Catch the Drift?

Robert G. Tompkins
Hochschule für Bankwirtschaft

Frankfurt MathFinance Workshop
Frankfurt, 27 March 2006

Background

Changing the measure from the "Real World" to the "Risk Neutral World" underlies the theory of option pricing.

For GBM, this is achieved by means of Girsanov, requiring only the drift of the Real World process to be modified. Leading to the famous drift adjustment:

$$d\ln S = \left(r - \frac{\sigma^2}{2} \right) dt + \sigma dw_Q$$

Background

Unfortunately, the "Real World" does not conveniently conform to a GBM process. What the exact process is, remains unknown.

Volatility is not constant nor time homogeneous, leading to the development of a wide variety of GARCH and Stochastic Volatility Models.

Even when such non-constant volatility factors are included in modelling the real world return process, often the returns remain non-normal. This has led to inclusion of alternative model specifications where the process driving either volatility or the underlying returns follow an alternative distribution to the usual Gaussian. Recently, Levy Processes such as the VG and NIG have been proposed.

Background

These models have been shown to adequately capture the dynamics of the "Real World" process, but the next step to arrive at an option pricing formula requires a measure change.

For many of these models, non-market-traded sources of risk are introduced and thus, markets are incomplete. There exist an infinite set of measure changes that do not allow arbitrage.

The objective of this research is to find the "correct" measure change for a given model. In this research, we will examine a modified GARCH (1,1) option pricing model, where the underlying return process is not GBM but rather an empirical Levy Process.

Thought Experiment

One of the cornerstones of no-arbitrage pricing is the fact that the price process must be a Martingale for all time points from t=0 to t=T. For the GBM process, this can be achieved by the change of measure previously presented.

How do you prove this result to the mathematically challenged?

Suppose, you set up a simple Monte Carlo simulation. You estimate prices in the usual way but leave out the risk neutral drift adjustment: $S_{t+1} = S_t e^{\sigma dZ}$

Then after 1000 paths, you take the natural log of the average price at each time point and draw a graph, what do you get?

Thought Experiment

Input Parameters for the Monte Carlo Simulation

- Starting Underlying Price, S_0 = 100 (ln S_0 = 4.6025)
- Annualised Volatility = 20%
- Interest Rate (dividend) = 0%
- Time steps 1 day (1/252)
- Estimation Horizon from 1 day to 180 days
- Risk neutral drift should be $-\frac{1}{2}\sigma^2 t$ or -7.93651E-05

