Undergraduate Economic Review

Volume 14 | Issue 1 Article 10

2017

Pricing Asian Options: Volatility Forecasting as a Source of Downside Risk

Adam T. Diehl
James Madison University, diehlat@dukes.jmu.edu

Pricing Asian Options: Volatility Forecasting as a Source of Downside Risk

Abstract

Asian options are a class of derivative securities whose payoffs average movements in the underlying asset as a means of hedging exposure to unexpected market behavior. We find that despite their volatility smoothing properties, the price of an Asian option is sensitive to the choice of volatility model employed to price them from market data. We estimate the errors induced by two common schemes of forecasting volatility and their potential impact upon trading.

Keywords

Asian Options, Volatility, Option Pricing, Risk Neutrality, GARCH

Cover Page Footnote

The author would like to thank Professor Vipul Bhatt for his comments and suggestions during the preparation of this manuscript.

INTRODUCTION

Asian options are a class of exotic derivatives that use an averaging procedure to generate their payoff. Typically, the average value of the stock price is calculated over the life of the option and compared to a fixed strike price at the time of exercise.[1] Both geometrically averaged and arithmetically averaged options are traded in the market, but arithmetic averaging is more common despite being analytically intractable (Lakhlani 2013). This has spurred the development of many procedures for estimating the price of arithmetic Asian options, but Monte Carlo simulation remains the benchmark without a closed form solution.

The popularity of Asian options in the over-the-counter (OTC) market stems from two characteristics of their payoff. First, because their payoff is structured by the average of the underlying price, investors reduce their exposure to the risk of price manipulation near option maturity. Second, the averaging procedure reduces the volatility in expected payoff compared to a vanilla European option. This combination of features makes these options particularly well suited for hedging exposure in currency markets or thinly traded commodity markets, where they are most popular (Mraovic and Zhang 2014). Conveniently, this smoothing effect also results in cheaper option premiums, a result demonstrated by Figure 1. The relative cheapness of these contracts also makes Asian options more attractive to corporate financiers.

While Asian options reduce the impact of localized volatility upon the payoff, the selection of a volatility model is relevant to pricing accuracy because these options are strongly path-dependent. Models of asset prices differentiate themselves based on how they handle volatility (Bollerslev et al. 1992). Not predicting a short period of particularly high or low volatility (such as in events like the “Flash Crash” of 2011) is unlikely to significantly affect the final payoff, but poor model specification is relevant across the life of the option. Errors accumulate and propagate, sending the average 'off-track,' resulting in poor pricing accuracy and potentially significant losses (Bollerslev and Mikkelsen 1999). Therefore, generating a more accurate forecast of volatility will produce a more accurate option premium.

The pricing accuracy of different models can be examined with historical data. Splitting the data into a training set and a validation set provides a benchmark for comparison. First, volatility forecasts are created to price the options, and then the option pricing models are re-run with the true underlying parameters observed over the period that the option would have been active. The goal of this estimation is to discover how significant pricing errors can be when using market data.

1 While "floating-strike" Asian options exist, they are less commonly traded (Alziary 1997).

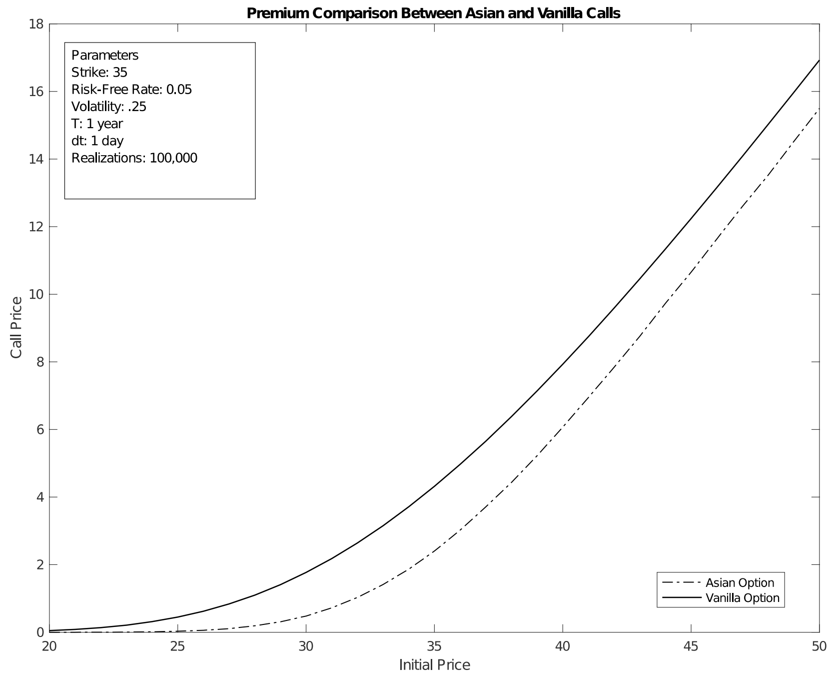

Figure 1: Premium comparison of Asian and vanilla options.

