NEKN34: Time Series Analysis
Spring Semester 2026

Lund University
School of Economics and Management

THE GARCH OF THE RISING SUN: ESTIMATING VOLATILITY IN YEN-USD

Alexander Eriksson Byström, Anselm Etzelmüller, Juan Sepulveda Odriozola

Replication Code: https://github.com/Alexerby/NEKN34-A2-Submission

1. Introduction

The Yen-Dollar currency pair is the second-most traded on in the world (Hinton (2023)) and has important implications for the world economy, some of which are presented in the following.

In the last decades, Japan has consistently maintained a lower interest rate than other countries. As Shi and Pande (2024) explain, this enabled investors to borrow Yen at interest rates close to zero, convert it to other currencies (e.g. the US-Dollar) and invest it in assets with higher yields. After realizing the returns, they are converted back into Yen to repay the loan. This is known as the Yen Carry Trade. If the Yen suddenly strengthens, global investors participating in the Carry Trade would have to sell their assets quickly to be able to repay their Yen loans, which could trigger massive sell-offs of assets worldwide.

Japan is also considered the world's largest creditor nation, due to Japanese investors holding much international debt and treasuries (Ito 2022). A particularly weak Yen could incentivize Japanese investors to pull their money out of foreign markets in order to invest it domestically, which could cause distress on markets in the US and other countries.

Furthermore, Japan is home to many important firms that export their products to the US and other countries, for instance car manufacturers such as Honda and Toyota, or producers of electronical devices, e.g. Sony and Nintendo. A weak Yen makes their products cheaper in foreign countries, while their prices would increase with a stronger Yen.

These global economic links put the Dollar-Yen exchange rate in the spotlight. To analyse its volatility is an important task for investors and economists, as it has important implications on current and future assets and goods prices. We add to the literature that is devoted to the heteroskedasticity of the Yen-Dollar exchange rate by conducting a wide replication of Tse (1998) and Tsui and Ho (2004). Both articles try to evaluate which time series model is best-fitting to the daily exchange rate of the two currencies. We aim to replicate their findings and extend them to a more current data sample.

The following sections are structured as follows: In Section 2, we present the data samples that we intend to use for the replication of Tse (1998) and Tsui and Ho (2004) as well as for our extended model, and provide descriptive statistics about them. Additionally, we provide the results of diagnostics tests about the underlying models of the mean. In Section 3, the model specifications for our replications and extension are specified. Section 4 presents the estimation results for all three datasets, hence the results of the replications and our extension. Section 5 concludes.

2. Data & Diagnostics

Our dataset is comprised of the daily exchange rate of the Japanese Yen to the US-Dollar (yen per dollar) from the 4^{th} of January 1971 to the 20^{th} of November 2023. This enables us to recreate the datasets used by Tse (1998) as well as Tsui and Ho (2004).

Dataset I replicates the dataset used by Tse (1998) and ranges from March 1978 to June 1994. Dataset II recreates the sample from Tsui and Ho (2004), thus ranging from March 1986 to February 2003. We thereafter extend their analysis to a data sample comprising of the remaining period of data, hence from March 2003 to November 2023, hereafter referred to as "Extended".

Table 1: Summary of Dataset Metadata and Sample Periods

ID	Start Date	End Date	Obs (T)
Dataset I	1978-01-04	1994-06-29	4140
Dataset II	1986-01-03	2003-02-21	4308
Extended	2003-01-03	2023-11-20	5637

Note 1: Start and end dates represent the available log-return series.

When working with exchange rate data, we expect time-varying volatility—specifically volatility clusters—and do not assume variance to remain constant over time, a phenomenon which is shown in Figure 1. This aligns with the empirical consensus that financial returns rarely are normally distributed, and that currencies can be considered financial assets. To check whether this actually applies to our data, we first examine the distributional characteristics in Table 2.

