Economics Bulletin

Volume 38, Issue 1

On modeling fossil fuel prices: geometric Brownian motion vs. variance-gamma process

Alejandro Mosiño
Universidad of Guanajuato

Alejandro Tatsuo Moreno-Okuno
Universidad de Guanajuato

Abstract

It is very common in the literature to assume that fossil fuels prices follow a geometric Brownian motion (GBM) process. However, the GBM is an imperfect process in the sense that it cannot capture correctly the frequent extreme movements in fossil fuel prices caused by the information generated within domestic and international markets. In this paper, we use fossil fuel index data to compare the performance of the GBM, which is based on normal density, with that of the variance-gamma (VG) process, which allows us to capture the skewness and the excess of kurtosis of price returns. We show that the VG process fits the data better than does the GBM, as indicated by several goodness-of-fit tests.

We thank the editor and three anonymous reviewers for their constructive comments, which helped us to improve this paper.
Citation: Alejandro Mosiño and Alejandro Tatsuo Moreno-Okuno, (2018) "On modeling fossil fuel prices: geometric Brownian motion vs. variance-gamma process", *Economics Bulletin*, Volume 38, Issue 1, pages 509-519
Contact: Alejandro Mosiño - Alejandro.Mosino@gmail.com, Alejandro Tatsuo Moreno-Okuno - atatsuo@hotmail.com.
Submitted: June 22, 2017. **Published:** February 27, 2018.

Economics Bulletin

Submission Number: EB-17-00495

On modeling fossil fuel prices: geometric Brownian motion vs. variance-gamma process

Alejandro Mosiño
Universidad of Guanajuato

Alejandro Tatsuo Moreno-okuno
Universidad de Guanajuato

Abstract

It is very common in the literature to assume that fossil fuels prices follow a geometric Brownian motion (GBM) process. However, the GBM is an imperfect process in the sense that it cannot capture correctly the frequent extreme movements in fossil fuel prices caused by the information generated within domestic and international markets. In this paper, we use fossil fuel index data to compare the performance of the GBM, which is based on normal density, with that of the variance-gamma process (VGP), which allows us to capture the skewness and the excess of kurtosis of price returns. We show that the VGP fits the data better than does the GBM, as indicated by several goodness-of-fit tests.

Submitted: June 22, 2017.

1 Introduction

Many authors in real options modeling have argued about the importance of choosing an adequate stochastic process to describe the price dynamics of the underlying asset. One of the favorites of real options theorists and practitioners is the geometric Brownian motion (GBM) process, which has been commonly used in financial derivatives valuation (Black and Scholes, 1973; Cox et al., 1979), and in corporate project valuation (Schwartz and Trigeorgis, 2001).

Since the seminal work by Pindyck (1980, 1984), the GBM process has also been largely used in natural resource applications. One example is Olsen and Stensland's optimal shutdown problem (Olsen and Stensland, 1988), which reveals one of the main advantages of using a GBM process: its mathematical simplicity. In particular, the GBM assumption results in differential equations whose analytical solutions are relatively easy to find and interpret. Another advantage of using a GBM is that its parameters are very easy to find by maximum likelihood estimation. This is particularly relevant to more empirical oriented models such as that of Detert and Kotani (2013).

Despite its advantages, the GBM lacks some important characteristics present in typical asset price data, and particularly those of special commercial commodities, such as oil. Specifically, typical data are usually leptokurtic and heteroscedastic, and could exhibit skewed distributions (Finlay, 2009). Several alternative processes have been used in the literature to circumvent some of these problems. These processes include the mean reverting —Ornstein-Uhlenbeck— process (MR), and the geometric mean reverting (GMR) process (Dixit and Pindyck, 1994). For instance, Gibson and Schwartz (1990), Dixit and Pindyck (1994), and Schwartz (1997) argue that, for oil and similar commodities, the MR would be more appropriated than the GBM, as it assumes that prices present random behavior in the short run, whereas, in the long run, they converge to equilibrium level, reflecting production marginal cost. As the GBM, the MR, and the GMR processes are very convenient, in the sense that they are (mathematically) very easy to handle and calibration requires no more than running an ordinary least squares regression (Dixit and Pindyck, 1994). There are other possibilities, such as processes combining MR and GBM with Poisson processes (Schwartz, 1997; Pindyck, 1999; Schwartz and Smith, 2000). However, these tend to be more complicated to use in theoretical and empirical applications.