This paper is organized as follows: Section 2 reviews the current literature concerning Asian option pricing procedures. Section 3 provides the theoretical framework used to generate the pricing algorithm. Section 4 details the simulation methodology and the data used to generate the results that are analyzed in Section 5. Concluding remarks are presented in Section 6.

LITERATURE REVIEW

The primary complication in pricing arithmetic Asian options is that the sum of finitely many log-normal random variables has no closed form solution (Alziary et al. 1997). In general, the sum of two probability distributions is calculated using the convolution integral of their respective characteristic functions, but a closed form representation for the characteristic function of the log-normal distribution is unknown (Fenton 1960). While infinite summation methods exist, these are only approximations because the averaging frequencies seen in the market are not continuous. Oil options, for instance, are typically averaged daily or monthly (Levy 1992). Furthermore, numerical integration of the sums is difficult due to the tail behavior of the log-normal distribution (Beaulieu and Rajwani 2004). While various lognormal sum approximation methods exist, there does not exist one method that is "best" across the entire distribution (Mehta et. al. 2007). Others have tried to directly approximate the option payoff, and the most common technique is to use some implementation of the Lévy lognormal approximation. In addition, P.D.E. methods and more exotic techniques exist, but they usually require extra restrictions to be placed on the options contract (Alziary et al. 1997). However, virtually every approximation method uses Monte Carlo simulation as its benchmark, so this paper will focus only on direct simulation. For a detailed comparison of the accuracy of these approximation methods in the context of Asian option prices, see Hsu and Lyuu (2011) and Nielson (2001).

THEORETICAL FRAMEWORK

Consider a continuous time economy where T is the maturity date of an option contract. This economy has at least two assets that do not pay dividends: a risk-free asset earning a return r, and a risky asset. Asian option contracts can be written on this risky asset with a fixed strike price of K, and these contracts can be exercised at time T. The market information available at the current time t is represented by the filtration $\mathcal{F}_t$, where $\mathcal{F}_t$ is a σ-algebra of subsets of the sample space Ω, which in this case is any possible price the asset could take on at a given point in time (Lawler 2014). Let the price of the risky asset be an integrable random variable S_t where S_t is $\mathcal{F}_t$ measurable. Let A_T be the average price over the life of the option (the averaging period). Now define $\psi_{t,i}$ to be a sequence of probability spaces (Ω, $\mathcal{F}_{t,i}, \mathrm{P}_{t,i}$), where i is a natural number indexing an individual realization of the time-path of an asset price (each probability space will be represented by a single price path in the simulation). The following diagram neatly represents the pricing process two representations of the pricing process.

$$\begin{tikzcd}
\Psi_{t,i} \arrow[r, "E[A_T|\mathcal{F}_t]"] \arrow[dr, "{E[E[A_T|\mathcal{F}_t]]-K}"'] & \Psi_i \arrow[d, "{E[A_T]-K}"] \\
& \Psi
\end{tikzcd}$$

Figure 2: Commutative diagram showing the pricing calculation and a two-stage decomposition of said process.

To prevent arbitrage, the price of the option today must be the market's expectation of the payoff at time T discounted back to account for the time value of money. This payoff is conditioned upon the future prices of the asset, which is information not yet present within the market. The best guess of this payoff is therefore $E[A_T|\mathcal{F}_t]$. However, the specific path that the asset could follow changes as new price information enters the market. Therefore, this conditional expectation is itself a random variable. The expected payoff of the option given what is known today can be found by taking the (unconditional) expectation $E[E[A_T|\mathcal{F}_t]]$ across all potential stock price paths, and subtracting the stock price.

This composition of expectations can be decomposed into a two-stage estimation of the option payoff. Consider $\psi_{i,t}$ as the output of a Monte Carlo simulation with N price paths (realizations) at τ time points arranged into an $N \times \tau$ matrix. The first stage is to map $\psi_{i,t} \rightarrow \psi_i$, which outputs an $N \times 1$ column vector by taking an average for each path across time. Taking the mean of the path averages (and subtracting K) is equivalent to defining a mapping $\psi_i \rightarrow X$ to generate a point estimate for the option payoff, which is adjusted for the time-value of money to price the option. The next stage is to generate the $N \times \tau$ output of the simulation.

