# A Work Project, presented as part of the requirements for the Award of a Master's degree in Finance from the Nova School of Business and Economics.

Banco Invest Field Lab on Option Volatility Models

- A Detailed Analysis of GARCH (p, q) Models -

Diogo Lemos Martins Granja - 29056

Tiago Fernandes da Cunha Ferreira - 29207

Tomás Macedo Lopes - 41033

Tomás Almeida e Silva Tavares - 23927

David Aleksidze - 42085

Work project carried out under the supervision of:

Professor João Pedro Pereira

#### Abstract

This project aims to analyze which volatility estimation model can better forecast volatility for Banco Invest. The compared models are the GARCH (1, 1), Exponentially Weighted Moving Average, Heston-Nandi GARCH, and two variations of the Heston stochastic volatility model. The model recommended for Banco Invest is the Heston, as it is the one that presents the closest results to the realized volatility and demonstrates the most stable estimates. Alternatively, if it is not Banco Invest's intention to use the implied volatility as an input when forecasting volatilities, the Heston-Nandi GARCH model should be taken into consideration.

# **Keywords**

Volatility, GARCH, EWMA, Heston-Nandi, Heston, Volatility Surface

# Acknowledgements

We would like to thank Paulo Monteiro, Humberto Pinto and Flávio Escada, as well as Banco Invest, for the opportunity and the support given throughout this project.

This work used infrastructure and resources funded by Fundação para a Ciência e a Tecnologia (UID/ECO/00124/2013, UID/ECO/00124/2019 and Social Sciences DataLab, Project 22209), POR Lisboa (LISBOA-01-0145-FEDER-007722 and Social Sciences DataLab, Project 22209) and POR Norte (Social Sciences DataLab, Project 22209)

# **Table of Contents**

1.	Intro	duction	3
2.	Literature review		5
	2.1	Implied Volatility and the Implied Volatility Surface	5
	2.2	GARCH (p, q) Model	7
	2.3	GARCH (1, 1) Model	8
	2.4	Exponential Weighted Moving Average (EWMA) Model	12
	2.5	Volatility Surface Regression	15
	2.6	Heston-Nandi GARCH Model	17
	2.7	Heston Model	20
3.	Methodology		29
	3.1	Data Collection	29
	3.2	GARCH (1, 1), EWMA and Heston-Nandi Models' Implementation	30
	3.3	Heston FFT and MCS Models' Implementation	31
	3.4	Methodology of Empirical Analysis	32
4.	Empirical Analysis		34
	4.1	GARCH (1, 1) Model	34
	4.2	Exponential Weighted Moving Average Model	42
	4.3	Exponential Weighted Moving Average Model - Fixed Lambda of 0,94	47
	4.4	Volatility Surface Regression	50
	4.5	Heston-Nandi GARCH Model	52
	4.6	Heston Model	57
	4.7	Limitations of the Empirical Analysis	65
5.	Conc	clusion	67
6.	Diogo Granja's Individual Contribution to the Project		71
	6.1	Autocorrelation Analysis for the GARCH (1, 1) and EWMA Models	71
	6.2	Possible Improvements for the GARCH (1, 1) Volatility Estimation	74
7.	List	of References	82
8.	Appe	endix	86

#### 1. Introduction

The following project, in partnership with Banco Invest, aims to maximize the precision of the volatility forecasts performed by the bank.

Banco Invest has diverse operations on the capital market, with this project aiming to assist the bank with structured products. This organization's department provides bundles of options to its customers where the products' payoff depends on the performance of the underlying assets of the options included. As an example, one of the products sold by the bank is an investment of a two-year period in a bundle, created by the bank, composed by five different stocks in the technology sector. The product is composed by stocks from ASML, Corning, Crown Castle, Vodafone and Qualcomm. The remuneration the customers receive on this investment depends on the returns of the bundle as a whole. For the aforementioned product, the customer would receive a payoff of 2,4% if all stocks have increased value at the due date, otherwise the yield would go down to 0%.

As an institution that guarantees a minimum payoff of 0%, as well as potential for positive returns, the bank cannot be exposed to market shocks, fluctuations and unpredictability. Therefore, Banco Invest hedges its position, to guarantee a consistent and controlled payoff, diminishing the risk of its operations.

As explained by Hull (2009, chapter 19), hedging a position, or delta hedging, implies buying or selling underlying asset stocks of the possessed option. Whether the position is long or short depends on the *delta* of the option, given by the change in the option price with respect to the price of the underlying asset. A position is considered hedged when delta equals zero. Maintaining delta at zero is a continuous process, as the position needs to be updated and adapted according to the market changes.

If the portfolio is properly hedged, the bank ensures that upon options' expiry, its portfolio of securities is balanced and has zero payoff. Banco Invest is then able to either yield positive returns to customers, if the spot prices go up, or incur no losses, if the spot prices don't go up.

However, while the payoff is guaranteed to be zero, losses can still be incurred. Constant purchase and sale of stocks involves substantial transaction costs. Moreover, many other Greeks such as Gamma, Vega and Rho should also be hedged to guarantee perfect hedging. Perfect hedging is not feasible in practice, due to accumulation of transaction costs (Hull 2009, chapter 19).

Given this, accurate expectations regarding the market considerably minimize the transaction costs associated with trading, thus maximizing the gains the bank can generate from its customers' investments. As the best estimate for how a stock, and consequently an option, is expected to behave in the future is the volatility of the underlying asset, volatility estimation plays a critical role on the hedging efficiency of the bank. An option price is severely influenced by the expected volatility and therefore, a competitive advantage exists when one institution has a better forecast on how the volatility is going to behave.

For years, Banco Invest has been using the Bloomberg implied volatilities, which are extrapolated from the price of the options. The bank has had positive results using implied volatilities but aims to enhance its option volatility accuracy even further. Moreover, due to the COVID-19 pandemic's repercussions, the implied volatilities did not always provide the best estimates, due to the unpredictability of the market. Since the implied volatilities have shown weakness in this period, the bank aims to find a model that will more effectively cover these weaknesses.

Consequently, the goal of this project is to comprehend not only how existing models are able to predict market volatilities, but also to assess the quality of the results proposed by the implied volatility model. To do so, the implied volatility results will be subjected to the same type of scrutiny as the other models' outputs.

The models to be implemented are the GARCH, EWMA, Heston-Nandi GARCH and Heston models. These models were selected based on their theoretical robustness and alignment with the needs proposed by the bank. At the same time, these models forecast volatility with different methods and different inputs, providing the bank a broader analysis of volatility models. The GARCH and the EWMA models only take into consideration historical volatilities. The Heston model, on the other hand, adds a stochastic volatility element to the initial input of the model. Finally, the Heston-Nandi GARCH incorporates both previously mentioned aspects, using the historical volatility in the parameterization process and the random component for the option pricing.

Each model will be evaluated not only on the fitness of the results when compared to the realized volatility but also their stability, as unstable models would increase transaction costs. The project also considers the applicability of the models, as very complex ones from a computation perspective may not be feasible. This represents a problem because Banco Invest updates twice a day, on average, 63 option volatilities for each of the over 70 stocks present in the portfolio. The bank has the need to have volatility estimates for options with seven different maturities, ranging from one month to two years, and nine distinct strikes, between 80% and 120% moneyness, for each stock of its portfolio.

#### 2. Literature review

# 2.1 Implied Volatility and the Implied Volatility Surface

Implied volatilities (IV) are derived from option prices in the market and are widely used among traders as the most reliable proxy for the market's volatility forecast. To compute such IV values, the Black-Sholes equation for pricing an option is used (Hull 2009, chapter 15). In this paper, the equation to price a call option will be given as an example, but the same rational

can be applied in the put option pricing equation. As such, equation 2.1.1 represents the starting point to properly derive IV estimates.

$$c = S_0 N(d_1) - K e^{-rT} N(d_2). (2.1.1)$$

Where  $N(d_1)$  and  $N(d_2)$  represent cumulative probability distribution functions for a variable with a standard normal distribution,  $d_1 = \frac{ln\binom{S_0}{K} + \binom{r-\sigma^2}{2}T}{\sigma\sqrt{T}}$  and  $d_2 = d_1 - \sigma\sqrt{T}$ .

To be feasible to derive the IV values from this equation, all variables, except the volatility parameter ( $\sigma^2$ ), must be known values. Therefore, it is necessary to have available underlying asset spot price at time zero ( $S_0$ ), the strike, maturity, and the value of the option (K, T and c, respectively), as well as the continuously compounded risk-free rate verified on the market (r). From this, it is possible to derive the intended IV following a trial-and-error approach, since it is not possible to invert equation 2.1.1 so that  $\sigma$  is expressed as a function of  $S_0$ , K, r, T, and c (Hull 2009, chapter 15).

This method has the advantage of directly computing the implied volatility of a given option as a function of its strike and maturity. This characteristic enables the creation of an implied volatility surface: a forecast of infinite implied volatilities in condition to the moneyness of the underlying asset and its time to maturity (Gatheral 2006, chapter 3). The Black-Scholes model (1973) predicts a volatility function curve for a certain fixed maturity across different strike prices that resembles a "smile", where the marker gradually decreases up until it reaches at-themoney point, and then increases again. With trading costs considered, the volatility curve will resemble a "smirk". This happens due to the traders being more concerned with hedging their own portfolios against market loss than capitalizing on huge gains, therefore, creating higher demand for out-of-the-money or at-the-money options.

A crucial assumption of this model is that the market is perfect when it comes to react to unexpected events and jumps in prices, that is, the implied volatility model assumes that the prices at which options are being traded in the market are correct and represent the true value of the asset. However, this scenario is not verified, since no market is 100% perfect, with options mispricing being a reality (Canina and Figlewski 1993).

In sum, it is clear that, although this method is widely used among option traders, the implied volatility model has its shortcomings, and it is adequate to try to find alternatives that can provide volatility estimates more aligned with the market's realized volatility.

# 2.2 GARCH (p, q) Model

The GARCH (p, q) (generalized autoregressive conditionally heteroskedastic) process was first introduced by Bollerslev (1986) and represents a generalization of the ARCH (q) (autoregressive conditionally heteroskedastic) model presented by Engle (1982). Both these methods allow "conditional variance to change over time as a function of past errors" (Bollerslev 1986), which represents a major difference when compared with other econometrics models where conditional variance is expected to remain constant.

Engle (1982) demonstrated that the variance of one period ( $\sigma_t$ ) is a random variable that is dependent of past volatilities ( $\sigma_{t-i}$ ), which led him to create a model that would forecast future variances as random variables that are dependent of past information. Given this, in both GARCH (p, q) and ARCH (q) the letters p and q stand for the number of lagged periods used to forecast volatility at a certain point in time. The difference between models is related with the existence of one more variable in the GARCH (p, q) process. To better understand this difference, both the ARCH (q) and the GARCH (p, q) equations must be studied.

The ARCH (q) model equation for the estimation of a given stock is

$$\sigma_n^2 = \gamma V_L + \sum_{i=1}^q \alpha_i \mu_{n-i}^2$$
 (2.2.1)

While the GARCH (p, q) method equation for the same forecast is

$$\sigma_n^2 = \gamma V_L + \sum_{i=1}^q \alpha_i \mu_{n-i}^2 + \sum_{i=1}^p \beta_i \sigma_{n-i}^2$$
 (2.2.2)

In both equations  $V_L$  stands for the long-run average variance,  $\mu_{n-i}^2$  is the squared return of the stock at the lagged period n-i, n is the sample size, while  $\alpha_i$  and  $\gamma$  are parameters of the models. In the GARCH (p, q) equation,  $\sigma_{n-i}^2$  is the conditional variance estimated for the lagged period n-i and  $\beta_i$  is a parameter of the model. The importance and meaning of the parameters will be discussed later in the paper.

According to Bollerslev (1986), "in the ARCH (q) process the conditional variance is specified as a linear function of past sample variances only, whereas the GARCH (p, q) process allows lagged conditional variances to enter as well", which enables the model to have an adaptive learning mechanism. Moreover, it is clear that if p = 0, then the GARCH (p, q) process is simply the ARCH (q).

Given this, in the following section of this thesis only the GARCH (p, q) process will be considered, as it represents a more complete approach to volatility forecast when compared with the ARCH (q) model.

#### 2.3 GARCH (1, 1) Model

The GARCH (1, 1) model is the simplest approach among all GARCH processes, since it considers only one lag period, making it computationally more feasible and applicable, when compared with other derivations of the model (Bollerslev 1986).

As it is Banco Invest's goal to have a computationally simple model as final output, and the GARCH (1,1) is often proven to derive reliable results (Bollerslev 1986), this method is going to be deeply studied and put to practice in the next chapters. Moreover, since the main objective of this project is to derive daily volatility forecasts for a given stock, from this point onwards  $\sigma_n^2$  represents the asset's daily variance,  $\sigma_{n-1}^2$  the daily variance lagged one period, that is, the conditional variance estimated for the day before and  $\mu_{n-1}^2$  the squared stock return in the day before.

The equation for the GARCH (1, 1) process is given by:

$$\sigma_n^2 = \gamma V_L + \alpha \mu_{n-1}^2 + \beta \sigma_{n-1}^2 \tag{2.3.1}$$

Where  $V_L$  stands for long-run average variance,  $\mu_{n-1}^2$  the day before squared returns of the underlying stock and  $\sigma_{n-1}^2$  the day before daily volatility. The Greek letters that are multiplying by each of the values are the parameters that allocate weight to each variable (Hull 2009, chapter 23). The three parameters must obey the following constraint:

$$\gamma + \alpha + \beta = 1 \tag{2.3.2}$$

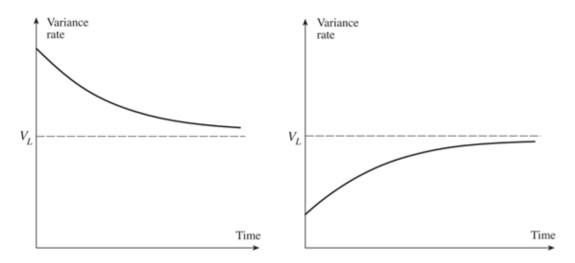
The model is often simplified to:

$$\sigma_n^2 = \omega + \alpha \mu_{n-1}^2 + \beta \sigma_{n-1}^2 \tag{2.3.3}$$

Where  $\omega$  is given by the product between  $\gamma$  and  $V_L$  in order to simplify the equation and the parameter estimation.

This model assumes the property of mean reversion aspect to the volatility forecasted, that is, the model anticipates that there will be a tendency for the daily volatility computed  $(\sigma_n^2)$  to get closer to the long term variance of the underlying asset as time passes. This characteristic of the model is illustrated in the following picture (Hull 2009, chapter 23):

Figure 2.3.1 – Illustration of the mean reversion property of the GARCH (1, 1) model



Furthermore, the understanding of the meaning of each parameter is imperative to understand how de model derives the volatility. There are three parameters that attribute wight to each component that GARCH (I, I) assumes to impact the volatility.

First, and the most relevant, is the  $\beta$ . This parameter defines how much of the volatility of the day n should be impacted by the volatility of the day before, n-1. As  $0 < \beta < 1$ , the higher the value the parameter the bigger correlation there will be between the volatility forecasted and the historical volatility. Moreover, because n-1 volatility is influenced by n-2, a higher value of  $\beta$  results in more past values being incorporated in today's volatility, resulting in increased stability of the model.

Another parameter,  $\alpha$ , is the one responsible for assuming a relation between the volatility and the asset's return, more precisely the squared returns. The squaring of the returns assures that, independently if the return was positive or negative, the impact in the volatility estimation will always be positive, assuring that the volatility predicted by this model cannot be negative.

Finally, the third parameter,  $\gamma$ , defines the weight attributed to the expected long-run volatility of the model. The higher the  $\gamma$  is, the stronger the mean reversion implied by the model will be. Opposing, if this parameter is very close to zero, the concept of the GARCH (1, 1) model being a mean-reverting one becomes virtually obsolete.

An important constraint of the model is that  $\alpha + \beta < 1$ , which consequently implies that  $\gamma$  is positive. If  $\alpha + \beta > 1$ , the model assumes a negative weight for the long-run variance, that is,  $\gamma$  would be lower than zero to respect equation 2.3.2. Moreover, as w would have the constraint to be larger than zero, VL would have to be negative in order to be in accordance with the equation  $w = VL \times \gamma$ , which is an impossible result as the volatility can never be negative. Under this scenario the model would be mean fleeing instead of mean reverting, thus violating a critical assumption of the GARCH model.

In order to estimate the correct parameters for this model, the maximum likelihood method must be performed. This process provides the values of the parameters that maximize the probability of the data occurring, and it is given by the expression:

$$\sum_{n=1}^{m} \left[ -\ln(\sigma_n^2) - \frac{\mu_n^2}{\sigma_n^2} \right]$$
 (2.3.4)

The maximization of this equation provides a result for  $\alpha$ ,  $\beta$  and  $\omega$ , parameters that can immediately provide the value for the long-term daily variance rate  $(V_L)$  for a given day (n) with the following equation:

$$V_L = \frac{\omega}{1 - \alpha - \beta} \tag{2.3.5}$$

Depending on the parameter results, weight will be attributed to the long-run variance, the previous returns of the asset and the volatility (Engle 2001).

The volatility result that is taken away from this model is a daily volatility that is going to be annualized assuming 252 business days in a year. Also, this volatility is not skewed by time to maturity or moneyness of an option. Banco Invest and Bloomberg assume that historical option volatilities are regarded as volatilities of options that have 100% moneyness. Moreover, the assumption was made that the yearly volatility will consist in a constant daily volatility that was predicted for the 252 business days. Therefore, based on the equation 2.3.3, the yearly volatility is given by the following equation:

$$\sigma_{annual} = \sqrt{252 \left(\omega + \alpha \mu_{n-1}^2 + \beta \sigma_{n-1}^2\right)}$$
 (2.3.6)

With this, the result from this model references the point of moneyness of 100% and maturity of one year in the volatility surface. Chapter 2.5 is going to extrapolate how this model can be incorporated to form a volatility surface.

#### 2.4 Exponential Weighted Moving Average (EWMA) Model

The exponentially weighted moving average (EWMA) represents an approach where the weights given to the volatilities of the dates n-t decrease exponentially through past observations. With this, the "decay" parameter of this model is given by  $\lambda$  and it is a constant between zero and one. This model provides a simple formula that computes today's volatility based on the previous days' volatilities, returns and on the value of  $\lambda$ , a parameter which is constant and used to define the weight given to each observation (Hull 2009, chapter 23). The formula of this model is therefore given by:

$$\sigma_n^2 = \lambda \sigma_{n-1}^2 + (1 - \lambda)\mu_{n-1}^2 \tag{2.4.1}$$

Although this formula appears to only use the volatility and the return of day n-1 to extrapolate the volatility for day n, it is important to mention that that is not the case. This can be explained by the fact that the day n-1 volatility is calculated from the return and volatility of day n-2. On its turn, the day n-2 value was based on a past observation, with this method being reproduced until the end of the time series. In sum, the volatility estimation for day n englobes all estimates previously performed for past periods. However, despite this process continuing until all data is used, it is important to refer that the weight given to a day at a certain point in the past becomes irrelevant and therefore this model does not require to store as much data as alternative models. Moreover, in addition to this positive aspect of the model, as the most recent day is incorporated in the model, the oldest data can be disregarded from the model as it is not relevant anymore. This allows the model to maintain its easiness to operate (Hull 2009, chapter 23). The formula that encompasses all the data in the model, from the first day of data until today, is given by:

$$\sigma_n^2 = (1 - \lambda) \sum_{i=1}^m \lambda^{i-1} \mu_{n-1}^2 + \lambda^m \sigma_{n-m}^2$$
 (2.4.2)

It is important to note that equation 2.4.1 and 2.4.2 are equivalent if the sample size (m) is large enough to make  $\lambda^m \sigma_{n-m}^2$  sufficiently small to be ignored (Hull 2009, chapter 23). Given this, the study of equation 2.4.2 is useful to better understand how the model behaves. This formula highlights how the exponential decline of the EWMA approach is given by the fact that the weight of today's observation is multiplied by the weight previously given to the result obtained in n-1, with the same rational being applied to the remaining sample size. From this, it is clear that, as the available data becomes older, its relevance exponentially declines.

