Modeling SPX Volatility to Improve Options Pricing

Jared Aiman, Vicente Iglesias, Sumit Sarkar

May 12, 2021

Dr. Kathy Ensor, Noah G. Harding Professor of Statistics

Dr. John A. Dobelman, Professor in the Practice of Statistics

Abstract

In this project, we develop a model to predict future stock market volatility and facilitate more accurate options pricing. The Black Scholes model gives an expected premium for an options contract; however, it uses an unknown fixed parameter referred to as volatility. We advance this by using a modified Glosten-Jagannathan-Runkle Generalized Autoregressive Conditional Heteroskedasticity (GJR-GARCH) model that uses previous returns, as well as the market's expectation of future volatility, to better predict future volatility. Additionally, we apply an Autoregressive Moving Average (ARMA) model to predict the value of future stock prices. We find that our model is able to model volatility better than using either the market volatility or a traditional GJR-GARCH model alone. This is particularly true due to our model's ability to capture the dependence between the S&P 500 returns and the changes in the market's expectation of volatility.

1. Introduction

Public companies sell shares of ownership in their company to help raise funds. These shares are referred to as stocks. After the initial sale, stocks are traded and can be bought and sold by anyone who wants to invest in the stock market. There are also associated derivatives traded on these stocks; a common type of derivative is the options contract, which gives the buyer the right to buy or sell 100 shares of an underlying asset at a pre-agreed price, called the strike price. Each of these terms is defined in Table A1 in Appendix A, along with other relevant terminology.

In addition to options on stocks, there are also options on indexes and Exchange-Traded Funds (ETF). An ETF is a portfolio of stocks and/or bonds that is created to mimic the overall performance of a standing index. For example, SPDR S&P 500 (SPY) is an ETF that tracks and mimics the performance of the Standard and Poor's 500 (SPX or S&P 500) - a collection of the top 500 largest publicly traded companies in the United States.

The CBOE Volatility Index (VIX) is an index that tracks the expected 30 day annualized volatility, or standard deviation of log returns, of the S&P 500 and is calculated using option prices with several different strikes. There are no tradeable VIX shares, but there are options traded on the VIX. According to Cao, Badescu, Cui, and Jayaraman (2020), volatility's movements are negatively correlated to price changes in the equity market. This is why the VIX is regularly referred to as a "fear index" - that is, the higher the VIX is, the more worried people are.

An investor can position themselves long or short volatility, where they make a profit if the volatility moves up or down respectively, through the use of options contracts. If the investor believes volatility will increase, they can buy options contracts. Similarly, a different investor may believe volatility is going to decrease and would thus position themselves short volatility by buying options contracts. In a study focused on volatility in oil pricing, Ensor, Han, Ostdiek, and Turnbull (2020) determined that negative news had a larger impact on volatility increase than positive news did. Additionally, negative news caused larger downward movement in price than positive news did for upward price movement. The research done by Ensor, Han, Ostdiek, and Turbull focuses on the specific difference between oil futures contracts (See Table C1, Appendix C) and US equities; however, the same dynamics likely may be applied to S&P 500 volatility.

The Black Scholes model is the best known model for calculating the value of an options contract (Black & Scholes 1973). Under certain assumptions, the Black Scholes model gives the value of the contract as a function of stock price, strike price, time until expiration, risk free rate, and the volatility (See Table C1, Appendix C). A key assumption of the model is that volatility is constant, and the logarithm of the stock prices moves as a normal distribution. It is well known that the Black Scholes model misprices options due to the assumption of a normal distribution (Corrado & Su 1996). However, the model is still commonly used to understand and derive options prices, as well as calculate implied volatility.

One popular theory states that the market is completely unpredictable and that prices are completely efficient. This is called the "Efficient Market Hypothesis," and it can be summarized by the belief that current stock prices are a reflection of all information that is publicly available. If this were true, it would be impossible to pick undervalued stocks without insider information (Malkiel & Fama 1970). Many people, such as William & Dobelman in 2020, have argued against the Efficient Market Hypothesis, and there has been substantial effort put into predicting the stock market.

The research in predicting the stock market has been done both to predict movement direction (Roondiwala et al. 2017) and also the uncertainty inherent in the move (Fleming et al. 1995). Despite the research on ways the stock market is not completely efficient, investing in an index fund generally generates better long-term results than almost all professional investors (Pisani 2019), implying that prediction is incredibly difficult to perform successfully. However, the potential payout for correct prediction continues to serve as a motivator for stock market modeling and prediction.

While knowing the uncertainty in future price moves does not necessarily lead to better profit when buying and selling stocks, it can lead to more profit on options contracts. Possessing this type of insight illuminates whether an option's price is currently underpricing the potential for large price movements in either direction. Options contracts are leveraged given that they are trading in moves from the strike price instead of zero like in stocks, so correctly pricing volatility would give investors and traders a high probability of increased returns on investment.

