

Corvinus University of Budapest
Faculty of Economics

Eötvös Lóránd University
Faculty of Science



Pricing Asian Options

A comparison of Laplace transform and Monte Carlo
simulation methods in the Black-Scholes framework

Author: Ákos Horváth

MSc in Actuarial and Financial Mathematics

Quantitative Finance Major

2013

Supervisor: Dr. Péter Medvegyev

Contents

1	Introduction	1
2	Valuation methodology	4
2.1	The Black-Scholes model	4
2.2	Risk-neutral valuation	5
2.2.1	Analytical methods	8
2.2.2	Numerical methods	8
2.2.3	Simulation methods	9
3	Pricing Asian options	12
3.1	Option types	12
3.2	Asian options	14
3.2.1	Geometric-average Asian options	15
3.2.2	Arithmetic-average Asian options	17
3.3	The Laplace transform method	21
3.3.1	Derivation of the analytical formula	21
3.3.2	The problem of inversion	24
3.4	The Monte Carlo simulation method	28
3.4.1	Basic method	29
3.4.2	Variance reduction by antithetic variates	30
3.4.3	Variance reduction by control variates	32
3.4.4	Implementation	34
4	Empirical analysis	35
4.1	Analytical results	35
4.1.1	Weakness of the Laplace transform approach	36
4.1.2	Asian option price tables	38
4.2	Simulation results	39
4.2.1	Discretization error	40
4.2.2	Simulation error	41
4.3	Comparison of methods	44
5	Conclusion	46

Appendices

List of Figures

1	Notional amounts of equity options globally (USD billion)	13
2	Comparison of analytical algorithms	35
3	Comparison of simulation methods ($m = 10$)	42
4	Correlation of proxy variates in variance reduction methods	43
5	Comparison of calculation times of different methods ($m = 10$)	45
6	Value of the Laplace transform using the Euler algorithm for $n = 20$.	52
7	Value of the Laplace transform using the Talbot algorithm for $M = 50$	53
8	Absolute error of Monte Carlo algorithms for different levels of m . .	54
9	Comparison of calculation times of different methods ($m = 10$)	55
10	Optimized memory management of simulation algorithms	55

List of Tables

1	Limits for the variance reduction method ($m = 10$)	44
2	Optimal method selection rules for pricing Asian options	45

Acknowledgement

Hereby, I would like to thank Dr. Péter Medvegyev for his kind counselling and assistance in the preparation of this work.

1 Introduction

What is the **value of a financial asset**? Is it the amount of money that you have to pay to acquire it? Or is it rather the amount of money for which you can sell it? According to the Financial Accounting Standards Board (FASB), which is responsible for the maintenance of the U.S. Generally Accepted Accounting Principles (GAAP), “fair value is the price that would be received to sell an asset or paid to transfer a liability in an orderly transaction between market participants at the measurement date”. (FASB, 2006, p. 2.) From this definition it follows that the (fair) value of an asset could be determined from orderly transactions. All we have to do then is simply check the price at which a recent transaction took place and we are done with valuation... as it turns out, in most cases this is not that simple. First of all, what if there are more transactions at different prices? Which one should we pick? What if the last but one transaction was for a far greater quantity than the last one? Also, what if there are no observable transactions, since transactions are private or don’t even happen at all? Is the asset worthless, then? By no means!

As we can see, and the recent (and still ongoing) global financial turmoil indicates, the above question is rather ill-posed and answering it could be a challenging job. However, we may pose a proxy question, which is much easier to handle: “How do market participants determine the value of an asset?”, or in other words: “What determines the market value of an asset?” Practice in financial markets instructs to first determine the expected value of future cash flows (be it positive or negative) from the asset and then calculate the present value of this expectation using a ‘proper’ discount rate. This approach, which Medvegyev and Száz (2010, p. 104.) calls the **Expected Value – Present Value principle**, breaks up the initial problem in two subproblems, which are

1. determining the probability distribution of future cash flows;
2. determining the discount rate properly suited to the probability distribution determined in the first step.

Moreover, without going into further details (for which refer to section 2.2), we postulate that by using a so-called “**risk-neutral probability measure**”, it is sufficient to determine only the (risk-neutral) probability distribution of future cash flows, since the same interest rate will apply for every asset. Therefore, if we have models describing the (risk-neutral) probability distribution of future cash flows and the dynamics of the (risk-free) interest rate, we can value the asset in question.

A simple but arguably the best-known model describing the price movements of financial instruments is the **Black-Scholes model**. In their original model, which we will introduce in section 2.1, Black and Scholes (1973) assumed that stock prices follow geometric Brownian motion and the risk-free interest rate is constant. Although being too simple to realistically describe the empirical properties (i.e. stylized facts) of observable market price movements, and thus we should expect only an approximation of actual derivative prices from it, this model is relatively easy to understand and is mathematically tractable, which make it a perfect starting-point both for empirical analysis and for more advanced model-building. That is why I will also conduct my research in this framework.

In Black and Scholes's footsteps, authors have successfully derived formulas for various derivative prices, probably the most renowned of which being the pricing formula for the European-exercise vanilla call option by Merton (1973), which is an efficient closed-form solution of the corresponding option valuation problem. However, at the dawn of "financial engineering", as new derivatives appeared in financial markets with "**bespoke**" or "**structured**" payouts (basically meaning complex payouts), it turned out that solving the Black-Scholes valuation problem is not that simple (to say the least) in many cases, especially in those cases where either the payout or the maturity (or both) of the derivative depends on the path taken by the underlying product's price (i.e. the derivative is "path-dependent"). This is a rather unnerving thought considering the size of the global option industry and the increasing popularity of complex-payout OTC derivatives.

A notable example for the above problem is the **valuation of Asian options**, which is the main subject of this thesis. The chief motivation behind choosing this type of option in particular is twofold:

1. from a financial perspective, Asian options are popular hedging and risk reduction instruments;
2. from a mathematical-computational perspective, the valuation problem of Asian options is challenging and its analytical solution is rather spectacular.

Exploring this type of option in section Pricing Asian options, we will also discuss what exactly an Asian option is. For now, suffice it to say that the valuation of this option type is an intriguing problem.

Once we have chosen our model describing the risk-neutral dynamics of assets and interest rates, the next step is valuation. As we will see in section 2.2, there are different approaches to do this. Although there exist numerous **computation techniques**, at the end of the day they fall into either of three categories:

- analytical methods
- numerical methods
- simulation methods

After briefly showing the difference between these approaches, we will move on to the valuation problem of Asian options and present two semi-analytical and simulation techniques. Also, we will see that the (not at all simple) calculation of the Laplace transform is only part of the solution, since we have to employ advanced numerical inversion algorithms so as to arrive at the desired option value. Despite the ever-increasing computational capacity of modern times, we will see that the implementation of neither the numerical inversion algorithms nor the simulation algorithms is easy as we quickly run out of available significant digits or memory resources in case the program code is not well-structured.

In section Empirical analysis, after overcoming the technical difficulties, we will compare the valuation algorithms presented earlier in terms of speed and robustness. Using **the outputs generated in MATLAB[®]**, I intend to make suggestions on which valuation approach is the most suitable in a given case. Hoping to improve on the results in a similar analysis of Fu, Madan and Wang (1999), I will strive to determine the weak-points of each algorithm as well as clarify their characteristics for different combinations of option parameters.

2 Valuation methodology

Before dealing with the valuation problem of Asian options, it is necessary to agree on the specific risk-neutral framework that we will be using. In the introduction, I pointed out that it is necessary to model the dynamics of asset prices and the risk-free interest rate before we could proceed with risk-neutral valuation. I also noted that my chosen model is the Black-Scholes framework.

In this section, first we will briefly recapitulate the assumptions and results of the Black-Scholes model, then we will introduce risk-neutral valuation and make a distinction between the main valuation approaches.

2.1 The Black-Scholes model

In their original model, Black and Scholes (1973, p. 4.) made the following assumptions about the behaviour of asset prices, interest rates and financial markets:

- the risk-free interest rate is constant over time (as opposed to models in which interest rates are deterministic or stochastic processes)
- the stock price (of any firm) follows geometric Brownian motion “with a variance rate proportional to the square of the stock price”
- the option is a “European-exercise” option (meaning that it can only be exercised at a predefined maturity)
- the stock pays no dividends (or other interim cash flows)
- there are no transaction costs of trading or borrowing
- there are no limitations to trading (including the ones concerning short-selling or the trading of fractional quantities)
- there are no limitations to taking out a loan or making a deposit at the risk-free interest rate

Although the last few assumptions are rather technical and only ensure that the functions and processes are well-behaved, the first three assumptions form the core of the model and define very specific stock price and interest rate behavior. Márkus (2012, p. 9.) calls this framework the “**bond–stock market**” as there are effectively three financial products in this setup:

- a risk-free bond B

- a stock S whose price follows geometric Brownian motion
- and a “European-exercise” derivative product f whose price is dependent on (is a real function of) the stock’s price

If we look at these products as processes, their respective price dynamics are

$$\begin{aligned} dB(t) &\doteq Br \, dt \\ dS(t) &\doteq \mu S(t) \, dt + \sigma S(t) \, dW(t) \\ df(t, S(t)) &\doteq \left(\frac{\partial f}{\partial t} + \frac{\partial f}{\partial S} \mu S(t) + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2(t) \right) dt + \frac{\partial f}{\partial S} \sigma S(t) \, dW(t) \end{aligned}$$

where W stands for the Wiener process. Furthermore, it can be easily shown that the logarithm of S is a normally distributed variate having the following dynamics

$$dX(t) \doteq d\ln(S(t)) = \left(\mu - \frac{1}{2} \sigma^2 \right) dt + \sigma \, dW(t) \quad (2.1)$$

which implies that **S is a lognormally distributed random variate** with the corresponding parameters.

With the help of the assumptions above, using a replication method called dynamic hedging, Black and Scholes (1973) managed to derive a partial differential equation for the derivative price. This is called the **Black-Scholes partial differential equation**

$$r f(t) = \frac{\partial f}{\partial t} + r \frac{\partial f}{\partial S} S(t) + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} \sigma^2 S^2(t) \quad (\text{BS PDE})$$

which is a notable result because in many cases it can be easily solved by means of standard parabolic PDE methods.

2.2 Risk-neutral valuation

How can we use the Black-Scholes model’s assumptions and results to price financial derivatives? In theory, by knowing the stock price’s dynamics and probability distribution, we can calculate the expected value of the derivative’s payout at maturity. All we have to do then is determine the present value of the expectation. What interest rate should we use for discounting? Should we use the same risk-free interest rate for a risky asset as for the risk-free bond? From our studies about the value of future cash flows and risk premium, we know that the discounting rate (i.e. expected return) should be an increasing function of the risk assumed. But how can

we determine the expected return of such a special asset as a derivative?

In this section we will see that by using the mathematical trick of changing probability measures, we will not need to. Before going into further explanation, let us expand our notations and definitions with those in Márkus (2012).

Notation 2.1 Let $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$ be a **filtered probability space**, where $\{\mathcal{F}_t\}_{t \geq 0}$ is the natural filtration generated by the Wiener process $W_{\mathbb{P}}$.

Definition 2.2 A $\pi(t) = (\beta(t), \gamma(t))$ process-pair is called **investment strategy** if both processes are predictable with respect to $\mathcal{F}_t$.

Definition 2.3 An investment strategy is **self-financing (SF)** if for all t : $B(t) d\beta(t) + S(t) d\gamma(t) = 0$ holds.

Notation 2.4 Let V_{π} be the **value process** of the portfolio (determined by a given π investment strategy) which contains $\beta(t)$ bonds and $\gamma(t)$ stocks at any given time. Accordingly, let $\tilde{V}_{\pi} = \frac{V_{\pi}}{B}$ be the **discounted value process** of the portfolio.

Now, let us explain the logic of risk-neutral valuation in the continuous-time Black-Scholes framework. Applying Girsanov's theorem, Márkus (2012) shows that by choosing

$$Z(t) \doteq \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{\mathcal{F}_t} = \exp \left\{ \left(\frac{\mu - r}{\sigma} \right) W_{\mathbb{P}}(t) - \frac{1}{2} \left(\frac{\mu - r}{\sigma} \right)^2 t \right\}$$

as an appropriate Radon-Nikodym-derivative process, the dynamics of the stock price will be

$$dS(t) = rS(t) dt + \sigma S(t) dW_{\mathbb{Q}}(t)$$

where $W_{\mathbb{Q}}(t) = W_{\mathbb{P}}(t) + \left(\frac{\mu - r}{\sigma} \right) t$ is the Wiener process with respect to probability measure $\mathbb{Q}$ ¹. From this result it can be derived that the discounted value process $\tilde{V}_{\pi}$ of any portfolio determined by a given π self-financing investment strategy will be a $\mathbb{Q}$ -local martingale, which implies

$$E_{\mathbb{Q}} \left(\tilde{V}_{\pi}(t) \mid \mathcal{F}_s \right) = \tilde{V}_{\pi}(s) \quad \forall \pi \in SF, \forall s \leq t$$

It follows that

$$V_{\pi}(0) \equiv \tilde{V}_{\pi}(0) = E_{\mathbb{Q}} \left(\tilde{V}_{\pi}(t) \right) \equiv e^{-rt} E_{\mathbb{Q}} (V_{\pi}(t)) \quad \forall \pi \in SF \quad (2.2)$$

which is worded in the following lemma.