The volatility that results from this model is largely dependent on the value of the parameter  $\lambda$ , since this value defines the weight given to previous periods. As this parameter gets lower, more prevalence is given to the most recent return. Moreover, as the decay becomes larger due to the lower values of the parameter, the sooner older data becomes obsolete and irrelevant for the estimation. On the opposite, as the parameter  $\lambda$  becomes greater, the model starts to react slower to changes in volatility, as it gives a greater weight to older data (Hull 2009, chapter 23), resulting in more stable estimations.

There are two possible ways to estimate the parameter to this model and both methods forecast quite different volatilities estimates. On one hand Hull (2009, chapter 23) proposes a process that is very similar to the GARCH (1, 1) model previously mentioned. Under this scenario, the parameter is estimated by maximizing the exact same likelihood function used in the GARCH process. The function that is to be maximized for this parameterization is then given by:

$$\sum_{n=1}^{m} \left[ -\ln(\sigma_n^2) - \frac{\mu_n^2}{\sigma_n^2} \right] \tag{2.3.4}$$

In which  $\sigma_n^2$  represents the volatilities estimate for a single period and  $\mu_n^2$  represents the return squared for the same period. By maximizing the sum of this equation for each day of the model the parameter ( $\lambda$ ) can be extracted. From an econometrical standpoint, this process is the

most correct to derive  $\lambda$ . However, when put to practice, it can result in worse estimates. This can be explained by the fact that the parameter is very sensitive to the data available, and therefore unusual shocks in the market or in that particular stock may result in extremely volatile  $\lambda$  values in short periods of time. Consequently, a volatile  $\lambda$  will originate less stable volatility forecasts, which is far from ideal given what is asked by Banco Invest (Hull 2009, chapter 23).

On the other hand, an empirical study performed by JPMorgan (1994) suggests fixing the parameter  $\lambda = 0.94$ . The study suggests that "across a range of different market variables, this value gives forecasts of the variance rate that come closest to the realized variance rate". Although this approach disregards the maximum likelihood equation, this method allows to have not only a more constant result for the volatility but also a result that often better represents the true market volatility. However, it is also clear to understand that this volatility estimation approach is less econometrically correct.

The EWMA model has historically been a trustworthy model from a theoretical base as well as from an empirical standpoint. The EWMA model attempts to produce similar results to the GARCH (1, 1) model with less parameters because it is using a similar methodology of declining the weights of historical values and therefore exponentially declining the usefulness of historical information, providing results more adequate to the current market circumstances (Jarešová 2010). Because it can theoretically arrive to similar results with a smaller data collection and number of parameters, this allows the model to be computationally less heavy.

However, the EWMA model does not incorporate the mean reversion of the GARCH (1, 1) model that is computed with the previously mentioned parameter w. The EWMA model is a singular case of the GARCH model where the parameter w is set as zero, consequently implying  $\alpha + \beta$  to equal one. From this, it is clear that the  $\lambda$  parameter in the EWMA has the same

function as the  $\beta$  parameter in the GARCH process, while the estimation of  $\alpha$  is ignored, being given by  $\alpha = 1 - \lambda$ .

The fact that this model does not account the mean reversion implies that it is theoretically less appealing. With this, by comparing these two models (GARCH and EWMA), there is a clear tradeoff between the theoretical quality of the two and the computational easiness of the two.

Once again, the volatility result that is taken away from this model is a daily volatility that is going to be annualized assuming 252 business days in a year. Therefore, based on the equation 2.4.1, the yearly volatility is given by the following equation:

$$\sigma_{annual} = \sqrt{252 \left[ \lambda \sigma_{n-1}^2 + (1 - \lambda) \mu_{n-1}^2 \right]}$$
 (2.4.4)

Moreover, this volatility is not skewed by time to maturity or moneyness of an option, therefore the result from this model references the point of moneyness of 100% and maturity of one year in the volatility surface.

# 2.5 Volatility Surface Regression

The results that are taken from the GARCH and EWMA models can be observed in the volatility surface as the annual volatility for an option at the money. Therefore, they will only provide results for the point T = 1 and K = 100% of the volatility surface (with T being the time to maturity and K the moneyness of the option). Since these models only provide a volatility forecast for a singular point of the surface, it is important to adapt this result to different moneyness and maturities. The following method would allow the extrapolation of results as a function of K and T while using the results of the models as initial inputs. Thus, the structure of the volatility surface will be similar to the implied volatility, but the values will differ from each model to another, depending on the input provided.

With the implied volatility being a function of K and T, one can use the ordinary least squares regression method to get the parameters that are able to mirror the implied volatility surface. Therefore, the coefficients can provide an estimation for a volatility for each time to maturity and moneyness (Nandi and Waggoner 2000).

The model provided by Dumas, Fleming and Whaley (1998) consists in:

$$\sigma(T,K) = a_0 + a_1K + a_2K^2 + a_3T + a_4T^2 + a_5KT$$
 (2.5.1)

With K being the strike, T the time to maturity and  $\sigma$  a function of both. It is expected that the parameters provided are a good fit. However, by adding more parameters and therefore creating a more elaborated model, it can be the case that the model becomes more precise but also more redundant due to statistically irrelevant parameters (Dumas Fleming and Whaley 1998). The quality of the fit is also heavily deteriorated by the inclusion of deep "in the money" or "out of the money" volatilities, as these options are characterized by higher volatilities and therefore higher deltas between volatility and moneyness (Andersen and Andreasen 2000). The estimation of the volatility surface that is going to be applied on the models previously mentioned will be conducted by minimizing the squared difference between the "real" implied volatility surface and the estimated one with the parameters, given by the equation 2.5.1.

Herewith, this regression method would construct an entire volatility surface by only inputting a singular volatility estimation at the point (T, K) = (1, 100%), as the ordinary least squared regression would provide a surface like the implied volatility one regarding the slope, while assuming the other models' volatility forecast. This regression will be applied to the results of the GARCH (I, I) and the EWMA models.

#### 2.6 Heston-Nandi GARCH Model

One of the conditions in the Black-Scholes model is the constant volatility, which is something that is not verifiable. One way to solve the problem is the time-varying volatility model of Heston (1993), that allows to relax this assumption. Steven L. Heston and Saikat Nandi developed a model in 2000, known as Heston-Nandi, which is a closed-form option valuation formula for a spot asset where the variance follows a GARCH (p, q) model, which can be done with the returns of the asset.

Previous researchers tried to create a GARCH option pricing models, like Duan in 1995, but no closed-form solution for the option pricing was available, meaning the usage of Monte Carlo simulations was needed. Heston and Nandi (2000) presented a closed-form GARCH model for the option pricing, solving the latter problem.

The Heston-Nandi derivation of the GARCH model is a very similar approach as the Duan option pricing model, which first introduced a mean-reverting GARCH model that eliminates the drawback of a singular and constant standard deviation. This model predicts the smile/smirk of the volatility curve which therefore provides adjusted results for different maturities and strikes (Kamiński 2013).

The closed-form model allows option valuation from Heston-Nandi to follow the GARCH variance as it is still based on past returns, while containing Heston's stochastic volatility model (Heston and Nandi 2000). With this, the option pricing model not only considers the path of the historical spot prices of the underlying asset but also the stochastic correlation between volatility and returns, which is explained by the leverage effect component of this model. The volatility for day *t* of this option pricing model is given by the equation:

$$h(t) = \omega + \sum_{i=1}^{p} \beta_i h(t - i\Delta) + \sum_{i=1}^{q} \alpha_i \left[ z(t - i\Delta) - \gamma_i \sqrt{h(t - i\Delta)} \right]^2$$
 (2.6.1)

Where h(t) is the past volatility, z is standard normal disturbance and  $\omega$ ,  $\alpha$ ,  $\beta$  and  $\gamma$  are parameters of the model. Omega ( $\omega$ ) is the weight assigned to the long-term volatility multiplied by the long-term volatility ( $\omega = x * \mathcal{V}_l$ ). Alpha ( $\alpha$ ) provides the weight assigned to the last observed return ( $u_{n-1}^2$ ). Beta ( $\beta$ ) is the weight assigned to the last variance observed ( $\sigma_{n-1}^2$ ) and Gamma ( $\gamma$ ) provides an important property that is disregarded in the basic GARCH (1, 1) model which is the leverage effect (Christoffersen, Jacobs and Ornthanalai 2013).

The leverage effect, first discussed by Black, is an asymmetry tendency regarding the different impacts in the volatility of positive and negative returns. This trend is present because of an increase of debt relative to equity when the spot price of a company decreases. With a decrease in equity and a stable debt, the Debt-to-Equity ratio increases which consequently increases the risk of the company. This explains why the volatility increase is more significant because of negative returns compared to positive ones when taking this parameter into consideration (Ait-Sahalia, Fan and Li 2013).

These parameters are estimated, like GARCH (1, 1), by maximizing a particular log likelihood function that not only includes the parameters previously mentioned but also a  $\lambda$ , the risk premium. This maximization equation is given by (Rouah and Vainberg 2007, chapter 6):

$$\log L = \sum_{t=1}^{n} \left( -\ln \sigma_t^2 - \frac{r_t^2}{\sigma_t^2} \right)$$
 (2.6.2)

A crucial part of the parameter calculation is the initial parameter values that are used in the model in order to estimate the optimized model. Because this model contains five parameters to be estimated, certain combinations of parameters do not create an equation that is able to be optimized, either because they tend to infinite or because the local maximum is far away from the optimum level of the function (Rouah and Vainberg 2007, chapter 6). Therefore, it is imperative for the initial variables to be relatively close to the correct optimum parameters.

These initial parameters used to maximize the model are parameters used in the Heston-Nandi paper (Heston and Nandi 2000).

With Rt being:

$$R_t \equiv \ln\left(\frac{S_t}{S_{t-1}}\right) = r + \lambda h_t + \varepsilon_t \tag{2.6.3}$$

The model assumes that h(t) represents the return premium, meaning the future returns of the underlying asset can be explained by a risk-free component and a component, proportional to the variance, that exceeds this riskless proportion of the return (Heston and Nandi 2000). Moreover, the model obeys two main assumptions, spot prices follow the path previously mentioned throughout time and that the call option with one period to expiration obeys the Black-Scholes-Rubinstein formula (Heston and Nandi 2000).

In order to price the option resulting from this variance, Heston-Nandi proposes a closed-form formula for European options. The formula allows to compute an option price for any given K and T (Byun 2011), in contrast with other models that require the use of a Monte Carlo to estimate the option price, which is more computationally intensive. The previously mentioned formula was firstly elaborated by Heston and Nandi (2000) and is given by:

$$C = \frac{1}{2}S_t + \frac{e^{-r(T-t)}}{\pi} \int_0^\infty Re \left[ \frac{K^{-i\phi}f(i\phi + 1)}{i\phi} \right] d\phi$$
$$-Ke^{-r(T-t)} \left( \frac{1}{2} + \frac{1}{\pi} \int_0^\infty Re \left[ \frac{K^{-i\phi}f(i\phi)}{i\phi} \right] d\phi \right)$$
(2.6.4)

With 
$$f(i\phi) = S(t)^{i\phi} exp \left[ A(t; T, i\phi) + \sum_{i=1}^{p} B_t(t; T, i\phi) h(t + 2\Delta - i\Delta) + \sum_{i=1}^{q} C_i(t; T, i\phi) \left[ z(t + \Delta - i\Delta) - \gamma_i \sqrt{h(t + \Delta - i\Delta)} \right]^2 \right]$$
(2.6.5)

$$A(t; T, i\phi) = A(t + \Delta; T, i\phi) + i\phi r + B_t(t + \Delta; T, i\phi) w$$
$$-\frac{1}{2} \log[1 - 2a_1B_1(t + \Delta; T, i\phi)]$$
(2.6.6)

$$B_{1}(t + \Delta; T, i\phi) = i\phi(\lambda + \gamma_{1}) - \frac{1}{2}\gamma_{1}^{2} + \beta_{1}B_{1}(t + \Delta; T, i\phi)$$

$$+ \left(\frac{\frac{1}{2}(i\phi - \gamma_{1})^{2}}{1 - 2a_{1}B_{1}(t + \Delta; T, i\phi)}\right)$$
(2.6.7)

The Heston-Nandi parameters previously mentioned are incorporated in equations 2.6.6 and 2.6.7, while " $f(i\phi)$  is the characteristic function of the logarithm of the spot price" (Heston and Nandi 2000).

With all the processes explained before, one obtains the European option price for any desired time to expiration as well as a variety of strikes. As previously mentioned, Heston-Nandi model follows the Black-Scholes-Rubinstein formula and therefore one can use this formula to extrapolate the local volatility (Heston and Nandi 2000) from the European option price obtained before. The Black-Scholes formula is used to reverse engineer the equation in order to extract the volatility, with inputs being the settlement date and maturity date, the strike of the option, risk free rate and finally the option price. With all the inputs previously mentioned, it is possible to extrapolate which volatility is implied in the option price estimated by the model, for a given maturity and strike price.

### 2.7 Heston Model

Heston (1993) developed a stochastic volatility option pricing model, which allowed variance to follow a random path. Heston assumes that the underlying asset follows a risk neutral continuous process:

$$dS_t = rS_t dt + \sqrt{V_t} S_t dW_t^1 (2.7.1)$$

$$dV_t = k(\theta - V_t)dt + \sigma\sqrt{V_t}dW_t^2$$
(2.7.2)

$$dW_t^1 dW_t^2 = \rho dt \tag{2.7.3}$$

This model takes nine main input parameters, divided into two groups. The first group consists of parameters used to price options with any model and that are observable in the

market, them being the stock price S, risk-free rate r, strike price K and maturity T. The second group can be obtained through a calibration process that will be explained below. It consists of initial volatility  $V_t$  of the asset, its long-term volatility  $\theta$ , the rate of mean reversion k between  $V_t$  and  $\theta$ , the volatility of volatility  $\sigma_t$  and the correlation between two random Wiener processes  $\rho$ , Wu (2019).

All components have specific effects on the resulting curve, where  $\rho$  affects the skewness of the distribution. As shown by Wu (2019) and Poklewski-Koziell (2012), the equation dictates that with positive  $\rho$ , the volatility of the asset and its returns will increase together. Conversely, with  $\rho < 0$ , the asset's volatility increases with a drop in log-returns. The parameter  $\sigma_t$ , the volatility of volatility, affects kurtosis which influences the heaviness of tails in general.

With lower kurtosis, tails will be heavier, while larger kurtosis signifies a distribution packed around the mean. Peak might also be changed, although not necessarily so, as  $\alpha$  affects the mean-reversion. The lack of mean-reversion would yield markets characterized by a considerable amount of assets with volatility either absurdly high or approaching zero, which are, in practice, fringe cases. As shown by Crisostomo (2014) or Poklewski-Koziell (2012), the Feller condition  $2a\theta \ge \sigma^2$ , from Feller (1951), is applied, which guarantees that all  $V_t \ge 0$ .

# **Closed Form Solution**

Heston (1993) suggested a closed form solution for the model:

$$C(S_t, \Pi_j, t, T) = S_t \Pi_1 - Ke^{-r(T-t)} \Pi_2, \qquad j = 1,2$$
 (2.7.4)

where,

$$\Pi_{1} = \frac{1}{2} + \frac{1}{\pi} \int_{0}^{\infty} Re\left(\frac{e^{-iuln(K)}\psi(u-i)}{i \cup \psi(-i)}\right) du = S^{*}$$
(2.7.5)

$$\Pi_{2} = \frac{1}{2} + \frac{1}{\pi} \int_{0}^{\infty} Re\left(\frac{e^{-iuln(K)}\psi l^{u}i}{iu}\right) du = P^{*}$$
(2.7.6)

Heston (1993) considers the characteristic function  $\psi$  of both probabilities to assume the form:

$$\psi_{xt}(u) = E[e^{iuxt}] = e^{\{C(u,t)\theta + D(u,t)V_0 + iu(\log(S_0) + rt)\}}$$
(2.7.7)

where,

$$C(u,t) = a \left[ r_{neg}(t) - \frac{2}{\sigma^2} \log \left( \frac{1 - ge^{-dt}}{1 - g} \right) \right]$$
 (2.7.8)

$$D(u,t) = r_{neg} \left[ \frac{1 - e^{-dt}}{1 - ge^{-dt}} \right]$$
 (2.7.9)

$$g := \frac{r_{neg}}{r_{pos}} \tag{2.7.10}$$

$$\frac{r_{neg}}{r_{pos}} = \frac{\beta \pm d}{2\gamma} \tag{2.7.11}$$

$$d = \sqrt{\beta^2 - 4\gamma\alpha} \tag{2.7.12}$$

$$\alpha = \frac{(-u^2 - iu)}{2} \tag{2.7.13}$$

$$\beta = k - \rho \sigma i u \tag{2.7.14}$$

$$\gamma = \frac{\sigma^2}{2} \tag{2.7.15}$$

Equation 2.7.4 is very similar to the Black-Scholes model. The first term represents the present value of the underlying asset, given optimal exercise and the second term is the present value of the strike.  $\Pi_1$  and  $\Pi_2$  have a similar role to  $d_1$  and  $d_2$  in the Black-Scholes framework, with the forward price and the strike price split into two different components. To compute the pricing of European option prices, equations 2.7.4, 2.7.5, 2.7.6 and 2.7.7 are approximated making use of various numerical integration techniques, as shown by Moodley (2005) and Poklewski-Koziell (2012).

However, Carr and Madan (1999) '...find that the use of the Fast Fourier Transform (FFT) is considerably faster than most available methods and furthermore, the traditional method

described in Heston, Bates, Bakshi and Madan, and Scott can be both slow and inaccurate...'. Kahl and Jäckel (2005) further improved on the work of Carr and Madan (1999) by the application of the rotation count algorithm to the original Heston formulas, thereby fixing the discontinuities in the integrals. This latter application of Fast Fourier Transform is the one followed in this paper.

### **Fast Fourier Transform**

Moodley (2005) demonstrates that Fast Fourier Transform is useful in solving equations in the following form:

$$w(k) = \sum_{j=1}^{N} e^{-i\frac{2\pi}{N}(j-1)(k-1)} x(j)$$
 (2.7.16)

This equation displays a similar form to that of equation 2.7.7, the characteristic function in the closed form solution of the Heston model. With this, the Fast Fourier Transform can now be applied:

$$\hat{f}(u) = \int_{-\infty}^{\infty} e^{-iux} f(x) \, d(x)$$
 (2.7.17)

The Inverse-Fourier Transform:

$$f(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{-iUx} f(u) d(u)$$
 (2.7.18)

This is the general framework of the Fourier Transforms. FFT is a methodology that greatly speeds up the Discrete Fourier Transform, through manipulation of the matrices in order to reduce the number of calculations the model would have to do. Given N number of different variables, without FFT, the number of calculations that would have to be run would proportionally be  $N^2$  and  $N[\log(n)]$ , respectively. So, the higher the number of calculations is, the larger the discrepancy between the two methodologies is.

In order to arrive at the solution and execute FFT, the characteristic function needs to be derived. This will be done on the log-underlying of any given S. The framework for the characteristic function for x is represented as:

$$\psi_{xt}(u) = E[e^{ixut}] \tag{2.7.7}$$

Where t is the part of the transform, E is the expectation and x is the variable that is being analyzed. The left side of the equation can be translated into:

$$\int_{-\infty}^{\infty} e^{-iutx} p(x) d(x) \tag{2.7.19}$$

Where  $e^{-itx}$  is multiplied by the probability density and then integrated from minus to plus infinity. The probabilities  $\Pi_1$  and  $\Pi_2$  are derived through reversing the characteristic function, as it had been laid out in the FFT framework. The result is obtained through running FFT on the characteristic function.

#### **Monte Carlo Simulation**

In addition to solving the closed form solution suggested by Heston (1993), Monte Carlo methods can be employed. Poklewski-Koziell (2012) mention that "Monte Carlo methods are used extensively in mathematical finance. They provide a convenient way of simulating stock price distributions and pricing options where closed form solutions are difficult to derive, or do not exist at all. For these reasons, the use of Monte Carlo methods is particularly useful ...". Monte Carlo simulations are ran based on equations 2.7.1 and 2.7.2.