In this paper, we propose using the variance-gamma (VG) process (Madan et al., 1998; Seneta, 2004). This process has been recently introduced in the financial literature, and it has proved to be superior to the GBM in fitting several asset prices (Finlay, 2009). Mathematically speaking, the VG process is very similar to the GBM, in that both have drift and volatility parameters. The former, however, is evaluated at a random time rather than being evaluated at a natural time.

Given its relative simplicity, we are able to compare the performance of the VG process to that of the classical GBM when modeling fossil fuel prices. To do this, we calibrate the IMF's fuel energy index as an example. This index summarizes the evolution of several fossil fuel indexes, such as crude oil, natural gas and coal. To compare both processes, we compute their goodness-of-fit by performing Chi-squared, Anderson-Darling, and Kolmogorov-Smirnov tests. We show that, in these tests, the null of VG cannot be rejected. We also compute the Value-at-Risk for both the GBM and VG process, and use it to measure the distribution adequacy for tail fits. We again show that VG process fits better the data that GBM. Our results are particularly relevant to real options modeling using exhaustible resources as underlying assets as, to our knowledge, not much work has been done on this subject.

The rest of the paper is as follows. After this introduction, we describe the essentials of the VG process in Section 2. In Section 3 we calibrate the IMF's fuel energy index and perform the above mentioned goodness-of-fit tests. We conclude the paper in Section 4.

2 The variance gamma process

In this section, we examine the main theoretical characteristics of the VG process. As with the GBM, the VG process is characterized by a deterministic linear trend and a stochastic component that captures random deviations of prices from its mean. The latter component is driven by a standard Wiener process. The key difference between the GBM and the VG process is that, in the latter, time t is replaced by the gamma process $\{g_t\}_{t\geq 0}$. The gamma process starts at zero, and has stationary and independent increments following the gamma distribution (Schoutens, 2003):

$$\Gamma\left(\frac{t}{\nu}, \frac{1}{\nu}\right), \quad \nu > 0.$$

Assume that $\{G_t\}_{t\geq 0} = g_t - g_q$, and that time period increments are given by $\Delta t = t - q$, with $t > q$. Then:

$$\mathbb{E}(G_t) = \Delta t$$
$$\mathrm{Var}(G_t) = \nu \Delta t.$$

The random time given by the VG process is usually interpreted as business time, rather than calendar time. It incorporates either the information flow or trading activity into the model: the more frenzied trading becomes, or the more information released to the market on a given day, the faster time flows (Finlay, 2009). The gamma process, like the Poisson process, is a pure jump process. This results in the VG process being a pure jump process with no diffusion component (Madan et al., 1998; Sullivan and Moloney, 2010).

Let $\{P_t\}_{t\geq 0}$ be a series of asset prices, and define $X_{\Delta t} = \ln(P_t) - \ln(P_q)$. Hence, the VG process can be defined as (Seneta, 2004):

$$X_{\Delta t} = \mu \Delta t + (g_t - g_q)\theta + (W_{g_t} - W_{g_q})\sigma. \tag{1}$$

In equation (1), $\mu, \theta, \sigma > 0$ are real constants, and W_t is the standard increment of a Wiener process. This is assumed to be independent of G_t. Note that the specification of the VG process allows us to capture more of the main characteristics shown by typical log returns data (Heyde and Liu, 2001; Finlay, 2009). As one can see from equation (1), $X_{\Delta t}$ is described by two drift parameters: a calendar-time drift parameter, μ, and another parameter that evolves in time g_t, θ. The latter parameter lets us take into account the occasional skewed distributions of typical log-returns' data. Also, in equation (1), σ measures the total volatility of the process, while ν is the variance of the gamma increment, G_t. This latter parameter enables the process to adjust the usual pronounced leptokurtic distribution of log-returns.