The price of a call and a put can be expressed as

$$c = e^{-r(T-t)} E[(\mathrm{A}_T - \mathrm{K})^+] \tag{1}$$

$$p = e^{-r(T-t)} E[(\mathrm{K} - \mathrm{A}_T)^+], \tag{2}$$

where A is the arithmetic average of the stock price, and E denotes the expectation under a risk-neutral measure (discarding negative payoffs). The task at hand is to generate S_t at discrete time nodes. According to geometric Brownian motion, the stochastic differential equation that models stock prices is

$$dS = \mu S dt + \sigma S dW_t, \tag{3}$$

where μ is the (annualized) expected return to the stock, σ is the (annualized) volatility, and dW_t is a Wiener process. It follows from Girsanov's theorem and the Novikov sufficient condition that a change of measure to a risk-neutral probability measure allows us to substitute μ with r to price the option (Steele 2001). Then by applying Ito's lemma and setting the relevant function of S to $\ln(S)$, the change in the log stock price is

$$\ln(S_t) - \ln(S_{t-1}) = \left(r - \frac{1}{2}\sigma^2\right) dt + \sigma dW_t. \tag{4}$$

This implies that in in discrete time

$$S_t = S_{t-1} \cdot e^{\left(r - \frac{1}{2}\sigma^2\right)\Delta t + \sigma z \sqrt{\Delta t}}. \tag{5}$$

THE GREEKS

Without a closed form pricing solution, the Greeks of an arithmetic Asian option must be computed numerically. In discrete time, the partial derivative is a partial difference. This paper focuses on two Greeks, Delta and Vega, as these highlight the primary difference between Asian and vanilla options.

The delta of an option measures its sensitivity to changes in the price of the underlying asset. Vanilla European options have a closed form expression for delta. The same cannot be said for (arithmetic) Asian options. A plot comparing the deltas of otherwise identical vanilla and Asian call options is below. The delta of an option is also a common proxy for the probability of exercise among options traders. The two deltas initially cross at the risk-free rate adjusted strike price, which is the at-the-money price, implying that they have the same probability of exercise here when the option is first written. The plot reveals an important concern for traders of Asian options, which is that the option is more sensitive to initial moneyness than its vanilla counterpart.

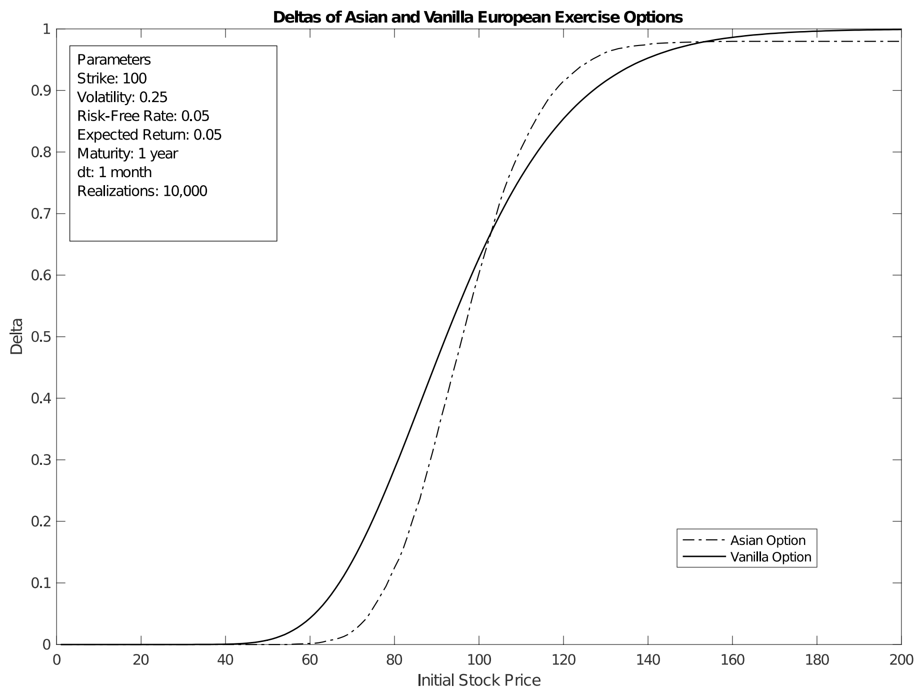

Figure 3: Deltas of two otherwise equivalent call options.