Kloeden and Platen, (1992, chapter 16), explore a robust simulation of stochastic differential equations applied to this context. Their work is of significance due to their derivation of Ito's Lemma which subsequently provided the basis for the Euler-Maruyama and Milstein schemes. Both Gatheral (2006, chapter 2) and Broadie and Kaya (2006) examines the application of Monte Carlo methods to the simulation of stochastic volatility models and the simulation schemes required to generate the variance and stock paths that makes a basis for the pricing of

European options. These schemes allow simulation of stock price processes by sampling from the exact distributions of the stock price and volatility process increments.

### **Euler-Maruyama and Milstein simulation schemes**

As firstly derived by Kloeden and Platen (1992, chapter 10) and secondly by Gatheral (2006, chapter 2), applying a simple Euler scheme to the Heston variance process equation 2.7.2 yields the following:

$$V_{t+1} = V_t + k(\theta - V_t)\Delta_t + \sigma\sqrt{V_t}Z_v\sqrt{\Delta_t}$$
 (2.7.20)

Given that  $Z_v$  is a standard normal random variable and  $\Delta_t$  is a step between moments. The Euler scheme does not guarantee non-negativity of output of volatility. To solve this, rules of absorption and reflection are employed. As shown by Higham and Mao (2005), absorption dictates that when  $V_t < 0$ , it gets set to  $V_t = 0$ . Deelstra and Delbaen (1998) demonstrated that reflection dictates that when  $V_t < 0$ ,  $V_t = -V_t$ . Lord, Koekkoek and Van Dijk (2010) propose this scheme to efficiently reduce the bias:

$$V_{t+1} = V_t + k(\theta - V_t^+)\Delta_t + \sigma \sqrt{V_t^+} Z_v \sqrt{\Delta_t}$$
 (2.7.21)

Given that  $V_t^+ = max(V_t, 0)$ , the asset paths are simulated through the following equation:

$$S_{t+1} = S_t + rS_t \Delta_t + S_t \sqrt{V_t^+} Z_s \sqrt{\Delta_t}$$
 (2.7.22)

Where  $Z_s$  is a standard normal random variable with correlation  $\rho$ , same as in the Heston equations, to  $Z_v$ . These variables can be rewritten into a function of independent standard random variables:

$$Z_{s} = Z_{1} \tag{2.7.23}$$

$$Z_{v} = \rho Z_{s} + \sqrt{1 - \rho^{2}} Z_{2} \tag{2.7.24}$$

The Milstein scheme expands on Euler-Maruyama scheme, adding an extra term from Ito's Lemma. This expansion reduces the number of scenarios when  $V_t$  becomes negative (Gatheral 2006, chapter 2).

$$V_{t+1} = V_t + k(\theta - V_t^+)\Delta_t + \sigma \sqrt{V_t^+} Z_v \sqrt{\Delta_t} + \frac{1}{4}\sigma^2 (Z_v^2 - 1)\Delta_t$$
 (2.7.25)

Now, applying the Milstein scheme to the stock simulation process:

$$S_{t+1} = S_t + rS_t \Delta_t + S_t \sqrt{V_t^+} Z_s \sqrt{\Delta_t} + \frac{1}{2} S_t V_t^+ (Z_s^2 - 1) \Delta_t$$
 (2.7.26)

Applying the Milstein scheme to equation 2.7.1 yields:

$$S_{j} = S_{j-1} + rS_{j-1}\Delta_{t} + S_{j-1}\sqrt{V_{t}^{+}}\sqrt{\Delta_{t}}z_{i} + \frac{1}{2}S_{j-1}V_{t}^{+}\left(\sqrt{\Delta_{t}}z_{i} - \Delta_{t}\right),$$

$$j = 1, ..., n \text{ and } i = 1, ...M$$
(2.7.27)

Where  $z_i$  is a random variable sampled from N(0,1), n is the number of steps and M is the number of sample paths.

#### **Risk Neutral Pricing**

As shown by Crisóstomo (2014), Monte Carlo option pricing assumes risk-neutral valuation. If q is the risk-neutral probability measure, the value of the option at time t can be calculated with the following equation:

$$V(S_t, t) = e^{-r(T-t)} E^q[h(S_t, T)]$$
(2.7.28)

In equation 2.7.28,  $h(S_t, T)$  is the expected payoff of the option, while r is the constant risk-free return on the asset. Under the risk neutral measure, the price of an option at time t is the sum of discounted payoffs. Since options are being priced at t=0, the previous option pricing equation can be simplified:

$$V(S_t, t = 0) = e^{-r(T)} E^q [\max(S(T) - K, 0]$$
 (2.7.29)

Where  $E^q[max\ (S(T)-K,0]]$  is the expected payoff of the call under the risk neutral measure. To expected payoff at maturity is estimated with the arithmetic mean of all the payoffs:

$$\bar{h}(S_T, T) = \frac{1}{M} \sum_{i=1}^{M} h(S_T^{i,n}, T)$$
 (2.7.30)

Where h is the option payoff and  $S_T^{i,n}$ , a simulated stock path of  $S_T$  over the  $i^{th}$  path, using n steps.

Similarly to the Heston-Nandi model, with all the processes explained before, one obtains the European option price for any desired time to expiration as well as a variety of strikes, for both Monte Carlo and FFT generated prices. The Black Scholes framework is then used to reverse engineer the volatility implied in the option price. It takes as inputs the time to maturity, the spot price, the strike, the risk-free rate, the dividend yield and the price of the option.

# **Parameter Estimation**

Lsqnonlin, a MATLAB predefined function, is a local optimization scheme usually used to calibrate the models. This is mainly due to lower computational demand the scheme presents. The alternative methods, Adaptive Simulated Anealing (ASA) and the Genetic Algorithm (GA) are both global optimization schemes. Only GA, out of the three, does not require initial parameters. As shown by Crisóstomo (2014), lsqnonlin relies on initial parameters more intensively than ASA does.

Poklewski-Koziell (2012), suggested that robustness and tradeoff in terms of speed must be assessed. While GA and ASA tend to produce more accurate results, *lsqnonlin* is faster. *Lsqnonlin* considers all the necessary variables in order to rigorously evaluate the option. Option behavior is affected by multiple different *greeks*, which dictate how its value is going to change with variance in time, interest rates, amongst other variables.

As demonstrated by Cristosomo (2014), lsqnonlin intuitively bases its output on reducing the nonlinear square errors. Suppose that N number of options are being evaluated.  $\phi$  represents the list of options from one to N.  $\phi_i$ , represents the market price of  $i^{th}$  option, either a call or a put option, while  $\widehat{\phi}_i$  will be the model price of the  $i^{th}$  option, according to the model parameter set given by  $\theta \in \mathbb{R}^n$ . The sum of the least squares of the errors will then be the squared difference between each option's market value and the estimate. The following equation describes that process:

$$SSD = \sum_{i=1}^{N} \omega_i \left( \Phi i \left( \sigma_i^{BS}, Ti, Ki \right) - \widehat{\Phi} i(\theta i, Ti, Ki) \right)$$
 (2.7.31)

In this equation, all but  $\theta$  represent the inputs that should be manually inputted.  $\omega_i$  is the weight that is assigned to a certain option. While it is possible to have all the weights equally assigned, the weights can be distributed according to options' liquidities. Liquidities are assessed based on the difference between the option's bid and ask prices. The scale to which the difference will matter does not have to be linear. The  $\sigma_i^{BS}$  component stands for the Black-Scholes volatility.  $T_i$  and  $K_i$  are the maturity and strike prices, respectively. The task of the function is finding the model parameters that minimizes the squared differences between market and model prices.

$$SSD(\theta *) = \min \sum_{i=1}^{N} \omega_i \left( \Phi i \left( \sigma_i^{BS}, Ti, Ki \right) - \widehat{\Phi} i (\theta i, Ti, Ki) \right)$$
 (2.7.32)

According to Bauer (2007) since the function is not linear and non-convex, the optimization technique is not as simple as it looks at the first glance. There might be several points where the function minimizes, creating a situation where multiple different outputs would be given. The initial parameter choice matters exactly for this reason.

The parameters estimated allow the computation of the previous methods and therefore the computation of a volatility estimate for a given moneyness and maturity.

# 3. Methodology

#### 3.1 Data Collection

For this project, the stock prices and the option prices from Qualcomm (QCOM), Bayer (BAYN) and SAP (SAP) were used. QCOM was the first underlying to be chosen as a suggestion of the bank, since it was a stock present in a structured product of Banco Invest and appeared to be relatively stable. Later, the bank suggested BAYN and SAP, as BAYN showed to be considerably unstable prior to the period in analysis and SAP showcased a very unexpected and extreme decline in spot prices during the last month of October. Both of these stocks are also part of the bank's portfolio.

The last two years of available data was the overall time range collected and all values were retrieved from Yahoo Finance's historical data. It is important to note that only adjusted closing prices were used in this research, thus eliminating misleading spikes in volatility when dividends are paid and, consequently, stock prices drop. Moreover, different time intervals were used when deriving each model.

First, for the EWMA volatility calculation, only one year of data was used as the model quickly diminishes the weight given to older data, making it obsolete. Second, the GARCH (1, 1) model includes two years of data since it is important to have more historical data to correctly compute the long-run volatility of the underlying asset. Third, for the Heston-Nandi model, only one year of data was used not only due to the computational heaviness of the model but also because smaller emphasis is given to the long-run variance, due to the existence of two additional parameters.

The Heston models, Heston FFT and Heston MCS, require several implied volatilities provided by Bloomberg with different strikes and maturities. The daily implied volatility surfaces of each of the stocks previously mentioned were provided by the bank, which were then used as an input by the Heston model. Finally, the implied volatility surfaces from

Bloomberg were also used to test the regression of the surface that allows the adjustments required for the EWMA and GARCH (1, 1) model.

The previously mentioned data collected is captured by each model to estimate the most appropriate parameters to forecast the volatility. More information provided to the model will theoretically result in more accurate estimations. A daily volatility forecast will be performed for each of the models using the updated parameters and the new market information, specifically the last day's return for the GARCH, EWMA and Heston-Nandi models and the last day's implied volatility surface for the Heston Model. The following daily results will be comprised between the dates of 09/10/2020 and 06/11/2020.

# 3.2 GARCH (1, 1), EWMA and Heston-Nandi Models' Implementation

The GARCH (1, 1) and EWMA models are based on the process described by Hull (2009, chapter 23). The processes consist on maximizing the likelihood function given by equation 2.3.4, in order to optimize the parameters. The parameter estimation will be performed on a daily basis which will then be used to forecast an annualized volatility estimation, with equation 2.3.6 for the GARCH (1, 1) model and equation 2.4.4 for the EWMA.

The Heston-Nandi GARCH model is performed using the excel template provided by Rouah and Vainberg (2007, chapter 6). The closed form version of this model was the one performed, instead of the Monte-Carlo one, to minimize the computational intensity of the model. Finally, the initial variables used for the parameterization of the Heston-Nandi model are, as previously mentioned, suggested in the original paper (Heston and Nandi, 2000). Since the output given by the Heston-Nandi model is an option price, the Black-Sholes equation is used to extract the volatilities.

#### 3.3 Heston FFT and MCS Models' Implementation

Both Heston models require inputs that can only be found by first calibrating the parameters to option prices fed to the model (usually these are the market prices). Once the calibration is performed, all the input parameters required for the Heston Model are achieved.

For the calibration of the parameters, the MATLAB function "Isnonquin", is used to minimize the errors between the market prices and the prices generated by one of the models, in this case, the Fast Fourier Transform prices. This is achieved by solving equations 2.7.4, 2.7.5, 2.7.6 and 2.7.7, generating initial FFT prices from the initial parameters  $\Omega_0 = (k, V_0, \theta, \sigma, \rho) = (2, 0.20, 0.25, 0.25, -1)$ , with upper and lower limits equal to  $u_{limit} = [0, 0, 0, 0, 0, -1]$ , respectively, as suggested by Crisóstomo (2014).

In addition, to avoid the volatility parameter reaching zero or negative values, the Feller condition,  $2k\theta > \sigma^2$ , is also applied to the calibration process, as shown by Crisóstomo (2014). These prices are then compared to the market prices given by the bank and equation 2.7.5 is minimized to obtain the optimal parameters  $\Omega_{optimal} = (k, V_0, \theta, \sigma, \rho)$ . These procedures were executed for the same time to maturity (T) and moneyness (M) pairs that Banco Invest provided as market options data. As this data was provided daily, the calibration procedure was also performed daily during the period of 09/10/2020 to 06/11/2020.

Once calibrated, both the Heston FFT and MCS prices can be computed. Heston FFT prices were computed by solving the same equations as before but using the optimized parameters  $\Omega_{optimal} = (k, V_0, \theta, \sigma, \rho)$  instead of the initial ones  $\Omega_0 = (k, V_0, \theta, \sigma, \rho) = (2, 0.20, 0.25, 0.25, -1)$ . Heston FFT prices were computed making use of same pricing code<sup>1</sup> used in the calibration procedure.

<sup>&</sup>lt;sup>1</sup> This code was written by Jonathan Frei, who based his work on various papers by Khal, Jäckel and Lord. This code can be found at <a href="https://github.com/jcfrei/Heston">https://github.com/jcfrei/Heston</a>

Heston MCS prices were generated by applying the Euler-Maruyama process to equations 2.7.1 and 2.7.2, generating the stock path simulating equation 2.7.26 and the variance path simulating equation 2.7.25, respectively. In addition, this method also makes use of two additional parameters, the number of time steps in the discretization n and the number of paths simulated M. (M,n) = (10000,250). The code works as follows, for each Monte Carlo simulation, two uniform and correlated random vectors are generated. The correlation is the estimated  $\rho$ . These vectors are standardized. Then, for each step of the discretization of the volatility path in equation 2.7.25, the next volatility is computed from the normal standard vector.

Afterwards, for each time-step in the discretization of the asset path in equation 2.7.26, the following asset value is computed using the previous asset value and the volatility value computed. For each of the 10000 Monte Carlo simulations, the payoffs of the options for all (M, n) pairs are stored and averaged and discounted using the risk-free rate, which was assumed to be r = 0% for both models. This algorithm produces the Monte Carlo option prices, under risk-neutral pricing. Heston MCS prices were also computed making use of MATLAB code<sup>2</sup>. To extract the Heston volatility from the prices, the Black-Scholes model is used.

# 3.4 Methodology of Empirical Analysis

The baseline of the empirical analysis will attempt to understand which model is better at predicting volatility. First, the estimations of each model will be compared against the implied volatility as this volatility is an appropriate benchmark for other volatility forecasts, not only due to its popularity, but also because it attempts to mirror market expectations regarding future volatility. The results derived from the models will also be benchmarked against the average of the fifteen and thirty days realized volatility of the underlying assets (Edwards and Lazzara

<sup>&</sup>lt;sup>2</sup> This code was written by Suhyun Kim. It can be found at <a href="https://www.mathworks.com/matlabcentral/fileexchange/34244-heston-simulation-using-monte-carlo">https://www.mathworks.com/matlabcentral/fileexchange/34244-heston-simulation-using-monte-carlo</a>

2016), which were computed with the formula proposed by Garman and Klass in 1980 and discussed by Brous, Ince and Popova (2010):

$$\delta_{GK}^{2} = 252 \left[ \ln \left( \frac{O_{n}}{C_{n-1}} \right)^{2} + \frac{1}{2} \ln \left( \frac{H_{n}}{L_{n}} \right)^{2} - (2\ln 2 - 1) \left( \ln \left( \frac{C_{n}}{O_{n}} \right) \right)^{2} \right]$$
(3.2.1)

 $H_n$  represents the high price of the underlying at day n,  $L_n$  is the low price of the same stock for the same day n, while  $O_n$  and  $C_n$  stand for opening and adjusted closing prices of the asset on day n.  $C_{n-1}$  denotes the adjusted closing price at the day n-l. This formula was suggested by Banco Invest, as it not only includes the daily standard deviation, with the closing price on n-l and the closing price of n, but also the volatility during the day captured by the  $H_n$  and the  $L_n$ .

Later, the stability of the models will be assessed, relative to other models and the Bloomberg implied volatilities. Standard deviation of the results will be used as a marker to assess stability. In order to minimize transaction costs, the model should exhibit a higher degree of stability than its competition. Although it is important for the models to capture volatility movements, unnecessary shocks will be deemed as a shortcoming of the volatility forecast. Unnecessary shocks are considered to be ones that result in sharp rises or declines of the volatility forecast, followed by an inverse movement of the estimation in the following days. The stability of the volatility estimations given by the models was a clear point of emphasis by Banco Invest throughout this project, which is measured by the standard deviation of the daily volatility estimations during the analyzed period.

To understand the comparisons between the several models and the realized volatility, the evaluation will be performed by comparing the volatility of options with one year to maturity and moneyness of 100%. Another point of the surface will be studied for the Heston models, the point (T, K) = (0.0833, 100%), due to the hypothesis that this model is less efficient when forecasting a volatility with shorter maturity (Wu 2019).

This analysis will mostly only be performed for the point (T, K) = (1,100%) because this is the original result for the GARCH (I, I) and the EWMA models. Since the volatility surface of this models is based on the slope of the implied volatility, the quality of results from another point of the volatility surface depends on the quality of the forecast for the volatility of (T, K) = (1, 100%). Therefore, the analysis of other points of the surface would yield similar results and conclusions for this model.

# 4. Empirical Analysis

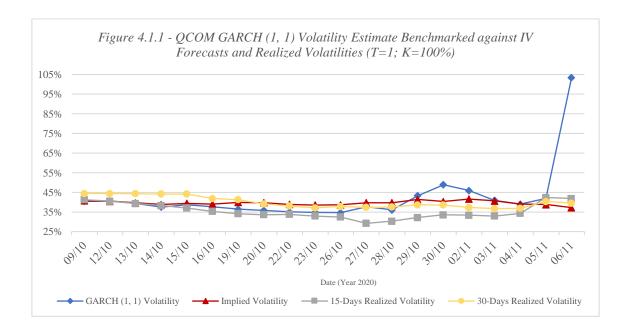
# 4.1 GARCH (1, 1) Model

When analyzing the quality of the GARCH (1, 1) volatility model compared to others, the theoretical robustness of this model did not assure stable and predictable results throughout the period in analysis. This is visible in the results of the QCOM stock, as this model seems to present results that are not aligned with the average realized volatility in several periods of the analysis.

In a first instance, the analysis will be performed on the proximity of GARCH results to the realized volatility and the implied volatility for the QCOM asset. The following figure provides a comparison between GARCH (1, 1) annual volatility estimates, IV and realized volatilities between 09/10/2020 and 06/11/2020. To estimate the GARCH (1, 1) model equation 2.3.6 is solved everyday between 09/10/2020 and 06/11/2020, with the parameters being optimized daily.

Moreover, each parameter optimization is done having into account two years of past QCOM stock prices, which enables the process to base its forecasts on past information. Regarding realized volatilities, it is important to note that equation 3.2.1 is solved fifteen and thirty times each day, for the fifteen and thirty days realized volatility, respectively, and the values presented in the figure are an average of those results. The implied volatility is the metric currently used

by Banco Invest and is directly extracted from Bloomberg. The analysis is performed for an atthe-money option, with one year to maturity.



As it is possible to see from the graph of annualized volatilities, despite displaying a similar result to the average realized volatility at some dates, it is clear that the volatility that results from the GARCH (1, 1) model are much more unstable than the other models. The shock of the 06<sup>th</sup> of November predicts an annual volatility of 103,42% due to an increase of the spot price from 04/11/2020 to 05/11/2020 of 128,97 to 145,41, a daily return of 12,75%. The model, with the current optimal parameters, is not able to correctly predict a stable annual volatility, which can be a problem regarding the hedging activity due to constant unnecessary transaction costs. It is important to mention that at the day 11/11/2020 the volatility forecast had already decreased to 52,09%, illustrating that this spike in volatility was unnecessary.

Looking at the average results throughout the period of the data of QCOM, present in table 4.1.1, it can be seen that both the implied volatility and the GARCH (1, 1) volatility are much closer to the average of thirty days of daily realized volatility than to the average of fifteen days. The numbers in table 4.1.1 represent the average of the results illustrated in figure 4.1.1. In.