Our project utilizes time series techniques to better predict volatility. We conduct our modeling using a combination of parameters which includes features such as past price and implied volatility. With this modeling, we are able to better predict future volatility when compared to common existing models.

2. Methodology

2.1 Model Background

The foundation of our work in modeling the stock market utilizes time series analysis to predict the future value and volatility of the stock market. To do this we combine two time series models. The first is the Autoregressive Moving Average (ARMA). Below the formula for an ARMA(1,1) model is given which has one lag in both the error term and response variable.

$$r_t = \alpha_0 + \alpha_1 r_{t-1} + \epsilon_t + \alpha_2 \epsilon_{t-1}$$

Here r_t is the return at time t, and the error term $\epsilon_t \sim$ skewed-$t(location = 0, scale = \sigma^2, skew, df)$. The distribution of the error term does not have to be a skewed-t distribution, but we used it to allow for asymmetry and fat tails in transitions. We allow for asymmetry and fat tails to better account for the likelihood of extreme outcomes. For example, at the beginning of the COVID-19 crisis, volatility increased drastically because of the frequent negative news surrounding COVID-19.

Currently, this type of event is a rare occurrence; however, the likelihood of similar rare occurrences may be underestimated by the normal distribution and the drastic increase in volatility observed shows why accounting for fat tails and asymmetry may give better results.

A Generalized Autoregressive Conditional Heteroskedasticity (GARCH) process models the volatility of the error term in the ARMA model for returns. The GARCH(1,1) models volatility as:

$$\sigma_t^2 = \beta_0 + \beta_1 \epsilon_{t-1}^2 + \beta_2 \sigma_{t-1}^2.$$

A Glosten-Jagannathan-Runkle GARCH (GJR-GARCH) model uses the GARCH model as a foundation but adds in β_3 term as a coefficient for a leveraged term for the squared residual. The leverage term takes into account asymmetry in how returns affect volatility. We apply a GJR-GARCH(1,1) for the volatility of the error process:

$$\sigma_t^2 = \beta_0 + \beta_1 \epsilon_{t-1}^2 + \beta_2 \sigma_{t-1}^2 + \beta_3 \epsilon_{t-1}^2 I_{t-1}$$

where $I_{t-1} = 1$ if $\epsilon_{t-1} < 0$. Both the GARCH and the GJR-GARCH can also be extended to have more lagged terms.

The ARMA(1,1)-GJR-GARCH(1,1) model combines these two models to represent both the returns of the underlying asset and the market implied volatility for the option. The ARMA(1,1) model estimates the returns, while the GJR-GARCH(1,1) models the variance of the error terms from the ARMA process. The ARMA-GJR-GARCH model serves as the basis for our modeling of both the returns of the underlying asset and the market implied volatility for the option.

Finally, we implement modifications that allow the implied volatility to impact the S&P 500 return volatility process and allow the implied volatility to be dependent on the S&P 500 returns. Our full model, referred to as the Modified ARMA-GJR-GARCH model, is shown in detail in Section 2.2.

2.2 Full Model

The model is stated in terms of S&P 500 with VIX, but the same approach can be taken using single stocks and the market implied volatility for that stock. For more information on definitions of variables and model parameters, see Table C1 in Appendix C.

The complete model is:

$$r_t - \alpha_0 = \alpha_1 (r_{t-1} - \alpha_0) + \epsilon_t + \alpha_2 \epsilon_{t-1} \tag{1}$$

$$\epsilon_t \sim \text{skewed-t}(location = 0, scale = \sigma_t^2, slant = \nu, df = \xi) \tag{2}$$

$$\sigma_t^2 = \beta_0 + \beta_1 \epsilon_{t-1}^2 + \beta_2 \sigma_{t-1}^2 + \beta_3 I_{t-1} \epsilon_{t-1}^2 + \beta_4 v_{t-1}^2 + \beta_5 v_{t-2}^2 \tag{3}$$

$$log(v_t / v_{t-1}) - \theta_0 = \theta_1 (log(v_{t-1}/v_{t-2}) - \theta_0) + \lambda_t + \theta_2 \lambda_{t-1} \tag{4}$$

$$\mu_{\lambda_t} = (\theta_3 \frac{\epsilon_t^2}{\sigma_t^2} + \theta_4 \frac{I_t \epsilon_t^2}{\sigma_t^2} + \theta_5 \frac{\epsilon_t}{\sigma_t} + \theta_6 \frac{I_t \epsilon_t}{\sigma_t}) * \chi_t \tag{5}$$

$$\lambda_t \sim \text{skewed-t}(location = \mu_{\lambda_t}, scale = \chi_t^2, slant = \iota, df = \eta) \tag{6}$$

$$\chi_t^2 = \phi_0 + \phi_1 \lambda_{t-1}^2 + \phi_2 \chi_{t-1}^2 + \phi_3 J_{t-1} \lambda_{t-1}^2 \tag{7}$$

Individual parameters in our variance process, shown in Equations (3) and (7), are not constrained, but the overall variance process is constrained to be positive.