Thought Experiment

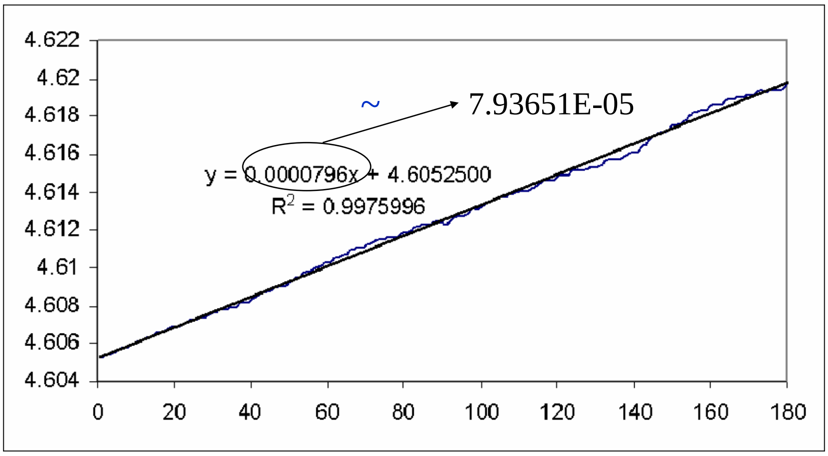

Where the regression is: $Ln(\bar{S}_t) = \alpha + \beta t$

Thought Experiment

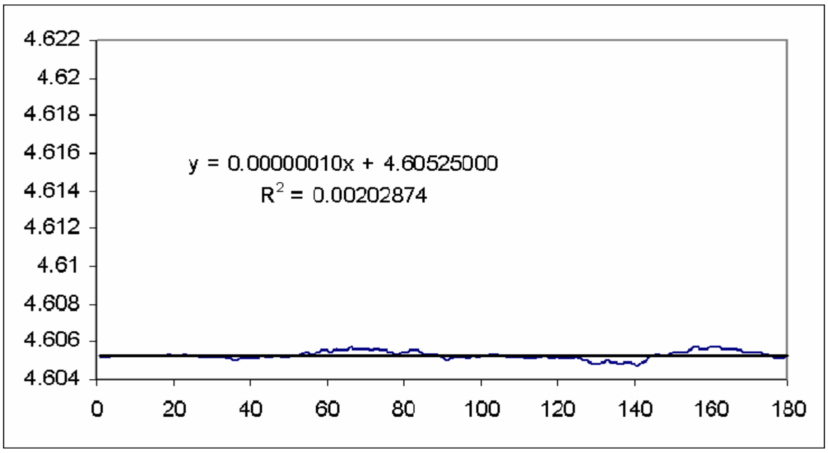

Including (the negative) of this estimated drift adjustment

Thought Experiment

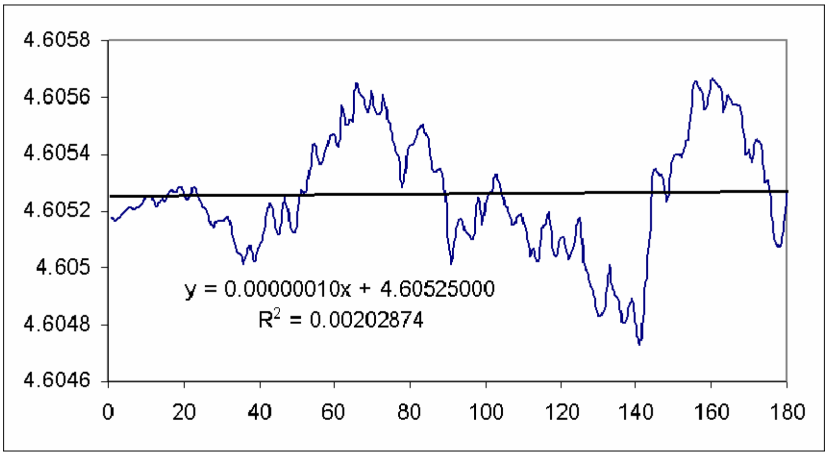

Smaller scale

Thought Experiment

Violá, we have for all time periods the natural logarithm of the average simulated future prices are equal to the initial starting value and we have the change of measure that makes the simulated price paths a Martingale.

Does this approach preclude arbitrage opportunities? The answer is yes! Duan and Simonato (1995) suggested an alternative approach to adjusting the drift using the “Empirical Drift Adjustment”. Like our approach, this precludes the existence of option arbitrage and as an experiment, using this same simulated data, we applied this approach.

Hochschule für Bankwirtschaft / HfB

Thought Experiment

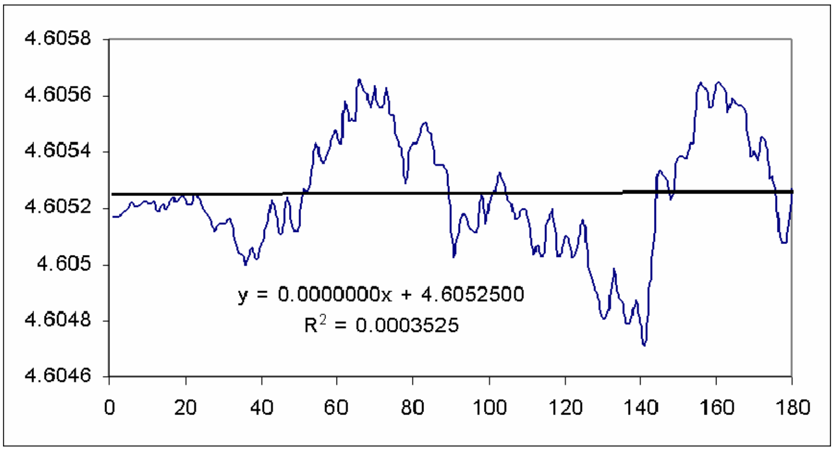

Using the Duan & Simonato (1995) Expirical Drift Adjustment

Hochschule für Bankwirtschaft / HfB

Alternative Real World Model

James and Colchester (2003) have shown that prices of currency options from 1992 to 2003 were equal to the average payout on these options.