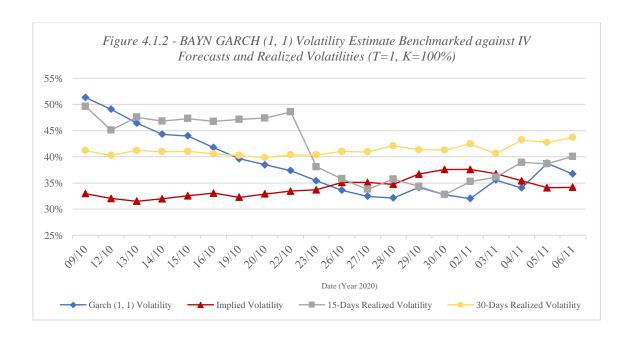
other words, an average of the daily results between 09/10/2020 and 06/1/2020 for the four metrics present in figure 4.1.1 was performed, and the values are displayed below.

Table 4.1.1 – QCOM average annualized volatility metrics during the 09/10/2020 – 06/11/2020 period

Volatility Forecasts (T= 1, K= 100%)		Realized Volatilitie	s (T= 1, K= 100%)
GARCH (1, 1) Implied Volatility		15-Days	30-Days
42,40% 39,58%		35,42%	39,94%

The results show that both the implied volatility and the GARCH (I, I) model, on average, forecast similar results to the thirty days' average realized volatility, slightly better than the fifteen days' average realized volatility. However, the quality the results between the implied volatility model and the GARCH (I, I) model are not similar if the standard deviation of the daily volatility forecasts of each model are taken into consideration. As it was mentioned before, it is important that the models do not showcase unnecessary volatility spikes due to the transaction costs associated for the bank. A lower standard deviation of the daily volatility forecast is a measure suggested by the bank to understand which models are the most stable. While the implied volatility was able to predict this result with a standard deviation of 1,06, the standard deviation for the same period of the GARCH (I, I) model results was 14,84.

The difference in the standard deviation between the two shocks seem to entail that, while the implied volatility predicts the market trajectory for the future, the GARCH (1, 1) model absorbs all the shocks from the market and reacts to them, therefore following the market trend, resulting in occasional unpredictable spikes in the volatility forecast. When analyzing the results from the BAYN stock, the model appears to better predict the market trend, which can be observed in the following graph:



While GARCH forecasts appear to tend to the levels of the average realized volatility, the implied volatility once again appears to predict much sooner how the realized volatility is going to behave as it is visible in the figure 4.1.2. On the opposite hand, GARCH (1, 1) appears to react constantly to the realized volatility, always following a similar trend to the fifteen-day average with a slight delay.

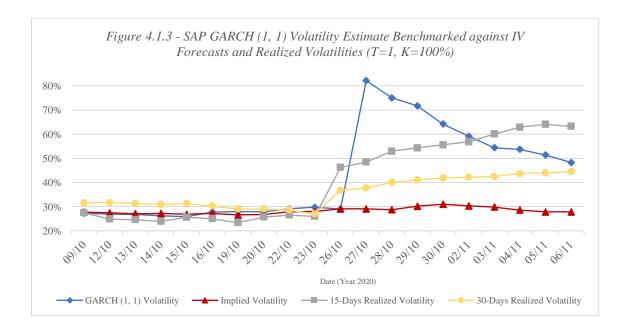
Table 4.1.2 - BAYN average annualized volatility metrics during the 09/10/2020 - 06/11/2020 period

Volatility Forecasts (T= 1, K= 100%)		Realized Volatilitie	es (T= 1, K= 100%)
GARCH (1, 1) Implied Volatility		15-Days	30-Days
38,49% 34,17%		41,34%	41,13%

The average results for the GARCH (1, 1) model are much closer to the realized volatility when compared to the implied volatility. While the implied volatility throughout the data expects the volatility to be stable, the GARCH model is reacting to the market fluctuations providing a more precise daily result. However, this daily precision of the GARCH (1, 1) model implies an increased instability of forecasts, with the standard deviation of the volatility results for GARCH being 5,86, while the implied volatility's standard deviation is 1,89.

Finally, regarding SAP, because this stock had such a shock due to profit forecasts resulting in a negative daily return of 21,94% on the 26/10/2020, this impact highlights the differences

between the implied volatility and the GARCH (1, 1) results. As a result of that unexpected drop, while the implied volatility model assumed that this negative return would not have a lasting impact in the yearly volatility and therefore had a very mild reaction, the GARCH model totally absorbed that shock resulting in an annualized volatility for that day of 82,14%, with a smooth decrease on the following days closer to the long-run volatility estimated by the model. These results can be observed in the figure 4.1.3:



Although the implied volatility model may underestimate the future impact of these shocks, it assumes that it will not have a preponderant impact in the yearly volatility. On the other hand, the GARCH (1, 1) estimates volatility spikes followed by a regression back to the long-run volatility estimated by the model, a characteristic that may present an obstacle on using this model to hedge the position on an option.

## GARCH (1, 1) parameters' analysis

Regarding the parameters of the GARCH (1, 1) model, emphasis will mostly be given to the QCOM parameters as well as the comparison to the BAYN ones, in order to understand how different parameterization results in different volatility forecast behavior. These two stocks

were chosen to be compared as they not only present very different parameters but also because the GARCH (1, 1) volatility estimations for the SAP stock provides results that are consistently overestimated, influencing the quality of the parameters. This analysis will be performed each parameter at a time.

The first parameter that is going to be analyzed is the  $\beta$ , which is responsible to analyze the impact of yesterday's volatility in the following day. For QCOM, the average value for this parameter throughout the length of the data is 0,512, a surprisingly low value when compared to the same parameter for the BAYN stock. This result provides a very important insight on why the forecast of the volatility for this asset is so inconsistent. A low value of this parameter results in a very high decay rate, meaning the past volatilities will have very low preponderance in today's prediction of volatility. This will therefore result in the very low weight given to n-2, n-3 volatilities, which would not occur if the  $\beta$  was higher, as there would be a higher stability given by the historical results. Comparing the QCOM result to BAYN, it is possible to observe that the optimal  $\beta$  of the latter is much higher, being 0,829. The differences in decay of the past volatility's impact are reflected in the following table:

Table 4.1.3 – Decay rate of past volatility's impact on tomorrow's forecast for the OCOM and BAYN assets

Day \ Stock	BAYN	QCOM
n-1	0,829273	0,512228
n-2	0,687694	0,262377
n-3	0,570286	0,134397
n-4	0,472923	0,068841
n-5	0,392183	0,035263

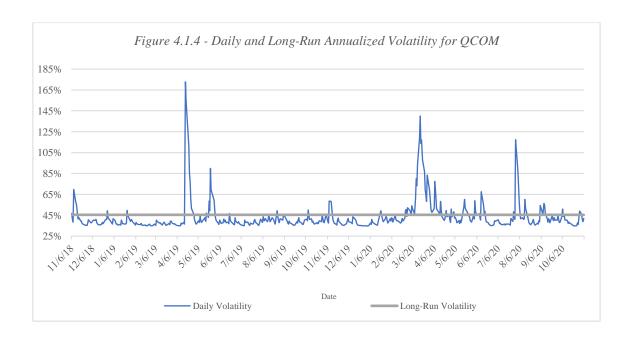
While the difference in the weights of n-1 is not abysmal, the exponential characteristic of the GARCH (1, 1) model results in a much steeper impact of past values when the initial value is lower. Therefore, two important results can be inferred from this table. First, while the n-5 impact in QCOM is almost insignificant, the BAYN n day volatility is still deeply affected by the day n-5 volatility. Second, while the weight of past volatilities are not directly attributed to

todays' volatility, the impact is embedded in the *n-1* data, which will therefore provide a much more stable result. This factor combined with much more stable returns has resulted in a smaller standard deviation on the volatility of BAYN, as the standard deviation for QCOM volatility forecast was 14,83 while for BAYN was 5,86.

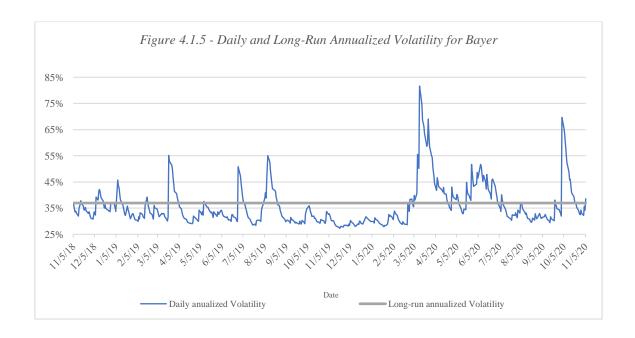
Regarding the parameter  $\alpha$ , which is responsible for the weight of the last squared returns in the daily volatility, it can also be explanatory for the differences in the way the two volatilities behave. In order to be in accordance with the equation given by  $\alpha + \beta + \gamma = 1$ , it is expected for QCOM to have a much higher  $\alpha$  parameter than BAYN, which indeed happens. The value of this parameter is 0,201 for the first, while it is 0,083 for the latter. This difference implies that the QCOM volatility is approximately 142% more affected by the squared returns of the asset compared to the impact in BAYN's volatility. The nature of this result can also explain why QCOM volatility is more prone to shocks resulting from unpredictable large results. While the returns are small, both volatilities are expected to remain stable. On the contrary, when the returns are large, QCOM is expected to have a shock in volatility, only caused by the returns, 142% larger than the BAYN shock.

Finally, regarding the  $\gamma$ , the parameter which is responsible to attribute a weight to the long-run volatility computed by the model, the result was also much larger in QCOM than BAYN, with the average of the first being 0,287 while the latter was 0,087. The parameter w is one which is optimally computed and is dependent not only on the  $\gamma$  value, but also on the VL, as  $w = \gamma \times VL$ . The w of QCOM is almost 400% larger than BAYN, as QCOM has not only a considerably higher  $\gamma$  but also has a larger Long-run volatility.

Comparing the two long-run volatilities when annualized, QCOM has an annual VL of 45,38% while the one from BAYN is 36,94%. To better understand how the VL related to the underlying asset volatility, the following graph shows how the QCOM daily annualized volatilities compares with the long-run volatility:

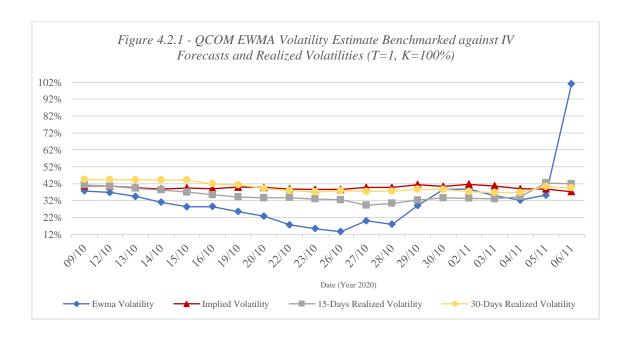


While the volatility predicted of the model is often around the long-run volatility predicted by the model, it is also visible that the model frequently forecasts spikes in volatility that are not really in accordance with the market, as the model quickly regresses to the long-run volatility that is expected. These spikes are results of unpredictable large returns, either positive or negative. These shocks, as well as low stability, might imply an increased difficulty in using this model in a hedging strategy. Moreover, the very constant result of the long-run volatility combined with the fact that the result is very close to the historical volatility of this asset, which is 43,42%, might also imply that this result is also not ideal to hedge, as it would not react immediately to market changes. Comparing to a similar graph regarding the BAYN stock, the following result shows that this asset has much less shocks, as well as less intense, than the QCOM one. Therefore, and as a result of the different parameterization between the two stocks, BAYN shows a much closer daily volatility to the long-run volatility, as it can be seen by the following graph:



# 4.2 Exponential Weighted Moving Average Model

The EWMA analysis follows a suggestion from Banco Invest to understand if the simplicity of the model would derive better volatility estimations. However, similar to the GARCH  $(I,\ I)$  process, the EWMA did not provide neither stable nor consistent volatility forecasts across the period under study. This statement is valid for the three stocks under analysis, but it becomes more relevant when addressing QCOM. The following graph entails not only how far the predictions of the model are when compared with the realized volatility of the underlying, but also how much more reliable and stable the implied volatility metric for this comparison. To estimate the EWMA model, equation 2.4.4 is solved everyday between 09/10/2020 and 06/11/2020, with the parameter being optimized daily. Moreover, each parameter optimization is done having into account one year of past QCOM stock prices, which enables the process to base its forecasts on past information.



It is crucial to analyze two-time frames present in the graph. First, it is important to understand the decrease in volatility expectations verified between 13/10 and 26/10. During these thirteen days the annual volatility forecast for QCOM decreased from 34,36% to 13,52%, a drop that was not verified in any other volatility metric. This situation was verified because, during this period, the average daily return of the stock was extremely low, 0,11%, which transmitted the incorrect idea to the model that the underlying was excessively stable. With an average optimal lambda of 0,792, the EWMA distributes much more weight to the last returns of the stock. With this, a period as the one verified between 13/10 and 26/10 exposes a great shortcoming of the process. In addition, after this stable period of returns, the spot prices of this asset became increasingly unstable, which immediately led the EWMA model to forecast much higher volatilities.

The second relevant time frame to be studied is related with the jump verified on 6/11. From the 4<sup>th</sup> to the 5<sup>th</sup> of November, QCOM stock increased its value by 12,75%, which led the model to predict the annual volatility of the underlying for the day after to be 101,05%, an impractical result for the bank's activities. Again, the value attributed to the optimal parameter is responsible for such unreliable forecast, combined with the large return and the nature of the

model. It is crucial to mention that, similarly to the GARCH (1, 1) model forecast, the estimated volatility immediately decreases in the next observations due to a stability of the returns after this date. This confirms that this spike in the volatility estimation was unnecessary, therefore highlighting a crucial shortcoming of the model.

As previously stated, the EWMA provides estimations that do not correspond with the volatilities verified on the market during the period under study. On the contrary, the implied volatility appears to be an extremely fair estimate when benchmarked against the average realized volatility of QCOM, regarding 30 days of returns. The following table contains all average values for the four volatility measures presented on figure 4.2.1 and confirms the previous statements.

Table 4.2.1 – QCOM average annualized volatility metrics during the 09/10/2020 – 06/11/2020 period

Volatility Forecasts (T= 1, K= 100%)		Realized Volatilitie	s (T= 1, K= 100%)
EWMA Implied Volatility		15-Days	30-Days
31,92%	39,58%	35,42%	39,94%

In addition, by analyzing figure 4.2.1, it becomes clear that the EWMA produces much more unstable results when compared to the implied volatility metric. In fact, the standard deviation of results of the EWMA model is 18,17, a value 17,21 times larger than the one derived from IV forecasts. This enormous difference in the standard deviation between the EWMA and the IV forecasts seem to entail that, similarly to the GARCH (1, 1), while the implied volatility predicts the market trajectory for the future, the EWMA approach simply reacts to the most recent market conditions.

Regarding the BAYN stock, similarly to the GARCH, the model appears to better predict the market trend, as it can be seen in the graph below.

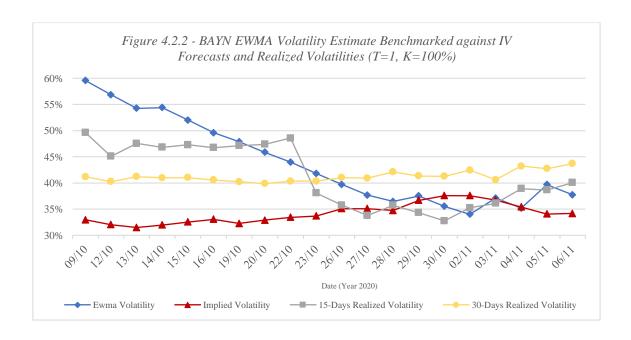


Figure 4.2.2 demonstrates that the EWMA estimation follows the trend of the fifteen days realized volatility and must be considered a good proxy for volatility from 19/10 onwards. In fact, the average volatility metrics for the period under study presented in the bellow table show that the EWMA was the model that better estimated volatility when compared with the implied volatility. In this case the implied volatility deviates about 7% from the average verified volatility of BAYN, while the model's estimation is less than 3% inaccurate.

Table 4.2.2 – BAYN average annualized volatility metrics during the 09/10/2020 – 06/11/2020 period

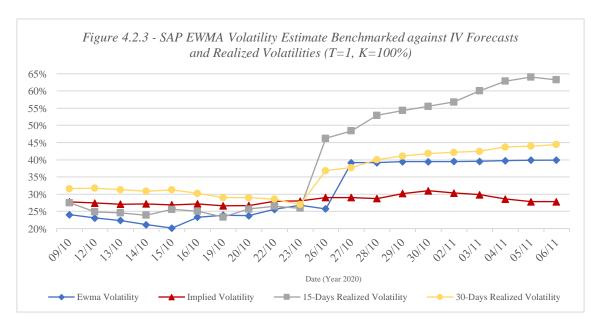
Volatility Forecasts (T= 1, K= 100%)		Realized Volatilitie	es (T= 1, K= 100%)
EWMA Implied Volatility		15-Days	30-Days
43,83%	43,83% 34,17%		41,13%

Such aligned results were only possible due to two main factors: first, Bayer's average optimal lambda for this period was 0,905, meaning that it was given less weight to the last observations, and consequently the forecast was able to incorporate more data. Second, unlike QCOM, this stock did not have periods of very unstable returns, being close to zero in some instances and being very high in others, which led to smoother volatility estimation.

Nevertheless, it is relevant to state that between 09/10 and 28/10 the annual volatility forecast for Bayer dropped from 59,56% to 36,46%, still evidencing the instability of results provided

by the EWMA approach. The standard deviation of results is once again much larger when compared with the value resulting from the implied volatility, being 8,11 for the first and 1,89 for the latter. This proves that, despite the fact of having a higher lambda than QCOM, the model is consistently much more sensitive to market oscillations when compared with the implied volatility metric. Such characteristic constitutes a drawback for successfully reducing unnecessary trading and consequently transaction costs on the hedging process of the bank.

The volatility estimation for the SAP stock by the EWMA method cannot be considered appropriate and exposes another shortcoming of the process. Although the following figure shows a close relationship amongst the forecast and the fifteen days average realized volatility for the period between 09/10 and 23/10, the values computed from 27/10 onwards are not reliable. As it was previously discussed, this underlying suffered a 21,94% spot price drop in value on the 26<sup>th</sup> of October, an unexpected price jump that the model was not able to incorporate in its forecasts. Given this, until the end of the period under study, the EWMA assumes a lambda of one, meaning that there is no estimation being performed, as the model simply assumes the first value of variance present in the data sample to be the daily variance for the entire data series. Figure 4.2.3 confirms these statements. Given these results, no other analysis regarding the EWMA approach will be performed on this security.



## 4.3 Exponential Weighted Moving Average Model - Fixed Lambda of 0,94

Due to the verified instability of the EWMA results, the following section studies how volatility forecasts behave if the parameter  $\lambda$  is fixed at 0,94, as suggested by JPMorgan. Considering this  $\lambda$  is larger than the optimal parameter for each the three underlying assets, this implies that the daily volatility forecast will account for more past days' volatilities. Similar to the GARCH process, as  $\lambda$  becomes larger, the decay rate of the past estimations becomes slower, often resulting in more stable results as the data in consideration increases. By following this approach, more past data will be incorporated in the estimation, which can lead to smoother results. In fact, the Figure 4.3.1 shows that this method provides not only much more stable results but also derives forecasts much more in line with the realized volatility of QCOM stock. To estimate the EWMA model with a fixed  $\lambda$  of 0,94 equation 2.4.4 is solved everyday between 09/10/2020 and 06/11/2020, but the parameter instead of being optimized is fixed at 0,94 during the entire period.