By using the properties of the standard Brownian motion, it is possible to conclude that the (conditional) distribution of $X_{\Delta t}$ is:

$$X_{\Delta t} - \mu \Delta t \mid_{G_t} \sim \mathcal{N}\left(\theta(g_t - g_q), (g_t - g_q)\sigma^2\right). \tag{2}$$

Also, we can compute the non-conditional density function of $X_{\Delta t}$ as (Brigo et al., 2009):

$$f_{X_{\Delta t}}(x) = \frac{2\exp\left(\theta\left(\frac{x-\mu}{\sigma^2}\right)\right)}{\nu^{\frac{\Delta t}{\nu}}\sqrt{2\pi}\sigma\Gamma\left(\frac{1}{\nu}\right)}\left(\frac{|x-\mu|}{\sqrt{\frac{2\sigma^2}{\nu}+\theta^2}}\right)^{\frac{\Delta t}{\nu}-\frac{1}{2}} \cdot B_{\frac{\Delta t}{\nu}-\frac{1}{2}}\left[\frac{1}{\sigma^2}|x-\mu|\sqrt{\left(\frac{2\sigma^2}{\nu}+\theta^2\right)}\right], \tag{3}$$

for $-\infty < x < \infty$, and $B(\cdot)$ being the Bessel function of the third class. Our knowledge of the non-conditional distribution of X_t, equation (3), under the VG assumption, lets us analyze how well this process fits the data compared to the GBM, whose log-returns follow a normal distribution, with a mean of $(\mu - 0.5\sigma^2)\Delta t$ and standard deviation equal to $\sigma\Delta t$.[1]

3 Fossil fuel prices as a variance-gamma process

We consider the IMF's fuel energy index, which summarizes the evolution of crude oil, natural gas and coal price indices. We have 304 monthly observations, starting from January 1992. We plot the data series in Figure 1.

Figure 1: Evolution of fuel energy index. Source: IMF

As can be seen in Figure 1, the fuel energy index exhibits sudden jumps. For instance, the sudden increase in the first half of 2008 is due to an oil prices' spike. Events that explain this behavior include (Smith, 2009):

[1] It is possible to verify that, if $(\mu + \theta) = (\mu - \frac{1}{2}\sigma^2)$, and σ is the same for both the GBM and the VG process, the difference between both models is the size of θ and ν. In the VG process, θ and the skewness are of the same sign. If $\theta = 0$, the skewness is also null, and $\frac{\nu}{\Delta t}$ is the excess of kurtosis compared to normal distribution (equal to 3).

- In February, Venezuela cut off oil sales to ExxonMobil during a legal battle over nationalization of the company's properties there.

- In late March, saboteurs blew up the two main export pipelines in the south, cutting approximately 300,000 barrels per day from Iraqi exports.

- On April 25, Nigerian union workers went on strike, causing ExxonMobil to shut down production of 780,000 barrels per day from three fields.

- On April 27, Scottish oil workers walked off the job, leading to closure of the North Forties' pipeline, which carries approximately half of the United Kingdom's North Sea oil production.

- On May 1, approximately 1.36 million barrels per day of Nigerian production was shut down, due to a combination of militant attacks on oil facilities, sabotage, and labor strife. At the same time, it was reported that Mexican oil exports had fallen sharply in April, due to rapid decline in the country's massive Cantarell oil field.

- On June 19, militant attacks in Nigeria caused Shell to shut down an additional 225,000 barrels per day.

- On June 20, just days before the price of oil reached its historic peak, Nigerian protesters blew up a pipeline, forcing Chevron to shut down 125,000 barrels per day.