Using the actual price path of the currencies, the prices of the options were compared to the average terminal payoffs.

We will re-examine their results using the price path of British Pound / US Dollar for the period from 1990 to 2004.

Alternative Real World Model

Tompkins & D'Ecclesia (2006) developed a new simulation approach, which both matches the empirical record and projects future price paths with exactly the same statistical characteristics as the historical record.

Historical return series are filtered using a GARCH (1,1) model. Simulated prices are generated using draws from this standardised set of historical "disturbances".

As opposed to "boot-strapping", the entire set of historical "disturbances" are selected rather than sampled with replacement. Price paths differ as the set of disturbances are mixed randomly for each simulation.

Alternative Real World Model

In that research, option prices for the S&P 500 were estimated based upon the actual empirical payoff. The mixed disturbances approach was found to match these empirical payoff values.

To express the payoffs in "risk neutral" terms in order to estimate implied volatilities, a naive drift adjustment was employed. Later tests showed that modifications were required to assure no arbitrage for the resulting option prices.

This research examines that problem.

The Unconditional Disturbances Model

Given a historical price series for the state variable S_t, for t = 0,…,T, returns are estimated using log price relatives:

$$r_t = \ln\left(\frac{S_t}{S_{t-1}}\right)$$

For the time period *0* to *T*, the unconditional mean μ, and standard deviation, σ, are estimated.

The Unconditional Disturbances Model

Normalizing the sequence of returns by these two moments yields:

$$u_t = \frac{r_t - \mu}{\sigma} \qquad \textbf{(1)}$$

where $\{u_t\}$ is the series of standardised "*disturbances*" from *0* to *T*. By design, the resulting disturbances have a mean of 0 and standard deviation of 1.

The Unconditional Disturbances Model

British Pound / US Dollar 1990-2004	Daily Unconditional Returns	Annualized Unconditional Returns	Daily Standardized Disturbances
μ	0.000027	0.00668	0.0
σ	0.005702	0.090525	1.0
skewness	-0.195489		-0.195489
kurtosis	5.538		5.538

Table 1. *Unconditional moments of the BP/US$ Returns, 1990 - 2004*

The Unconditional Disturbances Model

GARCH (1,1) Model

We introduce time varying volatility according to a GARCH (1,1) model.

$$\widehat{\sigma}_t(\eta_t) = \sqrt{\hat{\sigma}_t^{\,2}} = \sqrt{\omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2 + \eta_t^{\,2}} \qquad \textbf{(2)}$$

Using daily spot exchange rates for US Dollar /British Pound from 1990 to 2004, we estimated the following GARCH (1,1) parameters.

α = 0.075810072 Log Likelihood

β = 0.887753867 34157.82

ω = 1.18485E-06 Average volatility 9.0525%

The Unconditional Disturbances Model

GARCH (1,1) Model

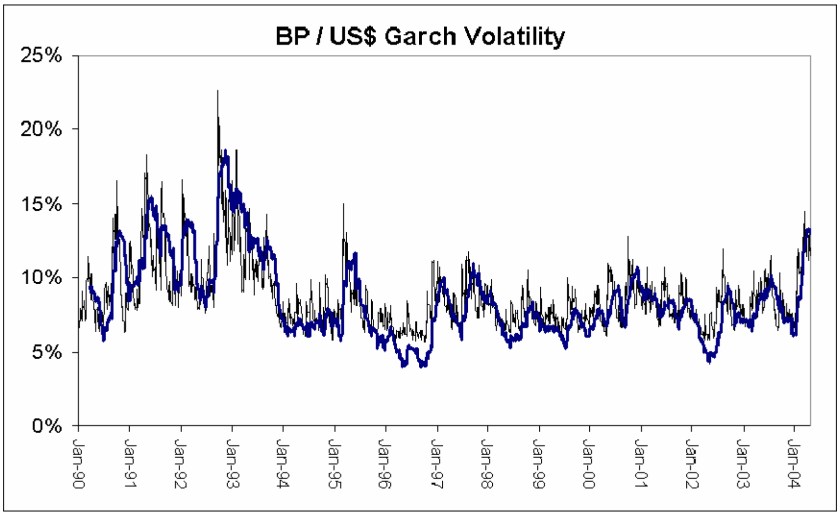

The Unconditional Disturbances Model

GARCH (1,1) Model

Following Barone Adesi et al. and Eberlein et. al. (2003):

the standardized disturbances, $\{u_t\}$, cannot be used directly to generate alternative price paths but must be "devolatised".