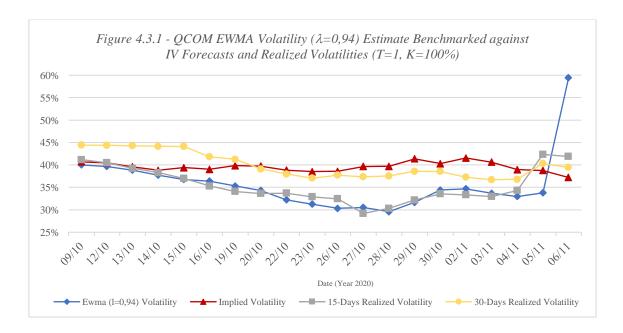
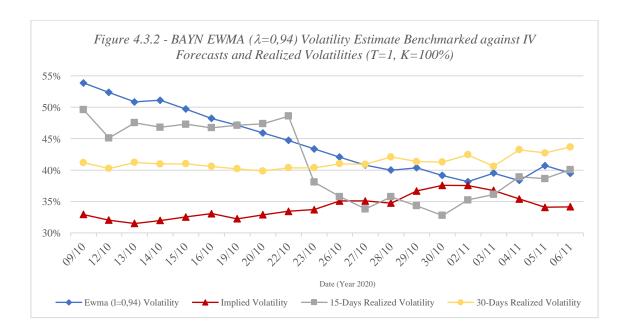


Figure 4.3.1 illustrates, on one hand, how precisely the model follows the pattern of the average fifteen days realized volatility of QCOM. On the other hand, it is clear that this EWMA approach is yet not able to provide smooth reactions to unexpected price jumps. Once again,

the model failed to derive consistent volatility forecasts when the 12,75% return occurred for the 6<sup>th</sup> of November, 1 day after the abnormal return. The model predicted an annual volatility of 59,41%. However, it is relevant to highlight the improvement of results when compared with the previous EWMA process. For the same day, 6<sup>th</sup> November, the optimized lambda value predicted an annual volatility of 101,05%. The stability of results is also improved by this method, as the standard deviation of forecasts drop to 6,38, a value 2,85 times lower than the one previously derived. However, such improvement was not sufficient to make this model as stable as the implied volatility metric. Similar to the previous models, the volatility estimation after the 6<sup>th</sup> of November exhibits a significant decrease.

Regarding the BAYN security, this approach was not responsible for relevant changes in the previous forecast, since the average optimal lambda previously computed is 0,905, very similar to 0,94. The following figure illustrates this situation, as it is clear that the EWMA for  $\lambda = 0,94$  line is very identical to the one presented in figure 4.2.2.

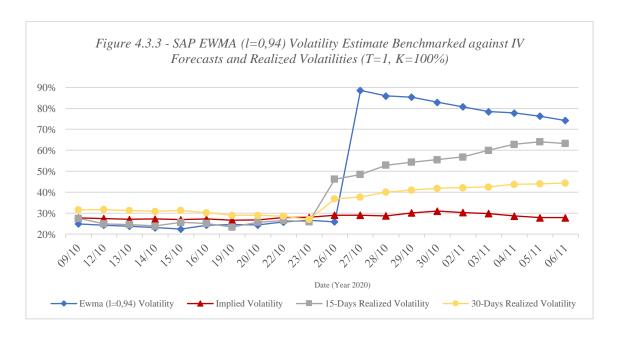


The average predicted volatility of this method, for the period under study, was 44,29%. This value is about 3% higher than the average realized volatility in both the fifteen and thirty days of results and, at the same time, is 0,45% higher than the average forecasted volatility by

the standard EWMA process. The standard deviation of results is 5,19, which represents an improvement when compared with the 8,11 value previously calculated.

Fixing the lambda at 0,94 did not mitigate the problems related with SAP stock. Despite the fact that there is no longer the case where lambda equals one after the shock verified on 26/10, the model currently goes from an annual forecast of 29,86% on the 26<sup>th</sup> October, to 88,60% after incorporating the unexcepted negative return. In the days before the model continues to follow the fifteen days average realized volatility, as it happened with the optimal lambda approach. This was only possible since, during the period before the abnormal result, the optimal lambda of SAP averaged 0,915, a value similar to 0,94.

However, the model's reaction to price jumps makes it inappropriate to use in hedging situations. As it can be seen on figure 4.3.3, after the spike in forecasts, on the 27<sup>th</sup> of October, the volatility estimation dropped to 74,23% on the 6<sup>th</sup> of November, which represents a 14% decrease in only seven business days. All in all, this process has a 29,10 standard deviation of results, an extremely high value when compared with the other studied metrics, as the implied volatility has a standard deviation of 1,30.



## 4.4 Volatility Surface Regression

When empirically analyzing the results provided by the regression of the implied volatility surface, the process of replicating this surface as a function of moneyness and maturity was performed every day of the data in order to correctly reproduce the surface daily. The results that are going to be studied will be from the volatility surface regression of the day 23/10/2020, a day where the forecast of volatility for the GARCH model is aligned with the realized volatility and the implied volatility.

First, as it was previously mentioned, the deep in the money and deep out of the money options' volatilities were disregarded to better reproduce a similar surface to the implied volatility one. The unpredictability of this set of volatilities have a significant impact in the quality of the model's output. With this, the difference in the error is considerable, as the inclusion of these values would increase the squared difference of the two volatility surfaces, the implied volatility one and the estimated one, from 12,56 to 42,38, an increase of the error term of 237% with only the increase of 14 observations to the surface regression. In addition to the initial 49 observations, this represents an increase of 28,6% on the number of observations, which leads to the conclusion that the error term increase is not proportional to the increment in observations.

Moreover, while the equation 2.5.1 previously mentioned assumed only five parameters in order to estimate the regression, the inclusion of more parameters more effectively estimated the surface which is empirically visible in the sum of the minimized differences between the real and estimated volatility surfaces. Therefore, while with the current seventeen parameters used the sum of the squared error is 12,56, the use of only five parameters consistently failed to achieve similar results as the regression with a lower number of parameters would not be able to adjust to specific local spikes of volatility for a given T and K. While using the seventeen parameters in the regression does not assure an absolute minimum error term, this

parameterization consistently achieves better error minimization. The inclusion of parameters that involved skewness and kurtosis, as well as more relations between the variables K and T allowed the model to achieve much closer results to the actual implied volatility surface. Following the process of the equation 2.5.1 suggested by Dumas, Fleming and Whaley (1998), , the squared error of each observation that resulted from this regression was:

*Table 4.4.1 – QCOM squared errors from volatility surface regression on 23/10/2020* 

T \ K	90%	95%	97,5%	100%	102,5%	105%	110%
0,083	0,11099	0,241416	0,27241	0,175684	0,036427	2,164327	0,153276
0,167	0,025316	0,170132	0,221761	0,188368	0,169007	0,161948	0,558036
0,25	0,456462	0,45831	0,284258	0,266992	0,24798	0,225125	0,093073
0,5	0,165145	0,120411	0,018502	0,79283	1,27799	0,627983	0,149814
1	0,187369	0,002637	0,031279	0	0,041699	0,078864	0,564677
1,5	0,094981	0,025937	0,346874	0,017773	0,023352	0,060302	0,114135
2	0,036593	0,032597	0,298814	0,001966	0,137571	0,471345	0,157746

Considering this, as the model is able to estimate how the volatility behaves as a function of strike and maturity, the model can now fit the outcomes of models such as the GARCH (1, 1) and the EWMA as a function of K and T, allowing these models to predict volatility for options with different maturity and moneyness.

By inputting one of this model's volatility result in the point (T, K) = (1, 100%), the model estimates a volatility for each position of the surface with a similar function to the implied volatility. By analyzing the result for an arbitrary position such as the (T, K) = (1.5, 105%), the outcome is:

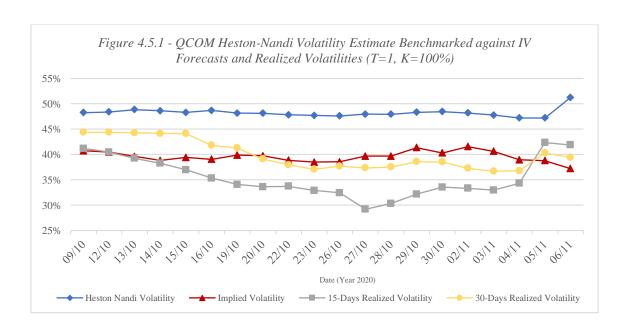
Table 4.4.2 – QCOM volatility estimates on 23/10/2020 for two different surface points

T/K	(T, K) = (1, 100%)	(T, K) = (1.5, 105%)	Differential
IV	38,51%	37,32%	3,18%
GARCH	34,79%	33,85%	2,78%
EWMA	15,32%	14,38%	6,56%
EWMA (λ= 0,94)	31,27%	30,33%	3,11%

As it is visible by the results, the model behaves very similarly for all different processes. However, the outcome of this model is deeply affected by the initial volatility input on the position (T, K) = (1, 100%). As it is evident by the standard EWMA results, if the initial input is flawed, the regression output will derive a similarly incorrect result. Therefore, the volatility regression model output is greatly influenced by the initial volatility estimation previously performed. It is visible by the differential column, meaning the relative difference from one point to the other, that the model assumes a similar slope to each volatility estimation model. The daily regression aims to constantly update the function of maturity and moneyness given by the implied volatility surface, in order to constantly provide the most accurate estimation of volatility for any given point of the surface.

#### 4.5 Heston-Nandi GARCH Model

When analyzing the QCOM stock volatility with the Heston-Nandi model, it is clear that this method's forecast differs considerably to the ones previously studied. The results, as it can be observed on the graph, appears to be stable, when compared with both the implied volatility and the realized volatilities. To estimate the Heston-Nandi model equation 2.6.4 is solved everyday between 09/10/2020 and 06/11/2020, with the parameters being optimized daily. Moreover, each parameter optimization is done taking into account one year of past QCOM stock prices, which enables the process to base its forecasts on past information. However, equation 2.6.4 outputs an option price, meaning that is necessary to use the Black-Sholes formula in order to extract the volatility estimates illustrated in figure 4.5.1.



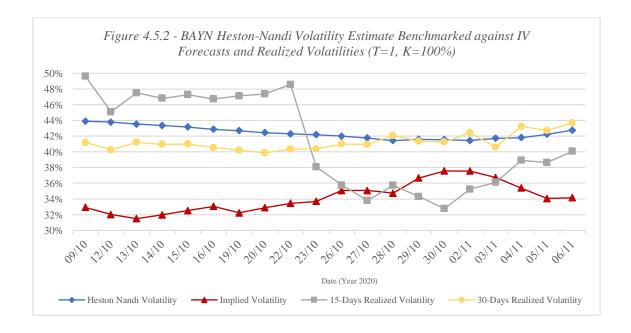
The Heston-Nandi model presented a standard deviation of the daily volatility forecasts of 0,83, a more stable result when compared to the implied volatility standard deviation of 1,06. This follows one of the major objectives of the bank and of this project, which is having a more stable model. However, the Heston-Nandi model is estimating a much higher volatility compared to the others. When the average volatility is compared between the models, Heston-Nandi stands out by at least eight percentage points, as it can be observed from the following table.

Table 4.5.1 – QCOM average annualized volatility metrics during the 09/10/2020 – 06/11/2020 period

Volatility Forecasts (T= 1, K= 100%)		Realized Volatilitie	es (T= 1, K= 100%)
Heston-Nandi Implied Volatility		15-Days	30-Days
48,25%	48,25% 39,58%		39,94%

Moreover, on one of the last days of the observation, 5th of November of 2020, the stock jumps from 128,97\$ to 145,41\$, a positive change of 12,75%. While the Heston-Nandi reacts to this jump on the next day, as the previous models also did, it is important to emphasize that the jump presented in this model is much milder than the shocks observed in the GARCH (1, 1) and EWMA volatility forecasts.

When analyzing the results from the BAYN stock, the model appears to do an appropriate prediction of the market while keeping the stability desired, which can be observed in the figure 4.5.2:



The Heston-Nandi model appears to be in line with the thirty days average realized volatility, despite showing less sensitivity to the stock returns' movement, since the thirty days average realized volatility moves a lot with the new returns. This impact is shown in the standard deviation, where the Heston-Nandi Model provides a value of 0,80 while the thirty days average realized volatility is 1,05.

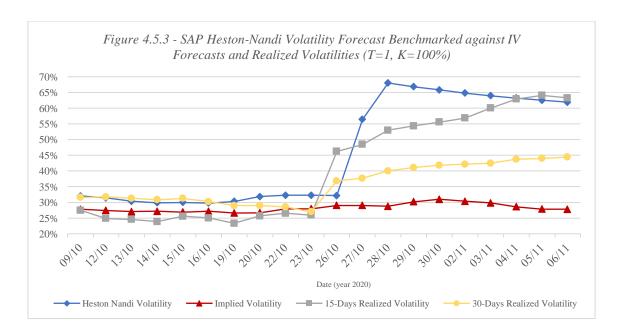
The implied volatility given by Bloomberg provides much lower volatility estimations when compared with both the Heston-Nandi and the thirty days average realized volatility, while the fifteen days average realized volatility is very unstable. When taking into consideration the average volatilities, observable in the table 4.5.2, the Heston-Nandi appears to, again, be very similar to the fifteen and thirty days average realized volatilities, while the implied volatility is about seven percentage points lower than the others.

Table 4.5.2 – BAYN average annualized volatility metrics during the 09/10/2020 – 06/11/2020 period

Volatility Forecasts (T= 1, K= 100%)		Realized Volatilities (T= 1, K= 100%)		
Heston-Nandi Implied Volatility		15-Days	30-Days	
42,42% 34,17%		41,34%	41,13%	

While the implied volatility expects the market to have a lower volatility in the future, the Heston-Nandi rejects that statement in parallel with the average annualized realized volatilities.

Finally, when analyzing the results given by SAP, which had a shock on the 26<sup>th</sup> of October that was previously mentioned, it can be observed a considerable spike in the Heston-Nandi volatility estimation, with the model changing its yearly forecast from 32% to 68% in only two days. Before this date, the model appears to be very stable and in line with the realized volatility, even though both being roughly 3% higher than the implied volatility.



When the price drop occurs, the implied volatility almost does not react to it, with a change of only 1%, staying in line with previous expectations. The Heston-Nandi model showcases a considerable reaction in the first days of the shock, slowly declining afterwards.

## Heston-Nandi parameters' analysis

The Heston-Nandi model is a derivation of the GARCH (1, 1) model for option pricing, which uses five parameters instead of three used by the GARCH. To understand the behaviour of these parameters, an analysis using QCOM and BAYN will be performed for each of the parameters.

First, considering Omega ( $\omega$ ), which is the weight assigned to the long-term volatility multiplied by the long-term volatility ( $\omega = x * \mathcal{V}_l$ ), it can be observed from the tables 4.5.3 and 4.5.4 that both stocks give very low weights to the long-term volatility, meaning their impact is minimal to the volatility forecast.

On the other hand, Beta ( $\beta$ ), which gives the weight assigned to the last volatility observed, is one of the most important parameters, with an average of 0,77 for QCOM and 0,91 for BAYN. Therefore, the impact of past volatilities is going to be more present for the BAYN forecast than the QCOM one. Moreover, as it is explained before in the GARCH (I, I) parameter analysis, the decay rate of the QCOM volatility forecast is going to be larger than the BAYN one, which results in less past volatilities being relevant for the volatility estimation.

Table 4.5.3 – Heston-Nandi average parameters' values during the 09/10/2020 – 06/11/2020 period for OCOM stock

	ω	α	β	γ	λ
Average Value	5,38 E-17	0,00021	0,77	3,57 E-10	3,57 E-12

Table 4.5.4 – Heston-Nandi average parameters' values during the 09/10/2020 – 06/11/2020 period for BAYN stock

	ω	α	β	γ	λ
Average Value	2,15 E-21	6,35 E-05	0,91	1,98 E-18	7,21 E-22

Regarding the parameter  $\alpha$ , which is responsible for the weight of the last squared returns in the daily volatility. The value of this parameter is 0,00021 for QCOM, while it is 0,000065 for BAYN. Even though both values are significantly low, QCOM's value is 329% higher than the BAY one, meaning that it is more affected by shocks resulting from large returns.

In the Heston-Nandi model, it is also possible to study the variance persistence of the stock, by calculating  $\beta + \alpha \gamma^2$ . If the value of the equation is lower than one, the stock is mean reverting, which is verified in both stocks, with QCOM verifying a value of 0,77 and BAYN of 0,91. The closer the value is to one, the more volatility will persist, while the closer it gets to zero, the faster volatility will revert to the long-run variance (Rouah and Vainberg 2007, chapter 6). Both stocks are much closer to one, especially BAYN, meaning the mean-reversion is minimal.

The gamma parameter is responsible for the leverage effect in the Heston-Nandi model, which is not included in the GARCH (I, I) model. This effect is more intense on QCOM then it is on BAYN, when considering that this parameter is larger for QCOM. However, it is important to emphasize that both models present a very low value for this parameter.

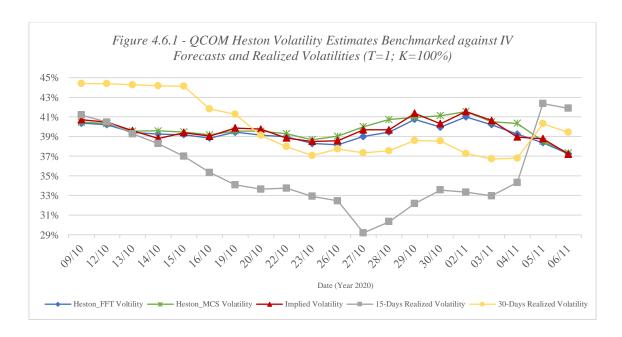
For the lambda parameter, even though it is estimated and used for the calculation of other parameters in the model, the value is almost insignificant. An approximation of -0,5 was performed as proposed by the first proposition of the Heston-Nandi paper (Heston and Nandi 2000), in order to include the risk-neutral process and achieve a better estimation and calculation of the option pricing. This approximation is used in the computations of the option pricing.

#### 4.6 Heston Model

Regarding the Heston models' empirical analysis, it will be not only based on the point (T, K) = (1, 100%) but also on the point (T, K) = (0.08333, 100%), in order to test the hypothesis previously mentioned that this model is not as efficient for options with shorter maturity. Considering this, the analysis of the realized volatility will only be performed for the point (T, K) = (1, 100%), due to the nature of this result.

Starting with the analysis of the point of the surface (T, K) = (1, 100%) with QCOM volatility forecast, as it was mentioned before, a high correlation with the implied volatility is evident

with both Heston models in analysis. By looking at the following figure, both the Heston FFT and the Heston MCS follow the trend of the implied volatility very closely. To estimate both Heston models, equations 2.7.1 and 2.7.2 are solved every day for Heston MCS and equations 2.7.4, 2.7.5, 2.7.6 and 2.7.7 for Heston FFT. This process is executed between 09/10/2020 and 06/11/2020, with all five parameters  $\Omega_0 = (k, V_0, \theta, \sigma, \rho)$  being optimized daily.



The fifteen days average realized volatility is considerably far away from the model's results, with the thirty days average being slightly closer. This occurs because the implied volatility is also deviating from the realized one. As both Heston processes have the implied volatility as an initial input, the disparity between the implied volatility and the realized volatility is associated with the Heston models forecasts deviance of the realized volatility. Regarding the actual results of QCOM, the average forecast for the data studied for all models was very similar, which is visible in the following results:

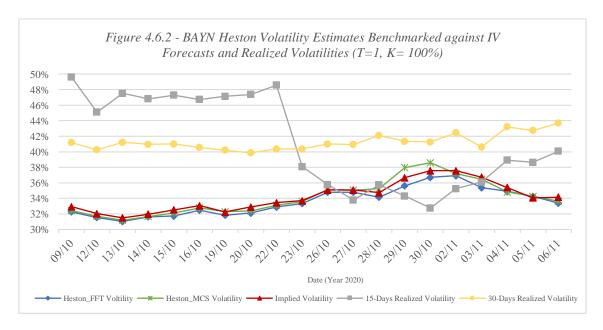
Table 4.6.1 – QCOM average annualized volatility metrics during the 09/10/2020 – 06/11/2020 period

Volatility Forecasts (T= 1; K= 100%)			Realized Volatilities	s (T= 1; K= 100%)
Heston MCS Heston FFT Implied Volatility		15-Days	30-Days	
39,70%	39,33%	39,59%	35,42%	39,94%

As it can be seen above, the average result from the period of study for the volatility forecast was very similar, not only between the Heston models and the implied volatility but also between the thirty days average realized volatility. While the Heston FFT expects, on average, a slightly lower volatility than the implied volatility, the Heston MCS forecasts the opposing way. The differences between these three averages are marginal.

Regarding the standard deviation, the Heston models presents a very important result. Both models are not only close to the implied volatility forecast but they also present a lower standard deviation. While the implied volatility presented a Standard deviation of 1,06, the standard deviation for the Heston FFT and for the Heston MCS is 0,92 and 1,00 respectively. This increased stability observable in this scenario is important for the bank in order to minimize its transaction costs involved with hedging. Moreover, it is also relevant to mention that both models correct the shortcoming of previous models regarding volatility jumps for this asset. As the Heston models are not dependent on spot prices to forecast the volatility, extremely positive or negative returns will have a very mild reaction in a long-term volatility forecast.

