λ_{GA} = ln G(T), and the corresponding threshold is $\lambda_{GA}(K)$ = ln K. A second conditioning variable which can be used follows from a first order approximation of A(T), as shown by Vanmaele et al. For t ≥ 0 the solution to the SDE in (2.1) easily follows as $S(t)=S(0)e^{(\mu-\frac{1}{2}\sigma^2)t+\sigma W(t)}$. Due to the convexity of the exponential, it is clear that the following first order approximation, which we will denote by Λ_{FA}, also acts as a lower bound for the arithmetic average:

$$
A(T) \geq \int_0^T S(0)e^{(\mu-\frac{1}{2}\sigma^2)t}\left(1+\sigma W(t)\right)\rho(t)\,dt \equiv \Lambda_{FA} \tag{4.7}
$$

The corresponding threshold value for this first order approximation is $\lambda_{FA}(K) = K$. □

We start with Rogers and Shi's lower bound, which is based on the following application of Jensen's inequality:

---

[5] There is no loss of generality in this. If Λ is not Gaussian and has F as its cumulative distribution function (cdf), $N^{-1}(F(\Lambda))$ is Gaussian and the corresponding threshold is $N^{-1}(F(\lambda(K)))$. Here N is the normal cdf.

$$\mathbb{E}\left[X^+\right] = \mathbb{E}\left[\mathbb{E}[X^+ \mid Y]\right] \geq \mathbb{E}\left[\mathbb{E}[X \mid Y]^+\right] \tag{4.8}$$

Rogers and Shi's lower bound then follows as:

$$\begin{aligned} \mathbb{E}[(A(T)-K)^+] &\geq \mathbb{E}[\mathbb{E}[(A(T)-K) \mid Z]^+] \\ &= \int_{-\infty}^{\infty} (\mathbb{E}[A(T) \mid Z = z] - K)^+ \, dF_Z(z) \equiv LB(Z) \end{aligned} \tag{4.9}$$

where F_Z is the cdf of the Gaussian Z. Curran's lower bound differs in that it conditions on Λ, and splits the expectation in two parts, depending on whether the conditioning variable is above or below its threshold:

$$\begin{aligned} c_A(T,K) &= \mathbb{E}[(A(T)-K)^+ 1_{[\Lambda<\lambda(K)]} + (A(T)-K)^+ 1_{[\Lambda\geq\lambda(K)]}] \\ &= \mathbb{E}[(A(T)-K)^+ 1_{[\Lambda<\lambda(K)]} + (A(T)-K) 1_{[\Lambda\geq\lambda(K)]}] \\ &\equiv c_1(T,K,\Lambda) + c_2(T,K,\Lambda) \end{aligned} \tag{4.10}$$

The special structure of Λ allows us to split the option price into two parts – c_1 and c_2. The second part can be calculated in closed-form by using the results from section 4.1. To this end define:

$$\begin{aligned} c_2(T,K,Z,z) &= \mathbb{E}\left[(A(T)-K) 1_{[Z\geq z]}\right] \\ &= \int_0^T \mathbb{E}[S(t) \mid Z = 0] \cdot \exp\left(\frac{\sigma_Z(t) \cdot \left(\mu_Z + \frac{1}{2}\sigma_Z(t)\right)}{\sigma_Z^2}\right) \cdot N\left(\frac{\sigma_Z(t) - (z - \mu_Z)}{\sigma_Z}\right) \rho(t) dt \\ &\quad - K N\left(\frac{\mu_Z - z}{\sigma_Z}\right) \end{aligned} \tag{4.11}$$

It follows that $c_2(T,K,\Lambda) = c_2(T,K,\Lambda,\lambda(K))$. A lower bound for c_1 is found once again by applying (4.8):

$$\begin{aligned} c_1(T,K,\Lambda) &= \mathbb{E}[(A(T)-K)^+ 1_{[\Lambda<\lambda(K)]}] = \mathbb{E}[\mathbb{E}[(A(T)-K)^+ 1_{[\Lambda<\lambda(K)]} \mid \Lambda]] \\ &\geq \mathbb{E}[\mathbb{E}[(A(T)-K) 1_{[\Lambda<\lambda(K)]} \mid \Lambda]^+] \\ &= \int_{-\infty}^{\lambda(K)} (\mathbb{E}[A(T) \mid \Lambda = \lambda] - K)^+ \, dF_\Lambda(\lambda) \end{aligned} \tag{4.12}$$

The resulting lower bound will be denoted by $LB(\Lambda)$. Rogers and Shi's and Curran's lower bound coincide when a conditioning variable Λ is used. Both approaches have in common that eventually we must numerically evaluate an integral with a discontinuous integrand, something which is undesirable. This problem is dealt with in the following section.

The reason why $LB(\Lambda)$ works so well is that the proposed choices of Λ resemble A(T) closely. The more information the conditioning variable contains about A(T), the lower the contribution of c_1 will be to the option price, and hence the smaller the error will be which we make in the lower bound. An estimate of this error can be made, and leads to Rogers and Shi's upper bound, recently sharpened by Vanmaele et al. [2006].

As a final note, it can be shown that $LB(\Lambda_{GA})$ is tighter than Vorst's lower bound. Vorst's lower bound, the value of the geometric average call, can be rewritten as:

$$\mathbb{E}[(G(T)-K)^+] = \mathbb{E}[(G(T)-K)1_{[\Lambda_{GA} \geq \lambda_{GA}(K)]}] \tag{4.13}$$

as $\{\Lambda_{GA} \geq \lambda_{GA}(K)\} = \{G(T) \geq K\}$. Hence, Vorst's lower bound is already smaller than $c_2(T,K)$ and therefore is clearly strictly smaller than $LB(\Lambda_{GA})$.

4.3. Closed-form expression for the lower bound

It has already been pointed out in the literature that when using Λ_{FA} or Λ_{GA} as conditioning variables the lower bound of Curran and Rogers and Shi can be calculated in closed form, see e.g. Thompson [1999a], Nielsen and Sandmann [2003] and Vanmaele et al. [2006]. The result on which it is based, stated generally in Lemma 1, is similar to the result used in Jamshidian's [1989] decomposition used for the pricing of options on a coupon-bearing bond.

Lemma 1:
Provided that $\sigma_Z(t)$, the covariance between the underlying asset and the conditioning variable Z, is strictly positive, there is a unique z'(K) such that:

$$\mathbb{E}[A(T) \mid Z = z'(K)] = K \tag{4.14}$$

In case $Z = \Lambda$, we know that $\lambda'(K) \in [-\infty, \lambda(K)]$.

Proof:
Due to the positivity of $\sigma_\Lambda(t)$ it is clear that the conditional expectation of A(T), built up of terms like (4.3), is strictly increasing and convex in z, and takes values from zero to infinity. This implies that there is a z'(K) such that (4.14) holds. If $Z = \Lambda$, the structure of Λ implies that the threshold $\lambda'(K) \in [-\infty, \lambda(K)]$. □

Remark 1:
For the conditioning variables Λ_{FA} and Λ_{GA} we indeed have that $\sigma_\Lambda(t)$ is strictly positive. This implies that the lower bound can be found by numerically searching for $\lambda'(K)$ in (4.14), and subsequently evaluating $LB(\Lambda) = c_2(T,K,\Lambda,\lambda'(K))$. □

Though the result of Lemma 1 is not restricted only to arithmetic Asian options, it is no longer applicable when there are negative correlations between the random variables and the conditioning variable. Therefore there is no closed-form solution for the lower bound when e.g.:

- we are using an arbitrary conditioning variable Z;
- we are considering a basket option with an arbitrary correlation structure between the assets.

In general one must therefore resort to (4.9) or (4.12), a problem Deelstra et al. [2004] also ran into. This involves a discontinuous integrand and as such is undesirable.

Luckily it turns out that even in the more general setup, the conditional expectation can only have a restricted number of forms. Though the derivations here will be based on the arithmetic Asian case, we stress that all results are fully general and apply to the situation where we have multiple underlyings as well. From (4.3) we can write the expectation of A(T), conditional upon a general Gaussian random variable Z, as:

$$f(z) \equiv \mathbb{E}[A(T) \mid Z = z] = \int_0^T \alpha(t) e^{\beta(t)\cdot z} dt \tag{4.15}$$

where $\alpha(t) \geq 0$ and $\beta(t) = \sigma_Z(t)/\sigma_Z^2$. We can write $Z = \int_0^T \eta(t) \ln S(t)\, dt$, where $\eta(t) \neq 0$ only if $\alpha(t) > 0$, and $\mathrm{Var}(Z) > 0$. It makes no sense for Z to have a component independent of all ln S(t) where $\alpha(t) > 0$. Let $\gamma(t) = \alpha(t)\beta(t)$. There are three possible situations:

1. $\gamma(t) \geq 0$: (4.15) will be monotone increasing;
2. $\gamma(t) \leq 0$: (4.15) will be monotone decreasing;
3. γ takes on both strictly positive and negative values.

The monotone increasing situation has been dealt with by Lemma 1. The monotone decreasing can be dealt with similarly. The situation in the third case is less clear. To analyse this assume that $\beta(t)$ is non-decreasing (otherwise we can rearrange the time-interval so that this holds true). Suppose s is the largest s for which $\beta(s) \leq 0$, so that for $0 \leq t \leq s$ we have $\beta(t) \leq 0$ (and $\beta(0) < 0$). For $s < t \leq T$ we have $\beta(t) > 0$. We now write:

$$f(z) = \int_0^s \alpha(t) e^{\beta(t) \cdot z} dt + \int_s^T \alpha(t) e^{\beta(t) \cdot z} dt \tag{4.16}$$

Given that γ takes on both strictly positive and negative values, f(z) is the sum of a monotone decreasing and a monotone increasing function. Furthermore, f(z) possesses a special structure for which we can prove that solving f(z) = K yields at most two solutions, making it particularly easy to calculate the lower bound numerically. Before showing this, we require the following lemma.

Lemma 2:
Consider f(z) in (4.15). In case there is a $t \in [0,s]$ for which $\gamma(t) < 0$, and there is a $t \in (s,T]$ for which $\gamma(t) > 0$, the first derivative of f(z) has only one zero.

Proof: See the appendix. □

In the case the lemma deals with, the conditional expectation f(z) has the property that:

$$\lim_{|z| \to \infty} f(z) = \infty \tag{4.17}$$

Since the derivative has only one zero (say z^*), z^* is the unique minimum of the function, so that f(z) will be decreasing until z^* and increasing afterwards. From a practical point of view it is perfectly valid to ask when this situation will occur. Given the general form of Z, assume there is a t such that $\eta(t)\mathrm{Cov}(Z, \ln S(t)) < 0$. Since $\mathrm{Cov}(Z,Z) = \mathrm{Var}(Z) > 0$, there must be an s such that $\eta(s)\mathrm{Cov}(Z, \ln S(s)) > 0$. Now, if we have either $\eta(u)$ nonnegative or nonpositive for all u, which is the case for the conditioning variables Λ_{FA} or Λ_{GA}[6], the existence of a negative covariance of Z with some ln S(t) implies the existence of a positive covariance of Z with some ln S(s), where both $\eta(t)$ and $\eta(s)$ are nonzero. We therefore may run into this situation when using Λ_{FA} or Λ_{GA} in a situation with multiple underlyings with some negative correlations between the underlyings. In this situation the conditional expectation will have the shape which is indicated below.

---

[6] For these conditioning variables we have $\eta(u) = \rho(u)$.

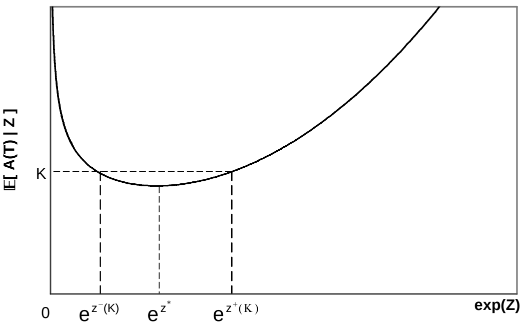

Graph 1: Possible shape of conditional expectation when negative correlations are present

Similarly, if η is not everywhere of the same sign, it is also possible to end up with a conditional expectation that has the shape indicated in Graph 1. Applying the knowledge gained from Lemmas 1 and 2 yields the following theorem, which provides a closed-form expression for the lower bound in case of an arbitrary conditioning variable Z.

Theorem 1:
Let f(z) be the conditional expectation of A(T) given Z, as defined in (4.15). If $\sigma_Z(t)$ is nonnegative or nonpositive for all t, let z'(K) denote the unique solution to f(z) = K. The lower bound can then be found as:

$$LB(Z) = \begin{cases} c_2\left(T, K, Z, z'(K)\right) & \forall_t\, \sigma_Z(t) \geq 0 \\ \mathbb{E}[A(T)] - K - c_2\left(T, K, Z, z'(K)\right) & \forall_t\, \sigma_Z(t) \leq 0 \end{cases} \tag{4.18}$$

If $\sigma_Z(t)$ is not of one sign, let z^* be the unique minimum of f(z). The equation f(z) = K has at most two solutions. If it has two solutions, we will denote the smallest as $z^-(K)$ and the largest as $z^+(K)$. The lower bound can then be found as:

$$LB(Z) = \begin{cases} \mathbb{E}[A(T)] - K & f(z^*) > K \\ c_2\left(T, K, Z, z^*\right) & f(z^*) = K \\ \mathbb{E}[A(T)] - K - c_2\left(T, K, Z, z^-(K)\right) + c_2\left(T, K, Z, z^+(K)\right) & f(z^*) < K \end{cases} \tag{4.19}$$

Proof: Follows immediately from Lemmas 1 and 2 and the results in section 4.1. □

Remark 2:
If we condition on a Λ we in addition know that we can restrict our searches for the solutions to the interval $(\infty, \lambda(K)]$, as $f\left(\lambda(K)\right) \geq K$ by definition of Λ. □

Theorem 1 allows us to write the lower bound for any option on a sum of correlated lognormal random variables as a closed-form expression, requiring at most three numerical searches. As the lower bounds in the arithmetic Asian case are already very tight, we do not consider optimising over the conditioning variable, something Deelstra et al. do in the case of multiple underlyings. The closed-form expression derived here would certainly facilitate such an optimisation greatly.