Events like those mentioned above prevents stochastic models, such as the GMB, from fitting the data well. However, as we explained in Section 2, the VG process could be more suitable. In the following section, we calibrate the parameters of both processes by maximum likelihood estimation. We next compare the fitting to real data.

3.1 Calibration

We first fix $\Delta t = 1$ and compute log-returns as in Section 2. We next run stationarity and autocorrelation tests. Our first finding is that the series are stationary as suggested by the very low (less than 1%) p-value of the augmented Dickey-Fuller test. A plot of the series (see Figure 2) also reveals that there are some heteroscedastic patterns and volatility clustering. This is a common characteristic of financial returns (Tsay, 2010). We also find less than 1% p-values of the Ljung-Box test with several lags, indicating that fossil fuel energy index displays autocorrelation in returns. This is a feature that characterizes many energy commodities, such as gas, oil and electricity (González-Pedraz et al., 2014). We also find that X_t^2 is autocorrelated. This last finding implies that the series are not iid and hence cannot be used to estimate the VG process parameters.

To deal with dependence of log-returns, we apply a GARCH filter (McNeil and Frey, 2000).[2] Let $X_t^* = \ln(P_t) - \ln(P_{t-1})$, $X_t = X_t^* - \overline{X}$, and

$$X_t = (\sigma_t)^f Y_t, \quad (\sigma_t^f)^2 = \alpha_0 + \alpha_1 X_{t-1}^2 + \beta_1 (\sigma_{t-1}^f)^2,$$

[2] For other filters refer to Lo et al. (2016).

i.e. a GARCH(1,1) model. $(\sigma_t^f)^2$ is estimated using quasi maximum likelihood. We next compute the filtered return series $\widehat{Y}_t = X_t/(\widehat{\sigma}_t)^f$, where $(\widehat{\sigma}_t)^f$ is the estimated GARCH(1,1) volatility. Autocorrelation test on $\widehat{Y}_t^2$ confirm that this is not autocorrelated. Hence, we can use $\widehat{Y}_t^2$ to calibrate the process.[3]

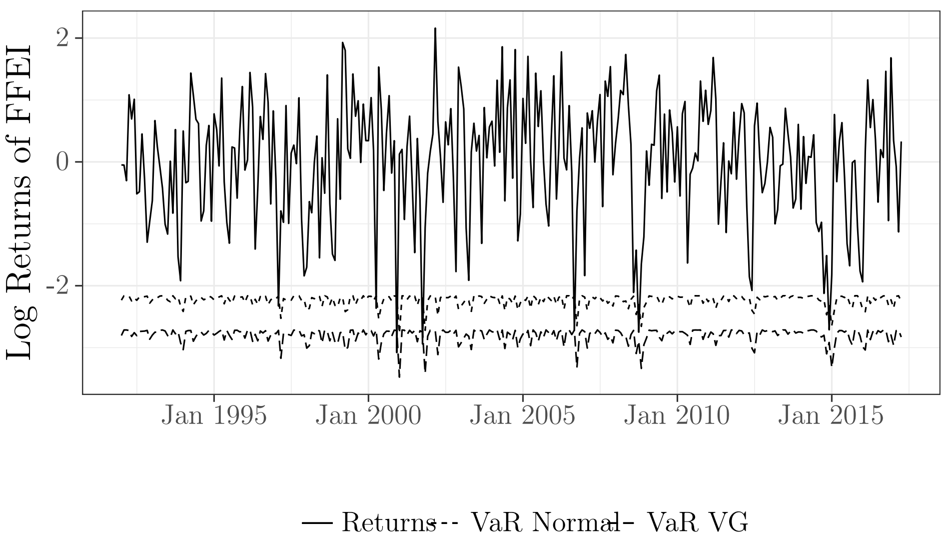

Figure 2: Evolution of log-returns of the fuel energy index and dynamic VaR estimation.