Each sample is negatively skewed and exhibits a kurtosis substantially higher than that of a normal distribution[1], thus indicating excess kurtosis and a leptokurtic distribution. This already indicates that our data samples are not normally distributed. To prove this formally, we run a Jarque-Bera test that tests for normality based on skewness and kurtosis where the outcome is a rejection of normality at the 1% level (Table 2). One effect of a non-normal distribution, especially excess kurtosis, is that t-tests may become inconsistent. However, large sample sizes—which we undoubtedly have—renders this concern moot. Having proved non-normality of the data, this provides a clear mandate for the use of non-Gaussian GARCH specifications which requires Quasi-Maximum Likelihood Estimation (QMLE) for the parameter estimates to remain consistent.

The next step is to assess whether conditional heteroskedasticity is present in the data. If no such effects are detected, the use of GARCH-type models is not justified. In that case, the variance can be treated as constant and estimated directly as σ^2 over the sample $\{t \in [-1, \ldots, -p]\}$, which can then serve as the forecast $\hat{\sigma}_t^2$.

To formally test for conditional heteroskedasticity, we employ Engle's (1982) test, commonly referred to as the ARCH-LM test. Prior to its introduction, traditional econometric models typically assumed

[1]That is $k = 3$.

a constant one-period-ahead forecast variance, an assumption that is generally relaxed in modern econometric practice.

Implementation of the test requires first specifying an appropriate autoregressive model in order to obtain residuals. The squared residuals $\hat{\epsilon}_t^2$ are then regressed on a constant and q of their own lagged values. In volatility modeling, such as in GARCH frameworks, the specification of the conditional mean primarily serves as a pre-filter rather than being the main object of interest.

One approach to selecting an appropriate model is to examine the ACF and PACF (see Figure 2), which indicate that an autoregressive model with approximately four lags may be suitable. However, adhering to the principle of parsimony, we instead opt for a more concise AR(1) specification as the estimated coefficient in the AR(1) model is found to be statistically significant based on a t-test (see Table 2). In addition, an augmented Dickey–Fuller test rejects the null hypothesis of a unit root at the 1% significance level, supporting the stationarity of the series.

As reported in Table 2, the Ljung-Box test fails to reject the null hypothesis of no residual autocorrelation for Dataset II and the Extended sample, confirming that the AR(1) filter successfully yields white-noise residuals. The choice of $Q(5)$ corresponds to one business week of trading activity, to ensure that the AR(1) filter has sufficiently captured the short-term linear dependencies in the return series. While the test is significant at the 5% level for Dataset I—suggesting some remaining linear dependence—we keep the AR(1) specification across all samples to ensure methodological consistency and parsimony, following the precedent established by Tse (1998) and Tsui and Ho (2004).

Having specified the AR(1) process for the conditional mean, we utilize the resulting residuals to perform the ARCH-LM test. The null hypothesis of homoskedasticity is strongly rejected at the 1% significance level for all datasets, confirming the presence of ARCH effects. This result indicates that the conditional variance is time-varying and justifies the use of GARCH-family models to capture the observed volatility clustering.

Table 2: Descriptive Statistics and Mean Model Diagnostics

ID	Mean	Std Dev	Skew	Kurt	AR(1)	JB-Stat	ARCH-LM	Q(5)
Dataset I	-0.021	0.663	-0.448	6.434	0.0356**	2173.03***	125.73***	12.06**
Dataset II	-0.012	0.708	-0.456	6.837	0.0385**	2791.76***	265.58***	9.21
Extended	0.004	0.597	-0.356	8.494	-0.0248*	7207.85***	243.67***	9.15

Note 1: *Note:* ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.
Note 2: Returns are defined as $r_t = (\ln S_t - \ln S_{t-1}) \times 100$.

3. Model Specifications

This section is organized into three parts. Section 3.1 outlines the specifications estimated for Dataset I, following the modelling framework of Tse (1998). Section 3.2 presents the corresponding specifications for Dataset II and evaluates the extent to which the original results can be replicated. Finally, Section 3.3 introduces our extended modelling framework, which builds upon the baseline specifications and incorporates long-memory dynamics.