5. Thompson's upper bound revisited

As the lower bound of Curran and Rogers and Shi was found to be a very accurate approximation of the arithmetic Asian option, subsequent research efforts have focused on upper bounds. As with the lower bounds, the first article to our knowledge which derives an upper bound for the value of an Asian option is again that of Vorst [1992]. It is found by adding the difference in forward value of the arithmetic average and the geometric average to the price of a geometric call. In Rogers and Shi [1995] an upper bound was derived via estimation of the error made in their lower bound. Nielsen and Sandmann [2003] sharpened this bound considerably when the conditioning variable is the geometric average. The analysis is however valid for any Λ which has the property that $\Lambda \geq \lambda(K)$ implies $A(T) \geq K$, as pointed out in Vanmaele et al. [2006]. We return to this upper bound in the next section.

Most other upper bounds can be reduced to an application of the following idea:

$$\mathbb{E}\left[(A(T)-K)^+\right] = \mathbb{E}\left[\left(\int_0^T (S(t)-Kf(t))\rho(t)\,dt\right)^+\right] \tag{5.1}$$

provided $\int_0^T f(t)\rho(t)\,dt = 1$, where f(t) may be stochastic. The upper bound is then derived as:

$$\begin{aligned}\mathbb{E}\left[\left(\int_0^T (S(t)-Kf(t))\rho(t)\,dt\right)^+\right] &\leq \mathbb{E}\left[\int_0^T ((S(t)-Kf(t))\rho(t))^+\,dt\right] \\ &= \int_0^T \mathbb{E}\left[(S(t)-Kf(t))^+\right]\rho(t)\,dt\end{aligned} \tag{5.2}$$

All upper bounds based on this form assume some function for f(t), be it deterministic or stochastic, and minimise the upper bound with respect to f(t). Note that the optimal[7] choice for f(t) is $1+\frac{1}{K}\left(S(t)-\int_0^T S(u)\rho(u)\,du\right)$, as this yields the same value as the Asian call. This does not help us, since we are searching for an upper bound which is easily calculable.

Thompson [1999a, 1999b] was the first to give this upper bound for stochastic f(t). In particular, he assumed a Gaussian form for f(t). It is this upper bound which will be the topic of the next few sections. Other upper bounds that appear in the literature are essentially either special cases of (5.1)-(5.2), or combine the above ideas with the conditioning techniques from the previous section. These bounds are documented in the following table.

Upper bound	**Idea**	**Paper**
Comonotonic upper bound (CUB)	Deterministic f(t)	Simon et al. [2000], Nielsen and Sandmann [2003]
Improved CUB (ICUB)	Condition on a Gaussian Z, and then apply the CUB	Dhaene et al. [2002], Vanmaele et al. [2006]
Partially exact CUB (PECUB)	Use Curran's idea and apply the ICUB to the c_1-part	Vanmaele et al. [2006]

Table 1: Various related upper bounds in the literature based on (5.1)-(5.2)

---

[7] The author would like to thank Antoon Pelsser for pointing this out. This is similar to a result from importance sampling, where a zero variance estimator can be derived if the value of the integral is known.

Though the results in Simon et al. [2000], Dhaene et al. [2002] and Vanmaele et al. [2006] rely heavily on comonotonicity theory, their upper bounds can be derived alternatively by using the calculus of variations, as Thompson does. Note that if we use the same conditioning variable Λ, we have PECUB(Λ) ≤ ICUB(Λ) ≤ CUB. As far as the ICUB goes, Vanmaele et al. found it to work best with $Z = W(T)$ as a conditioning variable, at least when applied to an arithmetic Asian option. We should mention that Nielsen and Sandmann [2003] also considered this upper bound when $\Lambda = \Lambda_{GA}$, and refer to it as $C_A^{u,G}$. However they did not succeed in finding an analytical expression for the optimal function f and resorted to numerical optimisation.

The following sections deal with Thompson's upper bound when f(t) has a general Gaussian form. Section 4.1 will address the issue of calculating the upper bound, given the parameters that enter f(t). Section 4.2 deals with the optimality conditions on f(t). The true optimal choice with respect to the deterministic part of f(t) can in general already not be found analytically, as this requires an analytic expression for the distribution of the sum of a lognormal and a normal random variable. The way Thompson approximates this is described in section 4.3, whereas section 4.4 discusses our proposition: a shifted lognormal approximation. Finally, section 4.5 considers a heuristic to approximate the optimal volatility parameter of f(t).

5.1. Calculation of the upper bound

In Thompson [1999a, 1999b] the following form is proposed for the strike-scaling f(t) which enters the upper bound in (5.1)-(5.2):

$$\begin{aligned} f(t) &= \mu(t) - \overline{\sigma} X(t) \\ X(t) &= \xi(t) \int_0^T W(u)\eta(u)\rho(u)\,du - W(t)\eta(t) \end{aligned} \tag{5.3}$$

The functions η, μ and ξ are deterministic, restricted to:

$$\int_0^T \mu(t)\rho(t)\,dt = \int_0^T \xi(t)\rho(t)\,dt = 1 \tag{5.4}$$

so that the restriction on f is satisfied. We will refer to this setup as Thompson's generalised Gaussian upper bound. It is slightly more general than the bound considered for Asian options in Thompson [1999a], although Thompson [1999b] suggests this setup for basket options. In Thompson [1999a] the scaled volatility $\overline{\sigma}$ was chosen equal to σ, and $\eta(t) = \xi(t) = 1$. Note that up to a different mean and scaling, f(t) is then equal to ln S(t) – ln G(T), hereby resembling the optimal choice. These restrictions will be relaxed in the following section, where we discuss Thompson's upper bound more thoroughly.

Given any choice of functions η, μ and ξ and the constant $\overline{\sigma}$, the upper bound can be written as a double integral. To evaluate the integral we must resort to numerical integration. We will now derive the formula for the upper bound. If we condition on W(t), we find that S(t) – Kf(t) is distributed as:

$$S(t) - K\mu(t) + K\overline{\sigma} \cdot \left(\mathbb{E}[X(t) \mid W(t)] + \sqrt{\mathrm{Var}(X(t) \mid W(t))} \cdot Z \right) \tag{5.5}$$

where $Z \sim N(0,1)$. Secondly, we know that the following holds:

$$\mathbb{E}[(a + bZ)^+] = aN\left(\tfrac{a}{b}\right) + b\,\varphi\left(\tfrac{a}{b}\right) \tag{5.6}$$

These are all the results we need to write down the expression for the upper bound. The expectation within the integral in (5.2) can be calculated as:

$$\begin{aligned}\mathbb{E}\left[(S(t)-Kf(t))^+\right] &= \int_{-\infty}^{\infty} \mathbb{E}\left[(S(t)-Kf(t))^+ \mid W(t)=\sqrt{t}\cdot z\right]\varphi(z)\,dz \\ &= \int_{-\infty}^{\infty} \left\{a(t,z)N\left(\tfrac{a(t,z)}{b(t,z)}\right)+b(t,z)\,\varphi\left(\tfrac{a(t.z)}{b(t,z)}\right)\right\}\varphi(z)\,dz\end{aligned} \tag{5.7}$$

where we defined:

$$\begin{aligned} a(t,z) &= S(0)e^{(\mu-\frac{1}{2}\sigma^2)t+\sigma\sqrt{t}\cdot z} - K\mu(t) + K\overline{\sigma}\,\mathbb{E}\left[X(t)\mid W(t)=\sqrt{t}\cdot z\right] \\ b(t,z) &= K\overline{\sigma}\sqrt{\operatorname{Var}(X(t)\mid W(t))}\end{aligned} \tag{5.8}$$

Simply substituting (5.7) into the expression for the upper bound in (5.2) yields the final formula for the upper bound:

$$\int_0^T\int_{-\infty}^{\infty} \left\{a(t,z)N\left(\tfrac{a(t.z)}{b(t,z)}\right)+b(t,z)\,\varphi\left(\tfrac{a(t.z)}{b(t,z)}\right)\right\}\varphi(z)\,\rho(t)\,dz\,dt \tag{5.9}$$

One must resort to numerical integration to evaluate this integral. In the next paragraph we consider the optimal choice for μ. From (5.7) it is clear that we can write Thompson's upper bound as a portfolio of delayed payment options, although we must create an artificial asset in order to achieve this. As we do not find a delayed payment structure in terms of options on the actual underlying itself, we do not pursue this further.

5.2. Optimality conditions

In the rest of this section we will take the functions η and ξ as given, and consider optimality conditions with respect to μ and the parameter $\overline{\sigma}$. Although it is possible to consider varying η and ξ, their optimality conditions are much more involved. The Lagrangian w.r.t. μ and $\overline{\sigma}$ is:

$$L(\lambda,\overline{\sigma},\{\mu(t)\}) = \int_0^T \mathbb{E}\left[(S(t)-Kf(t))^+\right]\rho(t)\,dt - \lambda\left(\int_0^T \mu(t)\rho(t)dt - 1\right) \tag{5.10}$$

The first order condition for the optimality of {μ(t)} follows from the calculus of variations, by means of the Euler-Lagrange equation. The condition can be rewritten as:

$$\mathbb{Q}(S(t)-Kf(t)\geq 0) = -\tfrac{\lambda}{K} \equiv \widetilde{\lambda} \tag{5.11}$$

for all t where ρ(t) is unequal to zero. Here $\widetilde{\lambda}$ is merely a constant. Defining Y(t) as:

$$Y(t) = S(t) + K\overline{\sigma}X(t) \tag{5.12}$$

we see that the condition in (5.11) can be rewritten as:

$$\mathbb{Q}(Y(t)\leq K\mu(t)) = 1-\widetilde{\lambda} \tag{5.13}$$

i.e. $K\mu(t)$ is equal to the $1-\tilde{\lambda}$ quantile of the probability distribution of $Y(t)$. Note that the random variable $Y(t)$ is equal to the sum of a lognormal and a normal random variable.

The optimal μ can be determined exactly when $\overline{\sigma}$ equals zero, for this we refer the reader to Thompson [1999a]. As mentioned earlier, this situation coincides with the CUB discussed in the previous section. For $\overline{\sigma}$ unequal to zero, μ is harder to determine analytically. This is due to the fact that there is no closed-form expression for the probability law of $Y(t)$. By approximating the probability law of $Y(t)$ we can obtain approximate solutions for the optimal choice of μ. In the next section we consider the particular approximation that Thompson chose, and improve upon it. The first order condition with respect to $\overline{\sigma}$ will be investigated at a later stage.

5.3. A shifted lognormal approximation

To circumvent the difficult probability law in condition (5.13), Thompson used the following first order approximation for $S(t)$:

$$S(t) = S(0)e^{(\mu-\frac{1}{2}\sigma^2)t+\sigma W(t)} \approx S(t)e^{(\mu-\frac{1}{2}\sigma^2)t}\big(1+\sigma W(t)\big) \tag{5.14}$$

which is valid for small values of $\sigma W(t)$. This approximation will typically only work well when $\sigma W(t)$ has a small variance (i.e. when $\sigma^2 t$ is small) or when the variance of the normally distributed component of $Y(t)$ is much larger than the variance of the lognormally distributed component. The approximation has the desired effect: the calculation of the quantile in (5.13) becomes easy. We refer the interested reader to Thompson [1999a] for the remaining details regarding the calculation of the upper bound.

Clearly, a better approximation for the probability law of $Y(t)$ will yield a tighter upper bound. We propose to approximate $Y(t)$ by a shifted lognormal random variable. A shifted lognormal random variable X can be written as follows:

$$X = \alpha + \exp(\mu + \sigma Z) \tag{5.15}$$

where $Z \sim N(0,1)$. This is a special case of the Johnson-I distribution, see Johnson [1949]. What we propose here is to fit a shifted lognormal random variable to the sum of a normal and a lognormal random variable by matching the first three moments. A reason for this is the following. Firstly, both the normal and the lognormal distributions are special cases of (5.15). When $\alpha = 0$ we have a lognormal distribution. To recognise the normal distribution, we reparametrise:

$$X = \beta - \tfrac{1}{q} + \tfrac{1}{q}\exp\big(q(\mu + \sigma Z)\big) \tag{5.16}$$

where q is restricted to be positive. This is also known as the q-model, considered e.g. by Khong-Huu [1999]. By letting q tend to zero, we find $\lim_{q \downarrow 0} X = \beta + \mu + \sigma Z$ so that X indeed approaches a normal random variable. Intuitively we may thus expect that the behaviour of X will be inbetween these two extremes.

Hill, Hill and Holder [1976] already considered how to fit a shifted lognormal random variable to an arbitrary distribution of which we know the first three moments. The parameters α, μ and σ can be determined analytically given these first three moments. The only condition required for it to be possible that a shifted lognormal random variable as in (5.15) is fit to an arbitrary distribution is that the distribution has a positive skewness or central third moment. This is proven in the following lemma.