¹Probability measure $\mathbb{Q}$ can be proven to be unique in the continuous-time Black-Scholes model as it defines a so-called “complete market”. For details refer to (Márkus, 2012, p. 35.).

Lemma 2.5 *The present value of any self-financing portfolio is equal to the expectation of the time- t value of the portfolio with respect to probability measure $\mathbb{Q}$ discounted by the risk-free interest rate.*

Applying the martingale representation theorem, it can be shown that for any European-exercise derivative there exists a corresponding self-financing investment strategy π^* replicating the value process of the particular derivative. In other words, the value of the derivative and the replicating portfolio will be equal at all times, and thus, applying equation (2.2), we arrive at

$$f(0, S(0)) \equiv V_{\pi^*}(0) = e^{-rt} E_{\mathbb{Q}}(V_{\pi^*}(t)) \equiv e^{-rt} E_{\mathbb{Q}}(f(t, S(t))) \quad (2.3)$$

In preparation for the final step, let us define a “risk-neutral stock” under probability measure $\mathbb{P}$:

$$dS^{(r)}(t) \doteq rS^{(r)}(t) dt + \sigma S^{(r)}(t) dW_{\mathbb{P}}(t) \quad (2.4)$$

noting that the crucial difference between S and $S^{(r)}$ is that μ in the drift coefficient has been swapped to the risk-free interest rate r .

Definition 2.6 (Risk-neutral logreturn) *Using (2.4) and the result in equation (2.1),*

$$dX^{(r)}(t) \doteq d \ln S^{(r)}(t) = \left(r - \frac{1}{2} \sigma^2 \right) dt + \sigma dW_{\mathbb{P}}(t) \quad (2.5)$$

describes the dynamics of the risk-neutral logreturn.

Consequence 2.7 *The distribution of $S^{(r)}(t)$ with respect to the $\mathbb{P}$ measure is the same as the distribution of $S(t)$ with respect to the $\mathbb{Q}$ measure.*

Consequence 2.8 *The distribution of $f(t, S^{(r)}(t))$ with respect to the $\mathbb{P}$ measure is the same as the distribution of $f(t, S(t))$ with respect to the $\mathbb{Q}$ measure.*

From these results and equation (2.3), we arrive at the following equation and theorem, which we will call the Expected Value – Present Value principle after Medvegyev and Száz (2010):

$$f(0, S(0)) = e^{-rt} E_{\mathbb{Q}}(f(t, S(t))) = e^{-rt} E_{\mathbb{P}}(f(t, S^{(r)}(t))) \quad (\text{EV-PV})$$

Theorem 2.9 (Expected Value – Present Value) *The present value of a European-exercise derivative equals to the (time- t) expectation of the derivative price with respect to probability measure $\mathbb{Q}$ discounted at the risk-free interest rate. This expectation can be calculated with regards to the $\mathbb{P}$ measure by assuming that the stock price has risk-neutral dynamics defined in equation (2.4).*

Having understood the Black-Scholes model and the risk-neutral valuation principle, we can now turn to introducing the different approaches which employ the above findings to solve the derivative valuation problem.

2.2.1 Analytical methods

Analytical methods aim at deriving a closed-form expression for the price of the particular derivative. This could be attempted by

- either solving the (BS PDE) with the appropriate boundary conditions given by the payout of the derivative at expiry
- or calculating the risk-neutral expected value of the derivative at expiry, using the **Expected Value – Present Value** principle outlined in section 2.2

Whichever method we choose, the existence of a closed-form expression is not guaranteed and depends solely on the payout of the derivative. From this point, the valuation problem is purely mathematical: the solution may be easy to find or it may not even exist.

Example 2.10 (European option) *A good example for an easy-to-derive closed-form expression is the formula for the European-exercise vanilla call option calculated in Merton (1973). In his famous work, Merton showed that the value of the option is given by the formula*

$$c(t) = c(S(t), K, T, \sigma, r) = S(t)N(d_+) - Ke^{-r(T-t)}N(d_-)$$

where d_+ and d_- are variables depending on the initial parameters and N denotes the cumulative distribution function of the standard normal distribution.

Although in some cases the derivation of the valuation formula is extremely challenging, financial mathematicians put considerable effort into finding it. The reason for this is that finding a closed-form expression for a particular financial derivative's price can be extremely rewarding as it yields an easy-to-calculate and exact value, which makes it superior to other (e.g. numerical or simulation) methods.

2.2.2 Numerical methods

Sometimes, a closed-form solution proves to be elusive or non-existent, but it is possible to derive a complex but deterministic (semi-analytical) equation. For instance,

we may arrive at a complex (single/double/triple) integral form which can be approximated by means of numerical methods (e.g. quadratures).

Example 2.11 (Laplace transform) *A notable example for the use of semi-analytical formulas is the case of Asian option valuation applying the Laplace transform, discussed in section Pricing Asian options.*

In other cases, financial mathematicians resort to the numerical solution of the (BS PDE). By substituting the appropriate derivatives of the unknown function f , we arrive at a partial "boundary value problem", the kind of which has extensive literature on numerical solutions.

2.2.3 Simulation methods

Finally, if we cannot make advances with the analytical solution, we can still turn to simulation methods known as "Monte Carlo methods", which is a code name² coined by John von Neumann for his secret operations during World War II. There exist numerous approaches but the basic principle is the same (Kalos and Whitlock, 2008):

1. using random number generating algorithms, generate outcomes for a random variate with known distribution;
2. applying the Law of Large Numbers, estimate the unknown expected value with the average of the generated outcomes;
3. when applicable, increase the sample size so as to constrain the estimate into the desired confidence interval.

Monte Carlo simulation methods are computer intensive but easy to understand and flexible algorithms in the sense that they are compatible with a broad family of (physical, mathematical or financial) stochastic models. From the perspective of financial derivative valuation, the flexibility of Monte Carlo simulation allows for complex stock price movements (e.g. jump diffusion) and derivative payouts (e.g. path-dependency or early-exercisability) (for further examples refer to Hull (2009)), which are otherwise extremely hard to handle by analytical methods.

²The name "Monte Carlo" stems from the Monte Carlo Casino in Monaco where the uncle of Neumann's colleague used to gamble.

Based on the results described earlier in sections 2.1 and 2.2, we can see that in the continuous-time Black-Scholes framework, Monte Carlo simulation is particularly simple and consists of the following steps:

1. applying that the logreturn is a normally distributed variate, whose risk-neutral dynamics are described in equation (2.5), generate risk-neutral logreturns up to the maturity of the derivative, cutting up the time to maturity as necessary;
2. taking the exponential of the generated returns, calculate the trajectories of (i.e. the paths taken by) the stock;
3. keeping in mind that the derivative's payout at maturity depends on the respective trajectory of the stock, calculate the payout for each outcome;
4. taking the average of outcomes, estimate the expected value of the derivative at maturity;
5. applying the **Expected Value – Present Value** principle, estimate the (present) value of the derivative by discounting the estimate for the expected value with the risk-free interest rate.

Without going into further details about the simulation algorithms that I will use in section Empirical analysis, we must point out the notoriously slow convergence of simulation methods. As I have noted earlier, these techniques rely heavily on the Law of Large Numbers, estimating the expected value by the average of simulated outcomes. Provided these outcomes are IID, the standard deviation of the estimate will be proportional to $n^{-0.5}$. Practically, this means that for every additional decimal digit precision, we have to make a hundred times greater (and longer-taking. . .) simulation. Observing how quickly the powers of 100 grow, it is not hard to see that the traditional Monte Carlo method becomes computationally intractable after the first 4-6 decimal digits of precision improvement.

For the above mentioned reasons, it is necessary to employ advanced simulation techniques so as to reduce the variance of the estimate. The umbrella term for these methods is “**variance reduction techniques**” and they fall into many categories. Without aiming at completeness, we list a few techniques here (for a detailed description refer to Kleijnen et al. (2010)):

- common random numbers (including antithetic variates)
- control variates

- conditioning and stratified sampling
- importance sampling and splitting
- quasi-Monte Carlo (not completely random) simulation

Remark 2.12 *Depending on the statistical properties of the estimated random variate, some techniques may be more suitable to a certain simulation problem than others. In the case of Asian option valuation, the use of control variates looks the most promising to me.*

At the end of section Pricing Asian options, I will describe the methods of antithetic random numbers and control variates in more detail, look into their statistical properties, and explain how they could be efficiently implemented in the case of option valuation.

3 Pricing Asian options

Having learnt the most important features of the Black-Scholes model and the fundamental logic behind risk-neutral valuation, we are ready to discuss the case of pricing Asian options. First, I will explain what makes an option “Asian” by making a distinction between the basic option types. Then, I will outline the valuation problem of Asian options and demonstrate how the Laplace transform could be applied in this particular case. Finally, we will turn to further investigating the opportunities in the simulation approach presented at the end of section 2.2.

3.1 Option types

The classification of various option contracts can sometimes be an ambiguous and rather confusing task, even for market participants. In order to create a solid foundation for a discussion of greater complexity, in the following I will introduce a suggested classification framework based on the information found in Hull (2009).

Apart from the direction of their market exposure (i.e. whether they are call or put options), we can classify option contracts along two fundamental dimensions. Based on their **exercise behavior**, market participants make a basic distinction between three different option types:

- European options: these options can be exercised only at the pre-agreed maturity date
- American options: these options can be exercised any time before the pre-agreed maturity date
- Exotic options: these options are neither European nor American, which means that there is something special in the way they can be exercised

Example 3.1 (Bermudan option) *The most common example for an exotic-exercise option is the so-called Bermudan option, where the holder of the option has the right to exercise at more than one pre-agreed (maturity) dates. Most exotic-exercise interest rate options are Bermudan options.*

The other dimension of classification is the rule which determines how the **payout at exercise** is calculated. Keeping the framework simple, I only distinguish two dimensions:

- Vanilla options: the payout of these options is calculated based on the underlying product’s price at the time of exercise


- Exotic options: the payout of these options is not “vanilla” in the sense that it depends not only on the exercise-date price of the underlying product but on the path taken by it

Remark 3.2 *The latter type of options (i.e. options with exotic payout calculation) are often referred to as “path-dependent” options.*

Example 3.3 (Barrier option) *A good example for “path-dependent” options is the (double) barrier option. At this type of option, the underlying product’s price must (not) cross (a) certain pre-agreed level(s), otherwise the payout is zero.*

Example 3.4 (Asian option) *Another excellent (if not the best) example for “path-dependent” options is the Asian option as the payout of this option depends on every single point of the path taken, by calculating the average of the underlying product’s prices during the option’s lifetime (i.e. from writing to maturity).*

Figure 1: Notional amounts of equity options globally (USD billion)



Source: BIS (2012)

Given the versatility and size of option markets (with reference to the chart above), now it must be clear to the reader why the valuation of financial options is a relevant topic. It is also apparent that **a significant portion of option contracts is traded OTC** (i.e. “over-the-counter”), where products are often customized and have complex payouts. This is why I referred to the European call option by the lengthy expression “European-exercise vanilla option” earlier in this thesis. For instance, it is very easy to get confused in the case of Asian options mentioned in example 3.4, since “Asian” looks like as though it referred to the option’s exercise behavior (just like “European”) but it is not! In fact, “Asian” refers to how the payout at exercise is calculated, and thus the proper name of an Asian option should

read “European-exercise Asian option”. Having said that, in the rest of this thesis we will call options at their usual market names (i.e. European, American, Asian, etc.), keeping in mind what the various prefixes refer to.