Regarding the BAYN volatility forecast, two important takeaways can be derived from the different estimations. The following graph presents the predictions of each model when compared against each other:

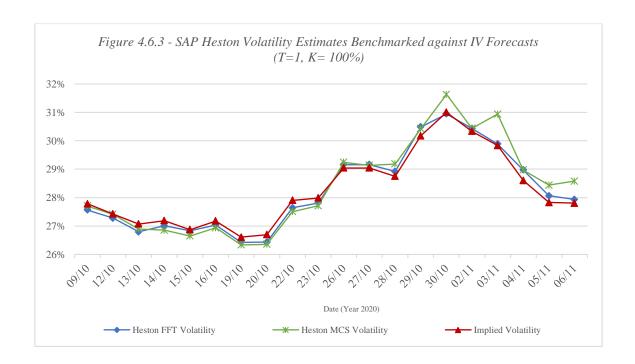


Firstly, similar to the QCOM analysis, both the Heston FFT and the Heston MCS appear to closely follow the implied volatility, throughout most of the period, and seem to deviate from the realized volatility estimations similarly to the implied volatility. The reason for this similarity was previously explained and also verified for this asset's analysis regarding the point (T, K) = (1, 100%).

Secondly, both Heston FFT and Heston MCS exhibit slightly lower average forecasts, in the period considered, than the implied volatility, 33.59% and 34.08%, respectively for Heston FFT and Heston MCS.

Where BAYN diverges from QCOM is in the standard deviations of the model. While Heston FFT still exhibits a lower standard deviation than implied volatility, 1.79 against 1.89, Heston MCS exhibits a higher standard deviation of 2.18. As a result, Heston FFT provides a slightly more stable model considering the marginal decrease in standard deviation when compared to the implied volatility and Heston MCS. This implicates that for this point of the surface and asset, Heston MCS would not be as efficient as the Heston FFT as it would unnecessarily increase the hedging cots for Banco Invest. A possible reason for this divergence in results is the number of stock paths simulated. The model was constructed assuming M = 10000 paths of the asset. Compared to the bank's own internal simulations, that assumed M = 100000 paths of the stock, Heston MCS's lower number of simulated paths might not adequately converge, making the model less useful than Heston FFT. This lower number of simulated paths was adopted in an effort to reduce the computational intensity of the Heston MCS model, making it faster to compute.

Regarding the SAP results, as the realized volatility provides very inflated results due to the spot price drop, the analysis of this forecast has additional value if those results are disregarded. Therefore, the following graph focuses only on the Heston FFT and Heston MCS when compared to the implied volatility:



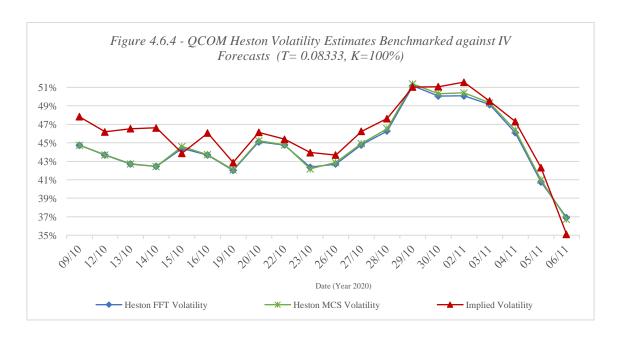
Similar to the results in the BAYN analysis, both the Heston MCS and the Heston FFT continue to present very similar results to the implied volatility ones for this point of the surface across all underlying assets that are being studied.

Contrarily to QCOM and BAYN, for this point of the volatility surface, the standard deviation of the model Heston FFT is 1,42, slightly larger than the implied volatility one, which is 1,30, which implies a slightly lower stability for this period of data. As was the case with BAYN, both models are producing a slightly more stable forecast than the model Heston MCS, which showcases a standard deviation of 1,60.

## Heston Model (T, K) = (0.0833, 100%)

Regarding the point of the surface (T, K) = (0.0833, 100%) a similar analysis will be performed, without the inclusion of the realized volatility, only comparing the implied volatility and both Heston Models forecasts. This point is important to analyze, since according to Wu (2019) shorter term maturity at-the-money options display a higher average error when compared with the rest of the surface.

Concerning the QCOM volatility forecast, as mentioned before, both models display a high degree of correlation with the implied volatility. In this specific situation, Heston MCS estimation, the green line in figure 4.6.4, is almost impossible to distinguish from the blue line marking the Heston FFT volatility, highlighting the closeness between forecasts.



Regarding the actual results of QCOM, the average forecast for the data studied for all models was very similar, with the following results of the volatility forecast:

Table 4.6.2-QCOM average annualized volatility forecasts and respective standard deviation during the 09/10/2020-06/11/2020 period

Volatility Forecasts (T= 0.0833, K= 100%)			Standard Deviation of Forecasts		
Heston MCS	Heston FFT	Implied Volatility	Heston MCS	Heston FFT	Implied Volatility
44,80%	44,69%	46,05%	3,56	3,47	3,70

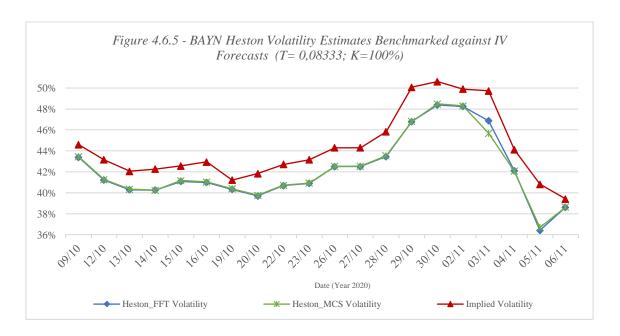
As demonstrated, the average result from the period of study for the volatility forecast were very similar, not only between the FFT, but also for the MCS. Unlike the (T, K) = (1, 100%) point, both Heston models expect, on average, a slightly lower volatility forecast than the implied volatility.

When it comes to the standard deviation, as has been the case consistently for reasons previously explained, Heston FFT tracks closer to the implied volatility forecasts than the

Heston MCS model. Heston FFT and Heston MCS presented standard deviations of 3,47 and 3,56, respectively, compared with 3,70 for the implied volatility.

Once again, the results seem to indicate that for this point of the surface and asset, Heston MCS would not be as efficient as Heston FFT, as its higher standard deviation would imply a higher frequency of hedging transactions, were this model to be implemented instead.

Regarding the BAYN volatility forecast, the following graph presents the predictions of each model, for the same point of the volatility surface, when compared against each other:



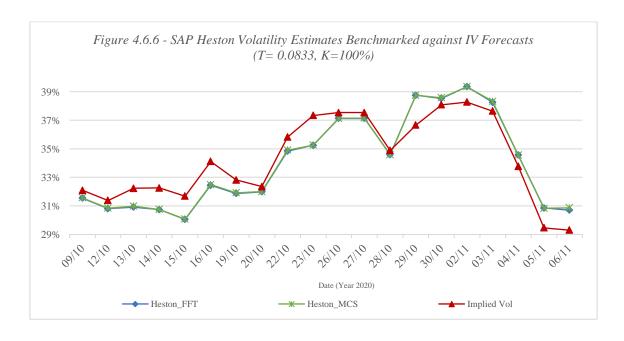
Both Heston FFT and Heston MCS, closely follow the implied volatility, as has been the case with every point of the asset for both models. The reason for this behavior is previously explained, in the QCOM analysis. Once again, the blue and green line present in figure 4.6.5 are almost impossible to distinguish.

Table 4.6.3 – BAYN average annualized volatility forecasts and respective standard deviation during the 09/10/2020 – 06/11/2020 period

Volatility Forecasts (T= 0.0833, K= 100%)			Standard Deviation of Forecasts		
Heston MCS	Heston FFT	Implied Volatility	Heston MCS	Heston FFT	Implied Volatility
42,23%	42,24%	44,28%	3,08	3,18	3,31

BAYN appears to exhibit the same behavior as QCOM, as both Heston models display a lower average volatility compared to IV, 42,24% for the FFT and 42,23% for the MCS, compared with 44,28% for the IV. The standard deviations for FFT and MCS are 3,18 and 3,08, respectively, compared with 3,31 for the IV. This contrasts with all previous cases as Heston MCS exhibits both a lower average volatility and standard deviation than Heston FFT and the implied volatility. The nature of this discrepancy is unknow, however, one possibility would be, as previously explained, a lack of simulated paths for the stock. This would make option pricing simulation more random, which could mean that for any given simulation, Heston MCS could exhibit values of the average volatility and standard deviation either lower or higher than its Heston FFT counterpart. What would appear constant is how closely both track the implied volatility of the market.

Regarding the SAP volatility forecast, the following figure presents the predictions of each model when compared against each other:



The scenario described in figure 4.6.6 is consistent with the ones presented in the two previous graphs. The two Heston volatilities derive marginally equal values during the period under study, with the green and blue lines being impossible to distinguish. Moreover, both

forecasts follow the trend of the implied volatility and present similar levels of stability, as it can be seen in table 4.6.4

Table 4.6.4 – SAP average annualized volatility forecasts and respective standard deviation during the 09/10/2020 – 06/11/2020 period

Volatility Forecasts (T= 0.0833, K= 100%)			Standard Deviation of Forecasts		
Heston MCS	Heston FFT	Implied Volatility	Heston MCS	Heston FFT	Implied Volatility
34,04%	34,01%	34,26%	3,20	3,21	2,93

For SAP, similarly to QCOM and BAYN, both Heston FFT and Heston MCS display a lower average volatility compared to the implied volatility. In this instance the difference between both models appears to be very marginal. The standard deviations, however, behave differently for SAP results. Not only are they much closer between Heston FFT and Heston MCS, 3,21 and 3,20 respectively, but they are both higher than the implied volatility, 2,93. In this case, although marginally, both models would decrease the stability of the estimate.

The conclusion reached by Wu (2019) that the Heston model would be a worse estimate than the implied volatility as option maturity decreases does not appear to be supported by the results, as both Heston approaches estimate stable values, averaging results in line with IV forecasts. However, it is important to note that the performed analysis is not able to prove that this model represents an improvement to the implied volatility, as the most correct technique of testing results is not possible to be implemented. This limitation is transversal to the whole empirical study and section 4.7 will provide details regarding the overall limitations of the entire empirical study.

## 4.7 Limitations of the Empirical Analysis

Although the previous analysis aims to recommend a specific volatility estimation model to Banco Invest, it is important to understand the characteristics of the project and therefore its limitations.

First, the time frame of the project does not allow the current analysis to fully capture the advantages and shortcomings of each model. As most of the volatility models required the implied volatility as an input, either for its computation or for the regression of the volatility surface, this analysis was only viable for the available period of data provided by the bank. Longer estimations must be implemented in order to correctly judge each of these volatility models. With this, considering the reputation and responsibility of the bank, the recommendations of this project are suggestions that should be deeply analyzed in order to ensure success in case of implementation.

Also, it is important to emphasize that the benchmarking performed in this paper is not the ideal analysis to properly comprehend the quality of the results. Although it is clear that models such as the GARCH (1, 1) and the EWMA may not be suited for the bank's activities, the analysis from the Heston-Nandi model and the Heston models are not sufficient to obtain definitive conclusions. The most accurate methodology would be to compare the past hedging payoffs and transaction costs of the bank, using the implied volatility, with the payoffs and costs resulting from the recommended models. However, given the reduced time frame of this project, such analysis was not feasible.

Finally, this paper also has limited statistical testing due to the nature of the results requested by Banco Invest, them being a comparison between the implied volatility, the volatility forecasted from the models being tested, and the realized volatility. Consequently, the project focuses mostly on the comparison of the volatility estimations, in order to distinguish which models are important for Banco Invest to further analyze.

#### 5. Conclusion

Following the empirical analysis previously performed, a cross-comparison of the results of each model is important to understand which approach has the potential to be used as the volatility forecast of Banco Invest hedging activities.

First, regarding the GARCH (1, 1) and its specific case, the EWMA, these models seem to be extremely sensitive to the returns of the underlying assets resulting in very unpredictable forecasts. Despite the volatility surface regression consistently performing well, correctly mirroring the implied volatility surface slope, the implementation of this surface is dependent on the quality of the volatility input. As it was previously mentioned, if the initial forecast of the GARCH (1, 1) or the EWMA is not appropriate, the remaining surface will turn out equally incorrect.

Despite the GARCH (1, 1) and the EWMA with  $\lambda = 0.94$  performing relatively well in the BAYN volatility forecast, the inconsistency of these models is evident in the analysis of QCOM and SAP's volatilities. The inconsistency of appropriate results imperatively denies the application of this models for Banco Invest, as a miscalculation of the volatility would make the bank incur in incorrect and unnecessary transactions.

Regarding the Heston-Nandi GARCH model, the quality of the forecasts of this model compared to the GARCH (1, 1) results seem to confirm the increased theoretical and practical robustness of the model.

On the one hand, while the Heston-Nandi GARCH model continues to work with the spot prices of the underlying assets, similarly to the models previously mentioned, the volatility outcomes of this model appear to be much more stable, having much lower standard deviations. Moreover, this model provides a closed-form of option pricing that can be translated to a volatility surface forecast (Heston and Nandi 2000). This allows the model to, by itself, compute a volatility for a specific moneyness and maturity without the assistance of the volatility surface

regression. The Heston-Nandi model also allows the bank to forecast the volatility for a particular point of maturity and moneyness without the computation of the entire volatility surface, minimizing the computational intensity involved with this process.

On the other hand, the Heston-Nandi GARCH model is still much more sensitive to intense shocks on the market than the implied volatility. The shortcomings of the GARCH (1, 1) that are constantly observable with frequent shocks in volatility are not totally fixed by this model. While unusual, the Heston-Nandi GARCH model fails to present stable results as it is evident in the analysis of the SAP volatility. However, it is relevant to emphasize that this shock is much smoother than the one forecasted by the GARCH (1, 1) and EWMA models.

Moreover, the initial parameterization of this model is not only very intensive but also arbitrary due to the various initial parameter attempts required to properly maximize the likelihood function, increasing the difficulty of implementation. Moreover, while stable, the Heston Nandi model presents a consistent bias on the volatility forecast, constantly presenting volatility predictions a few percentage points above the implied volatility estimation. As the company is currently using the implied volatility for its activities, the assumption of a volatility estimate higher than the current directly implies a safer approach of the bank, as the bank would start to expect a higher volatility than before.

Regarding the Heston models, it is evident by the empirical results that the output of these processes is very similar to the implied volatility one. The clear parallel between the Heston volatility forecast and the implied volatility one creates positive and negative qualities to the models. On one hand, because both Heston models use the implied volatility surfaces as input of option volatility forecast, this allows the models to incorporate assumptions that the previously mentioned models were not able to. As a result, the models do not present spikes in the volatility forecasts previously observable. Therefore, this process could be adopted by

Banco Invest without the significant miscalculations that were previously observable, making the Heston applicable for the bank's purpose.

Comparing the two Heston models, it is clear that the Heston FFT and the Heston MCS forecast very similar results between the two. However, while the Heston FFT is consistently similarly stable to the implied volatility or even more stable, the Monte-Carlo simulation method occasionally presents unpredictable spikes that are not necessarily in accordance with expectations, as volatility increases in one day followed by a decrease in the next one.

However, from the previous results, it is visible that there is a strong similarity between the Heston models forecasts and the implied volatility ones. Consequently, from a practical perspective, the application of this model would not drastically improve the volatility projection used by the bank. Nevertheless, the theoretical robustness of the Heston model in addition to the implied volatility input could yield slightly better volatility estimations that can only be fully comprehended with long-term comparison of the two model's hedging results.

The first recommendation is to disregard models such as the GARCH (1, 1) and the EWMA that are very unpredictable and unreliable for such purpose. Although the Heston-Nandi GARCH model provides some promising results, the shortcomings previously mentioned also propose implementation difficulties in a for Banco Invest.

As a result, one recommendation for the bank is to deeply analyze the feasibility of the implementation of the Heston model for the volatility estimation, more specifically the Heston FFT that presents more stable results than the Heston MCS. Because the bank operates with large investments, the slight improvement of the volatility estimation that is being proposed by the Heston model may yield considerable profits for Banco Invest. This would imply a constant parameterization of the model for a slight deviation of the implied volatility result, the one that Banco Invest already uses.

The incorporation of the implied volatility as an input for a better volatility forecast was crucial in order to generate a model without shocks in the estimation. It is important to emphasize that this recommendation is a suggestion for the bank to thoroughly investigate if this model yields better results in the long run. As mentioned before, the most precise way to do so is to compare the payoffs of the hedging activities from using the implied volatility and from using this model as an input of the expected volatility.

Because Banco Invest has shown some dissatisfaction towards the Implied volatility, an alternative recommendation is provided to the bank, one that may better fit the bank's purpose. If the bank is not interested in continuing to work with the implied volatility, the Heston-Nandi GARCH model is the best alternative model to be recommended. The previous results show that the Heston-Nandi GARCH, for the most part, can derive stable volatility forecasts that are aligned with the expectations. Therefore, an alternative suggestion is provided to Banco Invest in case the bank does not want to use the implied volatility as an input.

Concluding, two suggestions are provided to the bank from this project. On the one hand, if the bank is comfortable with the use of the implied volatility as an input, the parameterization and stochastic properties of the Heston FFT model combined with the stability of the implied volatility is expected to produce volatility estimations that are relevant to be better analysed in the future. As mentioned before, a slight improvement in the volatility estimation results in a competitive enhancement for Banco Invest. On the other hand, if the bank is not interested in a model that uses the implied volatility as an input, the Heston-Nandi model allows the bank to have a feasible alternative for a volatility model without the influence of this volatility estimation.

## 6. Diogo Granja's Individual Contribution to the Project

# 6.1 Autocorrelation Analysis for the GARCH (1, 1) and EWMA Models

Overall, the previous analysis concluded that the volatility estimations derived through the GARCH (1, 1) and EWMA approaches do not provide an appropriate proxy for the realized volatility verified on the market. The following section of this thesis will focus on trying to understand if this inability to correctly forecast volatilities is related to a possible poor econometric quality of the models. To achieve this goal, the models will be tested for autocorrelation, following the theory and process suggested by Hull (2009, chapter 23).

#### **6.1.1 Theoretical Framework**

According to Hull (2009, chapter 23), the squared returns of a stock (a proxy for its volatility) are highly positively autocorrelated, that is, an asset past returns influence its future returns. In other words, if a stock shows positive results in a past time series, it is expected to repeat a similar pattern of returns on a future time series, with the same rational being valid for periods where a given stock performs poorly. This characteristic transforms the squared returns in an inappropriate proxy for volatility.

The GARCH model, including the EWMA approach, assumes time-varying volatility, that should work in the following manner: when returns of a given day are highly positive (or negative), the returns of the following days are expected to follow the same pattern, making the volatility at this time period to be high. When the opposite occurs, that is, when returns are fairly stable for a given day, the GARCH model assumes that the same will occur in the following periods, resulting in low volatilities estimates. Given this, according to Hull (2009, chapter 23), if the GARCH model shows very little autocorrelation between a stock's squared return and its GARCH volatility estimate for the same day, then the model "has succeeded in explaining autocorrelations in the squared returns".

To test the econometric quality of the models, both Engle (2001) and Hull (2009, chapter 23) evaluate the variable  $\frac{\mu_n^2}{\sigma_n^2}$ , with  $\mu_n^2$  representing a stock's squared return at a given day, and  $\sigma_n^2$  it is the volatility estimate for the same day, for autocorrelation by performing the Ljung-Box test. This test was designed to evaluate the lack of fit in autoregressive-moving average models and has the following equation (Ljung and Box 1978):

$$Q = n(n+2) \sum_{k=1}^{k} \frac{\hat{p}_k^2}{n-k}$$
 (6.1.1.1)

Where *n* represents the sample size, *k* the number of lags being tested and  $\hat{p}_k^2$  the sample autocorrelation at lag *k*. By using this statistic with fifteen lags (*k*=15), Hull (2009, chapter 23) demonstrates not only that the squared returns of the S&P 500 are highly autocorrelated but also that the GARCH model fits its purpose and solves the autocorrelation problem. Assuming that *k*=15, if the result of the Ljung-Box test is a value below 25, then there is a 95% confidence interval that there is zero autocorrelation in the model (Burns 2002).