In line with the literature (Tse 1998; Tsui and Ho 2004) we will be specifying our variable of interest (JPY/USD) in terms of its log-transformation

$$r_t = \ln\left(\frac{S_t}{S_{t-1}}\right) \times 100 \tag{1}$$

where S_t denotes nominal returns in period t for the remainder of this assignment.

3.1 Dataset: I

Table 3 reports the model specifications from Tse (1998) that we seek to replicate. We employ the `arch` library by Sheppard (2025); however, its FIGARCH implementation does not support the leverage parameter γ and power parameter δ required for the asymmetric and power components of the FIAPARCH model. Consequently, we restrict our estimation to the symmetric fractionally integrated specifications, omitting FIAPARCH(1, d, 1) from our replication. This limitation extends to the fractionally integrated asymmetric specifications encountered in Section 3.2, where FIAGARCH(1, d, 1) and FIAPARCH(1, d, 1) are similarly excluded.

Table 3: Variance Model Specifications for Tse (1998) replication

Model Category	Specification Form (h_t or σ_t^δ)
Stable models	
GARCH(1,1)	$\eta + \alpha\epsilon_{t-1}^2 + \beta h_{t-1}$
APARCH(1,1)	$\eta + \alpha(\lvert\epsilon_{t-1}\rvert - \gamma\epsilon_{t-1})^\delta + \beta\sigma_{t-1}^\delta$
Fractionally integrated	
FIGARCH(1,d,1)	$\eta + [1-(1-\beta L)^{-1}(1-\phi L)(1-L)^d]\epsilon_t^2$
~~FIAPARCH(1,d,1)~~	$\cancel{\eta + [1-(1-\beta L)^{-1}(1-\phi L)(1-L)^d](\lvert\epsilon_t\rvert - \gamma\epsilon_t)^\delta}$

Note: Parameters are estimated via QMLE assuming a normal likelihood. Standard errors are calculated using the "robust" covariance estimator (White's estimator) provided by the `arch` package to ensure consistency under non-Gaussian distributions. Consistent with the Python implementation, η and η denote the variance intercepts, while ϕ and β represent the lag polynomials.

As Tse (1998) does not explicitly state the estimation method used for the FIAGARCH(1, d, 1) specification, we assume that it was estimated via QMLE, consistent with the other variance models considered. This choice also aligns with the robustness of QMLE under non-normality.

3.2 Dataset: II

Table 4 reports the model specifications from Tsui and Ho (2004) that we seek to replicate. As with the specifications from Tse (1998), we encounter limitations with the asymmetric GARCH specifications for fractionally integrated models, and therefore omit the FIAGARCH(1, d, 1) and FIAPARCH(1, d, 1) specifications from our replication exercise. The reasoning behind this decision is discussed in Section 3.1.

TABLE 4: VARIANCE MODEL SPECIFICATIONS FOR TSUI AND HO (2004) REPLICATION

Model Category	Specification Form (h_t or σ_t^δ)
Short-Memory (Stable)	
AGARCH(1,1)	$\eta + \alpha(\epsilon_{t-1} - \gamma)^2 + \beta h_{t-1}$
APARCH(1,1)	$\sigma_t^\delta = \eta + \alpha(\|\epsilon_{t-1}\| - \gamma\epsilon_{t-1})^\delta + \beta\sigma_{t-1}^\delta$
Long-Memory (Fractionally Integrated)	
~~FIAGARCH(1,d,1)~~	$\cancel{h_t = \frac{\eta}{1-\beta} + \lambda(L)(\epsilon_t - \gamma)^2}$
~~FIAPARCH(1,d,1)~~	$\cancel{\sigma_t^\delta = \frac{\eta}{1-\beta} + \lambda(L)(\|\epsilon_t\| - \gamma\epsilon_t)^\delta}$

Note: For all fractionally integrated models, the lag polynomial is defined as $\lambda(L) = 1 - (1 - \beta L)^{-1}(1 - \phi L)(1 - L)^d$. The parameter γ represents the asymmetry (leverage) effect, and δ represents the power transformation parameter.