3.2 Asian options

Having tackled the difficulties with classification and nomenclature, let us turn back to Asian options. We have mentioned that what makes these options “Asian” is the way their payout is calculated at exercise, namely that an average of the underlying product’s prices is taken instead of simply taking the underlying product’s price at exercise (as in the case of vanilla options). Using the notation of section Valuation methodology, the payout function of an Asian call option at exercise (expiry³) is

$$f(T, S(T)) = (Avg_{[0,T]}(S) - K)^+ \quad (3.1)$$

where K stands for the pre-agreed strike price and $Avg_{[0,T]}(S)$ stands for the average of the underlying’s prices from time 0 to time T (i.e. during the lifetime of the option). Now, let us apply the EV-PV principle for $t \equiv T$ (i.e. at expiry):

$$\begin{aligned} f(0, S(0)) &= e^{-rT} E_{\mathbb{Q}}(f(T, S(T))) = \\ &= e^{-rT} E_{\mathbb{P}}(f(T, S^{(r)}(T))) = e^{-rT} E_{\mathbb{P}} \left[(Avg_{[0,T]}(S^{(r)}) - K)^+ \right] \end{aligned} \quad (3.2)$$

This equation tells us that the final payout of the option will depend on the average of all the underlying’s prices during the option’s lifetime. In other words, the payout of the option is dependent on the entire path taken by the underlying, including every single point, which makes the Asian option, in a way, the “ultimate” path-dependent option. But what good is an option whose payout depends on the average price? Can a financial product like this survive the constant innovation pressure from financial engineering? Or is it just a passing fad of the market?

Well, according to a survey of 200 derivative-using non-financial U.S. firms by Bodnar et al. (1998), Asian options are the most popular exotic payout options chosen by non-financial firms for risk management. Vanilla options are (and will presumably remain) in the leading position in financial markets because of their simplicity, but it seems that Asian options have some added value to offer too. Based on Hull (2009) and Kemna and Vorst (1990), it seems that Asian options are the most attractive for those trading (hedging) in markets where

³Note that the time of (a possible) exercise and expiry coincides in the case of an European-exercise option, no matter whether it is “vanilla” or “Asian” in terms of payout calculation.

- either the volume is low (eg. in the corporate bond or commodity markets)
- or the volatility is high (eg. in the currency and interest rate markets)

In the first case, low trading volume allows for market manipulations and price jumps, which can be very unfavorable and generate rapid value changes for a European-exercise vanilla option holder close to maturity. In the second case, high volatility makes option strategies for hedgers too expensive to use. As opposed to what we are used to (i.e. price volatility is beneficial for option holders), in these cases market participants would like to avoid the inconvenience caused by hectic price fluctuations of the underlying product. As we can already see, **Asian options offer an efficient solution** for these and similar situations. By taking the average of market prices for a long period, adverse effects of market manipulation and price jumps are eliminated. Also, since the average of prices is more stable than the exercise-time price, Asian options are usually cheaper than their vanilla counterparts, which makes them a more efficient way of risk management.

The average shown in equation (3.1) can be calculated as **either the arithmetic mean or the geometric mean** of the underlying product's prices. Although practically there are no big differences in the payouts (i.e. both calculates some kind of a time-average of the underlying product's price), the theoretical option valuation is quite different for these variations. In the following, we are going to show how the risk-neutral valuation approach explained in section 2.2 can be applied to these two types of Asian options.

3.2.1 Geometric-average Asian options

Let us begin with Asian options where the average is calculated as the geometric mean of past underlying product prices. Based on Turnbull, McLean and Wakeman (1991), in the discrete case this means

$$Avg_{[0,T]}(S) = \sqrt[n+1]{\prod_{i=0}^n S(t_i)} \doteq G(T)$$

where we have partitioned the $[0, T]$ time interval into n subintervals.

$$\begin{aligned}
G(T) &= \sqrt[n+1]{\prod_{i=0}^n S(t_i)} = \exp \left\{ \frac{1}{n+1} \sum_{i=0}^n \ln(S(t_i)) \right\} = \\
&= \exp \left\{ \frac{(n+1) \ln(S(0)) + \sum_{i=1}^n (n+1-i) \ln(R(t_i))}{n+1} \right\} = \\
&= S(0) \exp \left\{ \frac{1}{n+1} \sum_{i=1}^n (n+1-i) X(t_i) \right\}
\end{aligned}$$

where X is the logreturn defined in equation (2.1). Since the X s are increments of a Wiener process, we know that they are $N\left(\frac{T}{n}(\mu - \frac{1}{2}\sigma^2), \frac{T}{n}\sigma^2\right)$ distributed IID random variates, and thus **$G(T)$ is lognormally distributed**. Therefore, switching to the risk-neutral framework introduced in section 2, the valuation problem of geometric-average Asian options means no greater challenge than the valuation of European-exercise vanilla options. Using Black's model for lognormally distributed underlying products (for more details, refer to Hull (2009)), Turnbull et al. (1991) derived the following pricing formula for geometric-average Asian call options:

$$\begin{aligned}
c_n(0) &= c_n(S(0), K, T, \sigma, r) = S(0)e^{-(r + \frac{n+2}{6(n+1)}\sigma^2)\frac{T}{2}} N(d_+) - Ke^{-rT} N(d_-) \\
d_+ &= \frac{\ln \frac{S(0)}{K} + \left(r + \frac{n-1}{6(n+1)}\sigma^2\right) \frac{T}{2}}{\sqrt{\frac{2n+1}{6(n+1)}}\sigma\sqrt{T}} \quad d_- = \frac{\ln \frac{S(0)}{K} + \left(r - \frac{\sigma^2}{2}\right) \frac{T}{2}}{\sqrt{\frac{2n+1}{6(n+1)}}\sigma\sqrt{T}}
\end{aligned} \tag{3.3}$$

For the continuous average formula, all we need to do is infinitely refine the partition of time interval $[0, T]$ by taking the limit of expression (3.3) in parameter n :

$$\begin{aligned}
c(0) &= \lim_{n \rightarrow \infty} c_n(0) = S(0)e^{-(r + \frac{1}{6}\sigma^2)\frac{T}{2}} N(d_+) - Ke^{-rT} N(d_-) \\
d_+ &= \frac{\ln \frac{S(0)}{K} + \left(r + \frac{\sigma^2}{6}\right) \frac{T}{2}}{\sigma\sqrt{T/3}} \quad d_- = \frac{\ln \frac{S(0)}{K} + \left(r - \frac{\sigma^2}{2}\right) \frac{T}{2}}{\sigma\sqrt{T/3}}
\end{aligned} \tag{3.4}$$

Remark 3.5 (Time- t price of path-dependent options) *At this point, it is important to note that the pricing formula for Asian options takes zero as the time argument. This is by no means a coincidence, since Asian options are path-dependent options regarding payout calculation (recall section 3.1). This implies that **it is impossible to find a formula for the time- t price of an Asian option without inputting the past prices** of the underlying product, since the average*

is not only dependent on the price level between time t and T (i.e. maturity) but it is also influenced by what happened between time 0 and t . However, this information is a random variate, changing from trajectory to trajectory, and thus cannot be omitted from a potential valuation formula.

3.2.2 Arithmetic-average Asian options

The payout of this kind of Asian options⁴ is determined by the arithmetic average of the underlying product's prices. In the continuous average case, let us define

$$Avg_{[0,T]}(S^{(r)}) = \frac{1}{T} \int_0^T S^{(r)}(t) dt \doteq \frac{A(T)}{T} \quad (3.5)$$

where $S^{(r)}$ is the risk-neutral underlying product process defined in equation (2.4). In theory, all we need to do now is substitute the time-average into equations (3.1) and (3.2), simply apply the EV-PV principle and calculate the value of an Asian call option:

$$C = e^{-rT} E_{\mathbb{P}} \left[(Avg_{[0,T]}(S^{(r)}) - K)^+ \right] = e^{-rT} E_{\mathbb{P}} \left[\left(\frac{A(T)}{T} - K \right)^+ \right] \quad (3.6)$$

As it turns out, solving this simple-looking equation hides several pitfalls haunting financial mathematicians for decades. The fundamental problem is caused by the unknown probability distribution of $A(t)$. More specifically, it can be shown that the sum of lognormally distributed random variates is not necessarily lognormally distributed, and thus the expected value in equation (3.6) is challenging to determine.

So much so that even such prominent authors as Kemna and Vorst (1990, p. 114.) claimed **“it is impossible to derive an explicit formula”** for the price of Asian options. Although based on a result of Mitchell (1968), who showed that the sum of lognormals is approximately lognormal, Turnbull et al. (1991) derived an approximation by calculating the option's first two moments and substituting back into the lognormal distribution's PDF, this still couldn't be used as a reliable pricing formula, and thus market participants resorted to simulation methods so as to determine the price of Asian options.⁵

⁴From now on we will refer to arithmetic-average Asian options, unless indicated otherwise.

⁵Simulation methods were introduced in section 2, and will be discussed in more detail in section 3.4.

The breakthrough can be credited to Geman and Yor (1993), based on whose seminal paper a series of authors, including Carr and Schröder (2003) as well as Dufresne (2004), commenced an exciting discussion over the validity and applicability of their approach⁶. Before proceeding with a more detailed explanation of their finding, let us do a crucial transformation of the integral in expression (3.5), using the notation of Medvegyev (2012) and applying the scaling-property of the Wiener process:

$$\begin{aligned}
A(T) &= \int_0^T S^{(r)}(t) dt = S(0) \cdot \int_0^T \exp \left\{ \left(r - \frac{\sigma^2}{2} \right) t + \sigma W_{\mathbb{P}}(t) \right\} dt = \\
&= S(0) \frac{4}{\sigma^2} \cdot \int_0^{\sigma^2 T/4} \exp \left\{ \left(r - \frac{\sigma^2}{2} \right) \frac{4}{\sigma^2} u + 2 \frac{\sigma}{2} W_{\mathbb{P}} \left(\frac{4}{\sigma^2} u \right) \right\} du = \\
&= S(0) \frac{4}{\sigma^2} \cdot \int_0^{\sigma^2 T/4} \exp \left\{ 2 \left(\nu \cdot u + \widetilde{W}_{\mathbb{P}}(u) \right) \right\} du
\end{aligned}$$

where $\widetilde{W}_{\mathbb{P}}$ is another Wiener process under probability measure $\mathbb{P}$, and

$$\nu \doteq \left(r - \frac{\sigma^2}{2} \right) \frac{2}{\sigma^2} = \frac{2r}{\sigma^2} - 1$$

Based on Geman and Yor (1993), let us define

$$A^{(\nu)}(T) \doteq \int_0^T \exp \left\{ 2 \left(\nu \cdot t + \widetilde{W}_{\mathbb{P}}(t) \right) \right\} dt \quad (3.7)$$

$$C^{(\nu)}(T, K) \doteq E_{\mathbb{P}} \left(A^{(\nu)}(T) - K \right)^+ \quad (3.8)$$

the latter being referred to as the “**normalized price of the Asian option**” by Carr and Schröder (2003, p. 3.).

Consequence 3.6 (Asian call option price) *Using the above definitions and substituting into equation (3.6), the value of the arithmetic-average Asian call option is given by:*

$$\begin{aligned}
C &= e^{-rT} E_{\mathbb{P}} \left(\frac{A(T)}{T} - K \right)^+ = e^{-rT} \cdot E_{\mathbb{P}} \left(\frac{S(0)}{T} \frac{4}{\sigma^2} A^{(\nu)}(\sigma^2 T/4) - K \right)^+ = \\
&= e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot C^{(\nu)} \left(\frac{\sigma^2 T}{4}, \frac{KT\sigma^2}{4S(0)} \right)
\end{aligned} \quad (3.9)$$

Notation 3.7 *Following Carr and Schröder (2003), in the rest of this thesis let us*

⁶Not incidentally, their work forms the cornerstone of this thesis too.

denote

$$h \doteq \frac{\sigma^2 T}{4} \quad \text{and} \quad q \doteq \frac{KT\sigma^2}{4S(0)}$$

The revolutionary finding of Geman and Yor (1993) was that there exists an analytical formula for the Laplace transform of $C^{(\nu)}(h, q)$ with respect to h . We will see the main steps of their derivation in the next section but before we proceed let us present three important findings in connection with Asian options.

Lemma 3.8 (Asian forward price) *The forward price of the arithmetic-average of stock prices for a given time interval $[0, T]$ is given by*

$$F \doteq e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot \frac{e^{2h(\nu+1)} - 1}{2(\nu+1)} \quad (3.10)$$

Proof. The above lemma is easily proven by substituting equation (3.7) into (3.10):

$$\begin{aligned} F &= e^{-rT} E_{\mathbb{P}} \left(\frac{A(T)}{T} \right) = e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot E_{\mathbb{P}} (A^{(\nu)}(h)) = \\ &= e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot E_{\mathbb{P}} \left(\int_0^h \exp \left\{ 2 \left(\nu \cdot t + \widetilde{W}_{\mathbb{P}}(t) \right) \right\} dt \right) = (\dots) \end{aligned}$$

Since the integrand is non-negative, we may apply Fubini's theorem and swap the order of the integrals:

$$\begin{aligned} (\dots) &= e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot \int_0^h E_{\mathbb{P}} \left(\exp \left\{ 2 \left(\nu \cdot t + \widetilde{W}_{\mathbb{P}}(t) \right) \right\} \right) dt = \\ &= e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot \int_0^h e^{2\nu t} E_{\mathbb{P}} \left(e^{2\widetilde{W}_{\mathbb{P}}(t)} \right) dt = e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot \int_0^h e^{2\nu t + 2t} dt = \\ &= e^{-rT} \frac{S(0)}{T} \frac{4}{\sigma^2} \cdot \frac{e^{2h(\nu+1)} - 1}{2(\nu+1)} \end{aligned}$$

■

Now, based on the previous lemma, we may derive the Black-Scholes price of the arithmetic-average Asian put option. The logic is the same as at the calculation of the vanilla option price, namely

$$(\xi - K)^+ - (K - \xi)^+ = (\xi - K) \quad \forall \xi$$

Regardless of the probability distribution of ξ , taking the expectation of both sides doesn't affect the equality, and thus we can apply equations (3.9) and (3.10) to arrive

at the desired Asian put option value.