In the presence of GARCH type effects, the standardized disturbances should be obtained using the volatility model of equation (2) with the following expression:

$$\{\hat{u}_t\} = [r_t - \mu] / \hat{\sigma}_t(0) \qquad \textbf{(3)}$$

The Unconditional Disturbances Model

Mixing Unconditional Disturbances – MUD Model

The devolatised unconditional disturbances $\{\hat{u}_t\}$ are used to generate alternative price paths:

1. **The disturbances are randomly mixed.**
2. **The volatility is estimated using the parameters from the GARCH (1,1) model**
3. **This entails drawing from the re-shuffling until all the disturbances are selected.**

Mixing Unconditional Disturbances Model

Mixed Paths (MUD Model) & GARCH (1,1) Volatility

Assuming the volatility of the process is measured using a GARCH(1,1) process and mixing the set of devolatised disturbances , using the following expression

$$\tilde{S}_{t+1} = \tilde{S}_t e^{(\hat{\sigma}_T(0)\varepsilon\hat{u}_t)} \qquad \textbf{(4)}$$

we get a dispersion of final terminal prices.

Hochschule für Bankwirtschaft / HfB

Mixing Unconditional Disturbances Model

Mixing the $\{\hat{u}_t\}$'s plus GARCH (1,1) Volatility

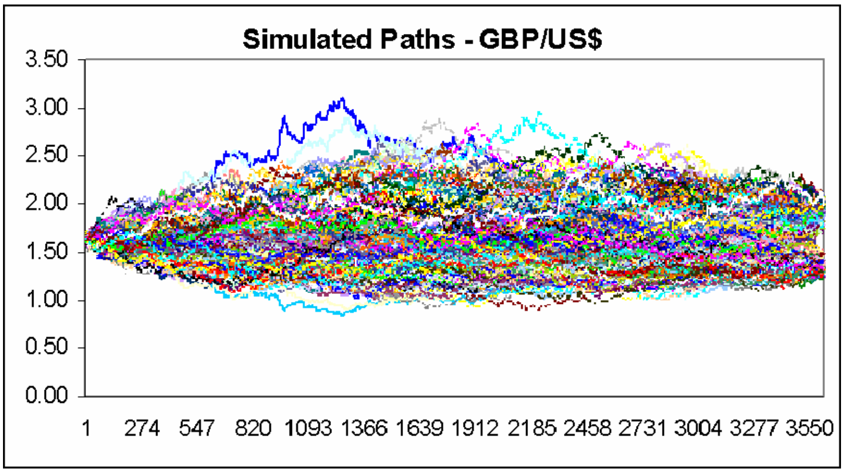

Hochschule für Bankwirtschaft / HfB

Mixing Unconditional Disturbances Model

MUD Simulated Moments vs. Original Path

	Original Price Series	Simulated Price Series
Mean	0.000027	0.00000
Standard Deviation	0.005702516	0.00564
Skewness	-0.195389482	0.00002
Kurtosis	5.5380	5.51082

Table 2. Statistical moments of simulated and original price series

Mixing Unconditional Disturbances Model

Input Parameters for the MUD Simulation

- Starting Underlying Price, S_0 = 1.612 (ln S_0 = 0.4775)
- Volatility = GARCH (1,1) average of 9.0525%
- Interest Rate (dividend) = 0%
- Time steps 1 day (1/252)
- Estimation Horizon from 1 day to 3619 days
- GBM drift would be $-\frac{1}{2}\sigma^2 t$ or 1.62593E-05

100 paths were estimated, ln(average price) per day was estimated and this was done 50 times.