Both return of S&P 500, r_t, and VIX, v_t, are modeled with an ARMA(1,1) process as shown in Equations (1) and (4), and the variance of the S&P 500 return and the VIX are modeled with a GJR-GARCH model shown in Equations (3) and (7). We modify this through the addition of parameters β_4 and β_5 in Equation (3). This allows the VIX to affect the volatility process of the S&P 500. Additionally, the error term of VIX returns can be correlated with the error (θ_5, θ_6) and squared error (θ_3, θ_4) of the S&P 500 return as shown in Equation (5). These additions reflect our team's intuition on how the stock market functions. If the market expects more volatility than the recent S&P500 returns would suggest, there is often a reason, and if there are very volatile returns it would probably also modify the market's expectation of volatility.

In addition to fitting this model, we also fit simpler baseline models for comparison. These consisted of the following: 1) A naive model where S&P 500 returns were normally distributed with a variance equal to the squared VIX on a daily scale 2) An ARMA(1,1)+GARCH(1,1) model for S&P 500 returns and VIX that assumed independence with a normal distribution. 3) An ARMA(1,1)+GJR-GARCH(1,1) model for S&P 500 returns and VIX that assumed independence with a normal distribution. We also fit our full model with different assumptions about the distribution of the error terms. These were independent normal, normal, and t, in addition to the skewed-t described here. The full specifications of these models are described in Appendix B.

Parameter estimation is done with the fGARCH (Wuertz et al. 2013) and ruGARCH (Ghalanos 2019) packages in R (R Core Team 2013) for the ARMA-GARCH and GJR-GARCH models respectively. Our modified GJR-GARCH models are fit using Maximum Likelihood Estimation with Broyden, Fletcher, Goldfarb, and Shanno optimization (Fletcher 1987). In addition, the sn package (Azzalini 2021) is used for the skewed-t distribution.

To ensure global optimality, models were fit to build on each other. For example, when fitting the modified ARMA-GJR-GARCH model with skewed-t-distribution, it was initially fit to optimize the skewness terms with the assumption that all of the other parameters would be equivalent to the parameters in the modified ARMA-GJR-GARCH model with t-distribution. After finding these skewness coefficients, the model was optimized over the entire parameter space using the earlier result as the initial parameters.

The model is trained on data from 2004-2015 and then tested on 2016-2020. For both the in-sample data and out-of sample data we compute daily predictions for return and squared return for both the S&P 500 and the VIX, and estimate the volatility process. We then compute the log-likelihood of the returns given our predictions and estimated volatility.

3. Data Description

We begin with the S&P 500. The S&P 500 tracks the performance of 500 of the biggest US companies and is generally the indicator used to measure how the stock market is performing. We analyze daily closing trading prices for the S&P 500 from 2004-2020 which were obtained from WRDS (Wharton Research Data Services 2021). The VIX data came from the Chicago Board of Exchange (CBOE) (2021) and covered the same time period. Our sample size is 4280 trading days for each dataset and all variables are recorded daily as the price at the end of each day where the market is open (trading day).

The daily values of both the VIX and S&P 500 are depicted in Figure 1. One can see two large spikes in VIX. The first occurred in 2008 during the global financial crisis and the second in 2020 due to COVID-19. Both of these events had corresponding drops in the S&P 500.