Lemma 3:
A sum of a normal and a lognormal random variable has a positive third central moment.

Proof: See the appendix. □

We conclude that a shifted lognormal variable can be fitted to the sum of a normal and a lognormal random variable. Returning to the problem at hand, we suggest to approximate Y(t) as:

$$Y_{SLN}(t) = \alpha(t) + \exp\big(\nu(t) + \omega(t)\cdot Z\big) \tag{5.17}$$

The shift, mean and volatility functions α(t), ν(t) and ω(t) will be determined for each point where ρ(t) is unequal to zero, by equating the first three moments of (5.17) to those of the original random variable Y(t). The same reasoning as before shows that condition (5.13) implies:

$$\frac{\ln\big(K\mu_{SLN}(t) - \alpha(t)\big) - \nu(t)}{\omega(t)} = \gamma_{SLN} \tag{5.18}$$

in turn yielding the following expression for $\mu_{SLN}(t)$:

$$\mu_{SLN}(t) = \tfrac{1}{K}\big(\alpha(t) + \exp\big(\nu(t) + \gamma_{SLN}\,\omega(t)\big)\big) \tag{5.19}$$

Finally, the constant γ_{SLN} is again determined by the condition in (5.4), but must now be determined via a numerical search.

Concluding, we expect that approximating Y(t) by a shifted lognormal random variable will yield better results than Thompson's approximation, at approximately the same computational cost. Thompson's approximation may be preferred when the variance of the normally distributed component is much higher than that of the lognormal component of the sum. This will be the case for high strike values or low volatilities. We now return to the Lagrangian and consider the optimality conditions for the scaled volatility $\overline{\sigma}$.

5.4. Optimal value of $\overline{\sigma}$

As stated before, in Thompson [1999a] only the choice of $\overline{\sigma}$ equal to σ and η(x) = ξ(x) = 1 was considered. This already gave great results for the situations he considered. For long maturities, high volatilities and high strike values, the effects of different choices for $\overline{\sigma}$, η and ξ will be considerable, as we will see.

We saw in the previous paragraphs that an approximately optimal choice for μ, given $\overline{\sigma}$, η and ξ, can be determined with relative ease. The first order conditions for η and ξ are more difficult and would typically create a large optimisation problem. For example, when we have N fixings we will have to determine 2N optimal values, subject to one constraint. Finding the optimal values will take a long time when the amount of fixings is large. One way to get rid of this problem is to parameterise η and ξ as functions of a small set of parameters. For our purposes however it seems that by only varying $\overline{\sigma}$, in conjunction with the shifted lognormal approximation to the optimal function μ, already yields results which outperforms all known upper bounds from the literature.

From the Lagrangian in (5.10) we see that the first order condition with respect to $\overline{\sigma}$ equals:

$$\begin{aligned}\frac{\partial L\left(\lambda,\overline{\sigma},\{\mu(t)\}\right)}{\partial\overline{\sigma}} &= \int_0^T \mathbb{E}\left[\left(-K\tfrac{\partial}{\partial\overline{\sigma}}\mu(t)+KX(t)\right)\cdot 1_{[S(t)-Kf(t)\geq 0]}\right]\rho(t)\,dt \\ &= \int_0^T K\,\mathbb{E}\left[X(t)\cdot 1_{[S(t)-Kf(t)\geq 0]}\right]\rho(t)\,dt = 0\end{aligned} \tag{5.20}$$

The Lagrange multiplier does not appear in this equation, since the condition in (5.4) implies:

$$\int_0^T \tfrac{\partial}{\partial\overline{\sigma}}\mu(t)\rho(t)\,dt = 0 \tag{5.21}$$

This property is also used in (5.20), in conjunction with the fact that the probability of each option being in-the-money is equal to $\widetilde{\lambda}$. As we use an approximation to derive the optimal function for μ, but subsequently calculate the upper bound exactly, equation (5.21) will not represent the optimality condition for our approximate model.

In practice we find that the upper bound is quadratic in a wide range around the optimal value of the scaled volatility. Outside this range the upper bound increases approximately linearly in the scaled volatility, the further we stray from the optimal value. Using this observation we can find a fast and accurate alternative to numerical optimisation for determining the optimal value of $\overline{\sigma}$. We propose the following algorithm:

```
1. Calculate the upper bound using $\mu_{SLN}$ for three carefully chosen values of $\overline{\sigma}$
2. Fit a quadratic function in $\overline{\sigma}$ to these values
3. Determine the value of $\overline{\sigma}$ in which the upper bound attains its minimum
4. Recalculate the upper bound in the approximately optimal $\overline{\sigma}$
```

Algorithm 1: Calculation of the SLNQuad upper bound

One could of course skip the last step and determine the minimum value of the upper bound by directly substituting the optimal value of $\overline{\sigma}$ in the quadratic function, though this procedure is not guaranteed to yield an upper bound. For this reason we recalculate the upper bound in step 4. In the first step the calculations can be sped up considerably by calculating the upper bound with less accuracy than in the last step. We will refer to the resulting upper bound as SLNQuad. In the final section it will be shown that this approximation gives results which are very close to the optimal value.

6. Partially exact and bounded approximations

A much-heard criticism of many analytical approximations is that the size of the error is not known beforehand. The majority of approximations are based on approximating the probability law of the arithmetic average by a distribution that is analytically tractable. The parameters of the approximating distribution are usually determined by matching the first couple of moments to that of the arithmetic average. Via an Edgeworth expansion it can be shown that the error made by approximating the pdf tends to zero when the number of moments that are matched tends to infinity, see Jarrow and Rudd [1982]. When the number of moments that are matched is finite, the error made will have to be determined via a numerical analysis. The approximations of e.g.

Turnbull and Wakeman [1991], Levy [1992], Curran[8] [1994], Posner and Milevsky [1998] and Milevsky and Posner [1998] have this shortcoming.

Two approximations for which the size of the error can be estimated are those of Vorst [1992] and more recently Vyncke, Goovaerts and Dhaene [2004]. Vorst's approximation essentially approximates the arithmetic average as a constant plus the geometric average, where the constant is determined such that the first moment coincides with that of the arithmetic average. As such it can be seen to lie between his lower bound and his upper bound.

Vyncke et al. [2004] propose two analytical approximations which by construction lie between a lower and an upper bound. Their idea is to take a convex combination of the LB(Λ) lower bound and the CUB or the ICUB(Z) upper bound. Interestingly, they show that there is a convex combination such that the first two moments of the approximation and the arithmetic average coincide. Therefore, though their approximation is a two-moment matching approximation, its accuracy can be gauged by the difference between the upper bound and the lower bound.

Inspired by these approximations, we will in the next section consider sufficient conditions for an approximation to lie between LB(Λ) and a recent sharpening of Rogers and Shi's upper bound due to Nielsen and Sandmann [2003] and Vanmaele et al. [2006]. The conditions give lead to a class of approximations, which we will name the class of partially exact and bounded (PEB) approximations. The error of these approximations also goes to zero when the strike price tends to zero or to infinity. The proof of this property will lead us to believe that Curran's approximation in fact diverges when the strike price tends to infinity. We show that this is indeed the case, before considering elements of the PEB approximations in the last subsection.

6.1. The class of partially exact and bounded approximations

This section will deal with the construction of an approximation which lies between LB(Λ) and $\mathrm{UB}_1(\Lambda)$, a sharpening of Rogers and Shi's upper bound due to Nielsen and Sandmann [2003] and Vanmaele et al. [2006]. The upper bound is based on the following inequality:

$$\begin{aligned} 0 &\le \mathbb{E}[X^+] - \mathbb{E}[X]^+ = \tfrac{1}{2}\big(\mathbb{E}\big[\,|X|\,\big] - |\mathbb{E}[X]|\big) \\ &\le \tfrac{1}{2}\mathbb{E}\big[\,|X - \mathbb{E}[X]|\,\big] \le \tfrac{1}{2}\sqrt{\mathrm{Var}(X)} \end{aligned} \tag{6.1}$$

The error made in the lower bound LB(Λ) can then be bounded from above by:

$$\begin{aligned} 0 &\le c_A(T,K) - \mathrm{LB}(\Lambda) \\ &= \mathbb{E}\Big[\mathbb{E}[(A(T)-K)^+ 1_{[\Lambda<\lambda(K)]} \mid \Lambda] - \mathbb{E}[(A(T)-K)1_{[\Lambda<\lambda(K)]} \mid \Lambda]^+\Big] \\ &\le \tfrac{1}{2}\mathbb{E}\Big[\mathrm{Var}\big(A(T)1_{[\Lambda<\lambda(K)]} \mid \Lambda\big)^{1/2}\Big] \equiv \varepsilon_1(\Lambda) \end{aligned} \tag{6.2}$$

leading to the upper bound:

$$\mathrm{UB}_1(\Lambda) = \mathrm{LB}(\Lambda) + \varepsilon_1(\Lambda) \tag{6.3}$$

---

[8] Note that this is only true for Curran's so-called "sophisticated" approximation. Curran's "naïve" approximation coincides with the lower bound of Rogers and Shi and as such lies between this lower bound and any upper bound. Throughout we will refer to the "sophisticated" approximation as Curran's approximation.

A bound that is slightly less tight, though easier to evaluate, can be found by bounding (6.2) further. We do not consider this bound here however. Note that Rogers and Shi's upper bound corresponds to the limit of $UB_1(\Lambda)$ when $\lambda(K)$ tends to infinity. As a consequence their original upper bound is independent of the strike.

To end up with an approximation that lies between $LB(\Lambda)$ and $UB_1(\Lambda)$ we could of course take a convex combination of both. A similar approach using the PECUB upper bound is employed in Vanmaele, Deelstra and Liinev [2004] to arrive at a bounded approximation. We will not follow this route, but will employ Curran's idea of decomposing the option price into an exact part and a part that has to be approximated. Recall that the part that has to be approximated is:

$$c_1(T,K,\Lambda) = \mathbb{E}[(A(T)-K)^+ 1_{[\Lambda<\lambda(K)]}] \tag{6.4}$$

The idea is to approximate A(T) by Ã(T), which has an analytically tractable law:

$$\tilde{c}_1(T,K,\Lambda) = \mathbb{E}[(\tilde{A}(T)-K)^+ 1_{[\Lambda<\lambda(K)]}] \tag{6.5}$$

so that the final approximation[9] of the forward option price is:

$$\tilde{c}(T,K,\Lambda) \equiv \tilde{c}_1(T,K,\Lambda) + c_2(T,K,\Lambda) \tag{6.6}$$

The following theorem supplies a sufficient condition for the resulting forward option price to be bounded below by $LB(\Lambda)$.

Theorem 2:
A sufficient condition that the approximating random variable Ã(T) must satisfy in order for the resulting approximation in (6.7) to be greater than or equal to $LB(\Lambda)$ is:

$$\mathbb{E}[\tilde{A}(T) \mid \Lambda = \lambda] \geq \mathbb{E}[A(T) \mid \Lambda = \lambda] \tag{6.7}$$

for $\lambda \in (-\infty, \lambda(K))$, i.e. its conditional mean given Λ must be greater than or equal to that of the arithmetic average.

Proof:

The following set of equations constitutes the proof:

$$\begin{aligned} \tilde{c}_1(t,K,\Lambda) &= \int_{-\infty}^{\lambda(K)} \mathbb{E}[(\tilde{A}(T)-K)^+ \mid \Lambda = \lambda] dF_\Lambda(\lambda) \\ &\geq \int_{-\infty}^{\lambda(K)} (\mathbb{E}[\tilde{A}(T) \mid \Lambda = \lambda] - K)^+ dF_\Lambda(\lambda) \\ &\geq \int_{-\infty}^{\lambda(K)} (\mathbb{E}[A(T) \mid \Lambda = \lambda] - K)^+ dF_\Lambda(\lambda) \end{aligned} \tag{6.8}$$

The first equation is merely another representation of (6.5). The first inequality follows from Jensen's inequality, whereas the second is an application of the condition imposed in (6.7). The last expression is simply the lower bound constructed for c_1 in $LB(\Lambda)$. This concludes the proof. □

---

[9] Typically the approximation will have to be calculated via numerical integration.

Theorem 2 supplies a nice condition to ensure that the resulting approximation in (6.6) is bounded below by a sharp lower bound. Any approximating random variable satisfying (6.7) can be used. Note that Curran's approximation also satisfies this condition. There are however other problems with this approximation, as we will see in the next subsection. We now supply sufficient conditions for the approximation to lie between LB(Λ) and $UB_1(\Lambda)$.

Theorem 3:
If we impose the following two conditions on the approximation random variable $\tilde{A}(T)$:

$$\begin{aligned}
&\mathbb{E}[\tilde{A}(T) \mid \Lambda = \lambda] = \mathbb{E}[A(T) \mid \Lambda = \lambda] \\
&\mathrm{Var}[\tilde{A}(T) \mid \Lambda = \lambda] \leq \mathrm{Var}[A(T) \mid \Lambda = \lambda]
\end{aligned} \tag{6.9}$$

for $\lambda \in (-\infty, \lambda(K))$, the resulting approximation in (6.6) lies between LB(Λ) and $UB_1(\Lambda)$.