3.3 Dataset: Extended

Our baseline model is the GARCH(1,1) of Bollerslev (1986), specified alongside the AR(1) mean equation as

$$r_t = \gamma r_{t-1} + \epsilon_t, \qquad \epsilon_t \mid \psi_{t-1} \sim N(0, h_t), \tag{2}$$
$$h_t = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta_1 h_{t-1}, \tag{3}$$

where h_t is the conditional variance. The GARCH(1,1) parsimoniously captures volatility persistence without the long lag structures required by pure ARCH models. Positivity of h_t is ensured by $\alpha_0 > 0$, $\alpha_1, \beta_1 \geq 0$, and the process is covariance stationary when $\alpha_1 + \beta_1 < 1$, under which the unconditional variance is $\sigma^2 = \alpha_0/(1 - \alpha_1 - \beta_1)$. While simple, it serves as a useful benchmark, though we expect it to be too short-sighted given the long-memory dynamics visible in Figure 2, as it assumes shocks decay exponentially rather than hyperbolically.

We therefore turn to FIGARCH(1, d, 1) (Baillie, Bollerslev, and Mikkelsen 1996), which replaces exponential with hyperbolic decay via fractional integration, allowing shocks to persist over much longer horizons. The $(1, d, 1)$ order is the most parsimonious specification that still allows for both an ARCH and a GARCH term alongside the fractional differencing parameter d, and is the specification used by both Tse (1998) and Tsui and Ho (2004) in the papers we replicate. The FIGARCH model was originally developed by Baillie, Bollerslev, and Mikkelsen (1996) using the DEM/USD exchange rate,[2] making it a natural fit for our JPY/USD setting.

We do not consider asymmetric specifications such as APARCH or FIAPARCH for our extension. In stock markets, negative shocks typically have longer-lasting effects than positive ones, but this asymmetric reaction generally cannot be confirmed for exchange rates (Ma and Kao 1990; Baillie, Bollerslev, and Mikkelsen 1996; Zhelyazkova 2024), as both an appreciation and depreciation of a currency can have positive and negative effects simultaneously[3]. As APARCH and FIAPARCH

[2]DEM is the currency code for the German mark ("Deutsche Mark").

[3]This notion was confirmed by Tse (1998). Tsui and Ho (2004) found evidence for small asymmetric effects – although associated with comparatively high standard errors – in some models, and none in others.

models specifically model asymmetric reactions to positive and negative shocks, their usage is not justified here.

Finally, our descriptive statistics in Table 2 show that all three samples exhibit excess kurtosis, suggesting that the normality assumption for ϵ_t may be too restrictive. We therefore estimate each model under three distributional assumptions: Gaussian, Student's t, and GED.

Model evaluation will thereafter be performed in two steps:

- **Information criteria:** We report AIC and BIC for all six specifications, as is standard in the GARCH literature.
- **Stability over time:** We compare our Gaussian GARCH(1,1) and FIGARCH(1, d, 1) estimates against the findings of Tse (1998) and Tsui and Ho (2004) to assess whether the volatility dynamics of the JPY/USD exchange rate have remained stable over time.

Table 5: Extended Dataset Model Specifications

Model	Distribution	Notes
Short-Memory		
GARCH(1,1)	Gaussian	Baseline
GARCH(1,1)	Student's t	Fat-tailed errors
GARCH(1,1)	GED	Fat-tailed errors
Long-Memory		
FIGARCH(1, d, 1)	Gaussian	Hyperbolic decay
FIGARCH(1, d, 1)	Student's t	Fat-tailed errors
FIGARCH(1, d, 1)	GED	Fat-tailed errors

Note: All models are estimated via QMLE alongside an AR(1) mean equation. Model selection is based on AIC and BIC. Asymmetric specifications are excluded on theoretical grounds, as exchange rates do not exhibit the asymmetric shock response documented in equity markets.