An option that begins significantly out of the money is less likely to be exercised, as it will likely accumulate subsequent "averaging days" that are also out of the money. This explains why the option contracts are cheaper; if they were not then arbitragers could make a profit offsetting their long positions by shorting Asian options that are significantly out of the money. A corollary of this effect is that the option is less sensitive to price manipulation; if counterparties seek to gain a profit by market manipulation, then it becomes significantly more expensive to engage in this manipulation because abnormal trading volumes must be sustained across the life of the option rather than just before exercise.

Turning now to volatility, vega denotes the sensitivity of the option price to changes in volatility in the underlying asset. Asian options demonstrate, ceteris paribus, a lower vega than their vanilla counterpart. This is demonstrated by the plot below of otherwise identical call options.

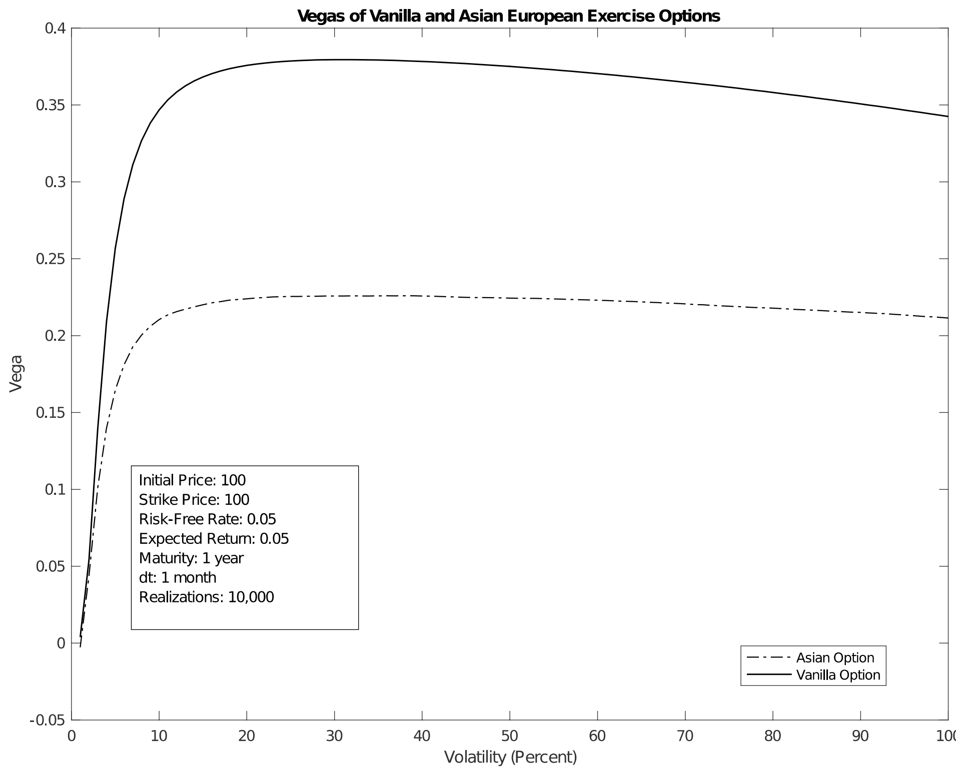

Figure 4: Vegas of two otherwise equivalent call options.

Outside of the low volatility range, vega tends to not have much curvature (particularly for Asian options), which is useful for vega hedging because even the presence of conditional heteroscedasticity is unlikely to significantly affect the option's vega. In the plot above, a 1% change in volatility will result in approximately a 22-cent change in the option premium whether the underlying's volatility is 20% or 60%. This implies that vega can be interpreted as the 'cost' of poor volatility forecasting. A higher vega implies a higher unit cost to an error in the volatility forecast used to price the option. However, even a low vega should not be underestimated, particularly in the highly-leveraged markets where Asian options are most popular. As a quick back of the envelope calculation, consider that, per these simulations, a reasonable vega for a currency near parity with USD is 0.001 per dollar of exposure on a 3 month, daily averaged contract. NASDAQ forex option lots are written on 10,000 dollars of the underlying currency and have a position limit of 600,000 contracts. On a position a tenth of this size, every percent error in the volatility forecast would be responsible for $600,000 worth of pricing error.