Proof:

The proof follows along the same lines as (6.2), which led to the construction of $UB_1(\Lambda)$:

$$\begin{aligned}
0 &\leq \tilde{c}(T,K,\Lambda) - LB(\Lambda) \\
&= \int_{-\infty}^{\lambda(K)} \left(\mathbb{E}[(\tilde{A}(T) - K)^{+} \mid \Lambda = \lambda] - (\mathbb{E}[A(T) \mid \Lambda = \lambda] - K)^{+} \right) dF_{\Lambda}(\lambda) \\
&= \int_{-\infty}^{\lambda(K)} \left(\mathbb{E}[(\tilde{A}(T) - K)^{+} \mid \Lambda = \lambda] - (\mathbb{E}[\tilde{A}(T) \mid \Lambda = \lambda] - K)^{+} \right) dF_{\Lambda}(\lambda) \\
&\leq \tfrac{1}{2} \int_{-\infty}^{\lambda(K)} \sqrt{\mathrm{Var}(\tilde{A}(T) \mid \Lambda = \lambda)}\; dF_{\Lambda}(\lambda) \\
&\leq \tfrac{1}{2} \int_{-\infty}^{\lambda(K)} \sqrt{\mathrm{Var}(A(T) \mid \Lambda = \lambda)}\; dF_{\Lambda}(\lambda) = \varepsilon_1(\Lambda)
\end{aligned} \tag{6.10}$$

It is clear that the first inequality holds, as the first condition in (6.9) implies via theorem 2 that the resulting approximation is greater than or equal to LB(Λ). The rest of the derivation is similar to the derivation in (6.2). It immediately follows that:

$$LB(\Lambda) \leq \tilde{c}(T,K,\Lambda) \leq UB_1(\Lambda) \tag{6.11}$$

which concludes the proof of theorem 3. □

We call the class of approximations for which condition (6.9) holds, the class of partially exact and bounded (PEB) approximations. Partially exact refers to the fact that they are constructed out of a part that is exact, whereas bounded refers to the fact that the approximation is known to lie between a sharp lower and an upper bound. By construction LB(Λ) is an element of this class of approximations. Other elements of this class will be considered in the last subsection of this section. In addition to being bounded, we find that the PEB class of approximations has some desirable properties which are stated in the following theorem.

Theorem 4:
The error made in the PEB approximations approaches zero when the strike price K approaches infinity. Furthermore, if K equals zero, the error is also equal to zero.

Proof:
Since the upper bound $UB_1(\Lambda)$ approaches a constant when K tends to infinity, we have to prove the first part in a different manner. Let us extend the approximation $\tilde{A}(T)$ in (6.9) to hold for all $\lambda \in \mathbb{R}$. For the approximating part of the PEB approximation we can then write:

$$\begin{aligned} \tilde{c}_1(T,K,\Lambda) &= \mathbb{E}[\left(\tilde{A}(T)-K\right)^+ 1_{[\Lambda<\lambda(K)]}] \\ &\le \mathbb{E}[\left(\tilde{A}(T)-K\right)^+] = \int_K^\infty \mathbb{Q}\left(\tilde{A}(T) > x\right)dx \\ &\le \int_K^\infty \mathbb{Q}\left(\left|\tilde{A}(T)\right| > x\right)dx = \int_K^\infty \mathbb{Q}\left(\tilde{A}(T)^2 > x^2\right)dx \\ &\le \int_K^\infty \tfrac{1}{x^2}\mathbb{E}[\tilde{A}(T)^2]dx = \tfrac{1}{K}\mathbb{E}[\tilde{A}(T)^2] \end{aligned} \tag{6.12}$$

where the Markov inequality is applied at the end. Since $\tilde{A}(T)$ has the same conditional expectation as the arithmetic average, and its conditional variance is smaller or equal to that of the arithmetic average, the same holds for the unconditional expectation and variance. Note that the first two moments of the arithmetic average are bounded; this in turn implies that the second unconditional moment of the approximating distribution is bounded. Finally, since the approximation $\tilde{c}_1(T,K,\Lambda) \ge 0$ and $\lim_{K\to\infty} \tfrac{1}{K}\mathbb{E}[\tilde{A}(T)^2] = 0$ we can invoke the sandwich theorem:

$$\lim_{K\to\infty} \tilde{c}_1(T,K,\Lambda) = 0 \tag{6.13}$$

The true value of the call option on the arithmetic average also approaches zero when the strike price approaches infinity. Since $0 \le c_2(T,K,\Lambda) \le c_A(T,K)$, we also have that c_2 approaches zero when K approaches infinity. Hence, the error made in the PEB approximation approaches zero when the strike price approaches infinity. We will now demonstrate that the error is zero when the strike price is zero. The Gaussian conditioning variable has the property that $\Lambda \ge \lambda(K)$ implies $A(T) \ge K$. The threshold value $\lambda(K)$ is the smallest value for which this holds. Since the arithmetic average is always larger than or equal to zero, we must have $\lambda(0) = -\infty$. Then:

$$\tilde{c}(T,0,\Lambda) = c_2(T,0,\Lambda) = \mathbb{E}[A(T)^+ 1_{[\Lambda>-\infty]}] = \mathbb{E}[A(T)] \tag{6.14}$$

which coincides with the true value of the call option when K = 0. □

Although this property seems rather trivial, it is not satisfied by all moment matching approximations. Let us first consider the case where the strike is zero. Since for distributions that can take on negative values $\mathbb{E}[\tilde{A}(T)^+] \ne \mathbb{E}[\tilde{A}(T)]$, the error does not equal zero in these cases. An example where this occurs is Posner and Milevsky's [1998] 4-moment matching approximation, which fits a Johnson Type-II distribution to the arithmetic average. The support of this distribution is $\mathbb{R}$, so that negative values can indeed be attained. Similar things can happen with their 3-moment matching approximation, which fits a shifted lognormal random variable to the arithmetic average.

6.2. Curran's approximation

Moment matching approximations that fit a distribution with a finite expectation and variance to the arithmetic average, satisfy the property that the error tends to zero when the strike price tends to infinity. Curran's approximation does not satisfy these properties. As Curran's approximation yields very accurate results for moderate strike prices, it is worthwhile to investigate why the approximation diverges for large strike values. In the last subsection we cure this undesirable property and obtain two very accurate PEB approximations.

From Curran [1994] it is clear that like our PEB approximations his approximation consists of an exact part (namely c_2 with the geometric average as the conditioning variable) and an approximating part, which satisfies:

Curran2M approximation:

$$\begin{aligned}
&\widetilde{A}(T) \mid G(T) = G(T) + \exp\left(\mu_{A|G}(K) + \sigma_{A|G}(K) \cdot Z\right) \\
&\mathbb{E}\left[\widetilde{A}(T) \mid G(T) = K\right] = \mathbb{E}\left[A(T) \mid G(T) = K\right] \\
&\operatorname{Var}\left(\widetilde{A}(T) \mid G(T) = K\right) = \operatorname{Var}\left(A(T) \mid G(T) = K\right)
\end{aligned} \tag{6.15}$$

for $G(T) \leq K$ and $Z \sim N(0,1)$. In words, the arithmetic average is approximated by a shifted lognormal random variable where the shift is equal to the geometric average. This ensures that the approximating quantity has the same natural restriction as the arithmetic average, namely that it is always larger than or equal to G(T). The parameters of the lognormal component, $\mu_{A|G}(K)$ and $\sigma_{A|G}(K)$, are constant and chosen such that the two moments of the approximating distribution coincide with that of the arithmetic average, conditional upon G(T) being equal to K. The rationale behind this assumption is that the largest contribution of c_1 to the option value will come from paths where the geometric mean is close to the strike price. Therefore, it is particularly important to fit the mean and variance of A(T) given G(T) close to the strike price. As such the approximation is a local two-moment fit at the strike price. To emphasise this fact we will refer to the approximation as the Curran2M approximation.

The Curran2M approximation is not a member of the PEB class, due to the fact that the conditional mean is not fit exactly in all points below the strike price. The fact that the approximation is strike-price dependent is what causes the approximation to diverge for large values of the strike price, which is proven in the next theorem.

Theorem 5:
The Curran2M approximation diverges when the strike tends to infinity.

Proof:
Consider the approximating part of the Curran2M approximation:

$$\begin{aligned}
&\mathbb{E}[\left(\widetilde{A}(T) - K\right)^{+} 1_{[G(T)<K]}] \geq \mathbb{E}[\left(\widetilde{A}(T) - K\right) 1_{[G(T)<K]}]^{+} \\
&= \left(\mathbb{E}[G(T) 1_{[G(T)<K]}] + \mathbb{Q}\left(G(T) < K\right) \cdot \left(\mathbb{E}[A(T) \mid G(T) = K] - 2K\right)\right)^{+}
\end{aligned} \tag{6.16}$$

where we used Jensen's inequality and property (6.15) to derive the last expression. For large values of K this lower bound will tend towards:

$$\left(\mathbb{E}[G(T)] + \mathbb{E}[A(T) \mid G(T) = K] - 2K\right)^+ \tag{6.17}$$

All that remains to be proven is that the conditional expectation of the arithmetic average at the strike price is superlinear in K, since then (6.17) clearly diverges for large K. To show that this indeed is the case, consider the covariance of ln S(t) with the geometric average G(T):

$$\sigma_{\Lambda_{GA}}(t) = \mathrm{Cov}\left(\ln S(t), \Lambda_{GA}\right) = \int_0^T \sigma^2 \min(t,u)\,\rho(u)\,du \tag{6.18}$$

This is non-decreasing in t. Suppose that s is the largest value in [0,T] for which ρ(s) > 0. The maximum covariance is then attained when t = s. In particular, using (4.3) we find:

$$\lim_{K\to\infty} \frac{\mathbb{E}\left[A(T) \mid G(T) = K\right]}{\mathbb{E}\left[S(s) \mid G(s) = K\right]\ \rho(s)} = 1 \tag{6.19}$$

The conditional expectation in the denominator is equal to:

$$\mathbb{E}\left[S(s) \mid G(s) = K\right] = \mathbb{E}\left[S(s) \mid G(s) = 1\right] \cdot K^{\sigma_{\Lambda_{GA}}(s)/\sigma^2_{\Lambda_{GA}}} \tag{6.20}$$

It is clear that the variance of Λ_{GA}, an average of the covariances $\sigma_{\Lambda_{GA}}(t)$, will be smaller than the largest covariance $\sigma_{\Lambda_{GA}}(s)$. Therefore (6.20), and also the conditional expectation in (6.17) is superlinear in K. This in turn causes the lower bound, and therefore also the Curran2M approximation itself, to diverge when the strike price tends to infinity. □

In practice we find that this is more noticeable for high volatility/long maturity situations. Though these situations may not be so important for options that just have an "Asian tail", it certainly is important when pricing rate-of-return guarantees which are embedded in many unit-linked insurance policies. These tend to have long maturities, and are in a Black-Scholes world equivalent to arithmetic Asian options, as has been shown by Schrager and Pelsser [2004]. Numerical examples of our finding will be shown in the last section. We finally note that the above analysis also holds for the approximations of Deelstra et al. [2004] that use constant coefficients.

6.3. Suggestions for partially exact and bounded approximations

From the manner in which the Curran2M approximation is constructed it should immediately be clear that we can cure the divergence issue by matching all two conditional moments exactly for those cases when Λ < λ(K). Hence it seems wise to consider the following approximations:

Conditional two-moment matching:

$$\tilde{A}(T) \mid \Lambda = \Psi(\Lambda) \quad \mathbb{E}[\tilde{A}(T) \mid \Lambda] = \mathbb{E}[A(T) \mid \Lambda] \quad \mathrm{Var}\left(\tilde{A}(T) \mid \Lambda\right) = \mathrm{Var}\left(A(T) \mid \Lambda\right) \tag{6.21}$$

for Λ < λ(K). Here $\Psi(\Lambda)$ is a non-negative random variable belonging to a distribution with at least two parameters, so that the first two moments can be fitted according to (6.21). Matching all first two conditional moments as in (6.21) implies that the first two unconditional moments are also matched. By doing this we may intuitively expect the resulting error to be smaller than when we only match the first two unconditional moments. Theorem 3 shows that a major benefit of doing this lies in the fact that we obtain a sharp analytical bound on the size of the error. Also, by virtue of theorem 4, we know that the error goes to zero when the strike price tends to zero or infinity. Regardless of the choice of Ψ, calculating $\tilde{c}_1$ is straightforward:

$$\begin{aligned}\tilde{c}_1(T,K,\Lambda) &= \mathbb{E}[\left(\tilde{A}(T)-K\right)^+ 1_{[\Lambda<\lambda(K)]}] \\ &= \int_{-\infty}^{\lambda(K)} \mathbb{E}\left[\left(\Psi(\lambda)-K\right)^+\right] dF_\Lambda(\lambda)\end{aligned} \tag{6.22}$$

In general we will have to resort to numerical integration to evaluate (6.22). We therefore need to calculate the conditional mean and variance of the arithmetic average in all points in which we evaluate the integrand. As Turnbull and Wakeman [1991] pointed out, the moments of the discretely sampled arithmetic average can be calculated very efficiently by using the property that the returns of the underlying are independent. Calculating the first m moments becomes a process of order O(mN), where N is the number of fixings. When we condition on Λ, this property is unfortunately forsaken; calculating the first m moments is now a process of order $O(N^m)$. Calculating the variance in each point of the algorithm may become too costly when N is very large. Theorem 3 provides a nice alternative: we can choose just to evaluate the variance in a number of points. As long as we ensure that the approximating conditional variance is lower than or equal to the true conditional variance, (6.10) is satisfied and we still have a PEB approximation. This also shows that if we can derive an analytical lower bound[10] on the variance of $A(T) \mid \Lambda$, we can use this expression as the fitted variance. In practice it may however be just as efficient to approximate the numerical integral in (6.22) with a smaller accuracy, requiring less evaluations of the conditional moments.