Consequence 3.9 (Asian put option price (put-call parity))

$$\begin{aligned}
P &= e^{-rT} E_{\mathbb{P}} \left(K - \frac{A(T)}{T} \right)^+ = \\
&= e^{-rT} \left[E_{\mathbb{P}} \left(\frac{A(T)}{T} - K \right)^+ - E_{\mathbb{P}} \left(\frac{A(T)}{T} - K \right) \right] = \\
&= C - F + e^{-rT} K
\end{aligned} \tag{3.11}$$

Remark 3.10 (Time-t value of the Asian call option) *Finally, let us provide an extension of the (3.9) option pricing formula for $t > 0$, based on Carr and Schröder (2003):*

$$C(t) = e^{-rT} \frac{S(t)}{T-t} \frac{4}{\sigma^2} \cdot C^{(\nu)}(h(t), q(t))$$

where

$$h(t) \doteq \frac{\sigma^2(T-t)}{4} \quad \text{and} \quad q(t) \doteq \frac{\sigma^2}{4S(t)} \left(K(T-t) - \int_0^t S(u) du \right)$$

Notice how past information about the trajectory appears in parameter q of the formula as I have already indicated in remark 3.5. This gives rise to a very interesting phenomenon which Carr and Schröder (2003, p. 4.) call “**pricing dichotomy**” in their work. According to the above time-t formula, there may exist a time point τ for which $q(t)$ becomes negative, which implies that

$$\begin{aligned}
C^{(\nu)}(h(t), q(t)) &= E_{\mathbb{P}} \left(A^{(\nu)}(h(t)) - q(t) \right)^+ = \\
&= E_{\mathbb{P}} \left(A^{(\nu)}(h(t)) - q(t) \right) = E_{\mathbb{P}} \left(A^{(\nu)}(h(t)) \right) - q(t) \quad \forall t \geq \tau
\end{aligned}$$

since $A^{(\nu)}(T)$ is by definition non-negative and $q(t)$ is a monotonous function of t . The above expected value is easy to calculate as we have already demonstrated at the calculation of the Asian forward price. What follows is that **the Laplace transform method is only required while $q(t)$ is still positive** (as it initially always is⁷), otherwise the valuation problem is easy to handle. This is because after time point τ , it is certain that the Asian call option will expire in-the-money, the only question remaining is how deeply.

⁷At the onset of the option $t = 0$, which implies that q is positive.

3.3 The Laplace transform method

Definition 3.11 (Laplace transform) *The Laplace transform $\mathcal{L}$ of a given function f is defined by the following equation:*

$$\mathcal{L}\{f(t)\}(\lambda) \doteq \int_0^{\infty} e^{-\lambda t} f(t) dt \quad (\text{LT})$$

assuming that function f is defined for every t in the interval $(0, \infty)$ and the integral is convergent.

According to Kano (2010, p. 3.), the Laplace transform is particularly useful in simplifying or restructuring complex time-dependent problems and systems. Hence, its popular use in physics, engineering and the study of stochastic processes. As it turns out, the application of the Laplace transform in the “complex time-dependent problem” of Asian option pricing is also a lucrative undertaking.

Theorem 3.12 (Geman-Yor) *The Laplace transform of the normalized price of the Asian option, defined in equation (3.8), with respect to h is*

$$\begin{aligned} \mathcal{L}\{C^{(\nu)}(h, q)\}(\lambda, q, \nu) &= \int_0^{\infty} e^{-\lambda h} C^{(\nu)}(h, q) dh = \\ &= \frac{(2q)^{1-\beta}}{2\lambda(\alpha+1)\Gamma(\beta)} \int_0^1 \exp\left\{-\frac{u}{2q}\right\} u^{\beta-2} (1-u)^{\alpha+1} du \end{aligned} \quad (\text{GY LT})$$

where

$$\gamma \doteq \sqrt{2\lambda + \nu^2}, \quad \alpha \doteq \frac{\gamma + \nu}{2}, \quad \beta \doteq \frac{\gamma - \nu}{2}$$

for every

$$\lambda > \max\{0, 2(\nu + 1)\}$$

3.3.1 Derivation of the analytical formula

Having stated the principal result of Geman and Yor (1993), in what follows I am going to present the two cornerstones of the proof, and then based on show how the particular Laplace transform is derived from these intermediate steps. The complete derivation of relation (GY LT) lies beyond the scope of this thesis. For a detailed proof, refer to Dufresne (2004) and Medvegyev (2012).

Definition 3.13 (Squared Bessel process) For every $\delta \geq 0$ and $x \geq 0$, the unique strong solution to the equation

$$X(t) = x + \delta t + 2 \int_0^t \sqrt{X(s)} dW(s)$$

is called the square of a δ -dimensional Bessel process started at x and is denoted by $BESQ^\delta(x)$.

Definition 3.14 (Bessel process) For every $\delta \geq 0$ and $a \geq 0$, the square root of $BESQ^\delta(a^2)$ is called the δ -dimensional Bessel process started at a and is denoted by $BES^\delta(a)$.

Remark 3.15 For $\delta \in \mathbb{Z}^+$ and $a \geq 0$, the Bessel process could be conceived of as the distance of the δ -dimensional Wiener process from zero.

Remark 3.16 For $\delta \geq 2$ and $a > 0$, it can be shown that $BES^\delta(a)$ is positive and satisfies the following equation:

$$R^{(\nu)}(t) = a + \left(\nu + \frac{1}{2} \right) \int_0^t \frac{1}{R^{(\nu)}(s)} ds + W(t)$$

where $\nu = \frac{\delta}{2} - 1 \geq 0$ is called the index of the process $BES^\delta(a)$.

The first key step to the proof is the following identity by Lamperti (1972).

Lemma 3.17 (Lamperti's identity) For $\nu \geq 0$

$$\exp \{W(t) + \nu \cdot t\} = R^{(\nu)} \left(\int_0^t \exp \{2(W(s) + \nu \cdot s) ds\} \right) = R^{(\nu)} (A^{(\nu)}(t))$$

where $R^{(\nu)}$ is a Bessel process with respect to the appropriate Wiener process.

Remark 3.18 (Time-changed Bessel process) In other words, Lamperti's identity tells us that the above exponential of the Wiener process can be conceived of as a time-changed Bessel process (started at 1).

The second key step could be derived by applying Lamperti's identity and calculating the transition density function of the Bessel process.

Lemma 3.19 *Let σ be an exponentially distributed random variate with parameter λ . For arbitrary ν , the distribution of $A^{(\nu)}(\sigma)$ is of the form $\xi/2\eta$, where $\xi \sim \beta(1, \alpha)$ and $\eta \sim \Gamma(\beta, 1)$ are independent random variates and*

$$\gamma \doteq \sqrt{2\lambda + \nu^2}, \quad \alpha \doteq \frac{\gamma + \nu}{2}, \quad \beta \doteq \frac{\gamma - \nu}{2}$$

The probability density function of $A^{(\nu)}(\sigma)$ is

$$f(x) = \frac{\lambda}{(2x)^{\beta+1}\Gamma(\beta+1)} \int_0^1 y^\beta (1-y)^{\alpha-1} \exp\left\{-\frac{y}{2x}\right\} dy \quad (3.12)$$

Having the above results in our hands, we can finish the derivation of the (GY LT). We are looking for

$$\begin{aligned} \mathcal{L}\{C^{(\nu)}(h, q)\}(\lambda, q, \nu) &\doteq \int_0^\infty e^{-\lambda h} \cdot C^{(\nu)}(h, q) dh = \\ &= \int_0^\infty e^{-\lambda h} \cdot E_{\mathbb{P}}(A^{(\nu)}(h) - q)^+ dh = \frac{1}{\lambda} \cdot E_{\mathbb{P}}(A^{(\nu)}(\sigma) - q)^+ = (\dots) \end{aligned}$$

where σ is an exponentially distributed random variate with parameter λ . Substituting the PDF of $A^{(\nu)}(\sigma)$ from equation (3.12)

$$(\dots) = \frac{1}{\lambda} \frac{\lambda}{\Gamma(\beta+1)} \int_0^\infty (x-q)^+ \frac{1}{(2x)^{\beta+1}} \int_0^1 y^\beta (1-y)^{\alpha-1} \exp\left\{-\frac{y}{2x}\right\} dy dx = (\dots)$$

then using that $\lambda = 2\alpha\beta$ (by definition from lemma 3.19)

$$(\dots) = \frac{(2\alpha\beta)2^{-\beta-1}}{\lambda\Gamma(\beta+1)} \int_q^\infty \int_0^1 (x-q) \frac{1}{x^{\beta+1}} y^\beta (1-y)^{\alpha-1} \exp\left\{-\frac{y}{2x}\right\} dy dx = (\dots)$$

where we can use $\Gamma(\beta+1) = \beta\Gamma(\beta)$ and apply Fubini's theorem and swap the integrals, since the integrand is non-negative for all u and x

$$(\dots) = \frac{\alpha 2^{1-\beta}}{2\lambda\Gamma(\beta)} \int_0^1 \int_q^\infty (x-q) \frac{1}{x^{\beta+1}} y^\beta (1-y)^{\alpha-1} \exp\left\{-\frac{y}{2x}\right\} dx dy = (\dots)$$

Finally, we substitute for $x = qy/u$ and rearrange the expression,

$$\begin{aligned}
(\dots) &= \frac{\alpha 2^{1-\beta}}{2\lambda\Gamma(\beta)} \int_0^1 \int_0^y \left(\frac{qy}{u} - q\right) \left(\frac{qy}{u}\right)^{-(\beta+1)} y^\beta (1-y)^{\alpha-1} \exp\left\{-\frac{u}{2q}\right\} \frac{qy}{u^2} du dy = \\
&\quad \frac{\alpha 2^{1-\beta}}{2\lambda\Gamma(\beta)} \int_0^1 \int_u^1 \left(\frac{qy}{u} - q\right) \left(\frac{qy}{u}\right)^{-(\beta+1)} y^\beta (1-y)^{\alpha-1} \exp\left\{-\frac{u}{2q}\right\} \frac{qy}{u^2} dy du = \\
&\quad \frac{\alpha(2q)^{1-\beta}}{2\lambda\Gamma(\beta)} \int_0^1 \int_u^1 \left(\frac{y}{u} - 1\right) u^{\beta-1} (1-y)^{\alpha-1} \exp\left\{-\frac{u}{2q}\right\} dy du = \\
&\quad \frac{\alpha(2q)^{1-\beta}}{2\lambda\Gamma(\beta)} \int_0^1 \exp\left\{-\frac{u}{2q}\right\} u^{\beta-1} \int_u^1 (1-y)^{\alpha-1} \left(\frac{y}{u} - 1\right) dy du = \\
&\quad \frac{\alpha(2q)^{1-\beta}}{2\lambda\Gamma(\beta)} \int_0^1 \exp\left\{-\frac{u}{2q}\right\} u^{\beta-1} \left(\left[\frac{(1-y)^\alpha}{\alpha} \left(\frac{y}{u} - 1\right) \right]_u^1 - \int_u^1 \frac{(1-y)^\alpha}{\alpha u} dy \right) du = \\
&\quad \frac{(2q)^{1-\beta}}{2\lambda\Gamma(\beta)} \int_0^1 \exp\left\{-\frac{u}{2q}\right\} u^{\beta-2} \left(- \left[\frac{(1-y)^{\alpha+1}}{\alpha+1} \right]_u^1 \right) du = \\
&\quad \frac{(2q)^{1-\beta}}{2\lambda\Gamma(\beta)(\alpha+1)} \int_0^1 \exp\left\{-\frac{u}{2q}\right\} u^{\beta-2} (1-u)^{\alpha+1} du
\end{aligned}$$

where $\beta - 2 > -1$ is necessary for the convergence of the integral, which is satisfied by the theorem's precondition $\lambda > \max\{0, 2(\nu + 1)\}$, implying $\beta > 1$. ■

Now that we understand how Geman and Yor (1993) derived the Laplace transform of the normalized price of the arithmetic-average Asian call option, we shall proceed with valuation as follows:

1. Calculate the Laplace transform given in (GY LT);
2. Invert the Laplace transform by means of numerical algorithms in order to get the normalized option value described in (3.8);
3. Substitute the normalized option value into equation (3.9) and calculate the option price.