Mixing Unconditional Disturbances Model

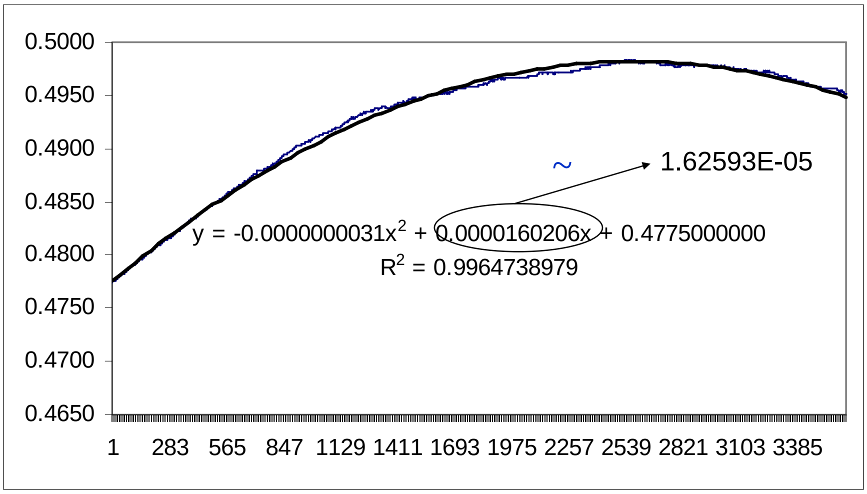

Risk Neutral Drift Adjustment – GBP/US$

As the previous graph indicates, the drift adjustment is not simplier a linear function of time but quadratic.

Of interest is that the first order term is extremely close to the GBM drift adjustment. But the second order term is negative and significantly different from zero.

What would we expect from a GARCH (1,1) type option pricing model.

Risk Neutral Drift Adjustment – GBP/US$

Heston and Nandi (2000) Garch Option Pricing Model Risk Neutral Dynamics:

$$\ln(S_t) = \ln(S_{t-\Delta}) + r - \lambda h_t + \sqrt{h_t} z_t$$

$$h_{t+\Delta} = \omega + \beta h_t + \alpha\left(z_t - \gamma\sqrt{h_t}\right)^2$$

$$\text{where} \quad z_t \sim N(0,1)$$

Risk Neutral Drift Adjustment – GBP/US$

Heston and Nandi (2000) Garch Option Pricing Model Risk Neutral Dynamics:

$$\ln(S_t) = \ln(S_{t-\Delta}) + r - \lambda h_t + \sqrt{h_t}\, z_t^*$$

$$h_{t+\Delta} = \omega + \beta h_t + \alpha\left(z_t^* - \gamma\sqrt{h_t}\right)^2$$

$$\text{where} \quad z_t \sim N(0,1),\; z_t^* = z_t + (\lambda + \tfrac{1}{2})\sqrt{h_t}$$

$$\text{when} \;\; \lambda = -\tfrac{1}{2},\; \gamma = \gamma' + \lambda + \tfrac{1}{2}$$

We are in a complete market and have a unique change of measure.

Risk Neutral Drift Adjustment – GBP/US$

In discrete time, the second order term does not disappear. Secondly, the underlying returns do not conform to lognormality. Therefore, this term should be related to the variance of the variance.

From Hodges and Tompkins (2002), the variance of the variance can be expressed as:

$$\tfrac{1}{2}\sigma^4(K-1)$$

where K is the unconditional kurtosis of the return series

In this case, this should be equal to -2.39937E-09

Risk Neutral Drift Adjustment – GBP/US$

To assess whether the coefficients of the quadratic regression on the Log of the Average simulated prices are significantly different from those in the Discrete Time non-GBM GARCH(1,1) Model, the following test was conducted.

For each of the 50 trials of the 100 simulated paths, the quadratic regression was estimated. From this, we estimated the standard error of the coefficients and conducted a t-test relative to the theoretical value of the Model.