It is not immediately clear how to choose Ψ, as the arithmetic average, conditional upon Λ, remains to be a sum of correlated lognormal random variables. For low to moderate strike values, the conditional variance of A(T) given G(T), when $G(T) \leq K$, will be smaller than the unconditional variance for low to moderate strike values. As such we expect the choice of Ψ may be less of an issue than when we are only matching the conditional moments. Since Curran's approximation works quite well in practice, we adapt his idea to obtain two PEB approximations:

Curran2M+ approximation:

$$\Psi(G(T)) = G(T) + \exp\left(\mu_{A|G}(G(T)) + \sigma_{A|G}(G(T))\cdot Z\right) \tag{6.23}$$

Curran3M+ approximation:

$$\Psi(G(T)) = \alpha_{A|G}(G(T)) + \exp\left(\mu_{A|G}(G(T)) + \sigma_{A|G}(G(T))\cdot Z\right) \tag{6.24}$$

The Curran2M+ functional form retains the natural restriction that $A(T) \geq G(T)$. By imposing (6.21) we extend the Curran2M approximation from a local two-moment fit at the strike price to a global two-moment fit. We note this approximation was also considered in Deelstra et al. [2004]. The Curran3M+ approximation sacrifices the natural restriction that $A(T) \geq G(T)$. It may however still be an improvement upon the Curran2M+ approximation, as the conditional skewness is now also matched exactly. The Curran3M+ approximation can also be seen as an extension of the 3-

---

[10] We were not able to derive a sharp lower bound.

moment matching approximation of Posner and Milevsky [1998]. Here, and also in Hill, Hill and Holder [1976] it is shown how to determine the coefficients α, μ and σ.

Both approximations ensure that the expectation within the integrand in (6.22) will be a Black-Scholes like expression, so that the evaluation of the integral is fairly straightforward. Note that when the amount of observations in the average is large, the extra computational effort required in the Curran3M+ approximation will be large, as we need to evaluate the third moment, a process of order $O(N^3)$, in each integration point.

The PEB approximations are found to be quite robust with regard to the specification of the conditioning variable and the approximation we use for the conditional law of the arithmetic average. Deelstra et al. [2004] have already demonstrated the first point, as they considered the Curran2M+ and Curran2M approximations for basket options. Using either Λ_{FA} or Λ_{GA} did not affect their results much. For the second point we mention that the Curran2M+ and 3M+ approximations seemed to yield the best results over a number of other choices we considered, though results for other distributions were close. To demonstrate the robustness with regard to the specification of the conditional law we include the 2M+Uniform approximation in our numerical results. This is obtained by replacing the lognormal random variable in (6.24) with a uniform random variable.

7. Numerical results and conclusions

In this final section we compare the new upper bound and the partially exact and bounded approximations to other bounds and approximations we mentioned in this document. We will focus on equally weighted, discretely sampled arithmetic Eurasian fixed strike calls. The stock underlying the arithmetic average will be dividend protected. In the first paragraph we briefly give implementation details for the bounds and approximations we reviewed, and compare calculation times. The next paragraph compares SLNQuad, and other choices one can make in Thompson's generalised Gaussian upper bound framework, to all other upper bounds. In the third paragraph we compare the performance of SLNQuad and the PEB approximations suggested in the previous section (Curran2M+ and Curran3M+) to the lower bounds and approximations we mentioned in sections 4 and 6 of this document. Several Greeks are calculated in order to assess whether the new bounds or approximations can be used to determine hedging positions. We end the paper with conclusions and recommendations.

Before starting the comparison, it is worthwhile to briefly review the outcomes of other studies. This will show us on what kind of examples we should focus our attention. The two studies we look at are those of Turnbull and Wakeman [1991] and those of Nielsen and Sandmann [2003]. Their examples also feature in other papers, e.g. in Vanmaele et al. [2006], which uses the examples from both studies to assess the performance of their upper bounds. In Table 2 on the next page the details are given about the examples in these studies. The first six columns describe the settings of the parameters for each of the examples. As before, N indicates the number of samplings in the average and r is the riskfree rate. The next three columns supply the spread between the smallest upper bound and the largest lower bound, for the at-the-money point. This is the point where the price of a call is equal to that of a put, i.e. when the strike is equal to the forward price of the arithmetic average. We measure the spread in this point as the at-the-money point is where many approximations already display large errors.

The upper bounds included in this comparison were UB_1, PECUB (with both Λ_{FA} and Λ_{GA} as conditioning variables), Thompson's original upper bound, SLNQuad and ICUB (with W(T) as the conditioning variable). The lower bounds we considered are $LB(\Lambda_{FA})$ and $LB(\Lambda_{GA})$. In the remainder of this document we will often measure the error made by an approximation or bound as a pricing error, which we define as:

Table 2: Spread between min. upper bound and max. lower bound, pricing error in Levy approximation for various studies

Studies	**Averaging frequency**	**Maturity**	**N**	**r**	**σ**	**Spread in bp. Between min. UB and max. LB (Relative error)**			**Pricing error (bp) in Levy approximation (Relative error)**
						Without Thompson and SLNQuad	*With Thompson, without SLNQuad*	*With Thompson, SLNQuad*	
Turnbull-Wakeman[11]	Daily	120d	30	8.62%	20%	0.57 (0.14%)	0.00 (0.00%)	0.00 (0.00%)	0.01 (0.00%)
	Daily	120d	30	8.62%	30%	1.27 (0.20%)	0.01 (0.00%)	0.01 (0.00%)	0.02 (0.00%)
	Daily	120d	30	8.62%	40%	2.22 (0.27%)	0.03 (0.00%)	0.03 (0.00%)	0.05 (0.01%)
Nielsen-Sandmann	Monthly	3y	36	4%	25%	20.63 (2.13%)	1.97 (0.20%)	1.72 (0.18%)	8.94 (0.92%)
	Monthly	10y	120	4%	25%	53.11 (3.40%)	14.53 (0.93%)	10.14 (0.65%)	47.96 (3.07%)
Our examples	Yearly	5y	5	5%	50%	134.59 (5.06%)	78.94 (2.97%)	34.20 (1.29%)	124.88 (4.70%)
	Yearly	30y	30	5%	25%	110.54 (5.78%)	73.74 (3.86%)	36.46 (1.91%)	135.69 (7.09%)

Table 3: Number of integration points required for an accuracy (in terms of pricing error) of $1\cdot10^{-3}$ bp; calculation times in brackets (in seconds, per 100 prices)

Maturity	**σ**	**UB$_1$**	**PECUB**	**ICUB**	**Thompson**	**SLNQuad**[12]	**Curran2M**	**Curran2M+**	**2M+Uniform**	**Curran3M+**
5y	50%	17 (0.08)	15 (0.38)	69 (0.77)	81 (0.29)	98 (0.64)	15 (0.03)	16 (0.09)	94 (0.33)	12 (0.16)
30y	25%	20 (2.00)	16 (2.86)	39 (3.22)	99 (2.01)	144 (4.62)	25 (0.17)	12 (1.10)	119 (10.34)	12 (14.21)

Table 4: Calculation times for other bounds and approximations (in seconds, per 100 prices)

Maturity	**σ**	**LB$_{GA}$**	**LB$_{FA}$/CUB**	**LB$_{FA}$/ICUB**	**2M**	**3M**	**4M**
5y	50%	0.02	0.05	0.81	0.01	0.01	0.02
30y	25%	0.12	0.81	4.18	0.03	0.04	0.06

---

[11] Note that the riskfree rate is equal to ln(1+0.09). Furthermore, the averaging occurs at the end of the contract, whereas in all other examples the averaging occurs throughout the life of the contract.

[12] To determine the optimal value of the scaled volatility we used 40 integration points.

$$\frac{e^{-rT}(\tilde{c}_A - c_A)}{S} \cdot 10000 \tag{7.1}$$

i.e. the discounted difference between the true forward price and the approximation (or bound), divided by the spot price, in basispoints (bp). It is a measure of the absolute error made. In the following we will mostly use this as a measure of error, as relative errors tend to explode for out-of-the-money calls. The spread between the upper bound and lower bound is simply the difference in their pricing error. In brackets the relative errors (in percentages) are supplied, where the Curran3M+ price acts as the true price. As we will see shortly, this price is very close to the true price; the error introduced hereby is thus rather small. The final column gives the pricing error of the Levy approximation.

The examples used by Turnbull and Wakeman have a relatively short maturity and a large number of averagings within this period. The volatility of the arithmetic average is therefore quite small, which is reflected in the spreads. There is little room for improvement left in this example, since even the simple (but widely used and quite effective in low volatility scenarios) Levy approximation prices the option with virtually no error. Nielsen and Sandmann already used two examples which are better test cases for the various bounds and approximations. Had they considered Thompson's upper bound, the error at the at-the-money point would have been reduced quite drastically. The SLNQuad bound performs even better. It is clear that the Levy approximation already displays larger errors for these examples.

For our own examples we chose to stress test the bounds and approximations even further than is done in Nielsen and Sandmann. In the first example we consider an Asian with yearly averaging and a maturity of 5 years. The volatility is an extremely high 50%. For our second example we lower the volatility to a moderate 25%, but increase the maturity to 30 years. The aforementioned unit-linked insurance policies, which contain embedded rate-of-return guarantees, tend to have long maturities. Although one is not advised to use a deterministic riskfree rate over such a long maturity (see e.g. Schrager and Pelsser [2004]), this example can still serve as a test case of how the bounds and approximations perform for long maturities.

7.1. Implementation details

Before letting the numerical results speak for themselves, we briefly discuss some implementation issues[13] for the various approximations and bounds. All calculations were carried out on a Pentium 4 2.66 GHz PC, in Visual Basic for Excel. Care was taken to optimise the code to allow for the fastest possible execution, given the programming language and processor. One should however not focus on the absolute time it took to calculate the various approximations and bounds, but on the relative difference in calculation times between the various methods.

To generate the true prices we numerically solved Večeř's PDE in (3.12) using a Crank-Nicolson finite-difference discretisation. Since in both examples the variance of the arithmetic average is quite high, due to either the high volatility or the long maturity, we found that a logarithmic transformation of the state space variable increased the numerical stability of the solutions. We transformed coordinates by defining $Y(t) = -\ln(1 - Z(t))$, where Z(t) is the value of the traded account divided by the stock price. Remember from section 3 that Z(t) can be related one-to-one to the strike of the Asian option under consideration. For the 50% volatility example we built a grid for strike prices ranging from 2.5 to 35000 (when S = 100). For the 25% volatility example the strike prices ranged from 0.5 to 55000. At the high strike prices we set the value of

---

[13] Naturally it goes too far to discuss all implementation details here. More details are available on request.

the option equal to zero; similarly, at low strike prices[14] we set the value equal to $e^{-rT}\left(\mathbb{E}[A(T)]-K\right)$. The ranges were chosen such that the error made was negligible.

For both examples we solved the PDE with three space/time points combinations; in the 5 year example the combinations were 500/5000, 1000/10000 and 2000/20000. We performed three Richardson extrapolations on the various outcomes. The largest spread (in bp) between the prices resulting from the second and third extrapolation was smaller than $5\cdot 10^{-4}$ bp for both examples. The first reduction mentioned in paragraph 3.2 indeed made the results more stable. The second reduction notably decreased the calculation time. Whereas the calculation time for the 1000/10000 space/time points combination was 47.41 seconds prior to using the reduction, it decreased to 38.08 seconds when using it, a reduction of about 20%.

We aim to compare the calculation times of the various methods in a fair manner. Since most approximations and bounds are based on a discretisation, we chose to find that discretisation where the pricing error (when compared to the true value for that approximation/bound) was just below $1\cdot 10^{-3}$ bp. In formulae, if the approximation uses a discretisation based on N points, we searched for the smallest N for which:

$$\frac{e^{-rT}\left|\tilde{c}_A(N)-\tilde{c}_A(\infty)\right|}{S}\cdot 10000 \le 1\cdot 10^{-3} \tag{7.2}$$

For the finite-difference solution of the PDE in the 50% volatility example, we found that performing one Richardson extrapolation on the 500/5000 and 1000/10000 space/time point combinations already yielded this accuracy. The combined calculation time for these two grids was 48.66 seconds. Calculation times were roughly the same for the 25% volatility example.

For the approximations and bounds that required numerical integration, we used Gauss-Legendre quadratures, see e.g. Press et al. [1996]. All upper bounds we considered require numerical integration. We must mention that the integral displayed for the ICUB and the PECUB in Vanmaele et al. [2006] was not numerically stable. Whereas their integrals were obtained by conditioning on a uniform random variable (i.e. $N(\frac{\Lambda-\mu_\Lambda}{\sigma_\Lambda})$ instead of Λ), we found it numerically much more stable to condition directly on a standardised version of Λ. The approximations requiring numerical integration are the recent approximations of Vyncke et al. [2004] (which we will refer to as the LB_{FA}/CUB and LB_{FA}/ICUB approximations), and the PEB approximations.

The search for the smallest number of integration points for which the error made in the discretisation was smaller than $1\cdot 10^{-3}$ bp was done over the set of strike prices we used to generate the results in the following sections. The results and corresponding calculation times are displayed in Table 3. For the UB_I and PECUB upper bounds the results displayed are only for the geometric average as conditioning variable. The calculation times and number of integration points required are however comparable when we use Λ_{FA} as the conditioning variable. In Table 4 calculation times are supplied for other bounds and approximations, namely the $LB(\Lambda_{GA})$ bound, which does not require any numerical integration, the approximations of Vyncke et al. and the moment matching approximations 2M, 3M and 4M. The 2M approximation is Levy's approximation; the other approximations stem from Posner and Milevsky's [1998] article. The 3M approximation fits a shifted lognormal distribution to the arithmetic average, whereas the 4M approximation fits a Johnson Type-II distribution to the first four moments of the average.