3.3.2 The problem of inversion

As it is described in Cohen (2007), the problem of inverting the Laplace transform is a nontrivial one, apart from a few special cases. Luckily, the Laplace transform is unique, in the sense that if $\mathcal{L}\{f\} = \mathcal{L}\{g\}$, then $f = g$ provided they are continuous

functions, which means that we need not worry about one approach leading to a different result than another. There have been various techniques invented, based on either of the following two inversion theorems.

Theorem 3.20 (Bromwich theorem) *If the Laplace transform $\mathcal{L}$ of a continuously differentiable function f exists, then*

$$f(t) = \mathcal{L}^{-1}\{\mathcal{L}(\lambda)\}(t) \doteq \frac{1}{i(2\pi)} \cdot \lim_{T \rightarrow \infty} \int_{\gamma-iT}^{\gamma+iT} e^{\lambda t} \mathcal{L}(\lambda) d\lambda, \quad (3.13)$$

where γ is a real number so that $\mathcal{L}$ has no singularities on or right from γ on the complex plane.

Theorem 3.21 (Post-Widder theorem) *If the Laplace transform $\mathcal{L}$ of a continuously differentiable function f exists, then*

$$f(t) = \mathcal{L}^{-1}\{\mathcal{L}(\lambda)\}(t) \doteq \lim_{n \rightarrow \infty} \frac{(-1)^n}{n!} \left(\frac{n}{t}\right) \mathcal{L}^{(n)}\left(\frac{n}{t}\right)$$

where $\mathcal{L}^{(n)}$ is the n -th derivative of the particular Laplace transform.

Although in contrast to the Bromwich formula the Post-Widder formula is restrained to the real line and thus looks easier to handle, its speed of convergence is very slow and the techniques based on them are prone to computational difficulties, as we can see in Fu et al. (1999). Therefore, I have chosen for the calculations two different algorithms based on the Bromwich theorem, serving as a check for each other. In what follows, I am going to present these algorithms described in Abate and Whitt (2006).

The Euler algorithm

Abate and Whitt (1995, p. 37.) named this algorithm the “Euler” algorithm because they employed Euler summation to accelerate convergence. Without going into further details, their basic idea was to look at the Bromwich integral of (3.13) as though it was a Fourier-series by substituting for $\lambda = a + iu$:

$$\begin{aligned} \mathcal{L}^{-1}\{\mathcal{L}(\lambda)\}(t) &= \frac{1}{i(2\pi)} \cdot \lim_{T \rightarrow \infty} \int_{\gamma-iT}^{\gamma+iT} e^{\lambda t} \mathcal{L}(\lambda) d\lambda = \\ &= \frac{1}{2\pi} \cdot \lim_{T \rightarrow \infty} \int_{-T}^T e^{(\gamma+iu)t} \mathcal{L}(\gamma + iu) du = \frac{e^{\gamma t}}{2\pi} \int_{-\infty}^{\infty} e^{iut} \mathcal{L}(\gamma + iu) du \end{aligned}$$

Their final result is the algorithm presented in Abate and Whitt (1995), which also Fu, Madan and Wang (1999) use in their calculations, being apparently a very (if not the most) efficient numerical algorithm in inverting the particular Laplace transform derived by Geman and Yor (1993).

Algorithm 3.22 (Euler) *Let us define*

$$s_n(t) \doteq \frac{e^{A/2}}{2t} \operatorname{Re} \left\{ \mathcal{L} \left(\frac{A}{2t} \right) \right\} + \frac{e^{A/2}}{t} \sum_{k=1}^n (-1)^k a_k(t)$$

and

$$a_k(t) \doteq \operatorname{Re} \left\{ \mathcal{L} \left(\frac{A + i(2\pi)k}{2t} \right) \right\}$$


Then the inverse Laplace transform could be approximated by the Euler series:

$$\mathcal{L}^{-1}\{\mathcal{L}(\lambda)\}(t) \approx \hat{f}(t, A, m, n) = \sum_{k=0}^m \binom{m}{k} 2^{-m} s_{n+k}(t)$$

where A , m and n are chosen parameters of the algorithm, for which Abate and Whitt (1995) recommends $A = 18.4$, $m = 11$ and $n = 15$ (increasing n as necessary).

The Talbot algorithm

Similarly to the Euler algorithm, this algorithm is also based on the Bromwich theorem. The idea of Talbot, described in Cohen (2007), was to mitigate the disturbing effects of oscillation as we move along the integration contour parallel to the imaginary axis by deforming (i.e. altering) the integration



Source: Kano (2010)

contour to a more appropriate one (let us call it Γ), which encloses all singularities of $\mathcal{L}$ in the complex plane. This idea is formally summarized in the equation

$$\mathcal{L}^{-1}\{\mathcal{L}(\lambda)\}(t) = \frac{1}{i(2\pi)} \cdot \lim_{T \rightarrow \infty} \int_{\gamma-iT}^{\gamma+iT} e^{\lambda t} \mathcal{L}(\lambda) d\lambda = \frac{ce^{\sigma t}}{i(2\pi)} \int_{\Gamma} e^{c\lambda t} \mathcal{L}(\sigma + c\lambda) d\lambda$$

Based on the logic above, Abate and Whitt (2006) elaborated the following numerical inversion algorithm.

Algorithm 3.23 (Talbot) *Let us define*

$$\delta_0 \doteq \frac{2M}{5}, \quad \delta_k \doteq \frac{(2\pi)k}{5} (\cot(k\pi/M) + i)$$

and

$$\gamma \doteq \frac{1}{2}e^{\delta_0}, \quad \gamma_k \doteq [1 + i(k\pi/M)(1 + \cot^2(k\pi/M)) - i \cot(k\pi/M)] e^{\delta_k}$$

Then the inverse Laplace transform could be approximated by the series

$$\mathcal{L}^{-1}\{\mathcal{L}(\lambda)\}(t) \approx \hat{f}(t, M) = \frac{2}{5t} \sum_{k=0}^{M-1} \operatorname{Re} \left\{ \gamma_k \mathcal{L} \left(\frac{\delta_k}{t} \right) \right\}$$

where M is a chosen parameter of the algorithm.

Implementation

In the Asian option valuation problem described in (3.9), we have

$$\mathcal{L}(\lambda) \equiv \mathcal{L}\{C^{(\nu)}(h, q)\}(\lambda, q, \nu)$$

which could now be inverted numerically with respect to λ to obtain

$$f(h) \equiv C^{(\nu)}(h, q)$$

The MATLAB[®] implementation of the introduced algorithms is presented in Appendix A: Euler algorithm and Appendix B: Talbot algorithm. For the sake of comprehension, I strove to keep the notation presented earlier in this thesis and insert as many comments into the code as possible. Without going into unnecessary technicalities, I would like to point out a couple of features and enhancements in the code:

- Both algorithms are based on substituting back into the Laplace transform $\mathcal{L}$ at different places on the complex plane. Most MATLAB[®] functions can take complex arguments, so this should not be a setback. One important hint to bear in mind though is the use of “1i” instead of “i” in the code so that MATLAB[®] immediately knows that we refer to the imaginary number and not a variable.
- A crucial step in calculating the Laplace transform presented in (GY LT) is complex integration on the $(0, 1)$ interval. Since solving this problem is not in the scope of this research, I used the built-in “integral” function for this purpose, which is based on “the global adaptive quadrature” according to the help in MATLAB[®].

- Based on my experience, during the implementation of either of the algorithms we have to work with quantities of very different magnitudes. More specifically, the integral is usually of magnitude 10^{-100} to 10^{-500} , whereas the fraction before is large enough to compensate it. This poses a big challenge as computers store numbers as so-called floating-point numbers⁸ and thus multiplication and division of such different size numbers can lead to the complete loss of precision. I have employed two techniques to mitigate this problem:

1. I took the (natural) logarithm of the whole expression in (GY LT) so as to exponentially reduce the magnitude of variables⁹. Although this step does not necessarily increase precision, it definitely earns us “space” in terms of magnitude.
2. I explicitly instructed MATLAB[®] to use its Symbolic Math Toolbox[®], which allows the user to choose an arbitrary level of precision for calculations. Regarding our particular problem, I was satisfied with a precision of 25 significant digits.

Having fully understood the semi-analytical valuation method using the Laplace transform of the Asian option’s normalized price, we are ready to price Asian options with two algorithms. But before starting our empirical analysis with respect to the precision and robustness of these algorithms, we will learn a couple of Monte Carlo simulation techniques, which will serve as an ultimate check for our calculations and make our option valuation independent of the Laplace transform approach.

3.4 The Monte Carlo simulation method

In what follows, I am going to present the simulation logic tailored to the specific valuation problem of arithmetic-average Asian options. Having already learnt the basics of the Monte Carlo method at the end of section 2.2, we know that this is a highly-flexible, robust approach, suitable to various frameworks and product types. Flexibility, however, has a price, which takes the form of variance in the case of simulation methods, meaning that the calculated price bears some uncertainty and fluctuates around the option’s true value. Striving to cut this variance to a minimum level, we are going to study in details how the variance reduction techniques

⁸The floating-point method represents numbers as $\{10^{exponent} \times significant\ digits\}$, using 32-bits or 64-bits for storing information according to the “single-precision” or “double-precision” IEEE 754 standard. This means about 6-9 or 15-17 significant digits of precision, respectively. (Wikipedia, 2013)

⁹This is a viable choice as the logarithm of the Γ function is a standard function in software packages.

of antithetic variates and control variates work in the case of financial derivative valuation.

3.4.1 Basic method

As the name suggests, this is the simplest form of Monte Carlo method. As we have already discussed earlier at the end of section 2.2, at this method we take the arithmetic-mean as an unbiased estimator of the expected value:

$$\widehat{E}(X) \doteq \frac{\sum_{i=1}^n X_i}{n} = \sum_{i=1}^n \frac{X_i}{n}$$

having a variance of

$$D^2(\widehat{E}(X)) = D^2\left(\sum_{i=1}^n \frac{X_i}{n}\right) = \sum_{i=1}^n \frac{D^2(X_i)}{n^2} = \frac{D^2(X)}{n}$$

where $X_i = f(\omega_i)$ represents the i^{th} outcome in our simulation, and they are IID random variates. The random function f , whose expected value we wish to estimate is the discounted payout function of the option as suggested by the EV-PV principle:

$$f(\omega) \doteq e^{-rT} \times \text{Payout}(\omega)$$

In the case of arithmetic-average Asian call options, this looks like

$$f(\omega) = e^{-rT} \left(\text{Avg}_{[0,T]}(S^{(r)}(\omega)) - K \right)^+$$

which has an expected value

$$E(f(\omega)) = e^{-rT} E_{\mathbb{P}} \left[\left(\text{Avg}_{[0,T]}(S^{(r)}(\omega)) - K \right)^+ \right]$$

This is precisely the risk-neutral value of the option according to equation (3.6). As we can see, it is this expected value that we are going to estimate through simulation of the underlying product's risk-neutral trajectory

$$\left\{ S_t^{(r)}(\omega) \right\}_{t \in [0,T]}$$

How will we generate one trajectory using the basic method?

In the case of European-exercise vanilla options (as explained in section 3.1), it is sufficient to generate the endpoint of the trajectory at time T , since the payout is only

dependent on the exercise-time price of the underlying product¹⁰. However, in the case of Asian options, every single point of the trajectory is taken into account when calculating the average of prices. Therefore, we will create an equidistant partition of time interval $[0, T]$ and then generate a particular trajectory by repeatedly applying the risk-neutral dynamics of the logreturn as described in equation (2.5). Hence, one trajectory will be constructed as

$$\left\{ S_{t_j}^{(r)}(\omega) \right\}_{j \in [0, m \cdot (T \times 365)]} \doteq \left\{ S(0) \cdot \exp \left\{ \sum_{i=0}^j X_{t_i}(\omega) \right\} \right\}_{j \in [0, m \cdot (T \times 365)]} \quad (3.14)$$

where

$$X_{t_i}(\omega) \doteq \begin{cases} 0 & \text{if } i = 0 \\ \left(r - \frac{\sigma^2}{2} \right) (t_i - t_{i-1}) + \sigma (W(t_i) - W(t_{i-1})) & \text{if } i > 0 \end{cases}$$

Sources of estimation error

Note that in the case of path-dependent derivatives simulation methods have two sources of estimation error:


1. The generated trajectories (outcomes) depend on the random event ω , and thus the calculated estimator will also be a random variate, fluctuating around the expected value that we are looking for. This is called **estimator variance**, and could be reduced by increasing n , which is the number of trajectories (outcomes) in our simulation.
2. The generated trajectories are discretized, which means they are only approximations of the true, continuous-time trajectories required for the calculation of the path-dependent derivative payout. In other words, we substitute the integral with a sum in the payout of the Asian option and no matter how great n and thus how accurate our estimator is, we will still have **discretization error**. This error could be reduced by increasing m , which is the number of fragments we break up each day of interval $[0, T]$ into.