We now estimate the VG parameters by following the procedure in Seneta (2004).[4] We compute the first four moments of the variance-gamma distribution by the method of moments. As $\Delta t = 1$ we get:

$$\widehat{\mu}^m = \overline{x} - \widehat{\theta}^m; \quad \widehat{\sigma}^m = s; \quad \widehat{\theta}^m = S\widehat{\sigma}^m; \quad \widehat{\nu}^m = \frac{K}{3} - 1, \tag{4}$$

where $\overline{x}$, s^2, S and K are mean, variance, skewness, and kurtosis sample estimates, respectively. Equation (3) is then maximized by using the results in equation (4) as an initial guess. Our resulting estimates are shown in Table I. To facilitate comparison, we also show the GBM estimated parameters in Table I.[5]

3.2 Goodness of fit

To analyze the performance of the VG process, we follow the recommendation in Göncü et al. (2013) and perform several goodness-of-fit tests. In particular, we perform Chi-squared, Anderson-Darling (A-D), and Kolmogorov-Smirnov (K-S) tests.[6]

[3] Autocorrelation functions of both X_t^2 and $\widehat{Y}_t^2$, as well as results from the augmented Dickey-Fuller test are available from the authors.

[4] See also Skindilias and Lo (2013) for calibration using the VG and other processes.

[5] To calibrate the GBM process see Dixit and Pindyck (1994), for instance.

[6] We also compared the GBM and VG process densities to the empirical density function. The VG distribution fits extremely well to data, specially in the right tail. QQ-Plots confirm this observation. Density plots and QQ-Plots are available from the authors.

Table I: Estimated parameters by using GBM and VG process (filtered data).

	Log-Returns	
	Normal distribution	*Variance-Gamma distribution*
μ	0.0015	2.3091
σ	0.9979	0.6181
θ	—	-2.3106
ν	—	0.1184

The Chi-squared goodness-of-fit test is a non-parametric test that is commonly used to compare the observed sample distribution with its expected probability distribution. It allows us to determine how well the theoretical distribution (in our case normal, or variance gamma) fits the empirical distribution. To perform the Chi-Square goodness-of-fit test, sample data is divided into K subintervals. Then, the number of points that fall into the k-th subinterval, for $k = 1, \ldots, K$, is compared with the expected number of points in that subinterval. Here, we follow Madan and Seneta (1987) and consider seven subintervals given by the bounds:

$$\{-\infty, -0.05, -0.025, -0.01, 0, 0.01, 0.025, 0.05, +\infty\}.$$

The relevant statistic is:

$$\chi^2_{K-1-m} = \sum_{i=1}^{K} \frac{(ob.f._i - ex.f_i)^2}{ex.f_i}$$

where $ob.f_i$ and $ex.f_i$ refer to the observed and expected frequencies at cell i, respectively. The expected frequency is defined as Np_i, where N is the sample size and p_i is the probability of a randomly drawn value to fall into the i cell; m is the number of parameters of the model. As usual, the value of the relevant statistic is compared with the critical value. If the calculated value of Chi-squared goodness-of-fit statistic is greater than the critical value, we reject the null hypothesis that there is no difference between the observed and the expected frequency at some chosen significance level. If the calculated value of Chi-squared goodness-of-fit statistic is less than the critical value, we do not reject the null hypothesis and conclude that there is no significant difference between the observed and the expected value.

Another non-parametric test to compare the observed sample distribution with its expected probability distribution is the K-S goodness-of-fit test. Contrary to the Chi-squared test, K-S is applied to unbinned data, i.e. it does not require discretization. In particular, the K-S test is based on the largest vertical difference between the empirical cumulative distribution and the hypothesized distribution. Given N ordered (from lowest to highest) data, $z_1, z_2, \ldots, z_N$, the K-S goodness-of-fit test is computed as:

$$KS = \max_{1 \leq i \leq n} \left\{ F(z_i) - \frac{i-1}{n}, \frac{i}{N} - F(z_i) \right\},$$

where $F(\cdot)$ is the (theoretical) cumulative distribution function. The hypothesis regarding the distributional form is rejected at some chosen significance level if the test statistic is greater than the critical value. However, as noted in Göncü et al. (2013), since we estimate the model parameters from the same initial data set, the critical values of the K-S test cannot be used. Nevertheless, the K-S test statistic can be used to quantify the fit of the VG model in comparison to the normal distribution.