The performance of the various bounds and approximations will be the topic of the next two paragraphs. When discussing their performance, we will reflect on their calculation times.

---

[14] Note that this condition is usually not invoked when the second reduction of section 3.2 is used.

7.2. Comparison of all bounds

In this paragraph we will discuss the performance of various parameterisations of Thompson's generalised Gaussian upper bound (5.2)-(5.4), relative to the other upper bounds known in the literature. The comparison will be on basis of our first example from Table 2, where we consider pricing a discretely sampled arithmetic Eurasian fixed strike call. The average is based on five yearly observations, $\sigma = 50\%$, and r is equal to 5%. Finally, $S(0) = 100$. The following varieties of Thompson's upper bound are considered:

- *Thompson*: Thompson's [1999a] upper bound, with $\eta(t) = \xi(t) = 1$ and $\overline{\sigma} = \sigma$;
- *Thompson($\overline{\sigma}$)*: same as *Thompson*, except that we numerically optimise over $\overline{\sigma}$;
- *SLNQuad*: same as *Thompson*, using a shifted-lognormal approximation to approximate the optimal μ; the optimal $\overline{\sigma}$ is approximated via a quadratic[15] approximation to the upper bound.
- *SLN*: same as *SLNQuad*, except that the optimal $\overline{\sigma}$ was found via numerical optimisation;
- $\mu, \overline{\sigma}$: Optimal values for μ and $\overline{\sigma}$, $\xi(t) = 1$;
- μ, η: Optimal values for μ and η, $\xi(t) = 1$;
- μ, $\xi, \overline{\sigma}$: Optimal values for μ, ξ and $\overline{\sigma}$;
- μ, ξ, η: Optimal values for μ, ξ and η; optimal upper bound of Thompson's generalised form;

The results are displayed in Graph 2. Note that the second picture is an enlargement of the first, to make it easier to distinguish between the numerous parameterisations. The horizontal axis denotes the moneyness of the option, which is an affine transformation of the strike price, defined as $K/\mathbb{E}[A(T)]-1$, i.e. the quotient of the strike and the forward of the arithmetic average, minus 1. Negative moneyness corresponds with in-the-money Asian calls, positive moneyness with out-of-the-money calls. A moneyness of zero indicates that the option is at-the-money. The range of strikes over which the graph is plotted was chosen such that the probability of the arithmetic average exceeding either side of the range was approximately equal to 5%.

Firstly considering Thompson's original parameterisation, it is clear that choosing $\overline{\sigma}$ optimally tightens the upper bound considerably. Using the shifted-lognormal approximation (cf. section 5.3) to approximate the optimal function μ has an even bigger impact, in particular for large strikes. The SLNQuad upper bound, which uses a quadratic approximation to approximate the optimal scaled volatility, is very close to the SLN upper bound, which uses the true value of the optimal scaled volatility. Furthermore, the upper bound obtained by optimally choosing both the function μ and $\overline{\sigma}$, is very close to both SLN and SLNQuad. These facts combined make that the SLNQuad upper bound is a fast and astonishingly accurate approximation to the upper bound found by optimally choosing μ and $\overline{\sigma}$. As far as the other parameterisations go, it is clear that by optimally choosing η and/or ξ, we can obtain even tighter upper bounds. Although this is theoretically interesting, it is far from practical as it involves numerically optimising over a large set of parameters when the number of fixings is large.

Comparing the various formulations of Thompson's upper bound to the other bounds included in Graph 2, we find that SLNQuad already outperforms all known upper bounds. This remains true when varying the parameters and the contract specifications. Reflecting on the calculation times required (Table 3), we notice that although the SLNQuad takes longer to evaluate than most bounds, it is not that much slower than the ICUB, which for large strike values is the second-best upper bound (and in fact faster for the 50% volatility example). Concluding, the SLNQuad upper bound is considerably sharper than all currently known upper bounds, and

---

[15] We evaluated the upper bound in $\overline{\sigma} = 0.5\sigma$, 0.75σ and σ and fit a quadratic function through these values to approximate the optimal scaled volatility.

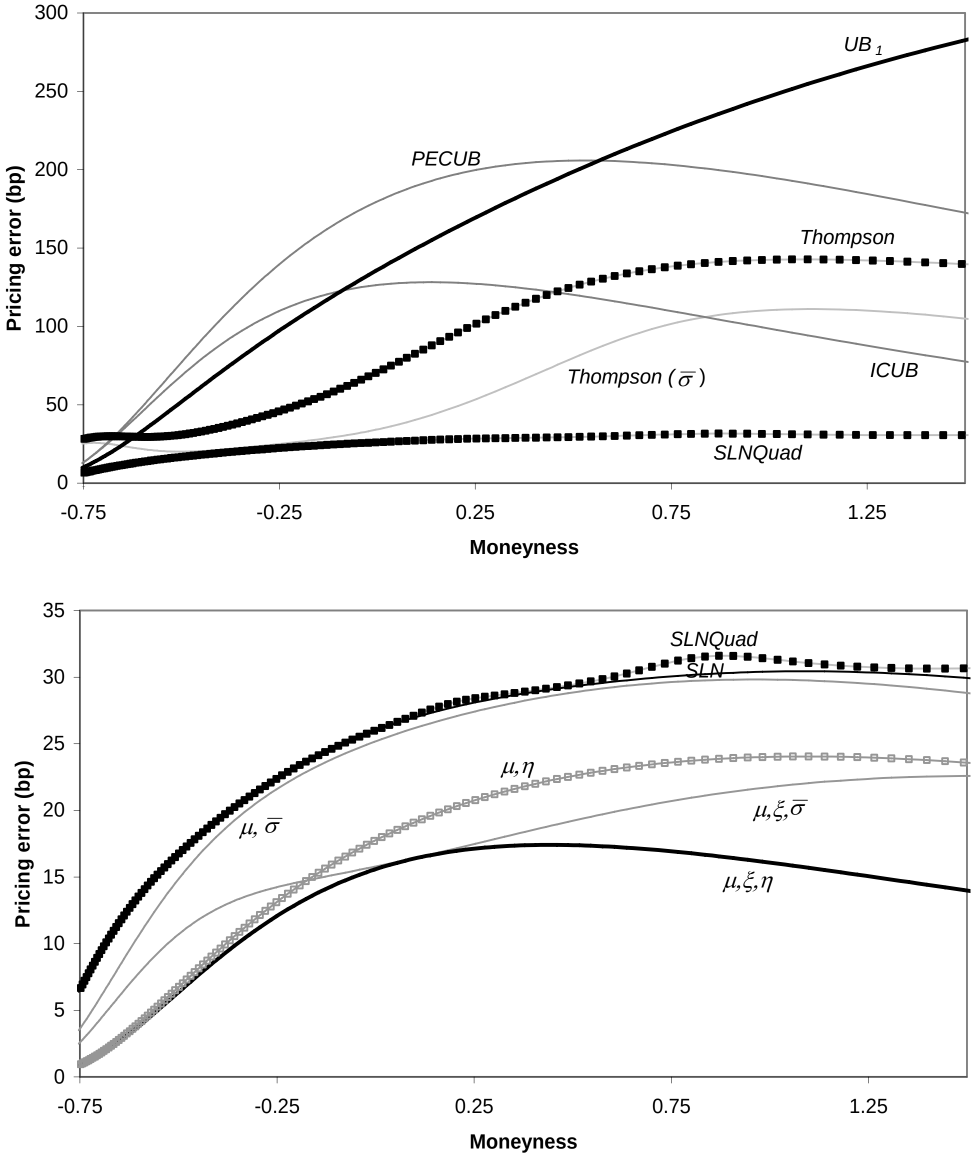

Graph 2: Parameterisations of Thompson's generalised Gaussian upper bound, $\sigma = 50\%$, 5y maturity

does not require a drastic increase in calculation time when compared to other upper bounds.

7.3. Comparison of all approximations

Having concluded that the SLNQuad upper bound outperforms all known upper bounds, and can be computed much faster than the even tighter variations of Thompson's generalised Gaussian upper bound, we now turn to testing the performance of the PEB approximations we introduced in section 6, and set them off against various other approximations. We also include the lower bound LB(Λ) in the comparison, with both Λ_{FA} and Λ_{GA} as conditioning variables. The approximations under consideration are Curran's approximation (Curran2M), the three PEB approximations (Curran2M+, 2M+Uniform and Curran3M+), as well as Levy's two-moment matching approximation (2M) and Posner and Milevsky's three and four-moment matching approximations (3M and 4M). Though we could have included many more approximations in this test, we chose to focus on the popular moment matching approximations. An approximation by Ju [2002], based on a Taylor expansion of the characteristic function of the arithmetic average, that is reported to work very well was found to be less accurate than Curran's and Rogers and Shi's lower bound in our examples, so that we excluded it from the comparison. Finally, we also include the SLNQuad upper bound in the numerical results, to set off its performance against the various approximations and bounds. The results are displayed on the following two pages.

As in section 7.2 the range of strikes over which the comparison is made was chosen such that the probability of A(T) exceeding either side of the interval was equal to 5%. Although this may seem like a very wide range of strike prices, we will show that the strike prices considered are not that extreme for the second example, where we consider an arithmetic average over 30 years. Suppose that we have a financial contract in which at each time t_0 (= 0) through t_{N-1} a fixed premium P is invested in the stock for the user; at the maturity date ($t_N = T$) the owner of the contract obtains the value of his investments at the maturity date. If this is lower than a guaranteed payment of K, the buyer of the contract receives this amount. The forward price of such a contract is equal to:

$$\mathbb{E}\left[\max\left(\sum_{i=0}^{N-1} P\tfrac{S(T)}{S(t_i)}, K\right)\right] = K + \mathbb{E}\left[\left(\sum_{i=0}^{N-1} P\tfrac{S(T)}{S(t_i)} - K\right)^+\right] \tag{7.3}$$

i.e. a guaranteed amount K, increased with a call option which gives the buyer the value of his investments above the guaranteed amount K. It is shown in Schrager and Pelsser [2004] that when we are in the Black-Scholes framework, the option in (7.3) is equivalent to an Asian option. Indeed, in this case we have:

$$\mathbb{E}\left[\left(\sum_{i=0}^{N-1} P\tfrac{S(T)}{S(t_i)} - K\right)^+\right] = \tfrac{NP}{S(0)}\mathbb{E}\left[\left(\sum_{i=1}^{N} \tfrac{1}{N}S(t_i) - \tfrac{S(0)}{NP}K\right)^+\right] \tag{7.4}$$

Suppose that we want to guarantee the user a (continuously compounded) return equal to g. We would then choose the strike such that $K = \sum_{i=0}^{N-1} Pe^{g(t_N - t_i)}$. In case of yearly averaging the strike of the Asian option in (7.4) would equal $\frac{S(0)e^{g}(e^{gN}-1)}{N(e^{g}-1)}$. Returning to our example with 30 years maturity, a guaranteed return of g = 10% would mean setting the strike of the Asian option equal to 668.5, a moneyness of 1.8. Although a rate-of-return guarantee of 10% is a bit on the high side, it is not completely unrealistic – the chosen guaranteed level will depend on the real-world behaviour of the underlying.

Returning to the numerical results, we first note that the SLNQuad upper bound and also the optimal choice of Thompson's generalised Gaussian upper bound are still not as tight as LB(Λ), which is just an application of something as seemingly straightforward as Jensen's inequality.