3.4.2 Variance reduction by antithetic variates

Keeping the logic and weak points of the basic method in mind, let us see how we could improve on our estimate of the expected value. First of all, let us try and reduce the variance of the estimator. A suitable technique is the method of

¹⁰This is exactly why Medvegyev and Száz (2010, p. 56.) call this kind of financial products “T-products”.

antithetic variates, described in Kleijnen et al. (2010). The basic idea here is using each generated random number twice, creating two trajectories from each logreturn sequence, where one is negatively correlated (antithetic) to the other.



The way these antithetic trajectories are created differs from simulation to simulation. It is usually a good idea to find a step in the simulation algorithm when we are dealing with symmetrically distributed random variates, and then simply “mirror” them about their expected value. In our particular case of logreturn sequences, what we are going to do is take the negative of each logreturn and this way create a “mirrored” trajectory. Hopefully, this new trajectory will be negatively correlated to the original one¹¹. Now let us create two samples: one from the original trajectories and another from the antithetic trajectories. Calling them sample A and sample B , let us create a new estimator:

$$\begin{aligned}\hat{E}_{anti}(X) &\doteq \frac{\hat{E}(X_A) + \hat{E}(X_B)}{2} = \\ &= \frac{\left(\sum_{i=1}^{n/2} X_i\right) / (n/2) + \left(\sum_{i=1}^{n/2} X_i^{(a)}\right) / (n/2)}{2} = \sum_{i=1}^{n/2} \frac{X_i + X_i^{(a)}}{n}\end{aligned}$$

where X_i and $X_i^{(a)}$ is the payout based on the i^{th} original and antithetic trajectory, and they are IID random variates, respectively. This has a variance of

$$\begin{aligned}D^2(\hat{E}_{anti}(X)) &= D^2\left(\sum_{i=1}^{n/2} \frac{X_i + X_i^{(a)}}{n}\right) = \\ &= \frac{\sum_{i=1}^{n/2} D^2(X_i) + \sum_{i=1}^{n/2} D^2(X_i^{(a)}) + \sum_{i=1}^{n/2} 2Cov(X_i, X_i^{(a)})}{n^2} = \\ &= \frac{D^2(X)/2 + D^2(X^{(a)})/2 + Cov(X, X^{(a)})}{n} \approx \\ &\approx \frac{D^2(X) + Cov(X, X^{(a)})}{n}\end{aligned} \tag{3.15}$$

¹¹Based on my empirical experience, there is a strong negative correlation between the payouts calculated from the original trajectories and those calculated from the antithetic ones. We are going to discuss this later in section 4.2.

where we have made use of the symmetric nature of the random walk, implying that the payout based on the antithetic trajectory has the same variance as the one based on the original trajectory.

The benefit of using this method is twofold

1. It is apparent from the above results that as long as the original and antithetic payouts have a negative covariance, **the variance of the estimator will be reduced compared to that of the basic method**, while the estimator remains unbiased. As we will see later in section Empirical analysis, the variance reduction achieved by this relatively simple method is substantial.
2. Since we use every random number twice, we only need to generate $n/2$ outcomes instead of the original n ones, which means that **our simulation has already become (almost) twice as efficient**.

3.4.3 Variance reduction by control variates

We have seen in the example of antithetic variates that it is possible to keep our estimator unbiased, while significantly reducing its variance. It is natural to be curious whether we can do better and cut down on the estimation variance further. As it turns out, it is possible to improve on our technique by using control variates instead of antithetic variates (Kleijnen et al., 2010). This method is based on the logic that if we add an expression to our original estimator that has zero expected value and is negatively correlated with the original estimator, then variance reduction could be achieved while our estimator remains unbiased. Let us denote the control variate with Y , while preserving our original variate as X . In this case, the estimator will be

$$\hat{E}_{control}(X) \doteq \hat{E}(X) + c\hat{E}(Y) = \sum_{i=1}^n \frac{X_i + cY_i}{n}$$

where c is an appropriately chosen constant, X_i and Y_i are the i^{th} outcomes of the original and control variates in our simulation, and they are IID random variates, respectively. It is easy to see that the variance will be

$$\begin{aligned} D^2(\hat{E}_{control}(X)) &= D^2\left(\sum_{i=1}^n \frac{X_i + cY_i}{n}\right) = \\ &= \frac{\sum_{i=1}^n D^2(X_i) + \sum_{i=1}^n D^2(cY_i) + \sum_{i=1}^n 2Cov(X_i, cY_i)}{n^2} = \\ &= \frac{D^2(X) + c^2 D^2(Y) + 2cCov(X, Y)}{n} \end{aligned} \quad (3.16)$$

Remark 3.24 *The way such a control variate is usually constructed is by picking a random variate η , whose expected value can be analytically calculated. Now, if we subtract the estimated and the true expected value from each other*

$$Y \doteq \eta - E(\eta) \quad (3.17)$$

then we have created a random variate Y with zero expected value. Since the subtraction of a constant term does not alter the variance (or the covariance with other random variates), all we need to pay attention to is picking a proper random variate η that is negatively correlated with our original variate.

Remark 3.25 *It can be easily shown (by derivation with respect to c) that the optimal value of the constant is $c^* = -\frac{Cov(X,Y)}{D^2(Y)}$ which minimizes the variance of the estimator.*

Kemna and Vorst (1990) showed that in the case of the arithmetic-average Asian option price, the use of the geometric-average Asian option price can be used as an effective control variate. In a later research, Fu, Madan and Wang (1999) compared using the discrete-time and continuous-time analytical formulas in equation (3.17) with interesting results. What they have found is that even though the appropriate formula for a discretized estimator would be the discrete-time formula, by using the continuous-time formula a greater variance reduction could be achieved. In other words, instead of following the best practice presented above in remark 3.24, by using a “biased” (more precisely, non-zero expected value) control variate, such as

$$\tilde{Y} \doteq \eta - E(\tilde{\eta})$$

we not only reduce the variance of our original estimator but also add an appropriate bias to it, which compensates for the above mentioned discretization error inherent in the basic simulation method. For this reason, I also decided to use this “discretization error adjusted version” of the control variate method in my calculations.

It is crucial to notice the advantages of this method:

1. It **reduces the variance** of our original estimator introduced in the basic method even more so than the method of antithetic variates¹²;

¹²Based on my experience, the correlation between the arithmetic-average Asian option price and the geometric-average Asian option price is very high (above 0.9), which indicates that a very significant variance reduction could be achieved. For further details refer to section Empirical analysis.

2. It effectively **eliminates some of the discretization error** inherent in the simulation technique by using a “biased” control variate as suggested by Fu, Madan and Wang (1999).

3.4.4 Implementation

Before proceeding with the empirical analysis of our valuation framework in the next section, let me add a few notes with regards to the simulation method discussed in this section.

1. I performed the simulation following the algorithms presented in this section. Although the true variances and covariances are unknown, I estimated them also from the sample, making the necessary adjustment of having $(n - 1)$ instead of n in the denominator.
2. The biggest challenge during simulation is to **avoid running out of memory**. Using today’s modern computers, this is usually not an issue when simulating only the endpoints of trajectories as in the case of European-exercise vanilla options. However, when the entire trajectory needs to be generated and stored in the memory, then we run out of resources alarmingly fast. Once this happens, the operation systems starts to use the virtual memory on the hard disk, which is by orders of magnitude slower than SDRAM memory, and the whole simulation process suddenly comes to a halt. Without going into further technical details, what I did to avoid this eventuality is to break each trajectory into small enough sections that the generated logreturn matrix fits nicely into the memory. Then I rolled over the trajectory and calculated the (quadratic) average step by step, efficiently reusing the available SDRAM memory space again and again as shown in Figure 10.


4 Empirical analysis

Having introduced various algorithms for the valuation of Asian options, now we are ready to generate outputs in MATLAB[®] and observe their efficiency and computational characteristics. First, we will examine how the two semi-analytical (i.e. Euler and Talbot) algorithms fare compared to each other as well as attempt to create an Asian call option pricing table by means of these algorithms. Then, we will contrast simulation outputs for different time partitions and sample sizes. Ultimately, we will compare the semi-analytical and the simulation algorithms and make suggestions on their appropriate usage.

4.1 Analytical results

At first glance, the semi-analytical algorithms presented in the previous section might look awkward but it turns out that **they are very efficient for a wide range of parameters** of the Asian call option price $C(S, K, r, T, \sigma)$. As Fu, Madan and Wang (1999, p. 21.) pointed out, “difficulties [...] seem to begin when $\sigma^2(T - t) < 0.01$ ”. My calculations confirm that the critical parameter is indeed $\sigma^2 T$, which I call “**volatility potential**”, since it fundamentally governs how far away the underlying product’s price can get from its current state during the lifetime of the option. This variable directly determines the place h at which the normalized option price $C^{(\nu)}(h, q)$, defined in equation (3.8), has to be calculated.

Figure 2: Comparison of analytical algorithms



Source: own calculations.

For the sake of simplicity, let us fix σ at one¹³ and observe what effect the changing

¹³Naturally, $\sigma = 1$ is an unrealistic scenario, but it makes easier for us to capture the changes in $\sigma^2 T$.

of T has on the calculation of an ATM ($S = K = 1$) Asian call option's price. As we have discussed in section 3.3.2, the main parameters of the Euler¹⁴ and Talbot algorithms, driving computational precision, are n and M , respectively. Now, the way I measured the efficiency of each algorithm is iteratively increase the precision parameters until the last four decimal digits of the option price is stabilized. The results are shown in Figure 2, where I depicted in log-log charts the times and parameter levels that each algorithm required to reach the desired (i.e. four-decimal) precision. From these charts, we may conclude the following:

- **The Euler algorithm is more robust** than the Talbot algorithm with respect to calculation time or parameter setting. In general, setting $n = 10$ produces an Asian option value of at least four significant digit precision.
- **The Talbot algorithm is more reliable** than the Euler algorithm in the sense that if m is set sufficiently high, then the precise Asian option value can be calculated for arbitrarily low $\sigma^2 T$ (given enough computation time). In contrast, the Euler algorithm fails to converge when $\sigma^2 T < 0.005$, no matter how high n is set.
- When $\sigma^2 T$ falls below 0.01 in the case of the Euler algorithm and below 0.1 in the case of the Talbot algorithm, we can see a rapid increase in both calculation times and parameter levels. However, when $\sigma^2 T$ stays out of the “danger zone”, **the semi-analytical method proves to be extremely efficient** and solve the inversion problem under 2-3 seconds.

Remark 4.1 *Based on the calculations in section 4.1.2, I found that the above results hold for different combinations of option parameters (i.e. the value of $\sigma^2 T$ fundamentally determines the difficulty of the inversion problem, regardless of the other parameters).*

4.1.1 Weakness of the Laplace transform approach

A question that remains to be answered is why the chosen semi-analytical algorithms fail to converge when $\sigma^2 T$ is low¹⁵. Considering that this is a phenomenon that also other authors, such as Geman and Eydeland (1995), Fu et al. (1999) and

¹⁴Fu et al. (1999) did an extensive research on tweaking of the Euler algorithm's other two parameters (i.e. A and m), concluding that the default parameter setting (i.e. $A = 18.4$ and $m = 11$) suggested by Abate and Whitt (1995) is indeed optimal. Thus, I left these parameters unchanged in my research.

¹⁵Which means that either the maturity of the Asian option is close or the volatility of the underlying product is low, implying that the option has little “time value” left (refer to Száz (2009)).

Craddock et al. (2000) found, it seems that as $\sigma^2 T \rightarrow 0$, inversion complications arise due to the inherent numerical characteristics of the Laplace transform, no matter which inversion algorithm we chose initially. Indeed, observing the charts presented in Appendix C: Behavior of the Laplace transform, we find that as $\sigma^2 T \rightarrow 0$, the Laplace transform “gets numerically out of hand”:

- In the case of the Euler algorithm, the Laplace transform oscillates and quickly tends to zero as k changes, making it unnecessary to perform the algorithm for $n > 20$ (and limiting the effect of our precision parameter at the same time). The problem is that as $\sigma^2 T \rightarrow 0$, the Laplace transform rapidly diminishes to orders below 10^{-10} , making the embedded numerical integration (refer to equation (GY LT)) increasingly hard to calculate with the desired precision.
- In the case of the Talbot algorithm, the Laplace transform shows some oscillating nature again as k changes but remains zero outside a given “location of oscillation”. The problem now is not that the transform diminishes to magnitudes which are hard to handle numerically but that this specific place of oscillation is unknown ex ante, moreover, it shifts along the k axis towards infinity as $\sigma^2 T \rightarrow 0$. This means that in order to reach the desired precision, we have to set k large enough that the “location of oscillation” is included in the final summation within the algorithm, which entails calculating the Laplace transform for an increasing number of times.