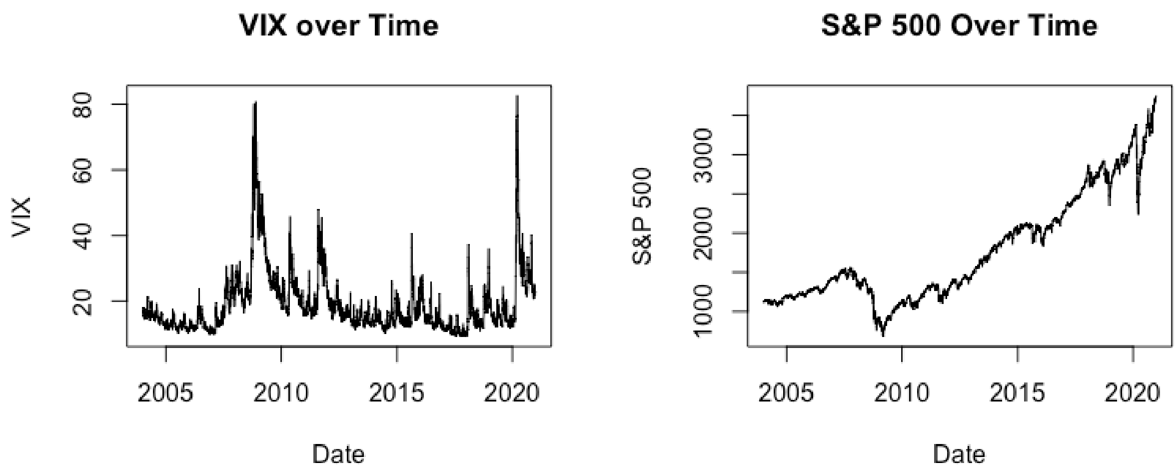

Figure 1: VIX (Volatility Index) and S&P 500 Daily Returns Over Time

4. Results

We computed the Root Mean Squared Error (RMSE) of returns, RMSE of squared returns, and log-likelihood of returns for our final model and the several specifications. These results are shown in Table 1. The models were not able to predict the direction of S&P 500 return successfully as all models do roughly as well as the naive model that assumes a constant mean return. Adding in a leverage term and incorporating both VIX and previous S&P 500 returns improved the RMSE of squared return for the S&P 500 in the training data. The benefit of including VIX for squared S&P 500 returns is less clear in the test data. Log-likelihood of S&P 500 returns went up noticeably in both the training and the test data when increasing model complexity other than adding the error dependency and switching to a skewed-t from a t-distribution. The largest improvements in results came from modeling the dependence of the S&P 500 and VIX error terms as the test RMSE for VIX returns went down by 45%.

The parameter estimates for the modified GJR-GARCH model with errors from a skewed-t distribution are given in Table C1 in Appendix C.

Looking at the model estimates, we find very large positive autocorrelation ($\alpha_1 = 0.965, \theta_1 = 0.874$) and very large negative moving average components ($\alpha_2 = -0.969, \theta_2 = -0.941$) for both the returns of the S&P 500 and the VIX.

The coefficient for new squared error in the GJR-GARCH model for S&P 500 is slightly negative ($\beta_1 = -5.92\text{e-}03$), but the leveraged coefficient for new squared error when error is negative, is positive and noticeably higher ($\beta_3 = 0.0662$). Variance has high autocorrelation ($\beta_2 = 0.895$). If one converts the VIX to be on a daily scale with 252 trading days in a year and in absolute terms instead of percentage (so a VIX of 20 would correspond to a daily standard deviation of $0.2/\sqrt{252}$), the coefficients for the lagged squared VIX in the GJR-GARCH for S&P 500 becomes 0.563 for a one day lag (β_4) and -0.526 for a two day lag (β_5), which are both strong effects.

	Training Data					
	SP500 Return			VIX Returns | SP 500 Return		
	RMSE	RMSE SQ	Log-Likelihood	RMSE	RMSE SQ	Log-Likelihood
Naive Model	0.0122	0.000461	9,804.7	0.0698	N/A	N/A
ARMA+GARCH	0.0121	0.000468	9,870.2	0.0689	0.01130	3,982.3
ARMA+GJR-GARCH	0.0121	0.000457	9,937.9	0.0689	0.01130	4,026.4
Normal, No Error Dep.	0.0122	0.000450	10,013.8	0.0689	0.01130	4,026.2
Normal	0.0122	0.000453	10,003.2	0.0355	0.00651	5,865.1
t-dist	0.0122	0.000456	10,010.7	0.0352	0.00654	5,914.3
Skewed-t	0.0122	0.000456	10,014.8	0.0350	0.00660	5,930.1
	Test Data					
	SP500 Return			VIX Return | SP 500 Return		
	RMSE	RMSE SQ	Log-Likelihood	RMSE	RMSE SQ	Log-Likelihood
Naive Model	0.0121	0.000623	4,228.1	0.0825	N/A	N/A
ARMA+GARCH	0.0121	0.000606	4,265.7	0.0823	0.0224	1,419.1
ARMA+GJR-GARCH	0.0121	0.000595	4,268.2	0.0823	0.0229	1,444.2
Normal, No Error Dep.	0.0121	0.000598	4,345.7	0.0824	0.0228	1,447.7
Normal	0.0121	0.000601	4,340.4	0.0452	0.0154	2,118.0
t-dist	0.0121	0.000601	4,376.6	0.0447	0.0151	2,204.2
Skewed-t	0.0121	0.000602	4,375.8	0.0448	0.0162	2,202.4

Table 1: Training Data Results using modeled prices without corrections. Results show Root Mean Squared Error (RMSE) for returns and squared returns. Additionally log-likelihood is calculated when trying to optimize function.

The relationship between the error for the VIX and S&P 500 is also shown in Figure 2. A low return is correlated with an increase in the VIX. This effect diminishes as the return becomes higher.