As for the approximations, let us first consider the Curran2M approximation. Although it performs very well for the 50% volatility example, the 25% volatility example shows that care should be taken when using this approximation. The effect we noticed in section 6.2, that the approximation diverges for large strike prices, is very noticeable in this example. As for the two, three and four-moment matching approximations, we notice that fitting more moments does not necessarily imply a smaller error. As anticipated earlier, we notice in both examples that indeed the error for the 4M approximation does not tend to zero when the strike price tends to zero. The second picture in both graph 3 and 4 is again an enlargement of the previous picture. Note that the LB_{FA}/CUB, LB_{FA}/ICUB, Curran2M+ and Curran3M+ approximations were not included in the

Graph 3: Performance of bounds and approximations, σ = 50%, 5y maturity, yearly averaging

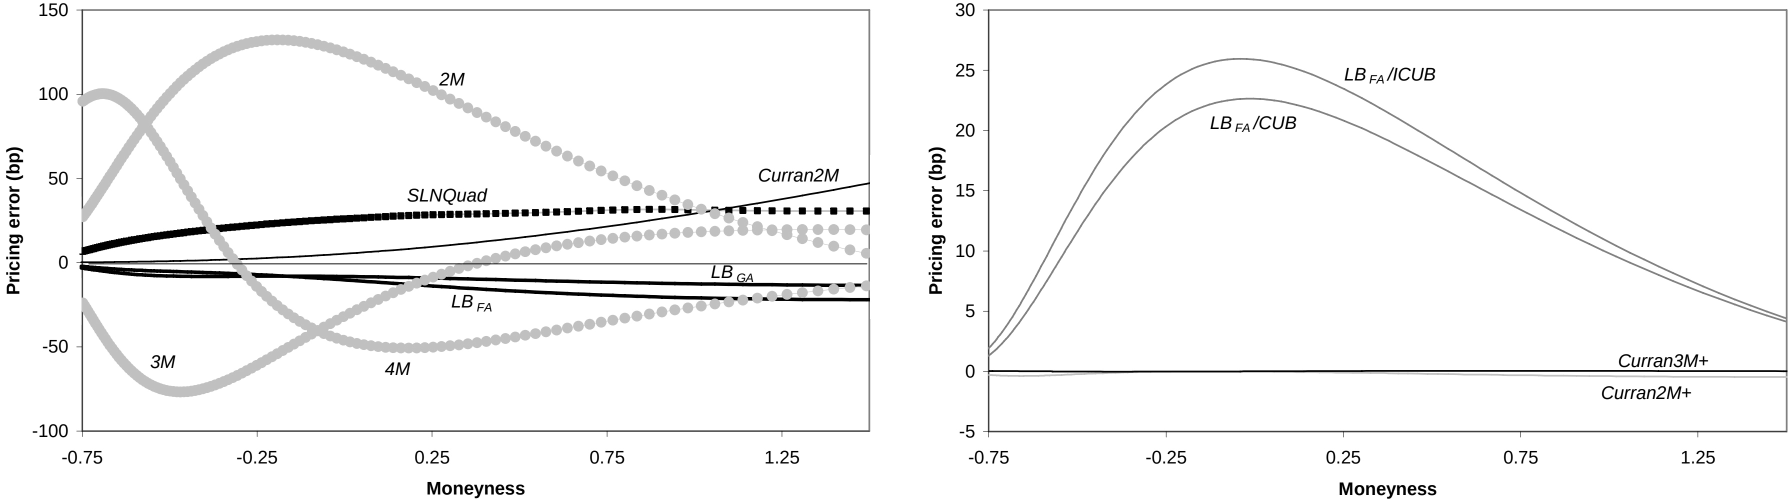

Table 5a: True price of Asian call and various approximations for three strike, σ = 50%, 5y maturity, yearly averaging

Strike	Moneyness	PDE	Curran2M+	2M+Uniform	Curran3M+	LB$_{FA}$/CUB	LB$_{GA}$	SLNQuad	Levy/2M
58.2370	-0.5	49.3944	49.3920	49.3967	49.3943	49.5103	49.3151	49.5617	50.3965
116.4741	0	26.5780	26.5778	26.5812	26.5781	26.8044	26.4962	26.8382	27.8268
174.7111	0.5	15.5342	15.5321	15.5436	15.5347	15.7073	15.4301	15.8286	16.3011

Table 5b: Pricing error (in bp)

Moneyness	Curran2M+	2M+Uniform	Curran3M+	LB$_{FA}$/CUB	LB$_{GA}$	SLNQuad	Levy/2M
-0.5	-0.24	0.23	-0.01	11.59	-7.93	16.73	100.21
0	-0.02	0.32	0.01	22.64	-8.18	26.02	124.88
0.5	-0.21	0.94	0.05	17.31	-10.41	29.44	76.69

Table 5c: Implied volatility (in bp)

Moneyness	Curran2M+	2M+Uniform	Curran3M+	LB$_{FA}$/CUB	LB$_{GA}$	SLNQuad	Levy/2M
-0.5	-1.10	1.06	-0.06	53.32	-36.70	76.83	449.36
0	-0.05	0.65	0.01	45.50	-16.43	52.29	251.98
0.5	-0.39	1.75	0.09	32.07	-19.29	54.53	142.00

Table 5d: True Delta, Gamma and Vega and those resulting from bounds and approximations; Vega is in percentages

	PDE			LB$_{GA}$			Curran3M+			SLNQuad		
Moneyness	*Delta*	*Gamma*	*Vega*	*Delta*	*Gamma*	*Vega*	*Delta*	*Gamma*	*Vega*	*Delta*	*Gamma*	*Vega*
-0.5	0.8164	0.0023	0.2166	0.8159	0.0023	0.2112	0.8164	0.0023	0.2166	0.8167	0.0022	0.2287
0	0.5733	0.0045	0.4981	0.5727	0.0045	0.4934	0.5733	0.0045	0.4980	0.5747	0.0045	0.5158
0.5	0.3876	0.0045	0.5396	0.3873	0.0045	0.5351	0.3876	0.0045	0.5396	0.3898	0.0045	0.5591

Graph 4: Performance of bounds and approximations, σ = 25%, 30y maturity, yearly averaging

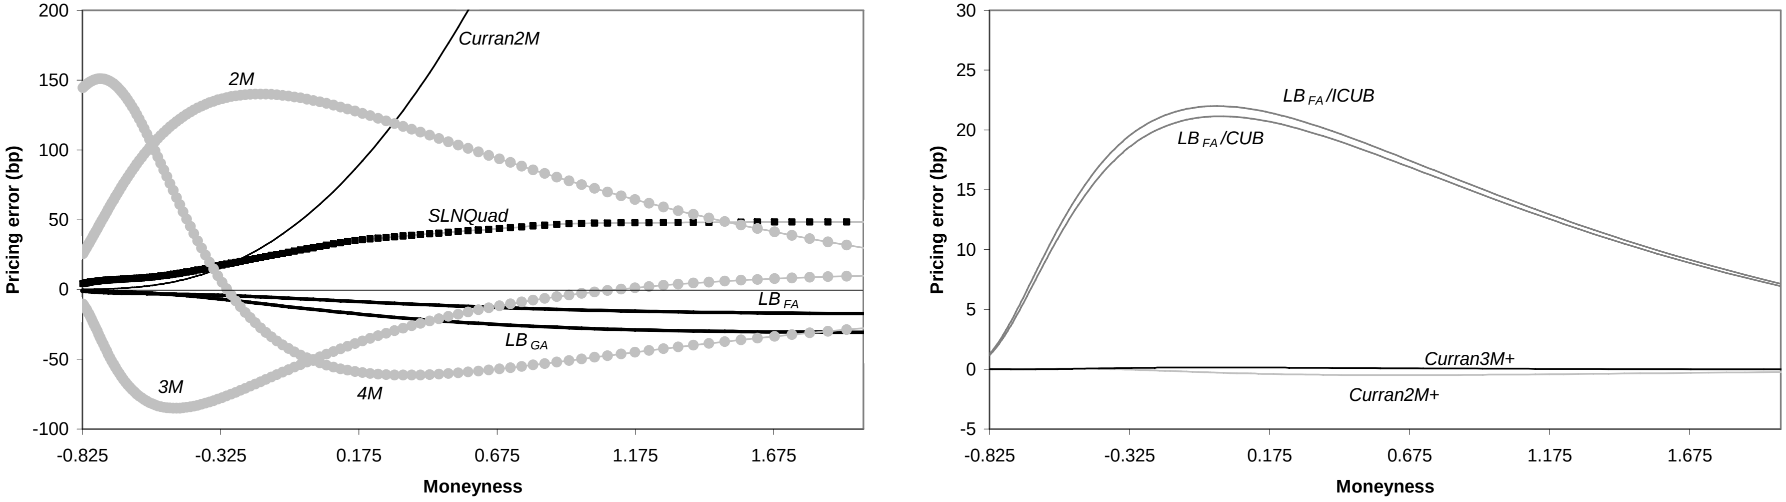

Table 6a: True price of Asian call and various approximations for three strike, σ = 25%, 30y maturity, yearly averaging

Strike	Moneyness	PDE	Curran2M+	2M+Uniform	Curran3M+	LB_{FA}/CUB	LB_{GA}	SLNQuad	Levy/2M
118.9819	-0.5	30.5153	30.5158	30.5246	30.5158	30.6554	30.4791	10.7417	31.6930
237.9638	0	19.1249	19.1220	19.1614	19.1263	19.3363	18.9845	29.4680	20.4832
356.9457	0.5	13.1168	13.1120	13.1676	13.1178	13.3010	12.8881	40.9490	14.1745

Table 6b: Pricing error (in bp)

Moneyness	Curran2M+	2M+Uniform	Curran3M+	LB_{FA}/CUB	LB_{GA}	SLNQuad	Levy/2M
-0.5	0.05	0.93	0.05	14.01	-3.62	10.74	117.77
0	-0.29	3.65	0.14	21.14	-14.04	29.47	135.83
0.5	-0.48	5.08	0.10	18.42	-22.87	40.95	105.77

Table 6c: Error in implied volatility (in bp)

Moneyness	Curran2M+	2M+Uniform	Curran3M+	LB_{FA}/CUB	LB_{GA}	SLNQuad	Levy/2M
-0.5	0.14	2.67	0.14	40.31	-10.45	30.93	334.94
0	-0.42	5.30	0.21	30.76	-20.36	42.91	199.75
0.5	-0.61	6.47	0.13	23.45	-29.07	52.18	135.13

Table 6d: True Delta, Gamma and Vega and those resulting from bounds and approximations; Vega is in percentages

	PDE			LB_{GA}			Curran3M+			SLNQuad		
Moneyness	*Delta*	*Gamma*	*Vega*	*Delta*	*Gamma*	*Vega*	*Delta*	*Gamma*	*Vega*	*Delta*	*Gamma*	*Vega*
-0.5	0.4655	0.0012	0.3467	0.4659	0.0012	0.3431	0.4655	0.0012	0.3467	0.4651	0.0013	0.3625
0	0.3503	0.0020	0.6884	0.3510	0.0020	0.6843	0.3503	0.0020	0.6886	0.3496	0.0020	0.7116
0.5	0.2663	0.0021	0.7858	0.2661	0.0021	0.7791	0.2663	0.0021	0.7860	0.2681	0.0020	0.8122

first picture, in order to not make the graph completely illegible. The results of the 2M+Uniform approximation were not included in any of the graphs, as its pricing error was quite close to that of the Curran2M+ approximation; it would therefore only clutter the graph.

The LB_{FA}/CUB approximation performs slightly better than the LB_{FA}/ICUB approximation. Comparing their calculation times in Table 2, we see that indeed the LB_{FA}/CUB should certainly be preferred, something which Vyncke et al. already concluded. These two approximations are quite a significant improvement over the 2M, 3M and 4M approximations, even though at heart they are also two-moment approximations.

With little extra calculation time however, the Curran2M+ method is virtually indistinguishable from a zero pricing error. For respectively the 50% and 25% volatility examples, its largest absolute pricing error (over the range of strike prices we considered) was 0.46 and 0.49 bp. The Curran3M+ method is even more accurate, with the largest absolute pricing errors equal to respectively 0.05 and 0.14 bp. This goes hand in hand with a significantly increased calculation time. As mentioned, calculating the third moment of the arithmetic average, conditional upon some Gaussian random variable Λ, is a process of the order $O(N^3)$, where N is the amount of fixings of the arithmetic average. This is particularly noticeable when N is large, as we see from Table 2.

As far as the robustness of the PEB approximations against the specification of the conditional law of the arithmetic average is concerned, we see that using a uniform distribution, which is totally different than the lognormal one, still yields a very reasonable approximation. The 2M+Uniform approximation still by far outperforms the LB_{FA}/ICUB and LB_{FA}/CUB approximations. Concluding, provided the conditioning variable is highly correlated with the arithmetic average, matching all first two conditional moments ensures that the approximation is very close to the true value of the Asian option, regardless of the distribution we choose to approximate the conditional law of the arithmetic average with. From table 1 we see that the calculation time of the 2M+Uniform approximation is significantly higher than that of the Curran2M+ approximation. This is caused by the nature of the uniform distribution: above a certain strike price, the option price is equal to zero, making the integrand in (6.22) discontinuous. Therefore we needed a much larger amount of integration points to obtain the same accuracy as the Curran2M+ approximation.

Finally, we note that UB_1 (see Graph 1) is the theoretical upper bound on the PEB approximations. For low to moderate volatilities and maturities this is one of the best upper bounds available, as is concluded in Vanmaele et al. For higher volatilities and maturities our examples show that it is outperformed by other upper bounds. It is clear that the error in the PEB approximations is much smaller than the theoretical error estimate we derived. Further research can hopefully derive a tighter upper bound for the PEB approximations.

In Tables 4a and 5a we supplied the current prices the best four approximations, the $LB(\Lambda_{GA})$ lower bound, the SLNQuad upper bound and Levy's approximation, so that the reader can judge the accuracy for herself. Pricing errors for these approximations and bounds are supplied in Tables 5b and 6b. An alternative measure of error is used in Tables 5c and 6c, where we show the error in implied volatility. The displayed results are deviations from the true volatility in bp. For example, the volatility we should have used in the PDE to obtain the same at-the-money price as the LB_{FA}/CUB approximation is equal to 25.4031% in the 25% volatility example. Whether this is an acceptable error will depend on the liquidity of the market.

In a trading environment it is not only important to have a good estimate of the price of a financial derivative. Calculating the correct hedging positions is equally, if not more, important. Although we have shown that the PEB approximations are extremely close to the true price, this in principle does not say anything about the accuracy of their Greeks. Nevertheless, as the approximations we propose are smooth and non-oscillating functions, it may be expected that the accuracy of the Greeks is comparable. Nielsen and Sandmann [2003] and Vanmaele et al. [2006]

supplied formulae for Delta, Gamma and Vega for their bounds. We will not do this here, as the Greeks of the PEB approximations and the SLNQuad upper bound are quite involved.

One of Nielsen and Sandmann's findings is that there are significant differences between the Vega of the LB(Λ_{GA}), the CUB and the UB$_2$(Λ_{GA}) bounds. As they supply the three Greeks for three volatility levels, the size of the Vega can be approximated numerically from their tables using e.g. a central difference. Doing this shows that their results for Vega cannot be correct, something which is also verified when recalculating their example. Furthermore, the Gammas they display are also incorrect. Recalculating their results shows that the three bounds they considered do in fact yield a reasonable estimate of Vega, invalidating their conclusions. This is also clear from our examples (Tables 5d and 6d). These examples show that the lower bound and the new SLNQuad upper bound already yield a good estimate on the size of the Delta, Gamma and Vega. As expected, the Curran3M+ approximation almost returns the true Greeks.