## **6.1.2 Empirical Results**

The empirical study was performed following the same technique as the one described by Hull (2009, chapter 23). First, the returns of each stock were tested for autocorrelation to confirm its existence. In this case, the variable  $\hat{p}_k$  present in equation 6.1.1.1 was  $\mu_n^2$ , evaluated at lag k. Then, the volatility estimates of each model were tested for autocorrelation to infer if the models correctly explained the autocorrelation present in the returns. In this scenario, the variable  $\hat{p}_k$  present in equation 6.1.1.1 was  $\mu_n^2/\sigma_n^2$ , evaluated at lag k.

The Ljung-Box test described in equation 6.1.1.1 was performed daily during the period between 09/10/2020 and 06/11/2020. It is important to note that for the GARCH (1, 1) model two years of daily data were put to test, whereas in the EWMA process one year of daily results

was the chosen time series. Moreover, like it was performed by Hull (2009, chapter 23), fifteen was the number of lags used in every computation of this statistical test.

The findings of the Ljung-Box statistic for the BAYN and QCOM securities were aligned with the theory and results presented by Hull (2009, chapter 23). In fact, appendix 6.1.2.1 illustrates that, for every day between 09/10/2020 and 06/11/2020, the squared returns of the two assets used to compute GARCH (1, 1) volatility forecasts are, most likely, autocorrelated. Appendix 6.1.2.2 demonstrates that the same situation happens when the data used to compute EWMA volatilities for the two stocks is tested for autocorrelation. Appendixes 6.1.2.3 and 6.1.2.4 show that the autocorrelation of squared returns is explained by both GARCH (1, 1), EWMA and EWMA with a fixed lambda of 0,94 in BAYN in QCOM securities, respectively. It is important to note that all the values present in appendixes 6.1.2.3 and 6.1.2.4 represent values below 25, which is sufficient to state, with 95% confidence, that there is no autocorrelation in the forecasts resulting from the models.

However, the findings resulting from the analysis of SAP are not aligned with the ones described above. First, it is crucial to understand why all test results after the 26<sup>th</sup> of October are extremely distinct from the ones derived until that date. From this day onwards all SAP related results present in appendixes 6.1.2.1, 6.1.2.2 and 6.1.2.5 indicate that it is highly unlikely that both the squared returns and volatility forecasts present autocorrelation in the time series. This situation was a direct result of the previously discussed 21,94% negative return that occurred on the 26<sup>th</sup> of October. This price movement was unprecedent in the time series under analysis, which made the Ljung-Box statistic to assume such low levels of autocorrelation for SAP.

Nevertheless, before this abnormal return SAP test results were still not completely aligned with what should be expected according to Hull's (2009, chapter 23) theory. As it is seen in appendix 6.1.2.1, the null hypothesis, that is, the presence of autocorrelation in squared returns,

is rejected with 95% confidence. When analyzing only one year of past returns (appendix 6.1.2.2) the null hypothesis must not be rejected.

Regarding the volatility forecasts of SAP, the ones where the maximum likelihood equation is applied appear to be free of autocorrelation, while the EWMA with a fixed lambda of 0,94 shows autocorrelation in its estimate. The results corroborating these statements can be found in appendix 6.1.2.5.

## **6.1.3 Key Takeaways**

In general, it is safe to assume that the problems related with both GARCH (1, 1) and EWMA forecasts are not a consequence of autocorrelated estimations. The models, when the maximum likelihood equation was used, appeared to correctly explain and remove, with a 95% confidence interval, the autocorrelation present in the squared returns.

Moreover, the Ljung-Box statistic was useful to understand that, while it can derive more stable volatility forecasts, the EWMA with the fixed lambda of 0,94 may result in autocorrelated estimates, which corroborates the previously idea that this approach is less correct from an econometrical standpoint.

# **6.2** Possible Improvements for the GARCH (1, 1) Volatility Estimation

The previous chapter made clear that the GARCH (1, 1) results were free of autocorrelation, indicating that the inaccurate forecasts are a result of other factors. This next section will attempt not only to understand if there is any computation error in the Excel template used to calculate the model but also describe alternative GARCH approaches that may result in more precise volatility estimations.

#### 6.2.1 An Alternative Program for Model Estimation: The use of Excel add-in NumXL

To guarantee the correctness of the GARCH (1, 1) forecasts derived, an alternative program was used, the Excel add-in *NumXL*. This tool was specially designed to "provide analytical and

statistical software tools and services for securities trading, risk management, and financial operations" (NumXL n.d.). The software gives the user the possibility to compute GARCH (p, q) parameters by simply inputting in the program the returns of the security under study and choosing the desired number of lags.

This exercise was performed, following the usual GARCH (1, 1) process, for the three stocks (QCOM, BAYN and SAP) during the sample period between 09/10/2020 and 06/11/2020. The average parameters' values computed can be seen in the following table, benchmarked against the previously derived parameters using Hull's (2009, chapter 23) GARCH (1, 1) model Excel template.

Table 6.2.1.1 – Average GARCH (1, 1) parameters' values for the 09/10/2020 – 06/11/2020 period. Comparison between the values derived by using Hull's (2009) Excel template and the Excel add-in NumXL

Parameter \ Stock	QCOM (Hull)	QCOM (NumXL)	BAYN (Hull)	BAYN (NumXL)	SAP (Hull)	SAP (NumXL)
ω	0,000234	0,000256	0,0000473	0,0000428	0,0000455	0,0000491
β	0,512	0,519	0,829	0,823	0,842	0,837
α	0,201	0,198	0,0833	0,0881	0,0496	0,0523

From the analysis of table 6.2.1.1 it is safe to conclude that both methods provide extremely similar values. The difference in results is residual and does not change the overall forecasts provided by the model.

All in all, the use of an alternative software demonstrated the correctness of the previously estimated parameters, showing that the lack of fit of the GARCH (1, 1) forecasts are not a result of computational mistakes.

## 6.2.2 Implementation of Higher Degrees of Lags: The GARCH (1, 2) Analysis

It has been proven that the GARCH (1, 1) does not provide appropriate volatility estimates for the securities under study. This happens mainly due to the instability present in the forecasted values. As such, the next step of this analysis will focus on understanding if, by considering a higher lag in the ARCH component of the model (Bollerslev 1986), the GARCH

(1, 2) process derives smother estimations. Given this, the equation used to estimate volatilities following the GARCH (1, 2) approach is the following

$$\sigma_n^2 = \omega + \alpha_1 \mu_{n-1}^2 + \alpha_2 \mu_{n-2}^2 + \beta \sigma_{n-1}^2$$
 (6.2.2.1)

As it can be seen above, adding a second lag of the ARCH component to the model represents attributing an alpha parameter not only the day before stock squared return but also to the day before that.

The volatility result that is taken away from equation 6.2.2.1 is a daily volatility that is going to be annualized assuming 252 business days in a year. Therefore, based on the equation 6.2.2.2, the yearly volatility is given by the following equation:

$$\sigma_{annual} = \sqrt{252(\omega + \alpha_1 \mu_{n-1}^2 + \alpha_2 \mu_{n-2}^2 + \beta \sigma_{n-1}^2)}$$
 (6.2.2.2)

However, it is important to note that GARCH (1, 1) forecasts usually derive extremely similar results when compared with higher orders of lags, thus being the most commonly used model derivation (Bollerslev 1986). Nevertheless, since in this case GARCH (1, 1) estimates failed to properly suit Banco Invest's objectives, it is worthwhile to explore alternatives.

With this being said, QCOM volatility estimates between 09/10/2020 and 06/11/2020 were put to test. Like in the previous analysis, two years of past data were considered to compute GARCH (1, 2) forecasts, and the volatility was computed for an at-the-money option, with one year to maturity. Moreover, the maximum likelihood equation used in this forecast is equal to the one used when performing GARCH (1, 1) estimates.

Table 6.2.2.1 displays the optimized GARCH (I, 2) average parameters' values during the period in analysis and compares the results with the parameters derived using the GARCH (I, 1) approach. First, it is relevant to note that the value attributed to the long term volatility is considerably higher in the GARCH (I, 2) process, which can result in more stable results. Regarding the  $\beta$ , it is clear that the weight given to the last volatility present in the time series

is much larger when following this approach. As a consequence, less weight is being given to past volatilities, meaning that less past information is being incorporated into the model. The alphas gained more relevance in this approach, averaging the same result during the entire period under study, which means that the weight given to the last two squared returns of QCOM was equally distributed.

Table 6.2.2.1 – GARCH (1, 2) average parameters benchmarked against GARCH (1, 1) average parameters, for the 09/10/2020 - 06/11/2020 period

Parameter \ Model	GARCH (1, 1)	GARCH (1, 2)	
ω	0,00023	0,00031	
β	0,512	0,285	
$\alpha_1$	0,201	0,186	
$\alpha_2$	NA	0,186	

The volatility estimates provided by the GARCH (1, 2) model, given an option with one year to maturity and 100% strike price, can be seen below, benchmarked against GARCH (1, 1) forecasts for the same option and period of time. To estimate the GARCH (1, 2) model equation 6.2.2.2 is solved everyday between 09/10/2020 and 06/11/2020, with the parameters being optimized daily. Moreover, each parameter optimization is done having into account two years of past QCOM stock prices, which enables the process to base its forecasts on past information.

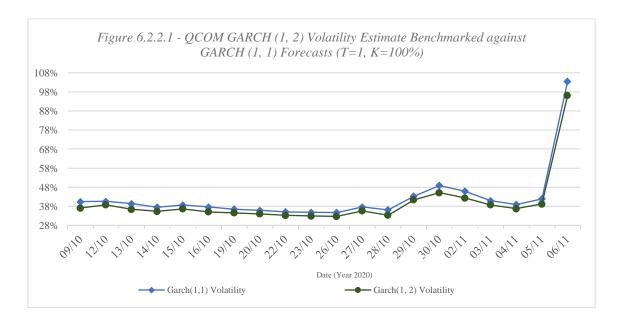
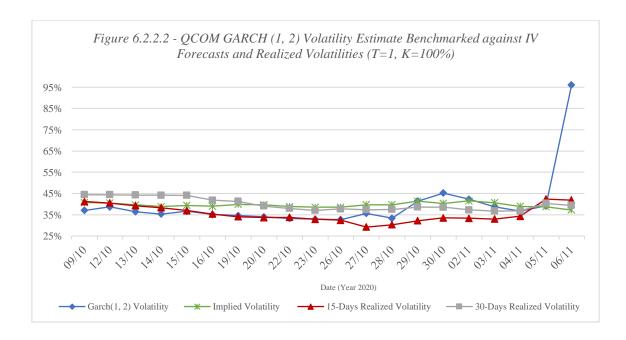


Figure 6.2.2.1 unequivocally demonstrates the similarity of results between the GARCH (1, 1) and the GARCH (1, 2) processes. The two lines follow the same pattern during the studied period, with the GARCH (1, 2) deriving relatively lower estimates. Given this, it is clear that a figure englobing IV forecasts, realized volatilities and GARCH (1, 2) estimates will be extremely similar to the one previously exhibited when the GARCH (1, 1) analysis was being performed. Figure 6.2.2.2 demonstrates such scenario.



Consequently, the average volatility forecasts derived by the two processes are similar. In fact, as it can be seen in table 6.2.2.2, GARCH (I, 2) average estimates are less than 3% lower than GARCH (I, 1) values. On average, the model derives roughly the same results as the IV metric and the thirty days realized volatility, while it presents results about 4% higher than the fifteen days realized volatility. Moreover, it is relevant to highlight that this process average forecasts were closer to the effectively verified volatility on the market than the estimates performed by the GARCH (I, I) model.

Table 6.2.2.2 – QCOM average annualized volatility metrics during the 09/10/2020 – 06/11/2020 period

Volatility	y Forecasts (T= 1, K	Realized Volatilitie	s (T= 1, K= 100%)	
GARCH (1, 2)	GARCH (1, 1)	Implied Volatility	15-Days	30-Days
39,76%	42,40%	39,58%	35,42%	39,94%

The major shortcoming of the GARCH (1, 2) process relies, similarly to the GARCH (1, 1) model, on the instability of its forecasts. Such feature can be seen not only in figure 6.2.2.2, but also by analyzing the standard deviation of results present in table 6.2.2.3. The values have a 13,67 standard deviation during the sample period, a value aligned with the one presented by GARCH (1, 1) forecasts, and 12,90 times higher than the standard deviation of IV estimates.

Table 6.2.2.3 – QCOM average standard deviation of results, for the point (T=1, K=100%), during the 09/10/2020 - 06/11/2020 period

	<b>Volatility Forecasts</b>	Realized V	Volatilities	
GARCH (1, 2)	GARCH (1, 1)	Implied Volatility	15-Days	30-Days
13,67	14,84	1,06	3,89	2,91

All in all, this analysis reinforced the idea that GARCH (p, q) models, while forecasting average volatilities aligned with the values verified on the market, are not suitable to perform stable volatility estimates, as the GARCH (1, 2) model is yet unable to provide a smooth reaction when unexpected price jumps occur.

# 6.2.3 Implementation of Higher Degrees of Lags: The GARCH (2, 1) Analysis

The previous section demonstrated that adding a second ARCH component to the model does not result in more stable volatility forecasts. Given this, chapter 6.2.3 will attempt to achieve the objective by adding a second lag of the GARCH component, that is, a GARCH (2, 1) analysis will be performed. This approach will assume the following equation (Bollerslev 1986):

$$\sigma_n^2 = \omega + \alpha_1 \mu_{n-1}^2 + \beta_1 \sigma_{n-1}^2 + \beta_2 \sigma_{n-2}^2$$
 (6.2.3.1)

Again, the volatility result that is taken away from equation 6.2.3.1 is a daily volatility that is going to be annualized assuming 252 business days in a year. Therefore, based on the equation 6.2.3.2, the yearly volatility is given by the following equation:

$$\sigma_{annual} = \sqrt{252(\omega + \alpha_1 \mu_{n-1}^2 + \beta_1 \sigma_{n-1}^2 + \beta_2 \sigma_{n-2}^2)}$$
 (6.2.3.2)

For consistency reasons, QCOM adjusted closing prices were used to perform GARCH (2, 1) forecasts. Again, two years of past information were used in all computations, and volatilities were forecasted for the period between 09/10/2020 and 06/11/2020 using equation 6.2.3.2. Like in previous analysis, the volatility was computed for an at-the-money option, with one year to maturity and the maximum likelihood equation used is equal to the one used when performing GARCH (1, 1) and (1, 2) estimates.

The parameters computed by the model can be found in appendix 6.2.3.1 benchmarked against GARCH (1, 1) values. The results are uniform between models, meaning that the added lag, when optimized, equals zero.

In conclusion, adding a second GARCH element to the process is, in this case, irrelevant, as the maximum likelihood equation continues to reach its higher value when  $\beta_2$  equals zero and the other parameters assume the values of the GACH (1, 1) estimation.

#### 6.2.4 Key Takeaways

First, it is relevant to note that, given the lack of fit of the results derived for QCOM, performing the analysis present in sections 6.2.2 and 6.2.3 for the remaining securities under study would be redundant. The further studies performed using different orders of lags demonstrated that GARCH (p, q) models are not appropriate to provide the smooth volatility estimates required by Banco Invest.

The standard deviation of forecasts remains excessively high when using the GARCH (1, 2) method, which would result in an increased and inefficient number of trades the bank would have to do. Such scenario would go against the main objective of the institution - reducing transactions and consequently trading costs. Moreover, GARCH (2, 1) forecasts ended up being equal to the GARCH (1, 1) values, thus resulting in no practical use.

All in all, the conclusions taken from this detailed study of GARCH (p, q) models are aligned with the ones found when studying the GARCH (1, 1) model. With this being said, it is clear that such type of methods must be disregarded and alternative models studied in order to achieve the objective of this thesis.

#### 7. List of References

- Ait-Sahalia, Yacine, Jianqing Fan, and Yingying Li. 2013. "The leverage effect puzzle: Disentangling sources of bias at high frequency." Journal of Financial Economics 109, no. 1: 224-249.
- Andersen, Leif, and Jesper Andreasen. 2000. "Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing." Review of derivatives research 4, no. 3: 231-262.
- Bollerslev, Tim. 1986. "Generalized autoregressive conditional heteroskedasticity." Journal of econometrics 31, no. 3: 307-327.
- Bouchaud, Jean-Philippe, Andrew Matacz, and Marc Potters. 2001. "Leverage effect in financial markets: The retarded volatility model." Physical review letters 87, no. 22: 228701.
- Broadie, Mark, and Özgür Kaya. 2006. "Exact simulation of stochastic volatility and other affine jump diffusion processes." Operations research 54, no. 2: 217-231.
- Brous, Peter, Ufuk Ince, and Ivilina Popova. 2010. "Volatility forecasting and liquidity: Evidence from individual stocks." Journal of Derivatives & Hedge Funds 16, no. 2: 144-159.
- Burns, Patrick. 2002. "Robustness of the Ljung-Box test and its rank equivalent." Available at SSRN 443560, no. 1: 1-17.
- Byun, Suk Joon. 2011. "A study on heston-nandi garch option pricing model." In 3rd International Conference on Information and Financial Engineering, no. 12: 467-471.
- Canina, Linda, and Stephen Figlewski. 1993. "The informational content of implied volatility." The Review of Financial Studies 6, no. 3: 659-681.
- Car, P. and Madan, D. B.. 1999. "Option Valuation Using the Fast Fourier Transform". Journal of Computational Finance, no. 2: 61-73.

- Christoffersen, Peter, Kris Jacobs, and Chayawat Ornthanalai. 2013. "GARCH option valuation: theory and evidence." The Journal of Derivatives 21, no. 2: 8-41.
- Crisóstomo, Ricardo. 2014. "An Analyisis of the Heston Stochastic Volatility Model: Implementation and Calibration Using Matlab.", CNMV Working papers No. 58.
- Deelstra, Griselda, and Freddy Delbaen. 1998. "Convergence of discretized stochastic (interest rate) processes with stochastic drift term." Applied stochastic models and data analysis 14, no. 1: 77-84.
- Dumas, Bernard, Jeff Fleming, and Robert E. Whaley. 1998. "Implied volatility functions: Empirical tests." The Journal of Finance 53, no. 6: 2059-2106.
- Edwards, Tim, and Craig J. Lazarra. 2016. "Realized Volatility Indices: Measuring Market Risk." McGraw Hill Financial.
- Engle, Robert F. 1982. "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation." Econometrica: Journal of the Econometric Society, no. 1: 987-1007.
- Engle, Robert. 2001. "GARCH 101: The use of ARCH/GARCH models in applied econometrics." Journal of economic perspectives 15, no. 4: 157-168.
- Gatheral, Jim. 2006. The volatility surface: a practitioner's guide. Vol. 357. John Wiley & Sons.
- Heston, Steven L., and Saikat Nandi. 2000. "A closed-form GARCH option valuation model." The review of financial studies 13, no. 3: 585-625.
- Higham, Desmond J., and Xuerong Mao. 2005. "Convergence of Monte Carlo simulations involving the mean-reverting square root process." Journal of Computational Finance 8, no. 3: 35-61.
- Homescu, Cristian. 2011. "Implied volatility surface: Construction methodologies and characteristics.", no. 1: 35-61.