4. Estimation Results

4.1 Dataset I

Table 6 reports the replication results for the dataset used in Tse (1998). For all our GARCH-specifications (`GARCH`, `APARCH`, `FIGARCH`) we have found close matches to the original findings in Tse (1998). The following are our key observations for all three specifications:

- **Magnitude matching:** All of our estimates match closely with the original results in Tse (1998). All of our estimates match to the third decimal except the mean (μ) for the mean estimating model ($AR(1)$) where we have a match to the second decimal.
- **Sign matching:** All estimates match in sign, except μ which does not match due to the low magnitude of the parameter.
- **Standard Errors:** All standard errors match to the third decimal.

Table 6: Replication Results for Dataset I (1978-1994)

	GARCH model	APARCH model	FIGARCH model
μ	-0.0112	-0.0101	-0.0098
	(0.0097)	(0.0097)	(0.0098)
ρ	0.0303	0.0307	0.0319
	(0.0177)	(0.0175)	(0.0181)
η	0.0280	0.0306	0.0605
	(0.0096)	(0.0115)	(0.0290)
α	0.0839	0.0891	—
	(0.0203)	(0.0220)	—
β_G	0.8540	0.8599	—
	(0.0364)	(0.0357)	—
γ	—	-0.0174	—
	—	(0.0754)	—
δ	—	1.7292	—
	—	(0.4234)	—
β_F	—	—	0.3770
	—	—	(0.1796)
d	—	—	0.2297
	—	—	(0.0531)
ϕ	—	—	0.2663
	—	—	(0.1669)

Note 1: Standard errors in parentheses are Bollerslev and Ole Mikkelsen (1996) robust standard errors.
Note 2: All models estimated via QMLE assuming a normal likelihood.

In the GARCH(1, 1) and APARCH(1, 1) specifications, we find the estimated persistence $(\alpha + \beta)$ to be approximately 0.9379 and 0.9490, respectively. These values, being close to unity, are consistent with the findings in Tse (1998) and suggest that the shocks to JPY/USD volatility are highly persistent. Our conclusions are that the series remains stationary ($I(0)$), but is better characterized by models that have longer memory—that is, the fractionally integrated ones.

4.2 Dataset II

Table 7 reports our replication results for Tsui and Ho (2004), where Dataset II is used. For the APARCH and AGARCH models the following are our main conclusions

- For APARCH, all of our estimates match the findings of Tsui and Ho (2004) to the second or third decimal. The sole notable difference is in the power parameter δ, which differs by 0.1304 relative to Tsui and Ho (2004). Given that our estimate remains statistically significant, we attribute this to differences in the optimizer used or the convergence criteria employed, rather than estimation error.

- For AGARCH, the core volatility parameters η, α, and β are closely replicated to the second or third decimal. However, the asymmetry parameter γ shows a substantial discrepancy of 0.2722, consistent with the pattern observed in the APARCH model. We attribute this to the same source—namely, sensitivity of the asymmetry parameter to optimizer choice and convergence criteria, as γ appears weakly identified given its large standard error relative to its point estimate. This is consistent with earlier results stating that reactions to exchange rate shocks are not asymmetric.

- Across both models, the discrepancies are concentrated in the asymmetry-related parameters

γ and δ, while the ARCH and GARCH components α and β are consistently well-replicated. This systematic pattern suggests optimizer sensitivity rather than random estimation noise as the primary source of deviation from Tsui and Ho (2004).

TABLE 7: REPLICATION RESULTS FOR DATASET II

	APARCH	AGARCH
μ	-0.0072	-0.0079
	(0.0100)	(0.0100)
ρ	0.0318	0.0314
	(0.0167)	(0.0167)
η	0.0143	0.0115
	(0.0067)	(0.0050)
α	0.0543	0.0421
	(0.0161)	(0.0141)
β_G	0.9313	0.9314
	(0.0198)	(0.0211)
γ	0.0771	0.0054
	(0.0934)	(0.0109)
δ	1.4759	—
	(0.3703)	—

Note 1: Standard errors in parentheses are Bollerslev and Ole Mikkelsen (1996) robust standard errors.
Note 2: All models estimated via QMLE assuming a normal likelihood.