The A-D goodness-of-fit test is a modification of the K-S test that attaches higher weights to the tails than does the K-S test (Göncü et al., 2013). This is computed as:

$$A^2 = -n - \sum_{i=1}^{n} \frac{2i-1}{n} \left(\ln F(z_i) + \ln(1 - F(z_{n+i-1})) \right),$$

where $F(\cdot)$ is, as in the K-S test, the cumulative distribution function, and z_i, with $i = 1, 2, \ldots, N$, are the ordered (from lowest to highest) log-returns. The A-D statistic can be compared to the critical values of the normal distribution in Stephens (1974).

Table II: Results from different goodness-of-fit tests. p-values are shown in parenthesis.

χ^2		**A-D**		**K-S**	
Normal	VG	Normal	VG	Normal	VG
18.594	14.659	1.6313	0.1666	0.0716	0.0304
(0.0199)	(0.0484)	(0.148)	(0.997)	(0.0882)	(0.9406)

In Table II we show the Chi-squared, K-S, and A-D tests values for both the normal and VG distributions. We also compute their associated p-values, these are shown in parenthesis. As we can see, Chi-squared test rejects both normal and VG distributions. A-D barely fails to reject the normal distribution and fails to reject VG. K-S rejects normal distribution, and fails to reject VG distribution. At least from the results of A-D and K-S test we have some evidence in favor of applying the VG process to model the fuel energy index.

3.3 Value at Risk

To further analyze the performance of the VG process, we compute the $(1-\alpha)\%$ dynamic Value-at-Risk (VaR). This is computed from the filtered data as:

$$\text{VaR}_{1-\alpha} = (\widehat{\sigma}_t)^f \widehat{F}^{-1}(\alpha),$$

where $\widehat{F}^{-1}$ is the quantile function corresponding to the estimated parameters of the VG distribution, and $(\widehat{\sigma}_t^f)$ is the estimated GARCH standard deviation of Section 3. VaR can be used to measure a distribution's level of adequacy for tail fits.

Table III: p-values for Kupiec test for several levels of α.

	p-values	
α	*Normal distribution*	*Variance-Gamma distribution*
0.01	0.0054148	0.9815645
0.05	0.2272968	0.4734198
0.1	0.4985621	0.7614667

We plot the dynamic VaR with $\alpha = 0.01$ in Figure 2. Figure 2 also gives us an idea of the relative frequency of VaR violations. To formally test the number of violations to the corresponding in-sample VaR level we compute the Kupiec likelihood ratio test (Kupiec, 1995). Resulting p-values for the Kupiec test for different levels are shown in Table III. These results imply that the VG distribution fits well the data, as Kupiec test fails to reject the hypothesis that VaR is correctly estimated.

4 Conclusion

In this paper, we model the fuel energy index as both a geometric Brownian motion (GBM) and a variance-gamma (VG) process. The former is more common in the literature of environmental resource economics, and it is mostly used by practitioners. Our analysis implies that, although not perfect, the variance-gamma distribution provides a better fit than does the normal distribution. This is because the VG process can capture the sudden extreme changes in the fuel energy index. Also, the goodness-of-fit tests we perform suggest that the null hypothesis of normality is usually rejected.

Our results are particularly relevant to real options models using exhaustible resources as underlying assets. One advantage of using the VG process is that its parameters are very easy to estimate. Additionally, because its structure is not too much different from that of the GBM, closed form solutions to purely theoretical natural resources real options models could eventually exist.

References

Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, 81(3):637–654.

Brigo, D., D'alessanddro, A., Neugebauer, M., and Triki, F. (2009). A Stochastic Processes Toolkit for Risk Management. *Journal of Risk Management for Financial Institutions*, 2(4):365–393.

Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263.