7.4. Conclusions and recommendations

In this paper we deal with the pricing of arithmetic Eurasian fixed strike options in the Black-Scholes framework. Since the underlying in the Black-Scholes framework is modelled as a geometric Brownian motion, the arithmetic average is a sum (or integral) of correlated lognormal random variables. As there is no closed-form expression for the probability of this sum, pricing these options is not trivial. Many research efforts to date have focused on deriving an accurate value for these types of options. We contribute to three areas of research.

Our first contribution is in the area of research that deals with the pricing of Asian options via the numerical solution of a PDE. We showed the link between the one-dimensional PDE of Rogers and Shi [1995], Andreasen [1999] and the recently derived PDEs by Hoogland and Neumann [2000a,b] and Večeř [2001]. For the latter PDEs, which Večeř found to be more numerically stable than the PDE of Rogers and Shi, we propose two reductions, which increase the numerical stability and reduce the calculation time.

Both the second and third contribution lie in the area of research that derives bounds on the value of an Asian option. We first showed how Rogers and Shi's lower bound can be evaluated in closed form for basket options, using at most three numerical searches. This replaces a numerical integral over a discontinuous integrand by a closed-form expression, which is practically relevant. Hereafter we considerably sharpened Thompson's [1999a,b] upper bound. This is important for the practically relevant case of options with long maturities. Numerical results show that the resulting upper bound (SLNQuad) is considerably tighter than recently introduced upper bounds in studies by Nielsen and Sandmann [2003] and Vanmaele et al. [2006].

Our final contribution deals with analytical approximations. We introduce a new class of analytical approximations, the partially exact and bounded (PEB) approximations, which can be proven to lie between Rogers and Shi's lower bound, and an upper bound recently derived by Nielsen and Sandmann. The error made in these approximations tends to zero when the strike price tends to zero or to infinity. We show that the latter property is violated by Curran's approximation; it in fact diverges when the strike price tends to infinity. Adapting Curran's idea to the class of PEB approximations however yields two approximations that almost return the true price, even for high volatilities and long maturities. Furthermore, the class of PEB approximations seems to be very robust to both the choice of approximating distribution for the conditional law of the arithmetic average, as well as to the conditioning variable, provided of course it is highly correlated with the arithmetic average. The approximations are found to outperform all of the current state-of-the-art bounds and approximations.

Bibliography

ANDREASEN, J. (1998). “The pricing of discretely sampled Asian and lookback options: a change of numeraire approach”, *Journal of Computational Finance*, vol. 1, no. 3, pp. 15-36.

BENHAMOU, E. (2002). “Fast Fourier Transform for discrete Asian options”, *Journal of Computational Finance*, vol. 6., no. 1, pp. 49-68.

CARVERHILL, A. AND L. CLEWLOW (1990). “Flexible convolution: valuing average rate (Asian) options”, *Risk*, vol. 4, no. 3, pp. 25-29.

CURRAN, M. (1992). “Beyond average intelligence”, *Risk Magazine*, vol. 5, no. 10.

CURRAN, M. (1994). “Valuing Asian and Portfolio Options by Conditioning on the Geometric Mean Price”, *Management Science*, vol. 40, no. 12, pp. 1705-1711.

DEELSTRA, G., LIINEV, J. AND M. VANMAELE (2004). “Pricing of arithmetic basket options by conditioning”, *Insurance: Mathematics and Economics*, vol. 34, no. 1, pp. 55-77.

DHAENE, J., DENUIT. M, GOOVAERTS, M., KAAS, R. AND D. VYNCKE (2002). “The concept of comonotonicity in actuarial science and finance: Theory”, *Insurance: Mathematics and Economics*, vol. 31, no. 1, pp. 3-33.

FALLOON, W. AND D. TURNER (1999). “The evolution of a market”, *Managing Energy Price Risk* (Second edition), Risk Publications.

GEMAN, H. AND M. YOR (1993). “Bessel processes, Asian options and perpetuities”, *Mathematical Finance*, vol. 4, no. 3, pp. 349-375.

GLASSERMAN, P., HEIDELBERGER, P. AND P. SHAHABUDDIN (1999). “Asymptotically optimal importance sampling and stratification for pricing path-dependent options”, *Mathematical Finance*, vol. 9, no. 2, pp. 117-152.

HENDERSON, V. AND R. WOJAKOWSKI (2002). “On the equivalence of floating and fixed-strike Asian options”, *Journal of Applied Probability*, vol. 39, no. 2, pp. 391-394.

HILL, I.D., HILL, R. AND R.L. HOLDER (1976). “Fitting Johnson Curves by Moments”, *Applied Statistics*, vol. 25, no. 2, pp. 180-189.

HOOGLAND, J.K. AND C.D.D. NEUMANN (2000a). “Asian and cash dividends: Exploiting symmetries in pricing theory”, CWI report MAS-R0019.

HOOGLAND, J.K. AND C.D.D. NEUMANN (2000b). “Tradable schemes”, CWI report MAS-R0024.

HULL, J.C. (2005). *Options, Futures, and Other Derivatives* (Sixth Edition), Prentice Hall.

HULL, J.C. AND A. WHITE (1993). “Efficient procedures for valuing European and American path-dependent options”, *Journal of Derivatives*, vol. 1, no. 1, pp. 21-31.

DEN ISEGER, P. AND E. OLDENKAMP (2006). “Pricing guaranteed return rate products and average price options”, Journal of Computational Finance, vol. 9, no. 3.

JAMSHIDIAN, F. (1989). "An exact bond option formula", *Journal of Finance*, vol. 44, no. 1, pp. 205-209.

JARROW, R. AND A. RUDD (1982). "Approximate option valuation for arbitrary stochastic processes", *Journal of Financial Economics*, vol. 10, no. 3, pp. 347-369.

JOHNSON, N.L.. (1949). "Systems of frequency curves generated by methods of translation", *Biometrika* vol. 36, pp. 149-176.

JU, N. (2002). "Pricing Asian and basket options via Taylor expansion", *Journal of Computational Finance*, Spring 2002.

KEMNA, A.G.Z. AND A.C.F. VORST (1990). "A pricing method for options based on average asset values", *Journal of Banking and Finance*, vol. 14, pp. 113-129.

KHUONG-HUU, P. (1999). "Swaptions with a smile", *Risk*, vol. 12, no. 8, pp. 107-111.

KLASSEN, T.R. (2001). "Simple, fast, and flexible pricing of Asian options", *Journal of Computational Finance*, vol. 4, no. 3, pp. 89-124.

LEVY, E. (1992). "Pricing European average rate currency options", *Journal of International Money and Finance*, vol. 11, pp. 474-491.

MILEVSKY, M.A. AND S.E. POSNER (1998). "Asian options, the sum of lognormals, and the reciprocal gamma distribution", *Journal of Financial and Quantitative Analysis*, vol. 33, no. 3, pp. 409-422.

NIELSEN, J.A. AND K. SANDMANN (2003). "Pricing bounds on Asian options", *Journal of Financial and Quantitative Analysis*, vol. 38, no. 2, pp. 449-473.

POSNER. S.E. AND M.A. MILEVSKY (1998). "Valuing exotic options by approximating the spd with higher moments", *Journal of Financial Engineering*, vol. 7, no. 2, pp. 109-125.

PRESS, W.H., TEUKOLSKY, S.A., VETTERLING, W.T. AND B.P. FLANNERY (1996). *Numerical Recipes in C* (Second Edition), Cambridge University Press.

ROGERS, L.C.G. AND Z. SHI (1995). "The value of an Asian option", *Journal of Applied Probability*, no. 32, pp. 1077-1088.

SCHRAGER, D.F. AND A.A.J. PELSSER (2004). "Pricing rate of return guarantees in regular premium unit linked insurance", *Insurance: Mathematics and Economics*, vol. 35, no. 2, pp. 369-398.

SIMON, S., GOOVAERTS, M.J. AND J. DHAENE (2000). "An easy computable upper bound for the price of an arithmetic Asian option", *Insurance: Mathematics and Economics*, vol. 26, no. 2-3, pp. 175-184.

THOMPSON, G.W.P. (1999a). "Fast narrow bounds on the value of Asian options", working paper, University of Cambridge.

THOMPSON, G.W.P. (1999b). "Topics in Mathematical Finance", PhD thesis, University of Cambridge.

TURNBULL, S.M. AND L.M. WAKEMAN (1991). "Quick algorithm for pricing European average options", *Journal of Financial and Quantitative Analysis*, vol. 26, no. 3, pp. 377-389.

VANMAELE, M., DEELSTRA, G. AND J. LIINEV (2004). "Approximation of stop-loss premiums involving sums of lognormals by conditioning on two variables", *Insurance: Mathematics and Economics*, vol. 35, no. 2, pp. 343-367.

VANMAELE, M., DEELSTRA, G., LIINEV, J., DHAENE, J. AND M.J. GOOVAERTS (2006). "Bounds for the price of discretely sampled arithmetic Asian options", *Journal of Computational and Applied Mathematics*, vol. 185, no. 1, pp. 51-90.

VEČEŘ, J. (2001). "A new PDE approach for pricing arithmetic average Asian options", *Journal of Computational Finance*, vol. 4, no. 4, pp. 105-113.

VORST, T. (1992). "Prices and Hedge Ratios of Average Exchange Rate Options", *International Review of Financial Analysis*, vol. 1, no. 3, pp. 179-193.

VYNCKE, D., GOOVAERTS, M.J. AND J. DHAENE (2004). "An accurate analytical approximation for the price of a European-style arithmetic Asian option", *Finance*, vol. 25, HS, pp. 121-139.

WILMOTT, P. (2006). *Paul Wilmott on Quantitative Finance* (Second Edition), John Wiley & Sons.

ZVAN, R., FORSYTH, P. AND K. VETZAL (1997/98). "Robust numerical methods for PDE models of Asian options", *Journal of Computational Finance*, vol. 1, no. 2, pp. 39-78.

Appendix – Proofs

Lemma 1:
Consider f(z) in (4.15). In case there is a $t \in [0,s]$ for which $\gamma(t) < 0$, and there is a $t \in (s,T]$ for which $\gamma(t) > 0$, the first derivative of f(z) has only one zero.

Proof:
The first derivative of f(z) can be written as:

$$f'(z) = \int_0^s \gamma(t)e^{\beta(t)\cdot z}dt + \int_s^T \gamma(t)e^{\beta(t)\cdot z}dt \tag{A.1}$$

where $\gamma(t) \leq 0$ for $0 \leq t \leq s$ and $\gamma(t) > 0$ for $s < t \leq T$. There is a $t \in [0,s]$ so that $\gamma(t) < 0$. Hence we have $\lim_{z\to-\infty} f'(z) = -\infty$ and $\lim_{z\to\infty} f'(z) \geq 0$. There has to be at least one zero. Now write:

$$f'(z) = 0 \Leftrightarrow -e^{\beta(s)\cdot z}\int_0^s \gamma(t)e^{(\beta(t)-\beta(s))\cdot z}dt = e^{\beta(s)\cdot z}\int_s^T \gamma(t)e^{(\beta(t)-\beta(s))\cdot z}dt \tag{A.2}$$

After dividing $e^{\beta(s)\cdot z}$ out, let us define the function on the left-hand side as g(z) and the function on the right-hand side as h(z). It is easy to verify that $\lim_{z\to-\infty} h(z) = 0$ and $\lim_{z\to\infty} h(z) = \infty$. Furthermore, since $\beta(t)$ - $\beta(s) > 0$ on (s,T], h'(z) will be strictly positive for $z > 0$. For g(z) we consider two situations. If $\beta(t) = \beta(s)$ for all $t \in [0,s]$ where $\gamma(t) < 0$, then g(z) is constant and larger than zero, and h(z) will intersect g(z) exactly once. Otherwise we have $\lim_{z\to-\infty} g(z) = \infty$ and $\lim_{z\to\infty} g(z) < \infty$. Due to g'(z) being negative we again only have one solution to g(z) = h(z). This in turn implies f'(z) has only one zero. □

Lemma 3:
A sum of a normal and a lognormal random variable has a positive third central moment.

Proof:
Let $X = \mu_1 + \sigma_1 Z_1 + \exp(\mu_2 + \sigma_2 Z_2)$ with Z_1 and Z_2 standard normal random variates with correlation ρ. Tedious calculations show that its third central moment equals:

$$\mathbb{E}\left[(X - \mathbb{E}[X])^3\right] = F^3\left(e^{3\sigma_2^2} - 3e^{\sigma_2^2} + 2\right) + 6\rho\sigma_1\sigma_2F^2\left(e^{\sigma_2^2} - 1\right) + 3F\rho^2\sigma_1^2\sigma_2^2 \tag{A.3}$$

where $F = e^{\mu_2 + \frac{1}{2}\sigma_2^2}$, the expectation of the lognormally distributed part. As this is a quadratic equation in ρ, we find that the minimum is attained for a value of ρ equal to $-\sigma_1^{-1}\sigma_2^{-1}F(e^{\sigma_2^2} - 1)$. Substituting this in (A.3) yields:

$$\mathbb{E}\left[(X - \mathbb{E}[X])^3\right] \geq F^3(e^{\sigma_2^2} - 1)^3 \tag{A.4}$$

which is positive for all possible parameter values. □