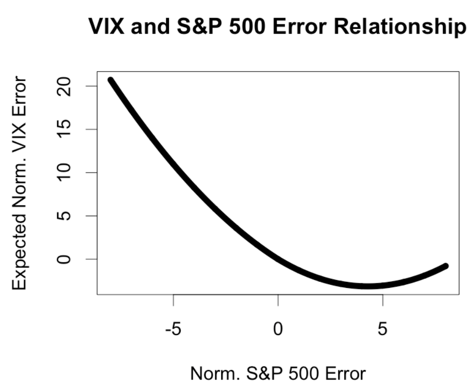

Figure 2: Mean for Skewed-t Distribution of Volatility Index (VIX) Error Given S&P 500 Error. S&P 500 error is divided by modeled daily volatility in S&P 500 and VIX is divided by modeled daily volatility in VIX

The GJR-GARCH model for VIX shows that an upward move in the VIX increases the volatility of the VIX ($\phi_1 = 0.0433$), but a large downwards move has a considerably smaller effect ($\phi_1 + \phi_3 = 0.027$). The volatility of the VIX also has strong autocorrelation ($\phi_2 = 0.766$), but not as strong as the autocorrelation of volatility for the S&P 500.

For the skewed t-distribution, we see that S&P 500 returns have a slightly positive skew ($\nu = 0.0015$) and VIX a negative skew ($\iota = -0.273$). Both have degrees of freedom of roughly 10 ($\xi = 10.7, \eta = 8.84$).

In Figure 3, we show our model's estimate for daily variance of S&P 500 returns and compare this to the daily VIX. The VIX and our variance process are closely related, but they occasionally diverge.

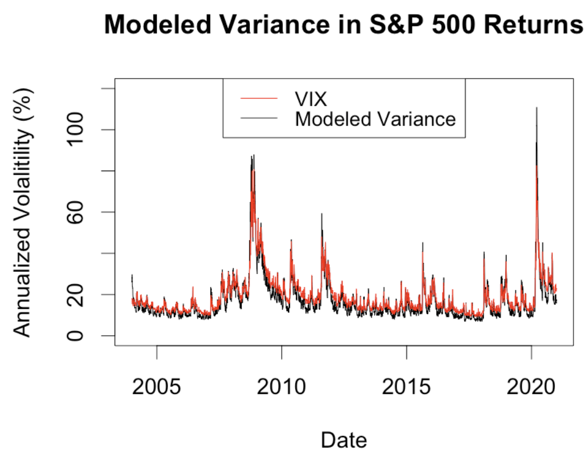

Figure 3: Modeled Variance of Daily Returns of S&P 500 vs. VIX Overlay

5. Discussion

Looking at the results of the different models on predicting returns of S&P 500 we see that none of the models improve from the naive model by more than 1% in RMSE. This fits with what we would expect to find if the Efficient Market Hypothesis were true. While the Efficient Market Hypothesis may not fully hold, it does help explain why using a simple ARMA process would be unable to generate large improvements in the RMSE. We do see increases in log-likelihood of S&P 500 from incorporating the VIX into the GJR-GARCH, but the further complications did not improve the likelihood by more than 1 in the training dataset and 31 in the test dataset, showing that our dependence modeling improved the estimates of the VIX, but not the estimates of the S&P 500, which is unsurprising given that the calculations of the VIX errors were conditional on knowing the S&P 500 returns and not the other way around. On the other hand, we see large effects of allowing for fat tails in the volatility. We also see, unsurprisingly, that the conditional distribution for VIX is improved by knowing the return. This is not information that is directly tradeable given that the current S&P 500 return is not known until the VIX return is known, but is still valuable in order to better understand how volatility changes over time.

Adding a leveraged term (β_3 and ϕ_3) produced improvements in the squared return RMSE values for the S&P 500 but not the VIX, despite increases in likelihood for both. We generally see strong increases in our likelihood estimations from our model when compared to less complex models. This implies that our model is better than naive models at predicting the full distribution of possible returns from the next trading day. This is pertinent information for traders looking to profit from options. Furthermore, our largest increases in performance came from including the dependence of the error terms. Clearly it is imperative to model the error terms dependently, as the S&P500 and VIX tend to move together over time, as highlighted in Figure 2.

The strong autocorrelation (α_1, θ_1) and moving average (α_2, θ_2) terms for returns of VIX and S&P 500 despite the lack of a strong RMSE from the ARMA+GARCH model makes it clear that despite the terms being high in magnitude, they do not add substantial predictive ability. This is actually incorporated into the model as the AR and MA components have opposite signs and largely cancel each other out.

We found a negative coefficient for new squared error of positive returns (β_1) in the GJR-GARCH model of the S&P 500. This result is surprising, but we believe that this can be explained with market overreaction. This means that when a large catalyst occurs, investors overreact and this shock is incorporated into the VIX which gets integrated into our variance model. If investors overreact, then having a negative coefficient for new squared error of positive returns can correct for the overreaction. On the other hand it is possible that large positive returns simply lead to lower volatility in the future. Similarly, we find that large downward movement in VIX does not change volatility of VIX, while large upwards moves tend to be directly correlated with increases in volatility of VIX. This shows that negative events in the stock market have a much larger impact on volatility than positive events.