4.3 Dataset: Extended

Considering the GARCH(1,1) model, both AIC and BIC choose GED as the best distributional assumption (see Table 8). The same applies to the FIGARCH(1, d, 1) model. Both criteria assign a marginally lower value to the GARCH-G model than to the FIGARCH-G model. Thus, the GARCH-G(1,1) is our prime model.

Now, we compare our results pertaining to "Extended" (see Table 8) to the findings of Tse (1998). Regarding the constant (μ) of the mean equation, the GARCH-G model on "Extended" produces very similar results to the GARCH model in Tse (1998). This implies that the average expected return of the exchange rate has not changed in the more recent time frame compared to the older time frame used in Tse (1998). We find a small negative autoregressive parameter ρ, indicating a small negative correlation between current and last period's exchange rate return that is only apparent in recent data.

Regarding the volatility equation, we cannot confirm the substantial positive constant in the volatility equation (η) found by Tse (1998), implying that the volatility in the exchange rate is lower in recent data. Additionally, we find a much lower α, indicating that the exchange rate is less volatile to recent shocks. The GARCH coefficient β, however, is comparable between the results for "Extended" and in Tse (1998). This means that the persistence of shocks is comparable in both datasets.

When it comes to the FIGARCH-G model, we can overall confirm the findings of Tse (1998) regarding the mean equation. "Extended" produces a much higher GARCH coefficient β though, implying a higher short-term persistence of volatility in recent times. Our estimates for the memory parameter d indicate a very slow decay of volatility shocks, suggesting that the JPY/USD exchange rate

Table 8: Estimation Results for Extended Dataset (2003–2023)

	GARCH-N	GARCH-t	GARCH-G	FIGARCH-N	FIGARCH-t	FIGARCH-G
μ	0.0125	0.0112	0.0010	0.0118	0.0112	0.0010
	(0.0070)	(0.0056)	(0.0003)	(0.0068)	(0.0056)	(0.0008)
ρ	-0.0172	-0.0336	-0.0111	-0.0151	-0.0323	-0.0107
	(0.0145)	(0.0126)	(0.0009)	(0.0146)	(0.0128)	(0.0058)
η	0.0024	0.0016	0.0014	0.0088	0.0086	0.0012
	(0.0009)	(0.0007)	(0.0007)	(0.0043)	(0.0031)	(0.0048)
α	0.0512	0.0478	0.0484	—	—	—
	(0.0092)	(0.0093)	(0.0089)	—	—	—
β_G	0.9439	0.9512	0.9503	—	—	—
	(0.0099)	(0.0096)	(0.0092)	—	—	—
β_F	—	—	—	0.6654	0.6637	0.9508
	—	—	—	(0.0968)	(0.0594)	(0.1562)
d	—	—	—	0.4341	0.4273	1.0000
	—	—	—	(0.0911)	(0.0535)	(0.4986)
ϕ	—	—	—	0.2830	0.2864	0.0000
	—	—	—	(0.0445)	(0.0368)	(0.3395)
AIC	9162.96	8625.45	8542.33	9156.32	8618.63	8544.40
BIC	9196.15	8665.27	8582.15	9196.15	8665.08	8590.86

Note 1: Standard errors in parentheses are Bollerslev and Ole Mikkelsen (1996) robust standard errors.
Note 2: All models estimated via QMLE alongside an AR(1) mean equation.
Note 3: Column suffixes denote the error distribution: -N (Gaussian), -t (Student's t), -G (GED).

has become even more persistent in the 21$^{\text{st}}$ century. Similar to Tse (1998), the long-term memory coefficient d is statistically significant, while the short-term memory coefficient ϕ is not.

To compare our "Extended" results to those of Tsui and Ho (2004), we need to rely on a comparison of our FIGARCH results and their FIAGARCH results. Their FIAGARCH model produces a significant d but of much smaller magnitude than our FIGARCH model. This could imply that the d-variable in our FIGARCH model suffers from an omitted variable bias in the sense that it misinterprets asymmetry as persistence, thus indicating that asymmetric reactions play a much more pronounced role for exchange rates than previously claimed. Tsui and Ho (2004)'s AGARCH model likewise finds a significant role of asymmetry. That the remaining variables of their AGARCH model and our GARCH model are similar in significance and magnitude implies that the asymmetry is suppressed by the sharp decay in these kinds of models.