Detert, N. and Kotani, K. (2013). Real options approach to renewable energy investments in Mongolia. *Energy Policy*, 56:136–150.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.

Finlay, R. (2009). *The Variance Gamma (VG) Model with Long Range Dependence: A model for financial data incorporating long range dependence in squared returns*. VDM Verlag.

Gibson, R. and Schwartz, E. S. (1990). Stochastic Convenience Yield and the Pricing of Oil Contingent Claims. *The Journal of Finance*, 45(3):959–976.

Göncü, A., Karahan, M. O., and Kuzubas, T. U. (2013). Fitting the Variance-Gamma Model: A Goodness-of-Fit Check for Emerging Markets. *Bogazici Journal of Economics and Administrative Sciences*, 27(2):1–10.

González-Pedraz, C., Moreno, M., and Peña, J. I. (2014). Tail risk in energy portfolios. *Energy Economics*, 46:422–434.

Heyde, C. C. and Liu, S. (2001). Empirical realities for a minimal description risky asset model. The need for fractal features. *Journal of the Korean Methematical Society*, 38(5):1047–1059.

Kupiec, P. H. (1995). Techniques for Verifying the Accuracy of Risk Measurement Models. *The Journal of Derivatives*, 3(2):73–84.

Lo, C. C., Skindilias, K., and Karathanasopoulos, A. (2016). Forecasting Latent Volatility through a Markov Chain Approximation Filter. *Journal of Forecasting*, 35(1):54–69.

Madan, D. B., Carr, P. P., and Chang, E. C. (1998). The Variance Gamma Process and Option Pricing. *Review of Finance*, 2(1):79–105.

Madan, D. B. and Seneta, E. (1987). Chebyshev Polynomial Approximations and Characteristic Function Estimation. *Journal of the Royal Statistical Society. Series B (Methodological)*, 49(2):163–169.

McNeil, A. J. and Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7(3-4):271–300.

Olsen, T. E. and Stensland, G. (1988). Optimal shutdown decisions in resource extraction. *Economics Letters*, 26(3):215–218.

Pindyck, R. S. (1980). Uncertainty and Exhaustible Resource Markets. *The Journal of Political Economy*, 88(6):1203–1225.

Pindyck, R. S. (1984). Uncertainty in the Theory of Renewable Resource Markets. *The Review of Economic Studies*, 51(2):289–303.

Pindyck, R. S. (1999). The long-run evolution of energy prices. *The Energy Journal*, 20(2):1–27.

Schoutens, W. (2003). *Lévy Processes in Finance: Pricing Financial Derivatives*. Wiley.

Schwartz, E. and Smith, J. E. (2000). Dynamics in Commodity Prices. *Management Science*, 46(7):893–911.

Schwartz, E. and Trigeorgis, L., editors (2001). *Real options and investment under uncertainty*. MIT Press, Cambridge, Mass. [u.a.].

Schwartz, E. S. (1997). The Stochastic Behaviour of Commodity Prices: Implication for Valuation and Hedging. *The Journal of Finance*, 52(3):923–973.

Seneta, E. (2004). Fitting the Variance-Gamma Model to Financial Data. *Journal of Applied Probability*, 41(2004):177–187.

Skindilias, K. and Lo, C. C. (2013). *Energy Security: Stochastic Analysis of Oil Prices*, pages 155–178. Springer London, London.

Smith, J. (2009). The 2008 oil price shock: Markets or mayhem?

Stephens, M. A. (1974). EDF Statistics for Goodness of Fit and Some Comparisons. *Journal of the American Statistical Association*, 69(347):730–737.

Sullivan, C. O. and Moloney, M. (2010). The Variance Gamma Scaled Self-Decomposable Process in Actuarial Modelling.

Tsay, R. S. (2010). *Analysis of Financial Time Series: Third Edition*.

Venegas-Martínez, F. (2006). *Riesgos Financieros y Económicos: Productos Derivados y Decisiones Económicas Bajo Incertidumbre*. Thompson.