We also see that the skewed t-distribution for S&P 500 returns has right skew (ν) and has fat tails (ξ). Seeing right skew is quite surprising, given that the skewness of the modeled error terms is negative and intuition generally suggests returns are left skewed. Despite this, the model still gives a better likelihood than using an unskewed t-distribution in the training data. This effect is marginal and does not hold for the test data. It is still somewhat surprising to see that the modeled skew parameter for the S&P 500 is not a strong negative. This may be due to differences in the skew between the tails and the rest of events. In general, it appears that modeling skew makes a smaller difference than we may intuitively expect.

Future research would work on making sure that our model is well calibrated at predicting the distribution of longer term returns and on incorporating our model into an options pricing formula similar to the Black Scholes model and building a profitable trading strategy with that.

In conclusion, our modified GJR-GARCH model provides valuable insight into the challenges of volatility prediction. To generalize, we find that our model is not able to predict the direction of the S&P 500 returns, however, we are able to observe that large negative moves in the S&P 500 returns increase both the VIX and true volatility of returns. Most importantly, our model better captures how the market moves over time and is a useful research tool for options traders who wish to estimate future volatility.

References

Azzalini, A.(2015). Package 'sn'. The skew-normal and skew-t distributions, 1-3. URL https://cran.r-project.org/web/packages/sn/sn.pdf.

Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. Journal of Political Economy, 81(3), 637-654.

Cao, Hongkai & Badescu, Alex & Cui, Zhenyu & Jayaraman, Sarath. (2020). Valuation of VIX and Target Volatility Options with Affine GARCH Models. Journal of Futures Markets.

CBOE. (2021). CBOE Global Markets Data. URL https://www.cboe.com/market_data_services/.

Corrado, C. J., & Su, T. (1996). Skewness and kurtosis in S&P 500 index returns implied by option prices. Journal of Financial Research, 19(2), 175-192.

Ensor, K.B., Han, Y., Ostdiek, B. et al. (2020). Dynamic jump intensities and news arrival in oil futures markets. Journal of Asset Management 21, 292–325.

Ensor, K. B., & Koev, G. M. (2014). Computational finance: correlation, volatility, and markets. Wiley Interdisciplinary Reviews: Computational Statistics, 6(5), 326-340.

Fleming, J., Ostdiek, B., & Whaley, R. E. (1995). Predicting stock market volatility: A new measure. Journal of Futures Markets, 15(3), 265-302.

Fletcher, R. (1987). Practical methods of optimization. John Wiley & Sons, New York.

Ghalanos, A. (2019). Package 'rugarch'. URL https://cran.r-project.org/web/packages/rugarch/rugarch.pdf.

Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. The Journal of Finance, 48(5), 1779-1801.

Kanniainen, J., Lin, B., & Yang, H. (2014). Estimating and using GARCH models with VIX data for option valuation. Journal of Banking & Finance, 43, 200-211.

Lin, Y. N., & Chang, C. H. (2010). Consistent modeling of S&P 500 and VIX derivatives. Journal of Economic Dynamics and Control, 34(11), 2302-2319.

Malkiel, B. G., & Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. The Journal of Finance, 25(2), 383-417.

Pisani, B. (2019). Active fund Managers trail the S&P 500 for the ninth year in a row in triumph for indexing. CNBC.

Roondiwala, M., Patel, H., & Varma, S. (2017). Predicting stock prices using LSTM. International Journal of Science and Research (IJSR), 6(4), 1754-1756.

R Core Team (2013). R: A language and environment for statistical computing.

Wang, Y. H. (2009). Nonlinear neural network forecasting model for stock index option price: Hybrid GJR–GARCH approach. Expert Systems with Applications, 36(1), 564-570.

Wharton Research Data Services (2021). WRDS. URL https://wrds-www.wharton.upenn.edu/.

Williams, E. E., & Dobelman, J. A. (2020). Random Walk To Nowhere, A: How The Professors Caused A Real 'Fraud-on-the-market'. World Scientific.

Wuertz, D., Setz, T., Chalabi, Y., Boudt, C., Chausse, P., & Miklovac, M. (2013). Package 'fGarch'. Technical report, working paper/manual, 09.11. 2009. URL http://cran.r-project.org/web/packages/fGarch/fGarch.pdf.

Appendix

Appendix A: Terminology

Term	Definition
Annualized Volatility	The variation in an asset's value over the course of a year (used as a risk measure for an investment).
Expiration (for options)	The last day you can choose to exercise your right to buy (call options) or sell (put options) an underlying asset at the strike price.
Futures	Futures are an example of a derivative and represent an agreement between a buyer and seller to transact an asset at a predetermined future date and price. Usually, these contracts are for commodities such as oil, corn, sugar, gas, etc.
Exchange Traded Fund (ETF)	A fund that is designed to follow rules to track a certain group of underlying investments
Implied Volatility (IV)	The market's expectation of a stock of ETFs volatility over the next year for a 1 standard deviation move.
Options Contracts	An agreement between a buyer and a seller to give the buyer an opportunity the right to buy or sell the underlying asset at a strike price on or before a specific date.
Risk Free Interest Rate	The theoretical rate of return of an investment with 0 risk (the minimum return an investor expects for any investment).
Standard & Poors 500 index (a.k.a the S&P, S&P 500, or SPX)	A market-cap weighted index of the largest 500 publicly traded companies.
Strike Price	The price at which a put or call option can be exercised.
Volatility Index (VIX)	A market index representing the market's expectations for volatility in the next 30 days.