5. Conclusion

This study successfully replicated the findings of Tse (1998) and Tsui and Ho (2004) for the JPY/USD exchange rate. Importantly, our replication of the asymmetric models from Tsui and Ho (2004) (Dataset II) demonstrated that the asymmetry parameters (γ) were statistically insignificant, hinting that they are weakly identified in the original paper. This finding provided the empirical justification for omitting asymmetric effects, along with coherent argumentations in previous literature (e.g. Baillie, Bollerslev, and Mikkelsen (1996)), and utilizing more parsimonious symmetric GARCH and FIGARCH specifications in our extended analysis. That said, the comparison of our results and Tsui and Ho (2004)'s casts a doubt on the notion of non-asymmetric reactions in exchange rates.

The analysis of the Extended sample (2003–2023) shows a structural shifts in the currency pair's dynamics. While average return expectations remain stable around zero, a negative autoregressive parameter indicates oscillating behaviour of returns. More importantly, our estimates show higher short-term persistence (β) and a memory parameter (d) that indicates a very slow hyperbolic decay of shocks in recent times compared to earlier periods. Finally, the GED distribution was found to provide the most accurate fit across all specifications, underscoring the non-normal, leptokurtic nature of modern exchange rate returns.

References

Baillie, Richard T., Tim Bollerslev, and Hans Ole Mikkelsen (1996). "Fractionally integrated generalized autoregressive conditional heteroskedasticity". In: *Journal of Econometrics* 74.1, pp. 3–30. ISSN: 0304-4076. DOI: https://doi.org/10.1016/S0304-4076(95)01749-6. URL: https://www.sciencedirect.com/science/article/pii/S0304407695017496.

Bollerslev, Tim (1986). "Generalized autoregressive conditional heteroskedasticity". In: *Journal of Econometrics* 31.3, pp. 307–327. ISSN: 0304-4076. DOI: 10.1016/0304-4076(86)90063-1. URL: https://www.sciencedirect.com/science/article/pii/0304407686900631.

Bollerslev, Tim and Hans Ole Mikkelsen (1996). "Modeling and pricing long memory in stock market volatility". In: *Journal of Econometrics* 73.1, pp. 151–184. ISSN: 0304-4076. DOI: https://doi.org/10.1016/0304-4076(95)01736-4. URL: https://www.sciencedirect.com/science/article/pii/0304407695017364.

Engle, Robert F. (1982). "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation". In: *Econometrica* 50.4, pp. 987–1007. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1912773 (visited on 02/26/2026).

Hinton, Caleb (Aug. 24, 2023). *How do currency pairs work?* Expert Analysis. URL: https://www.currencytransfer.com/blog/expert-analysis/how-do-currency-pairs-work (visited on 02/26/2026).

Ito, Hiroyuki (Aug. 2022). *An Impact of Exchange Rate Fluctuations*. URL: https://www.rieti.go.jp/en/columns/a01_0702.html (visited on 02/26/2026).

Ma, Christopher K. and G. Wenchi Kao (June 1990). "On Exchange Rate Changes and Stock Price Reactions". In: *Journal of Business Finance & Accounting* 17.3, pp. 441–449. DOI: 10.1111/j.1468-5957.1990.tb01196.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1468-5957.1990.tb01196.x.

Sheppard, Kevin (Oct. 2025). *bashtage/arch: Release 8.0*. Version v8.0.0. DOI: 10.5281/zenodo.593254. URL: https://doi.org/10.5281/zenodo.593254.

Shi, Rui and Prashant Pande (Dec. 2024). *Understanding Currency Carry Trades: The Yen Carry Trade and Its Impact on ASEAN+3 Economies*. Analytical Note. ASEAN Macroeconomic Research Office (AMRO). URL: https://amro-asia.org/wp-content/uploads/2024/12/20241219-Analytical_Note_Carry_Trade.pdf.