Risk Neutral Drift Adjustment – GBP/US$

Across 50 trials	Second Order	First Order
average	-3.1040E-09	1.6021E-05
stdev	6.3694E-10	2.0159E-06
avg(ln)	-3.1E-09	1.60206E-05
Theory	-2.39937E-09	1.62593E-05
t-test	-7.77812989	-0.837424746
Null (0)	-34.415	56.194

Risk Neutral Drift Adjustment – GBP/US$

The first order drift adjustment is not significantly from the GBM drift of $-\frac{1}{2}\sigma^2 t$

But the second order term is significantly more negative than that suggested by the modified GARCH (1,1) Model.

However, both the first and second order drift adjustments are similar in orders of magnitude and sign of the model.

Hochschule für Bankwirtschaft / HfB Mixing Unconditional Disturbances Model

MUD simulation versus Bootstrapping

As this MUD simulation approach is new and relatively untried, a natural question would be how would the estimated regression coefficients compare if a traditional Bootstrapping method were used.

For this approach, the unconditional disturbances were sampled with replacement and revolatised in exactly the same manner as the MUD simulation.

Mixing Unconditional Disturbances Model

MUD simulation versus Bootstrapping

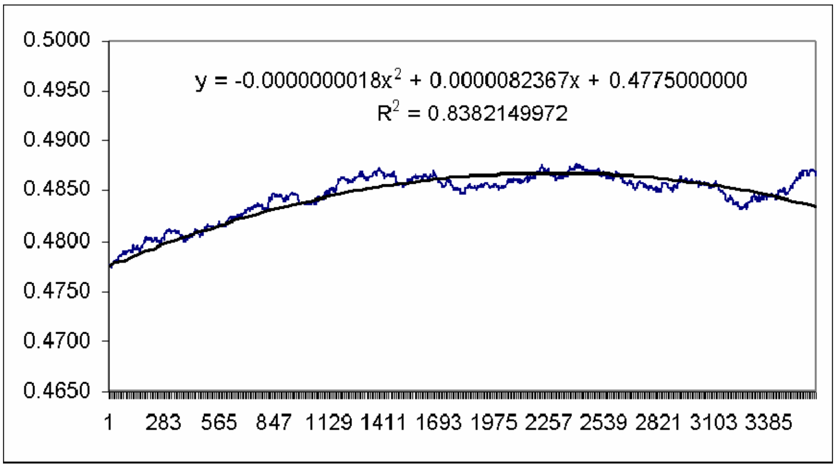

Mixing Unconditional Disturbances Model

MUD simulation versus Bootstrapping

Across 50 trials	Second Order	First Order
average	-1.808E-09	8.604E-06
stdev	5.242E-09	1.994E-05
avg(ln)	-1.8E-09	8.2367E-06
Theory	-2.3994E-09	1.6259E-05
t-test	0.80845509	-2.8450761
Null (0)	-2.428	2.921

MUD simulation versus Bootstrapping

The first order drift adjustment is now significantly from the GBM drift of $-\frac{1}{2}\sigma^2 t$

However, the second order term is not significantly different from the variance of the variance adjustment previously proposed.

However, the boot strapping method introduces substantial variation. The standard error of the regression coefficients is 8 times higher for the second order term and 10 times higher than the first order term.

MUD simulation versus Bootstrapping

Therefore, it would appear that the MUD simulation approach yields more consistent results with the GARCH (1,1) model [for the first order term] with much lower variability in the coefficient estimates compared to bootstrapping.

It remains a puzzle what the second order term in the drift adjustment implies. Simply adjusting the drift for higher order terms does not seem to adequately explain the magnitude of this adjustment. However, the sign and relative size of the adjustment are close to what would be theorertically expected

Mixing Unconditional Disturbances Model

Further Research

The next steps in the research will be:

1. To build a "black box" simulation programme and compare theoretical option prices to actual option prices.
2. Determine hedging ratios from the pricing model and estimate hedging performance versus the standard Garman Kohlhagen (1983) model and the Heston (1993) model.
3. Evaluate the pricing and hedging of Exotic options.
4. Incorporate leverage effects to evalute options on stocks and stock indices.

Hochschule für Bankwirtschaft
Professor Robert G. Tompkins
Sonnemannstrasse 9 - 11
60314 Frankfurt am Main
Phone (+49) 69 / 154 008 - 718
Fax (+49) 69 / 154 008 - 728
tompkins@hfb.de
http://www.hfb.de