Table A1: Important Terminology Introduced in this Paper.

Appendix B: Models

Naive

$$
\begin{aligned}
r_t - \alpha_0 &= \epsilon_t \\
\epsilon_t &\sim N(location = 0, scale = \frac{(\frac{v_{t-1}}{100})^2}{252}) \\
log(v_t/v_{t-1}) &= \lambda_t
\end{aligned}
$$

ARMA+GARCH

$$
\begin{aligned}
r_t - \alpha_0 &= \alpha_1(r_{t-1} - \alpha_0) + \epsilon_t + \alpha_2\epsilon_{t-1} \\
\epsilon_t &\sim N(location = 0, scale = \sigma_t^2) \\
\sigma_t^2 &= \beta_0 + \beta_1\epsilon_{t-1}^2 + \beta_2\sigma_{t-1}^2 \\
log(v_t/v_{t-1}) - \theta_0 &= \theta_1(log(v_{t-1}/v_{t-2}) - \theta_0) + \lambda_t + \theta_2\lambda_{t-1} \\
\lambda_t &\sim N(location = 0, scale = \chi_t^2) \\
\chi_t^2 &= \phi_0 + \phi_1\lambda_{t-1}^2 + \phi_2\chi_{t-1}^2
\end{aligned}
$$

ARMA+GJR-GARCH

$$
\begin{aligned}
r_t - \alpha_0 &= \alpha_1(r_{t-1} - \alpha_0) + \epsilon_t + \alpha_2\epsilon_{t-1} \\
\epsilon_t &\sim N(location = 0, scale = \sigma_t^2) \\
\sigma_t^2 &= \beta_0 + \beta_1\epsilon_{t-1}^2 + \beta_2\sigma_{t-1}^2 + \beta_3 I_{t-1}\epsilon_{t-1}^2 \\
log(v_t/v_{t-1}) - \theta_0 &= \theta_1(log(v_{t-1}/v_{t-2}) - \theta_0) + \lambda_t + \theta_2\lambda_{t-1} \\
\lambda_t &\sim N(location = 0, scale = \chi_t^2) \\
\chi_t^2 &= \phi_0 + \phi_1\lambda_{t-1}^2 + \phi_2\chi_{t-1}^2 + \phi_3 J_{t-1}\lambda_{t-1}^2
\end{aligned}
$$

Normal No Error Dependence

$$
\begin{aligned}
r_t - \alpha_0 &= \alpha_1(r_{t-1} - \alpha_0) + \epsilon_t + \alpha_2\epsilon_{t-1} \\
\epsilon_t &\sim N(location = 0, scale = \sigma_t^2) \\
\sigma_t^2 &= \beta_0 + \beta_1\epsilon_{t-1}^2 + \beta_2\sigma_{t-1}^2 + \beta_3 I_{t-1}\epsilon_{t-1}^2 + \beta_4 v_{t-1}^2 + \beta_5 v_{t-2}^2 \\
log(v_t/v_{t-1}) - \theta_0 &= \theta_1(log(v_{t-1}/v_{t-2}) - \theta_0) + \lambda_t + \theta_2\lambda_{t-1} \\
\lambda_t &\sim N(location = 0, scale = \chi_t^2) \\
\chi_t^2 &= \phi_0 + \phi_1\lambda_{t-1}^2 + \phi_2\chi_{t-1}^2 + \phi_3 J_{t-1}\lambda_{t-1}^2
\end{aligned}
$$

Normal

$$
\begin{aligned}
r_t - \alpha_0 &= \alpha_1(r_{t-1} - \alpha_0) + \epsilon_t + \alpha_2\epsilon_{t-1} \\
\epsilon_t &\sim N(location = 0, scale = \sigma_t^2) \\
\sigma_t^2 &= \beta_0 + \beta_1\epsilon_{t-1}^2 + \beta_2\sigma_{t-1}^2 + \beta_3 I_{t-1}\epsilon_{t-1}^2 + \beta_4 v_{t-1}^2 + \beta_5 v_{t-2}^2 \\
log(v_t/v_{t-1}) - \theta_0 &= \theta_1(log(v_{t-1}/v_{t-2}) - \theta_0) + \lambda_t + \theta_2\lambda_{t-1} \\
\mu_{\lambda_t} &= (\theta_3\frac{\epsilon_t^2}{\sigma_t^2} + \theta_4\frac{I_t\epsilon_t^2}{\sigma_t^2} + \theta_5\frac{\epsilon_t}{\sigma_t} + \theta_6\frac{I_t\epsilon_t}{\sigma_t}) * \chi_t \\
\lambda_t &\sim N(location = \mu_{\lambda_t}, scale = \chi_t^2) \\
\chi_t^2 &= \phi_0 + \phi_1\lambda_{t-1}^2 + \phi_2\chi_{t-1}^2 + \phi_3 J_{t-1}\lambda_{t-1}^2
\end{aligned}
$$