Remark 4.2 *Although in case of the Euler-algorithm we quickly run into a “dead-end street”, since the embedded integral cannot be calculated numerically to an arbitrary precision, the deformation of the Bromwich-contour in case of the Talbot algorithm seems to leave some chance for improvement. If we could somehow (either numerically or analytically) pinpoint the “location of oscillation” without having to calculate all those zeros before the oscillation unnecessarily, we could significantly improve on the algorithm’s performance. For now, I leave this research direction open.*

Remark 4.3 *Facing the serious numerical difficulties of Laplace transform inversion, we fully understand why Craddock et al. (2000, p. 79.) claimed that “numerical inversion of Laplace transform is still more of an art than a science. [...] Even for a specific problem, such as pricing an Asian option, no method deals satisfactorily with all possible subcases that we may reasonably expect to arise in practice.”*

4.1.2 Asian option price tables

After solving a relatively complicated financial mathematical problem, the question naturally arises how to make the results available and usable for practitioners. For example, it is a well-known result that according to the Black-Scholes-formula for vanilla call options, the percentage price (i.e. c/S) of the option is governed by only two parameter combinations:

- “moneyness”, which is K/S
- “volatility potential”, which is $\sigma^2 T$

This implies that we can create a 2D table showing the percentage prices of vanilla call options along these two parameter combinations, and this table can be readily used by market practitioners without having to use or even know the pricing formula. Given that the “volatility potential” was our critical parameter at the inversion of the Laplace transform, I was intrigued whether a similar table could be created for Asian call options. Looking at pricing formula (3.9) and calculating the (initial) percentage price of the Asian call option

$$\begin{aligned} \frac{C(0)}{S(0)} &= e^{-rT} \frac{4}{\sigma^2 T} \cdot C^{(\nu)} \left(\frac{\sigma^2 T}{4}, \frac{K}{S(0)} \frac{\sigma^2 T}{4} \right) = \\ &= e^{-rT} \frac{4}{\sigma^2 T} \cdot E_{\mathbb{P}} \left(\int_0^{\sigma^2 T/4} \exp \left\{ 2 \left(\underbrace{\left(\frac{2\sigma^2}{r} - 1 \right)}_{\nu} t + \widetilde{W}_{\mathbb{P}}(t) \right) \right\} dt - \frac{K}{S(0)} \frac{\sigma^2 T}{4} \right)^+ \end{aligned}$$

we can realize (in a rather disheartened way) that there are four parameter combinations governing the option’s price:

- “moneyness”, which is $m \doteq K/S$
- “volatility potential”, which is $h \doteq \sigma^2 T$
- “risk-free bond price”, which is $P \doteq e^{-rT}$
- “nu”, which is $\frac{2\sigma^2}{r} - 1$

Luckily, ν could be expressed in terms of the risk-free bond price and the volatility potential

$$\nu = -2 \left(\frac{\ln P}{h} - 1 \right)$$

which means that m , h and P are sufficient to determine the percentage price of an Asian call option. This results in a **3D table, which I calculated**¹⁶ for different levels of m and various combinations of h and P . The price of the Asian put option can also be easily determined from these tables, by using the put-call parity relationship presented in (3.11), for which the knowledge of m , h and P is sufficient.

4.2 Simulation results

We have investigated how well the two numerical algorithms perform for different combinations of option parameters. Now, before comparing them with simulation methods, let us look into the latter on their own and examine the three Monte Carlo algorithms that we introduced in section 3.4. As we have explained earlier, in the case of path-dependent derivatives (such as Asian options) simulation techniques have two governing parameters:

- n , which is the number of trajectories simulated
- m , which is the number of fragments each day is partitioned into

where n determines the estimation error, while m determines the discretization error. Striving to **isolate the effects of these parameters** as much as possible, first we will examine what level of m leads to a satisfactory approximation of the continuous Asian option price, and only then will we turn to adjusting n in order to attain an estimate with desirable precision. Just as in the case of semi-analytical methods, we will restrain our analysis to ATM Asian call options, where $S = K = 1$ and $\sigma = 1$. The benefit of this restriction is threefold:

- the simulation results will be comparable with the numerical ones
- the relationship between T and various simulation errors will serve as a proxy for the relationship between $\sigma^2 T$ and various simulation errors¹⁷
- observable (market) volatilities fall below 100%, and thus if we manage to fine-tune our simulation algorithms so as to produce satisfactory estimates even under these dire circumstances, then the simulated errors will be an upper-bound of those which we could possibly face in real-world situations

¹⁶Because of the strict format restrictions of this work, I had to omit these tables even from the appendix due to a lack of available space.

¹⁷This makes perfect sense as in our framework the logreturns are modeled as independent normally distributed random variates with variance σ^2 . Hence, simulation variances should be proportional to $\sigma^2 T$.

However, before proceeding with the fine-tuning of the above parameters and measuring the performance of the different simulation algorithms, let us have a word about what level of error we consider “desirable”. So that we can compare the performance of simulation methods with the performance of semi-analytical methods, we should aim at the same (i.e. roughly four-decimal digit) precision. All we have to do then is to set parameters n and m so that the confidence interval¹⁸ around the estimated option value is sufficiently small to ensure the four-decimal precision. Since the probability distribution of the Asian option price (i.e. the random variate whose value we strive to estimate) is hard to determine, let us apply Chebyshev’s inequality to give a rough approximation for the confidence interval by means of the multiple of the standard deviation:

$$P(|E(X) - X| \geq k\sigma) \leq \frac{1}{k^2}$$

which tells us that choosing $k = 4$, we can be about 95% certain that the true Asian option value falls between $\hat{E}(X) - 4\sigma$ and $\hat{E}(X) + 4\sigma$. In other words, σ must be of order 10^{-5} so that the desired four-decimal precision is satisfied.

4.2.1 Discretization error

Keeping in mind what we have just concluded about simulation precision, let us determine what level of m ensures a good approximation of the underlying product’s continuous trajectory. If we set $n = 100,000$, which is relatively large, we expect to receive simulation results with low simulation error. This implies that by changing m , we should be able to discern changes in discretization error, only slightly disturbed by any residual simulation error in the output. Applying this logic, I ran the simulation for different maturities and gathered the output in Appendix D: Charts to Monte Carlo simulation displaying the absolute error, which I defined as the absolute difference of the simulated option value and the true (semi-analytical) option value. From these charts, we can deduce the following:

- The control variate method has lower (discretization) error than the other two methods by about one order of magnitude.
- There is a rapid decrease in discretization error when using $m = 10$ instead of $m = 1$. However, the discretization error reduction is not significant when m is further increased to $m = 100$.

¹⁸For this we should also pick a suitable significance level. For now, let this be 95%.

- Interestingly, the antithetic variate method does not reduce but rather amplifies discretization error. On second thought, this phenomenon is not surprising at all, as in the case of the antithetic variate method we use each generated random number twice so as to cut down on simulation error and time. However, this also implies that the discretization error (if any) in a given simulated trajectory is emphasized as there are less independent trajectories to offset it.
- By taking a closer look at the absolute errors for different levels of m , we might notice patterns, such as peaks or troughs, especially in case of the first two methods. This is not an accident, as I used the same random number series¹⁹ for each simulation, which means that longer simulations (i.e. where T is larger) will include the same random number sequences as their shorter counterparts and hence any discretization error will also be preserved. This “slowly decaying bias” in our simulation comes in very handy now, making it easier for us to identify those simulation methods which effectively reduce discretization error. For example, it is very hard to discern patterns in the chart of the control variate method, which implies that this method does exactly what it professes, namely that it reduces discretization error inherent in the simulation of the arithmetic-average by offsetting it with the one inherent in the simulation of the geometric-average.

From the above results, we can conclude that it is best to choose the control variate method, as it effectively reduces discretization error. Also, it is unnecessary to set m higher than 10, as creating more refined trajectories significantly increases simulation time but does not bring significant discretization error reduction to the system. Hence, from now on we will work with $m = 10$.


4.2.2 Simulation error

From a perspective of discretization error reduction, the method of control variates clearly is the best performing one. Now, let us make a comparison between simulation methods based on the simulation error of their outputs, measured by the standard deviation of the estimate.

The estimated standard errors produced by each simulation method for $n = (\#) \cdot 10^3$ trajectories are shown on a log-log scale in Figure 3, from which we could draw the following inferences:

¹⁹This could easily be achieved by initializing the random number generator with the same so-called “seed value” before each simulation.

Figure 3: Comparison of simulation methods ($m = 10$)



Source: own calculations.


- The antithetic method achieves only a slight reduction in simulation variance compared to the basic Monte Carlo method.
- The control variate method achieves an extreme reduction in simulation variance compared to the basic Monte Carlo method. The variance reduction is so efficient that we may even think of using only $n = 10,000$ trajectories.
- The standard error from regular and antithetic methods grows only at a rate $\sim x^{0.55}$, whereas the control variate method's standard error grows at a faster rate $\sim x^{1.05}$. It looks as though the significant discretization error reduction that we could observe earlier has a price, and thus the standard error of the control variate method catches up with the standard error of other methods as $\sigma^2 T$ increases.

The question naturally arises whether we could quantify the variance reduction by the antithetic and the control variate methods. For this, we need to estimate the correlation of antithetic variates and control variates, respectively. As we can see in Figure 4, these correlations are not constants but they depend on the volatility potential.

Now, from equations (3.15) and (3.16) we know that

$$D^2(\hat{E}_{anti}(X)) \approx \frac{D^2(X) + Cov(X, X^{(a)})}{n} \approx \frac{D^2(X) \left(1 + \rho(X, X^{(a)})\right)}{n}$$

Figure 4: Correlation of proxy variates in variance reduction methods



Source: own calculations.

and

$$\begin{aligned}
 D^2(\hat{E}_{control}(X)) &= \frac{D^2(X) + c^2 D^2(Y) + 2cCov(X, Y)}{n} \approx \\
 &\approx \frac{D^2(X) - \frac{Cov^2(X, Y)}{D^2(Y)}}{n} \approx \frac{D^2(X) (1 - \rho^2(X, Y))}{n}
 \end{aligned}$$

provided $D^2(X) \approx D^2(X^{(a)})$ and $D^2(X) \approx D^2(Y)$. Comparing these simulation errors with the variance of the basic Monte Carlo method, which is

$$D^2(\hat{E}(X)) = \frac{D^2(X)}{n}$$

we can measure the variance reduction. In the case of the antithetic method, $\rho \in (-0.4, -0.2)$ ensures a 20-40% variance reduction, which is roughly 10-20% standard error reduction. In contrast, $\rho \approx 1$ in the case of the control variate method results in a close to 100% standard error reduction. Also, by looking at the formulas, where $D^2(X)$ is multiplied by $(1 + \rho)$ in the antithetic and by $(1 - \rho^2)$ in the control variate case, we can imagine why the simulation error of the latter method grows more faster as $\rho \rightarrow 0$.

From the results presented, we may conclude that **the control variate method effectively reduces both discretization error and simulation error** even for $m = 10$ and low values of n . Using this method, standard error remains below the desired 10^{-5} if

Table 1: Limits for the variance reduction method ($m = 10$)

$\sigma^2 T$	n
< 0.025	10,000
< 0.05	100,000
< 0.1	1,000,000

Source: own calculations.

4.3 Comparison of methods

Having compared the different numeric and simulation methods separately and having identified the optimal parameter values for each, let us compare the two approaches in the search for an answer to the research question of this work, namely “What method (algorithm) should we use for pricing Asian options?”

We have seen that the Euler and Talbot algorithms are powerful tools for pricing Asian options, though being rather sensitive to low values of the “volatility potential”, and have also concluded that the control variate method is able to pinpoint the price of Asian options very accurately by efficiently reducing discretization and simulation variance. The only question that remains is how fast these algorithms are compared to one another. Is there a preference order we could set, maybe even as a function of $\sigma^2 T$? To answer this question, let us have a look at Figure 5, where we can see what time it takes for


- the Euler algorithm to calculate the Asian option price to the desired (i.e. four-decimal) precision
- the control variate algorithm to run a simulation for a given n , where each curve is calculated until the desired precision is maintained

Let us note that while the time required by numeric algorithms is fundamentally a function of $\sigma^2 T$, simulation times depend basically on T (i.e. how long trajectories need to be generated). This implies that for each σ , the relative position of timings change, as we may observe the differences between Figure 5 and Figure 9 (in Appendix D: Charts to Monte Carlo simulation), calculated for $\sigma = 20\%$ and $\sigma = 10\%$, respectively²⁰.