Tse, Y. K. (1998). "The Conditional Heteroscedasticity of the Yen-Dollar Exchange Rate". In: *Journal of Applied Econometrics* 13.1, pp. 49–55. ISSN: 08837252, 10991255. URL: http://www.jstor.org/stable/223265 (visited on 02/25/2026).

Tsui, Albert and Kin-Yip Ho (2004). "Conditional heteroscedasticity of exchange rates: further results based on the fractionally integrated approach". In: *Journal of Applied Econometrics* 19.5, pp. 637–642. URL: https://EconPapers.repec.org/RePEc:jae:japmet:v:19:y:2004:i:5:p:637-642.

Zhelyazkova, Slaveya (Dec. 2024). "Application of Multivariate FIGARCH and FIAPARCH Models in the Analysis of Volatility of Exchange Rates, Gold and Oil Prices". In: *Izvestia Journal of the Union of Scientists - Varna Economic Sciences Series* 13.2, pp. 131–142. DOI: 10.56065/IJUSV-ESS/2024.13.2.131. URL: https://doi.org/10.56065/IJUSV-ESS/2024.13.2.131.

A. Figures

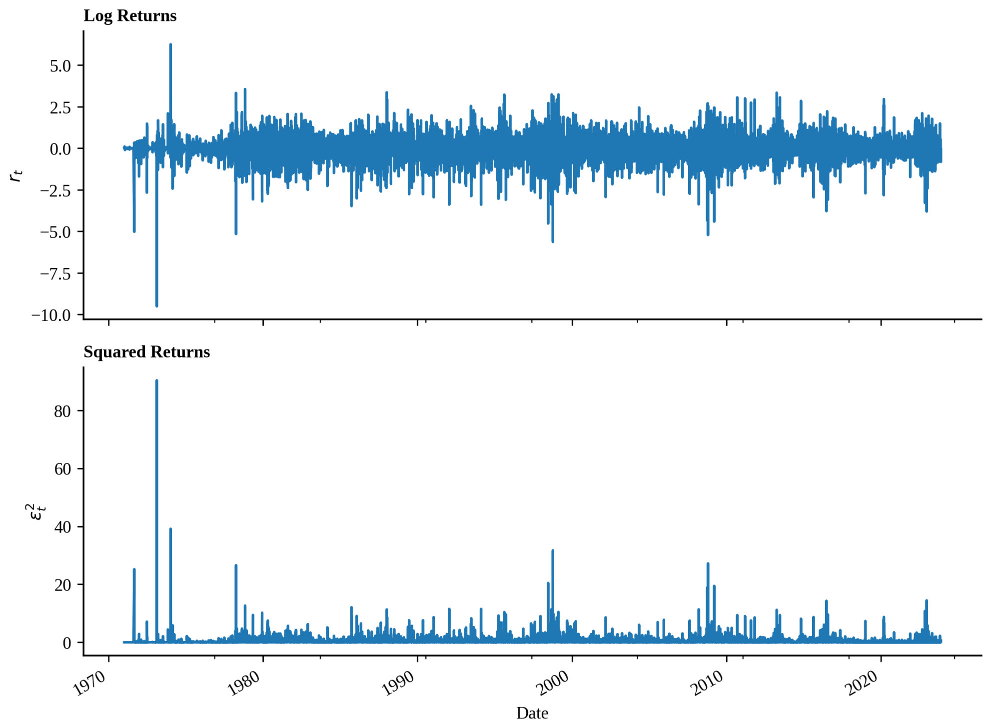

Figure 1: Daily Log Returns and Volatility Persistence (JPY/USD). Upper panel displays daily log return series (r_t). Lower panel displays squared returns (r_t^2) as proxy for conditional variance.

Figure 2: Autocorrelation and Partial Autocorrelation Analysis.
Note: The vertical axis has been capped at 0.2 to enhance the visibility of higher-order lags; the first lag (lag 0) is equal to 1.0 by definition.

B. AI Statement

During the completion of this assignmment, ChatGPT and Gemini assisted us with writing Python code and debugging LaTeXissues.