t

$$
\begin{aligned}
r_t - \alpha_0 &= \alpha_1(r_{t-1} - \alpha_0) + \epsilon_t + \alpha_2\epsilon_{t-1} \\
\epsilon_t &\sim t(location = 0, scale = \sigma_t^2, df = \xi) \\
\sigma_t^2 &= \beta_0 + \beta_1\epsilon_{t-1}^2 + \beta_2\sigma_{t-1}^2 + \beta_3 I_{t-1}\epsilon_{t-1}^2 + \beta_4 v_{t-1}^2 + \beta_5 v_{t-2}^2 \\
log(v_t/v_{t-1}) - \theta_0 &= \theta_1(log(v_{t-1}/v_{t-2}) - \theta_0) + \lambda_t + \theta_2\lambda_{t-1} \\
\mu_{\lambda_t} &= (\theta_3\frac{\epsilon_t^2}{\sigma_t^2} + \theta_4\frac{I_t\epsilon_t^2}{\sigma_t^2} + \theta_5\frac{\epsilon_t}{\sigma_t} + \theta_6\frac{I_t\epsilon_t}{\sigma_t}) * \chi_t \\
\lambda_t &\sim t(location = \mu_{\lambda_t}, scale = \chi_t^2, df = \eta) \\
\chi_t^2 &= \phi_0 + \phi_1\lambda_{t-1}^2 + \phi_2\chi_{t-1}^2 + \phi_3 J_{t-1}\lambda_{t-1}^2
\end{aligned}
$$

C: Additional model parameters and definitions

Symbol	Description	Estimate
t	Day in the Dataset. Only Includes Trading Days	N/A
r_t	S&P500 log Return on Day t	N/A
v_t	VIX on Day t	N/A
ϵ_t	Error Term for Daily S&P500 Return	N/A
σ_t^2	Daily Variance in Error Term for Daily S&P500 Return	N/A
λ_t	Error Term for Daily VIX Return	N/A
χ_t^2	Daily Variance in Error Term for Daily VIX Return	N/A
I_t	Indicator Function that Equals 1 if $\epsilon_t < 0$	N/A
J_t	Indicator Function that Equals 1 if $\lambda_t < 0$	N/A
μ_{λ_t}	Conditional Location for VIX Residual	N/A
α_0	Constant for Daily S&P500 Return	2.79e-04
α_1	Autoregressive Parameter for Daily S&P500 Return	0.965
α_2	Moving Average Parameter for Daily S&P500 Return	-0.969
β_0	Constant in GJR-GARCH for Daily S&P500 Return	2.03e-07
β_1	New Squared Error Component in GJR-GARCH for Daily S&P500 Return	-5.92e-03
β_2	Former Variance Component in GJR-GARCH for Daily S&P500 Return	0.895
β_3	Coefficient for Indicator Parameter in GJR-GARCH for Daily S&P500 Return	0.0662
β_4	Coefficient for Previous Day's Squared VIX in GJR-GARCH for Daily S&P500 Return	2.23e-07
β_5	Coefficient for Squared VIX from Two Days Earlier in GJR-GARCH for Daily S&P500 Return	-2.09e-07
θ_0	Constant for Daily VIX Return	-3.80e-05
θ_1	Autoregressive Parameter for Daily VIX Return	0.874
θ_2	Moving Average Parameter for Daily VIX Return	-0.941
θ_3	Coefficient for Today's Squared S&P500 Error for Daily VIX Return	0.171
θ_4	Coefficient for Indicator times Today's Squared S&P500 Error for Daily VIX Return	-0.0387
θ_5	Coefficient for Today's S&P500 Error for Daily VIX Return	-1.46
θ_6	Coefficient for Indicator times Today's S&P500 Error for Daily VIX Return	-0.0694
ϕ_0	Constant in GJR-GARCH for Daily VIX Return	1.31e-04
ϕ_1	New Squared Error Component in GJR-GARCH for Daily VIX Return	0.0433
ϕ_2	Former variance Component in GJR-GARCH for Daily VIX Return	0.766
ϕ_3	Coefficient for Indicator parameter in GJR-GARCH for Daily VIX Return	-0.0406
ν	Skew in Error Term for Daily S&P500 Return	1.50e-03
ξ	Degrees of Freedom Parameter in Error Term for Daily S&P500 Return	10.7
ι	Skew in Error Term for Daily VIX Return	-0.273
η	Degrees of Freedom Parameter in Error Term for Daily VIX Return	8.84

Table C1: Model Parameter Definitions and Estimates.