Therefore, we may claim that **the control variate algorithm generally beats numerical algorithms in terms of calculation time** for usual values of σ and

²⁰In usual market scenarios stock price volatility lies between these two values, and thus Figure 5 and Figure 9 show the extremities for the σ parameter.

Figure 5: Comparison of calculation times of different methods ($m = 10$)



Source: own calculations.

T . However, we should also incorporate another factor in our analysis before making a conclusion, namely calculation precision. As we could observe in Figure 3 and Table 1, even when using the control variate method, **the standard error of the estimate exceeds our desired level of precision when $\sigma^2 T$ is above a given limit**, indicating that it might be worth to switch to numeric algorithms to maintain precision. Based on our findings, we may set up the following decision rules to optimize the Asian option pricing methodology, where the choice between the Euler algorithm and control variate algorithm depends on the specific value of σ , as indicated earlier.

Table 2: Optimal method selection rules for pricing Asian options

$\sigma^2 T$	optimal method
< 0.025	control ($n = 10k$)
< 0.05	control ($n = 100k$) or Euler ($n = 10$)
< 0.1	control ($n = 1000k$) or Euler ($n = 10$)
	Talbot ($M = 10$)

Source: own calculations.

5 Conclusion

Having explained the basics of “risk neutral valuation” in section Valuation methodology, we classified the various option types and expanded on the different approaches to Asian option pricing in section Pricing Asian options. Then, we saw in section Empirical analysis how the two semi-analytical (numerical) methods and the three simulation methods that this work focuses on fare for different combinations of option parameters.

Especially, we have pointed out the importance of $\sigma^2 T$ in the solution of the Asian option pricing problem and presented how it influences the choice of valuation method. Also, we have seen that there is an interesting dichotomy between the usage of numerical and simulation methods as one group of methods tends to excel when the other shows poor performance and vice versa, depending on the value of $\sigma^2 T$ as shown in Table 2. This is good news in the sense that the two approaches seem to compensate each other so that **we will always have an algorithm at hand which is able to price an Asian option within 10 seconds with a precision of four decimal digits**²¹.

Of course, one could argue that either when a higher decimal precision is required or the scale parameter (i.e. the stock price S) of the stock price is larger, the timings in Figure 5 are not realistic. This is true, but before making new charts or improvements to the algorithms, we should consider whether it is sensible to calculate derivative prices to five or six decimal digits precision at all, as these financial products are very sensitive to the value of their parameters and there is also a significant estimation error in the input parameters (such as σ or S).

The only problematic case is when market volatility is extremely (unrealistically) low (i.e. below 5%). In such cases, the value of σ^2 is very small and numerical algorithms fall into the “danger zone” for usual values of T . This means that unless T is very large, we are constrained to using simulation methods, which maybe a computationally demanding task.

Finally, with respect to the limitations of the Black-Scholes framework, further research could be done into the relaxation of its assumptions, such as

- the incorporation of dividend flow into the model

²¹In my perception, this speed and precision is sufficient for most market circumstances.

- the relaxation of the constant volatility assumption, where time-dependent (temporal) or place-dependent (local) volatility structures could be applied
- the use of more general Lévy processes instead of the Wiener process to model the logarithmic returns of the stock price)

Though mathematically challenging, these research directions are auspicious and likely to lead to a deeper insight into the behavior and valuation of Asian options.

References

- Abate, J. and Whitt, W. (1995): Numerical inversion of Laplace transforms of probability distributions. *ORSA Journal on Computing* 7: 36–43.
- Abate, J. and Whitt, W. (2006): A unified framework for numerically inverting Laplace transforms. *INFORMS Journal on Computing* 18: 408–421, doi:10.1287/ijoc.1050.0137.
- BIS (2012): Quarterly Review. Technical report, Bank for International Settlements, URL http://www.bis.org/publ/qtrpdf/r_qt1212.htm.
- Black, F. and Scholes, M. (1973): The pricing of options and corporate liabilities. *Journal of Political Economy* 81: 637.
- Bodnar, G. M., Hayt, G. S. and Marston, R. C. (1998): Survey of financial risk management by U.S. non-financial firms. *Financial Management* : 70–91.
- Carr, P. and Schröder, M. (2003): Bessel processes, the integral of geometric Brownian motion, and Asian options. *Rossiiskaya Akademiya Nauk. Teoriya Veroyatnostei i ee Primeneniya* 48: 503–533, doi:10.1137/S0040585X97980543.
- Cohen, A. M. (2007): *Numerical methods for Laplace transform inversion*, Numerical Methods and Algorithms 5. New York: Springer.
- Craddock, M., Heath, D. and Platen, E. (2000): Numerical inversion of Laplace transforms: A survey of techniques with applications to derivative pricing. *Journal of Computational Finance* 4: 57–82.
- Dufresne, D. (2004): *Bessel Processes and a Functional of Brownian Motion*. Centre for Actuarial Studies, Department of Economics, University of Melbourne.
- FASB (2006): *SFAS 157: Fair Value Measurements*. Financial Accounting Standards Board.
- Fu, M. C., Madan, D. B. and Wang, T. (1999): Pricing continuous Asian options: A comparison of Monte Carlo and Laplace transform inversion methods. *Journal of Computational Finance* 2: 49–74.
- Geman, H. and Eydeland, A. (1995): Domino effect. *Risk* 8: 65–67.
- Geman, H. and Yor, M. (1993): Asian options, Bessel processes and perpetuities. *Mathematical Finance* 2: 349–375.

- Hull, J. C. (2009): *Options, futures, and other derivatives*. Pearson.
- Kalos, M. H. and Whitlock, P. A. (2008): *Monte Carlo methods*. Wiley-Blackwell, Weinheim, 2nd ed., doi:10.1002/9783527626212.
- Kano, P. O. (2010): Numerical Laplace transform inversion and selected applications. URL http://math.arizona.edu/~brio/WEEKS_METHOD_PAGE/NLAPcolloquium2010.pdf, presentation.
- Kemna, A. G. Z. and Vorst, A. C. F. (1990): A pricing method for options based on average asset values. *Journal of Banking & Finance* 14: 113–129.
- Kleijnen, J., Ridder, A. and Rubinstein, R. (2010): Variance reduction techniques in Monte Carlo methods .
- Lamperti, J. (1972): Semi-stable Markov processes. I. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* 22: 205–225.
- Medvegyev, P. (2012): Pénzügyi matematika. Budapesti Corvinus Egyetem, egyetemi jegyzet.
- Medvegyev, P. and Száz, J. (2010): *Meglepetések jellege a pénzügyi piacokon*. Budapest: GT-Print Kft.
- Merton, R. C. (1973): Theory of rational option pricing. *Bell Journal of Economics and Management Science* 4: 141–183.
- Mitchell, R. L. (1968): Permanence of the log-normal distribution. *JOSA* 58: 1267–1272.
- Márkus, L. (2012): Eszközár folyamatok modellezése, valamint európai és amerikai opciók árazása. Eötvös Lóránd Tudományegyetem, egyetemi jegyzet.
- Száz, J. (2009): *Devizaopciók és részvényopciók árazása*. Bankárképző központ.
- Turnbull, S. M., McLean, S. and Wakeman, L. M. (1991): *A quick algorithm for pricing European average options*. Cambridge University Press.
- Wikipedia (2013): Floating point. URL http://en.wikipedia.org/wiki/Floating_point.

Appendix A: Euler algorithm


```
1 function x = Euler(n,P,S,K,T,sigma)
2 %% Recommended parameters by Abate-Whitt (1995).
3 m = 11;
4 A = sym(18.4);
5
6 %% Parameters of the normalized option price.
7 h = sym((T*sigma^2) / 4);
8 q = sym((K*T*sigma^2) / (4*S));
9 nu = sym(-2*log(P) / (sigma^2*T) - 1);
10
11 %% Lambda calculation (Euler algorithm).
12 arg = sym(zeros((n+m) + 1,1));
13 for k = 0:(n+m)
14     arg(1+k) = (A + 2*k*pi*1i) / (2*h);
15 end
16
17 if sum(single(abs(arg)) < max(0, 2*(single(nu)+1))) ~= 0
18     x = -999; % An easy to identify numeric error output.
19     return
20 end
21
22 try
23 %% Substitution into the Laplace transform.
24 mu = sqrt(2*arg + nu^2);
25 alpha = (mu + nu)/2;
26 beta = (mu - nu)/2;
27
28 integrand = @(alpha,beta,q,u) u.^(double(beta)-2).* ...
29     (1-u).^(double(alpha)+1).*exp(-u/(2*double(q)));
30 a = sym(integral(@(u)integrand(alpha,beta,q,u), 0, 1, ...
31     'ArrayValued', true, 'AbsTol', 0, 'RelTol', 1));
32 a = log(a) + log(2*q)*(1-beta) - log(2.*arg) ...
33     - log(alpha+1) - mfun('lnGAMMA',double(beta));
34 a = real(exp(a));
35
36 %% Summation (Euler algorithm).
37 s = (exp(A/2)/h).*((-1).^(0:(n+m)))'.*a;
38 s(1) = s(1)/2;
39 for k = 1:(n+m)
40     s(1+k) = s(1+k) + s(k);
41 end
42
43 x = sym(0);
44 for k = 0:m
45     x = x + nchoosek(m,k)*2^(-m)*s(1+(n+k));
46 end
47
48 x = x*P*(S*4)/(T*sigma^2);
49 x = double(x);
50
51 catch
52     x = double(-999);
53 end
54
55 end
```

Appendix B: Talbot algorithm


```
1 function x = Talbot(M,P,S,K,T,sigma)
2 %% Parameters of the normalized option price
3 h = sym((T*sigma^2) / 4);
4 q = sym((K*T*sigma^2) / (4*S));
5 nu = sym(-2*log(P) / (sigma^2*T) - 1);
6
7 %% Lambda calculation (Talbot algorithm).
8 d = sym(zeros(M,1));
9 d(1) = 2*M/5;
10 for k = 1:(M-1)
11     d(k+1) = 2*k*pi/5 * (cot(k*pi/M) + 1i);
12 end
13
14 g = sym(zeros(M,1));
15 g(1) = 1/2 * exp(d(1));
16 for k = 1:(M-1)
17     g(k+1) = (1 + 1i*(k*pi/M)*(1 + (cot(k*pi/M))^2) ...
18             - 1i*cot(k*pi/M))*exp(d(k+1));
19 end
20
21 arg = d/h;
22
23 if sum(single(abs(arg)) < max(0, 2*(single(nu)+1))) ~= 0
24     x = -999; % An easy to identify numeric error output.
25     return
26 end
27
28 try
29     %% Substitution into the Laplace transform.
30     mu = sqrt(2*arg + nu^2);
31     alpha = (mu + nu)/2;
32     beta = (mu - nu)/2;
33
34     integrand = @(alpha,beta,q,u) u.^(double(beta)-2).* ...
35         (1-u).^(double(alpha)+1).*exp(-u/(2*double(q)));
36     a = sym(integral(@(u)integrand(alpha,beta,q,u), 0, 1, ...
37         'ArrayValued', true, 'AbsTol', 0, 'RelTol', 100));
38     a = log(a) + log(2*q)*(1-beta) - log(2.*arg) ...
39         - log(alpha+1) - mfun('lnGAMMA',double(beta)) + log(g);
40     a = real(exp(a));
41
42     %% Summation (Talbot algorithm).
43     x = sym(0);
44     for k = 0:(M-1)
45         x = x + a(k+1);
46     end
47
48     x = x*2/(5*h)*P*(S*4)/(T*sigma^2);
49     x = double(x);
50
51 catch
52     x = double(-999);
53 end
54
55 end
```

Appendix C: Behavior of the Laplace transform


Figure 6: Value of the Laplace transform using the Euler algorithm for $n = 20$



(a) $\sigma^2 T = 1$




(b) $\sigma^2 T = 0.1$




(c) $\sigma^2 T = 0.01$

Source: own calculations.


Figure 7: Value of the Laplace transform using the Talbot algorithm for $M = 50$



(a) $\sigma^2 T = 1$



(b) $\sigma^2 T = 0.1$




(c) $\sigma^2 T = 0.01$


Source: own calculations.

Appendix D: Charts to Monte Carlo simulation


Figure 8: Absolute error of Monte Carlo algorithms for different levels of m



(a) Basic method




(b) Variance reduction by antithetic variates



(c) Variance reduction by control variates


Source: own calculations.

Figure 9: Comparison of calculation times of different methods ($m = 10$)



Source: own calculations.

Figure 10: Optimized memory management of simulation algorithms



Source: own calculations.