METHODOLOGY

Monte Carlo simulation does not require a closed form solution for the pricing function, which makes it a natural choice to deal with sums of lognormal random variables. Unfortunately, simulations can only provide estimates of the true value for the option price. For a point estimate provided by a run of the simulation, the standard error of the option price is on the order of $\frac{1}{\sqrt{N}}$ where N is the number of simulations (Mraovic and Zhang 2014). As $\sqrt{N}$ grows more slowly than N, large numbers of simulations are required to generate accurate results. The variance in the option price is approximately one to two cents when using 100,000 price paths. As the Monte Carlo simulation imposes no restriction on the nature of the option contract, it is assumed that the averaging period and the maturity period of the option are identical (as opposed to forward and backward starting options (Alziary 1997)), that options can only be exercised at maturity, and that all options are at the money at the beginning of the maturity period. All contracts have a maturity and averaging period of three months, with daily averaging. Option contracts are priced assuming a lot size of a single unit of the underlying stock.

Historical financial data was used to calculate the relevant option parameters. This paper utilizes three years of daily closing stock prices, beginning on January 2, 2001 and ending on January 2, 2004. This stock data was sourced from a Bloomberg terminal for a variety of randomly chosen stocks. All stocks during this period were screened to ensure the presence of ARCH effects, and the three-month period after January 2, 2004 did not exhibit any notable regime changes in the market. After screening the stock tickers, seven suitable stocks remained: CS (Credit Suisse), M (Macy's), MDT (Medtronic), NCR (NCR Corporation), POT (Potash Corporation of Saskatchewan), WFC (Wells Fargo), and WFT (Weatherford International). These companies all trade on the New York Stock Exchange but represent a variety of different industries, from financial services (CS and WFC) to computer hardware (NCR) to chemicals (POT). The risk-free rate is defined here as the average of the daily one month Treasury constant maturity rate, with data sourced from FRED.

Since the Black Scholes model assumes volatility is constant, three constant volatility estimates are generated from this data. The first is historical volatility, which is the most straightforward means of estimating volatility. Let ζ_t represent the time series of daily returns for a stock over the three-year training period. Then the historical volatility is

$$\sigma_h = \sqrt{252}\sqrt{\mathrm{Var}(\zeta_t)}. \tag{6}$$

Unfortunately, historical volatility is a backwards looking estimate, which provides cause for concern when attempting to predict future asset volatility. A common method of generating a forward-looking volatility estimate is to use GARCH (1,1), which conveniently also can output a constant volatility estimate. The GARCH volatility equation is

$$\sigma_t^2 = \alpha_0 + \alpha_1 X_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \tag{7}$$

where X_t is a white noise process (Posedel 2005). Restricting $0 \leq \alpha_1, \beta_1 < 1$ implies that σ_t is integrable and positive semi-definite. Furthermore, a constant estimate of volatility can be estimated by computing

$$\frac{\alpha_0}{1 - \beta_1 - \alpha_1}, \tag{8}$$

which is used in this paper. Finally, the "true" volatility is computed in the same manner as historical volatility, but instead uses the three months of data held back from the training set.

RESULTS

Differences in volatility models resulted in severe mismatches in option premium estimates, as shown by the table of results below. Two stocks, Credit Suisse (CS) and Wells Fargo (WFC), highlight the dangers of poor volatility forecasting performance. For Credit Suisse, the historical volatility estimate resulted in option pricing errors upwards of $0.80 per unit of stock. Options written on WFC that were priced using historical volatility were nearly double the cost of options priced using GARCH (1,1), which offered a substantially better prediction of future volatility. For some stocks, neither historical volatility nor GARCH offered accurate volatility predictions. Volatility forecasting error for Macy's (M) resulted in approximately $0.40 pricing error for both forecasting methods, with historical volatility overestimating the option premium (and therefore volatility), and GARCH underestimating the options premium. M also highlights the risk of underestimating the impact of vega. Vega for this stock is relatively low, at only 3 cents per 1% error in volatility (for reference, the vega on an otherwise identical vanilla option is just over 5 cents). However, the historical volatility estimate is 39.4%, the GARCH volatility estimate is 13.5%, while the true volatility observed over the life of the option was 27.1%. Such large mismatches in volatility estimates explains most of the pricing error when taking even this low vega into account.

These results suggest that investors interested in pricing Asian options should carefully review risk management practices that relate to sources of model error. While GARCH(1,1) will be sufficient for some stocks (here consider WFC and WFT), these results clearly demonstrate that when volatility models fail, the cost of failure is high. While in these experimental conditions there is a true benchmark for volatility estimates, this not the case when trading, so errors will only be realized after the fact.