- Hull, John. 2009. Options, futures and other derivatives/John C. Hull. Upper Saddle River, NJ: Prentice Hall,
- Jarešová, Lucia. 2010. "EWMA historical volatility estimators." Acta Universitatis Carolinae. Mathematica et Physica 51, no. 2: 17-28.
- Kahl, Christian, and Peter Jäckel. 2005. "Not-so-complex logarithms in the Heston model." Wilmott magazine 19, no. 9: 94-103.
- Kamiński, Szymon. 2013. "The pricing of options on WIG20 using GARCH models". no. 1: 2006-13.
- Kloeden, Peter E. and Eckhard Platen. 1992. "Numerical Solution of Stochastic Differential Equations", Springer-Verlag
- Ljung, Greta M., and George EP Box. 1978. "On a measure of lack of fit in time series models." Biometrika 65, no. 2: 297-303.
- Lord, Roger, and Christian Kahl. 2007. "Optimal Fourier inversion in semi-analytical option pricing.", Journal of Computational Finance, no. 10(4): 1–30.
- Lord, Roger, Remmert Koekkoek, and Dick Van Dijk. 2010. "A comparison of biased simulation schemes for stochastic volatility models." Quantitative Finance 10, no. 2: 177-194.
- Moodley, N. 2005. "The Heston Model: A Practical Approach", Honours Thesis: The University of the Witwatersrand.
- Morgan, J. P. 1996. "RiskMetricsTM--Technical Document, Morgan Guaranty Trust Companies." Inc. New York.
- Nandi, Saikat, and Daniel F. Waggoner. 2000. "Issues in hedging options positions." Economic Review-Federal Reserve Bank of Atlanta 85, no. 1: 24.
- NumXL. 2020. "About Us". 2020. https://numxl.com/about-us/

- Poklewski-Koziell, Warrick. 2013. "Stochastic volatility models: Calibration, pricing and hedging." PhD diss., University of the Witwatersrand, Faculty of Science, School of Computational and Applied Mathematics.
- Rouah, Fabrice D., and Gregory Vainberg. 2007. Option pricing models and volatility using Excel-VBA. Vol. 361. John Wiley & Sons.
- Stephane, Moins. 2002. "Implementation of a simulated annealing algorithm for \Matlab", no. 1: 1-34.
- Wolters, Jürgen and Uwe Hassler. 2005. "Unit root testing", Diskussionsbeiträge, No. 2005/23, Freie Universität Berlin, Fachbereich Wirtschaftswissenschaft, Berlin
- Wu, Hsin-Fang. 2019, "From Constant to Stochastic Volatility: Black-Scholes Versus Heston Option Pricing Models". *Senior Projects Spring 2019*. 163.
- Yahoo Finance. 2020. "Bayer Aktiengesellschaft (BAYN.DE)." 2020. https://finance.yahoo.com/quote/BAYN.DE
- Yahoo Finance. 2020. "QUALCOMM Incorporated (QCOM)." 2020. https://finance.yahoo.com/quote/QCOM
- Yahoo Finance. 2020. "SAP SE (SAP.DE)." 2020. https://finance.yahoo.com/quote/SAP.DE

# 8. Appendix

Appendix 1 - Bayer comparison between all the different models

Dates	Price	Returns	Imp. Vol	EWMA 0,94	Garch(1,1)	EWMA Param	HN-Garch	Heston_FFT	Heston_MCS	Real.15-Days	Real. 30-Days
09/10/2020	46,73		32,94%	53,88%	51,33%	59,56%	43,90%	32,26%	32,43%	49,62%	41,18%
12/10/2020	46,36	-0,79%	32,02%	52,35%	49,07%	56,85%	43,79%	31,53%	31,71%	45,10%	40,25%
13/10/2020	44,75	-3,47%	31,49%	50,85%	46,41%	54,27%	43,52%	31,03%	31,19%	47,53%	41,21%
14/10/2020	44,33	-0,95%	31,97%	51,12%	44,30%	54,38%	43,35%	31,62%	31,66%	46,81%	40,97%
15/10/2020	44,11	-0,49%	32,53%	49,70%	43,97%	51,99%	43,16%	31,73%	32,11%	47,29%	41,00%
16/10/2020	44,81	1,59%	33,05%	48,22%	41,75%	49,58%	42,87%	32,48%	32,77%	46,73%	40,56%
19/10/2020	44,33	-1,07%	32,24%	47,16%	39,62%	47,84%	42,69%	31,83%	32,28%	47,13%	40,21%
20/10/2020	43,62	-1,60%	32,88%	45,91%	38,46%	45,84%	42,44%	32,11%	32,39%	47,37%	39,86%
22/10/2020	42,33	-2,96%	33,44%	44,70%	37,35%	43,96%	42,29%	32,89%	33,13%	48,57%	40,35%
23/10/2020	42,42	0,22%	33,70%	43,37%	35,41%	41,82%	42,18%	33,33%	33,59%	38,08%	40,35%
26/10/2020	42,56	0,31%	35,08%	42,06%	33,59%	39,70%	41,99%	34,80%	35,12%	35,76%	41,00%
27/10/2020	41,91	-1,52%	35,08%	40,79%	32,44%	37,67%	41,77%	34,80%	35,05%	33,80%	40,93%
28/10/2020	40,71	-2,88%	34,73%	39,99%	32,13%	36,46%	41,42%	34,15%	35,25%	35,73%	42,08%
29/10/2020	40,79	0,20%	36,66%	40,35%	34,14%	37,51%	41,60%	35,62%	37,96%	34,31%	41,34%
30/10/2020	40,36	-1,04%	37,57%	39,13%	32,76%	35,52%	41,56%	36,71%	38,59%	32,76%	41,26%
02/11/2020	41,81	3,58%	37,56%	38,15%	31,99%	34,00%	41,45%	36,92%	37,21%	35,25%	42,43%
03/11/2020	41,96	0,37%	36,74%	39,52%	35,55%	37,08%	41,72%	35,38%	36,44%	36,11%	40,60%
04/11/2020	43,75	4,27%	35,40%	38,35%	34,04%	35,19%	41,80%	34,90%	34,81%	38,91%	43,22%
05/11/2020	43,79	0,10%	34,06%	40,71%	38,74%	39,67%	42,23%	34,26%	39,35%	38,64%	42,73%
06/11/2020	42,74	-2,40%	34,16%	39,47%	36,70%	37,73%	42,76%	33,37%	33,64%	40,07%	43,68%
Average			33,75%	44,66%	40,18%	44,64%	42,96%	32,50%	34,08%	43,32%	42,34%
Stand. Dev			1,78	4,74	6,47	7,42	1,15	2,81	2,22	6,02	2,68

Appendix 2 - Qualcomm comparison between all the different models

Dates	Price	Returns	Imp. Vol	EWMA 0,94	Garch(1,1)	EWMA Param	HN-Garch	Heston_FFT	Heston_MCS	Real.15-Days	Real. 30-Days
09/10/2020	124,87		40,72%	40,04%	40,42%	37,67%	48,25%	40,36%	40,46%	41,20%	44,40%
12/10/2020	126,69	1,46%	40,47%	39,65%	40,59%	36,73%	48,41%	40,20%	40,33%	40,46%	44,39%
13/10/2020	127,46	0,61%	39,60%	38,85%	39,36%	34,36%	48,85%	39,42%	39,55%	39,27%	44,28%
14/10/2020	129,88	1,90%	38,81%	37,74%	37,54%	30,94%	48,63%	39,25%	39,57%	38,28%	44,17%
15/10/2020	128,58	-1,00%	39,40%	36,75%	38,68%	28,27%	48,31%	39,20%	39,46%	36,98%	44,11%
16/10/2020	129,03	0,35%	39,04%	36,40%	37,62%	28,39%	48,69%	38,85%	39,18%	35,33%	41,81%
19/10/2020	128,42	-0,47%	39,86%	35,32%	36,44%	25,39%	48,15%	39,45%	39,57%	34,08%	41,28%
20/10/2020	128,30	-0,09%	39,75%	34,29%	35,85%	22,77%	48,12%	39,15%	39,55%	33,64%	39,12%
22/10/2020	128,38	0,06%	38,84%	32,25%	35,07%	17,64%	47,85%	38,98%	39,29%	33,73%	37,98%
23/10/2020	128,88	0,39%	38,51%	31,27%	34,79%	15,32%	47,70%	38,28%	38,68%	32,91%	37,07%
26/10/2020	126,20	-2,08%	38,57%	30,35%	34,68%	13,52%	47,62%	38,15%	39,04%	32,45%	37,72%
27/10/2020	125,91	-0,23%	39,68%	30,52%	37,65%	19,97%	47,97%	39,00%	40,00%	29,19%	37,34%
28/10/2020	121,58	-3,44%	39,69%	29,60%	36,09%	17,87%	47,94%	39,46%	40,74%	30,33%	37,54%
29/10/2020	126,44	4,00%	41,36%	31,66%	43,28%	29,08%	48,32%	40,76%	40,96%	32,18%	38,59%
30/10/2020	123,36	-2,44%	40,31%	34,41%	48,89%	38,66%	48,46%	39,95%	41,12%	33,55%	38,54%
02/11/2020	123,97	0,49%	41,55%	34,68%	45,94%	38,79%	48,21%	41,00%	41,54%	33,33%	37,27%
03/11/2020	125,45	1,19%	40,63%	33,68%	40,93%	34,73%	47,78%	40,20%	40,50%	32,96%	36,71%
04/11/2020	128,97	2,81%	38,96%	32,98%	38,95%	32,21%	47,21%	39,24%	40,35%	34,32%	36,79%
05/11/2020	145,41	12,75%	38,79%	33,79%	41,89%	35,14%	47,22%	38,37%	38,53%	42,36%	40,32%
06/11/2020	145,01	-0,28%	37,20%	59,41%	103,42%	101,05%	51,25%	37,23%	37,32%	41,87%	39,43%
Average			39,16%	41,37%	46,06%	42,53%	49,34%	38,77%	39,99%	39,33%	40,22%
Stand. Dev			1,15	10,63	15,00	23,67	1,90	1,21	2,88	7,25	2,53

Appendix 3 - SAP comparison between all the different models

Dates	Price	Returns	Imp. Vol	EWMA 0,94	Garch(1,1)	EWMA Param	HN-Garch	Heston_FFT	Heston_MCS	Real.15-Days	Real. 30-Days
09/10/2020	133,06		27,78%	24,89%	27,29%	24,01%	32,06%	27,57%	27,70%	27,50%	31,57%
12/10/2020	134,38	0,99%	27,42%	24,22%	26,88%	23,06%	31,40%	27,28%	27,41%	24,87%	31,72%
13/10/2020	134,08	-0,22%	27,07%	23,80%	26,70%	22,32%	30,35%	26,79%	26,90%	24,56%	31,29%
14/10/2020	134,32	0,18%	27,19%	23,09%	26,25%	21,10%	29,86%	27,01%	26,85%	23,87%	30,85%
15/10/2020	130,56	-2,80%	26,87%	22,40%	25,87%	20,14%	29,96%	26,83%	26,65%	25,59%	31,25%
16/10/2020	132,84	1,75%	27,18%	24,29%	27,75%	23,31%	29,76%	27,03%	26,94%	25,03%	30,25%
19/10/2020	131,00	-1,39%	26,61%	24,51%	27,96%	23,87%	30,27%	26,42%	26,33%	23,35%	28,99%
20/10/2020	127,54	-2,64%	26,69%	24,37%	27,92%	23,73%	31,84%	26,43%	26,36%	25,68%	28,98%
22/10/2020	124,50	-2,38%	27,90%	25,76%	29,09%	25,61%	32,27%	27,64%	27,50%	26,52%	28,55%
23/10/2020	124,90	0,32%	27,98%	26,64%	29,75%	26,75%	32,26%	27,81%	27,71%	25,93%	26,98%
26/10/2020	97,50	-21,9%	29,04%	25,86%	28,94%	25,74%	32,17%	29,16%	29,24%	46,23%	36,81%
27/10/2020	96,95	-0,56%	29,04%	88,60%	82,14%	39,17%	56,39%	29,16%	29,14%	48,44%	37,69%
28/10/2020	92,24	-4,86%	28,75%	85,93%	75,00%	39,16%	68,01%	28,92%	29,18%	52,95%	40,01%
29/10/2020	93,26	1,11%	30,18%	85,43%	71,72%	39,46%	66,83%	30,48%	30,43%	54,32%	41,09%
30/10/2020	91,49	-1,90%	31,00%	82,94%	64,22%	39,47%	65,80%	30,94%	31,63%	55,55%	41,84%
02/11/2020	90,18	-1,43%	30,33%	80,75%	59,14%	39,51%	64,82%	30,43%	30,44%	56,81%	42,16%
03/11/2020	93,95	4,18%	29,83%	78,49%	54,40%	39,53%	63,94%	29,89%	30,94%	60,04%	42,46%
04/11/2020	96,75	2,98%	28,60%	77,81%	53,72%	39,76%	63,18%	28,97%	28,97%	62,89%	43,72%
05/11/2020	98,21	1,51%	27,83%	76,33%	51,35%	39,87%	62,53%	28,06%	28,44%	64,06%	43,97%
06/11/2020	97,25	-0,98%	27,80%	74,23%	48,22%	39,89%	61,88%	27,94%	28,58%	63,26%	44,42%
Average			28,10%	52,64%	43,22%	31,96%	49,75%	28,09%	29,88%	43,05%	36,53%
Stand. Dev			1,28	27,92	17,89	8,29	15,49	1,40	7,35	17,26	6,43

Appendix 6.1.2.1 - Ljung-Box test results for the squared returns present in the time. series used to compute GARCH (1, 1) forecasts

Date \ Stock	BAYN	QCOM	SAP
09/10/2020	36,470	70,059	20,427
12/10/2020	36,358	70,061	20,615
13/10/2020	35,681	70,165	20,755
14/10/2020	33,556	70,301	20,934
15/10/2020	33,968	70,397	21,112
16/10/2020	33,153	70,399	21,016
19/10/2020	33,100	70,488	21,009
20/10/2020	33,134	70,490	20,446
22/10/2020	34,547	70,645	20,391
23/10/2020	34,483	70,743	20,345
26/10/2020	34,571	70,952	20,686
27/10/2020	38,404	71,369	3,495
28/10/2020	38,423	71,390	3,593
29/10/2020	38,422	71,288	4,110
30/10/2020	38,492	71,182	2,929
02/11/2020	38,566	71,148	2,640
03/11/2020	38,461	71,140	2,618
04/11/2020	38,550	71,142	2,565
05/11/2020	38,235	71,227	2,513
06/11/2020	38,246	67,183	2,248

 $Appendix\ 6.1.2.2\ -\ Ljung\ -Box\ test\ results\ for\ the\ squared\ returns\ present\ in\ the\ time\ series\ used\ to\ compute\ EWMA\ forecasts$ 

Date \ Stock	BAYN	QCOM	SAP
09/10/2020	35,445	141,264	44,670
12/10/2020	35,193	140,997	44,658
13/10/2020	34,165	141,021	45,737
14/10/2020	30,206	140,906	52,629
15/10/2020	30,166	140,863	53,473
16/10/2020	28,797	140,767	52,699
19/10/2020	28,515	140,822	87,610
20/10/2020	28,242	140,964	87,388
22/10/2020	27,730	141,426	86,673
23/10/2020	27,753	141,703	86,192
26/10/2020	27,680	142,058	86,143
27/10/2020	27,759	142,243	6,098
28/10/2020	27,726	142,847	5,910
29/10/2020	27,604	142,625	6,083
30/10/2020	27,674	142,238	3,662
02/11/2020	27,627	142,292	2,896
03/11/2020	27,519	142,740	2,840
04/11/2020	27,807	143,059	2,313
05/11/2020	27,370	143,423	2,032
06/11/2020	27,353	122,376	1,749

Appendix 6.1.2.3 - Ljung-Box daily test results for volatility forecasts of BAYN stock

EWMA EWMA Date \ Stock GARCH (1, 1) (optimized  $\lambda$ )  $(\lambda = 0.94)$ 09/10/2020 7,417 22,083 21,081 7,375 22,339 12/10/2020 20,968 7,371 22,430 13/10/2020 20,721 14/10/2020 6,132 10,689 12,605 6,202 10,751 12,731 15/10/2020 16/10/2020 5,885 9,549 10,461 19/10/2020 5,938 9,354 10,343 20/10/2020 5,989 9,159 10,288 22/10/2020 6,253 7,454 8,108 23/10/2020 6,068 7,602 8,246 26/10/2020 6,022 7,590 8,255 27/10/2020 6,374 8,274 7,596 28/10/2020 6,377 7,597 8,280 6,392 7,607 8,291 29/10/2020 30/10/2020 6,398 7,647 8,345 02/11/2020 6,379 7,648 8,348 03/11/2020 6,463 7,662 8,344 6,521 7,690 8,384 04/11/2020 05/11/2020 6,446 7,663 8,337 06/11/2020 6,504 8,335 7,665

Appendix 6. 1.2.4 - Ljung-Box daily test results for volatility forecasts of QCOM stock

Date \ Stock	GARCH (1, 1)	EWMA (optimized λ)	EWMA (λ=0,94)
09/10/2020	0,564	7,124	3,679
12/10/2020	0,563	7,094	3,682
13/10/2020	0,560	7,023	3,670
14/10/2020	0,548	7,102	3,701
15/10/2020	0,552	7,100	3,713
16/10/2020	0,553	7,110	3,725
19/10/2020	0,553	7,028	3,701
20/10/2020	0,557	7,038	3,714
22/10/2020	0,569	7,221	3,775
23/10/2020	0,565	7,323	3,816
26/10/2020	0,570	7,472	3,859
27/10/2020	0,595	7,113	3,896
28/10/2020	0,580	7,105	3,908
29/10/2020	0,590	7,003	3,875
30/10/2020	0,587	7,288	3,901
02/11/2020	0,583	7,506	3,927
03/11/2020	0,569	7,137	3,853
04/11/2020	0,560	7,314	3,873
05/11/2020	0,572	7,523	3,939
06/11/2020	0,622	7,514	3,529

Appendix 6.1.2.5 - Ljung-Box daily test results for volatility forecasts of SAP stock

Date \ Stock	GARCH (1, 1)	EWMA (optimized λ)	EWMA (λ=0,94)	
09/10/2020	3,523	8,569	13,902	
12/10/2020	3,454	11,070	14,464	
13/10/2020	3,421	11,229	15,377	
14/10/2020	3,356	13,857	20,977	
15/10/2020	3,373	13,806	21,095	
16/10/2020	3,397	13,784	20,012	
19/10/2020	3,388	19,748	26,845	
20/10/2020	3,075	19,905	27,155	
22/10/2020	3,093	20,431	27,608	
23/10/2020	3,062	20,586	27,913	
26/10/2020	3,124	20,944	28,382	
27/10/2020	1,281	6,098	5,596	
28/10/2020	1,136	5,910	5,470	
29/10/2020	1,042	6,083	5,285	
30/10/2020	0,775	3,662	4,343	
02/11/2020	0,723	2,896	3,466	
03/11/2020	0,689	2,840	3,491	
04/11/2020	0,547	2,313	3,351	
05/11/2020	0,451	2,032	1,287	
06/11/2020	0,471	1,749	1,035	

Appendix 6.2.3.1 – GARCH (2, 1) parameters benchmarked against GARCH (1, 1) parameters for QCOM stock, during the 09/10/2020 – 06/11/2020 period

		GARC	CH (2, 1)		GARCH (1, 1	')	
	ω	α	$oldsymbol{eta_1}$	$oldsymbol{eta}_2$	ω	α	$oldsymbol{eta_1}$
09/10/2020	0,000241	0,189	0,514	0	0,000241	0,189	0,514
12/10/2020	0,000240	0,190	0,513	0	0,000240	0,190	0,513
13/10/2020	0,000243	0,192	0,507	0	0,000243	0,192	0,507
14/10/2020	0,000240	0,192	0,509	0	0,000240	0,192	0,509
15/10/2020	0,000240	0,194	0,509	0	0,000240	0,194	0,509
16/10/2020	0,000241	0,195	0,507	0	0,000241	0,195	0,507
19/10/2020	0,000240	0,197	0,507	0	0,000240	0,197	0,507
20/10/2020	0,000238	0,196	0,510	0	0,000238	0,196	0,510
22/10/2020	0,000232	0,196	0,517	0	0,000232	0,196	0,517
23/10/2020	0,000231	0,199	0,515	0	0,000231	0,199	0,515
26/10/2020	0,000234	0,206	0,505	0	0,000234	0,206	0,505
27/10/2020	0,000231	0,211	0,506	0	0,000231	0,211	0,506
28/10/2020	0,000230	0,210	0,508	0	0,000230	0,210	0,508
29/10/2020	0,000232	0,209	0,508	0	0,000232	0,209	0,508
30/10/2020	0,000232	0,209	0,508	0	0,000232	0,209	0,508
02/11/2020	0,000233	0,208	0,509	0	0,000233	0,208	0,509
03/11/2020	0,000230	0,204	0,516	0	0,000230	0,204	0,516
04/11/2020	0,000226	0,201	0,522	0	0,000226	0,201	0,522
05/11/2020	0,000219	0,198	0,534	0	0,000219	0,198	0,534
06/11/2020	0,000230	0,224	0,520	0	0,000230	0,224	0,520