Table 1: Asian option prices under different volatility forecasts

Stock	Volatility	Call Price	Put Price
CS	Historical	2.233762	2.123826
CS	GARCH	1.412262	1.283331
CS	Observed	1.484792	1.371772
M	Historical	1.261422	1.176934
M	GARCH	0.4524793	0.3787000
M	Observed	0.8754793	0.7969439
MDT	Historical	1.933642	1.803678
MDT	GARCH	1.710804	1.552222
MDT	Observed	1.882623	1.746523
NCR	Historical	0.5855573	0.5550229
NCR	GARCH	0.4186111	0.3832460
NCR	Observed	0.3005722	0.2738922
POT	Historical	0.1774972	0.1606975
POT	GARCH	0.1766955	0.1619346
POT	Observed	0.1977078	0.1808683
WFC	Historical	0.9482303	0.8612777
WFC	GARCH	0.5145017	0.4248920
WFC	Observed	0.5456992	0.4521035
WFT	Historical	0.5451017	0.5180468
WFT	GARCH	0.4069263	0.3823108
WFT	Observed	0.4086320	0.3802554

Source: own calculations.

Figure 5: Table of results displaying simulation output from 100,000 realizations per item.

CONCLUSION

Losses from model error can be significant; banks across the globe have historically lost upwards of 80 million dollars due to the modeling errors from even a single trader (Economist 1997). While Asian options hedge local volatility, they remain sensitive to the quality of one's forecasts. The risks are higher for OTC options trading because it tends to deal in significantly higher volumes than retail traders can afford, which means that even small errors can quickly scale up. Path dependent options are particularly vulnerable to this source of error as volatility modeling controls price-path behavior of the asset, quickly sending estimated prices off course. The risk is especially relevant for Asian options as they are the most commonly traded in currency and commodity markets where leverage amplifies minor pricing errors. These results show that major losses can be experienced even in the relative simplicity of a Black Scholes world, which cannot account for even larger sources of loss that could be exposed by models incorporating regime switching and stochastic volatility. Therefore, the first step for market actors seeking to price Asian options should be to select a robust and trusted volatility model.

References

B. Alziary, J. Décamps, and P. Koehl. "A P.D.E. Approach to Asian Options: Analytical and Numerical Evidence." *Journal of Banking and Finance* 21 (1997): 613–40.

N. Beaulieu, and F. Rajwani. "Highly Accurate Simple Closed-Form Approximations to Lognormal Sum Distributions and Densities." *IEEE Communication Letters* 8, no. 12 (December 2004).

T. Bollerslev, R. Chou, and K. Kroner. "ARCH Modeling in Finance." *Journal of Econometrics* 52 (1992): 5–59.

T. Bollerslev, and H. Mikkelsen. "Long-Term Equity Anticipation Securities and Stock Market Volatility Dynamics." *Journal of Econometrics* 92 (1999): 75–99.

L. Fenton. "The Sum of Log-Normal Probability Distributions in Scatter Transmission Systems." *IRE Transactions on Communications Systems*, March 1960, 57–67.

W. Hsu, and Y. Lyuu. "Efficient Pricing of Discrete Asian Options." *Applied Mathematics and Computation* 201 (2011): 9875–94.

V. Lakhlani. "Pricing and Hedging Asian Options." Master's Thesis, Utah State University, 2013.

E. Levy. "Pricing European Average Rate Currency Options." *Journal of International Money and Finance* 11 (1992): 474–91.

N. Mehta, J. Wu, and J. Zhang. "Approximating a Sum of Random Variables with a Lognormal." *IEEE Transactions on Wireless Communications* 6, no. 7 (July 23, 2007). doi:10.1109/TWC.2007.051000.

A. Mraovic, and Q. Zhang. "Valuation of Asian Options with Levy Approximation." Master's Thesis, Lund University, 2014.

L. Nielsen. "Pricing Asian Options." Master's Thesis, Aarhus University, 2001.

P. Posedel. "Properties and Estimation of GARCH(1,1) Model." *Metodološki Zvezki* 2, no. 2 (2005): 243–57.

M. J. Steele, *Stochastic Calculus and Financial Applications*. 2nd ed. Stochastic Modelling and Applied Probability 45. Springer, 2001.

"NatWest Holed." *The Economist*, March 20, 1997, sec. Finance and Economics. http://www.economist.com/node/146045. Accessed June 9, 2017.