Normality of Daily Stock Returns

Evidence from Apple Microsoft and NVIDIA
January 2020 to October 2025

Abstract

This study examines whether daily stock returns for Apple (AAPL), Microsoft (MSFT), and NVIDIA (NVDA) follow a normal distribution. The sample window is January 1, 2020 to October 22, 2025. It contains 1,460 closing-price observations and 1,459 daily logarithmic returns per company, calculated from the supplied dataset's Close fields. Descriptive statistics, histograms, normal Q-Q plots, box plots, and the Shapiro–Wilk and Jarque–Bera tests are used to assess normality. Pearson kurtosis is 9.27 for AAPL, 10.84 for MSFT, and 7.05 for NVDA, compared with 3 for a normal distribution. Both tests report p-values below 0.001 for every company, and the Q-Q plots show substantial departures in both tails. Taken together, the evidence indicates that a single normal distribution is an inadequate description of these pooled daily returns. The conclusion is limited to the selected stocks, sample period, and price fields; it does not identify the best alternative distribution or establish that returns are independent over time.

1 Introduction

Stock prices generate a large quantity of financial data, but the statistical behavior of those data is not necessarily simple. One of the most familiar probability models is the normal distribution, which is symmetric, centered around its mean, and characterized by a relatively low probability of extreme observations. Because of its mathematical simplicity, the normal distribution has long been used as a benchmark in financial statistics and quantitative finance.

The central question of this paper is whether this benchmark describes actual daily stock returns. The issue matters because estimates based on a normal distribution may understate the frequency of unusually large gains and losses. Early studies by Mandelbrot (1963) and Fama (1965) challenged the Gaussian description of financial price changes. This paper examines the same distributional question in a recent sample of technology stocks.

This study focuses on three large U.S. technology companies: Apple, Microsoft, and NVIDIA. The research question is: Do daily stock returns follow a normal distribution? The analysis uses daily logarithmic returns over a fixed sample window, January 1, 2020 to October 22, 2025. It combines descriptive statistics, graphical diagnostics, and formal tests to assess a single normal distribution for each company over the full period.

The rest of the paper is organized as follows. Section 2 reviews relevant literature on financial return distributions and the normality assumption. Section 3 describes the dataset and the construction of daily logarithmic returns. Section 4 explains the statistical methods. Section 5 presents the empirical results. Section 6 interprets those results in relation to prior research. Section 7 discusses limitations, and Section 8 concludes.

2 Literature Review

The normal distribution as a financial benchmark

The normal distribution is attractive in financial modeling because its behavior is mathematically well understood and can be summarized using only a mean and a variance. A normal distribution is symmetric, has zero skewness, and has Pearson kurtosis equal to 3. These properties make normality a useful benchmark for asking whether observed financial returns display unusual asymmetry or tail behavior.

A distributional model must distinguish symmetry from normality. A symmetric distribution can still assign much more probability to extreme observations than a normal distribution. It is also necessary to distinguish the distribution pooled over several years from the distribution conditional on information available on a particular day. These distinctions guide the interpretation of the evidence reviewed below.

Early evidence on return distributions

Mandelbrot (1963) studied speculative price changes, including cotton prices, and argued that large movements occurred too often for a Gaussian model. He proposed a stable-Paretian family of distributions as an alternative. His study established an influential approach to describing extreme financial price movements.

Fama (1965) extended the empirical study of price changes to stocks and reported substantial departures from a normal distribution. Together, these early studies motivate examining the tails of returns as well as their average and standard deviation. They provide historical evidence against treating normality as an assumption that can be accepted without testing.

Volatility clustering and conditional distributions

Cont (2001) reviews recurring properties of financial returns, including heavy tails and volatility clustering. Volatility clustering means that large absolute returns tend to be followed by other large absolute returns, while quieter observations also occur together. It concerns the magnitude of movements, rather than a sequence of returns with the same sign.

Bollerslev (1986) develops the generalized autoregressive conditional heteroskedasticity model, or GARCH, in which conditional variance depends on past squared disturbances and past conditional variances. Such a model allows volatility to persist over time. A process with conditionally normal innovations can have a non-normal distribution when observations with different variances are pooled. Consequently, a rejection of full-sample normality alone does not identify the source of the departure.

Asymmetry and the interpretation of moments

Peiró (1999) examines daily returns in eight international stock markets and three exchange rates. He finds that symmetry tests based on sample skewness can be misleading when their reference distribution assumes normality. His findings support using several diagnostics together: a small skewness estimate does not establish normality, and a non-normal distribution need not be strongly asymmetric.

Return distributions during COVID 19

Baker et al. (2020) document an unusually high frequency of large daily U.S. market moves during the early COVID-19 crisis. Wątorek et al. (2021) examine return distributions across several asset classes and contracts for differences during 2017–2020, reporting that the pandemic disruption affected market dynamics and tail behavior. Their instruments and time scales differ from this paper's daily individual-stock sample, so their findings provide context rather than a direct test of the three companies studied here.

Alternatives to the normal distribution

The stable distributions considered by Mandelbrot (1963) and the asymmetric Student-t family developed by Zhu and Galbraith (2010) provide alternatives that can accommodate more extreme observations than a Gaussian model. A standard Student's t distribution is symmetric and has a tail parameter controlling the frequency of extremes; asymmetric extensions also allow different behavior on the two sides. These models offer candidate explanations for distributional shape, but each requires fitting and evaluation on the data of interest.

Scope of the present study

The present study applies established methods to AAPL, MSFT, and NVDA over a common recent period. Its contribution is a focused comparison using an explicitly defined price field and reproducible calculations. The analysis assesses the pooled return distributions; it does not estimate a volatility model, attribute movements to particular events, or compare fitted alternative distributions.

3 Data

The analysis uses 3_Companies_OHLCV.csv from the public stock-ohclv-dataset repository (Sakibul786, n.d.). The repository describes daily Open, High, Low, Close, and Volume data for AAPL, MSFT, and NVDA covering January 2015 to October 22, 2025. The sample is restricted to January 1, 2020 through October 22, 2025, inclusive. The first available trading date is January 2, 2020.

After date filtering and sorting, the sample contains 1,460 price observations per company. There are no duplicate dates, missing Close values, or non-positive Close values. The first price observation has no earlier observation within the selected window, so it contributes no return. This leaves 1,459 returns per company, dated January 3, 2020 to October 22, 2025. No prices are filled in for weekends or holidays, and no observed returns are removed or winsorized.

The price variables are exactly AAPL_('Close', 'AAPL'), MSFT_('Close', 'MSFT'), and NVDA_('Close', 'NVDA'). The CSV contains no separate Adj Close columns. Open, High, Low, and Volume are not used to compute the returns.

Under Yahoo Finance's historical-data convention, Close is adjusted for stock splits, while Adj Close is adjusted for both splits and dividend distributions (Yahoo Finance, n.d.). Split adjustment puts earlier prices on a comparable share basis and avoids treating a mechanical change in price per share as a loss. Dividend adjustment additionally accounts for distributions to shareholders. A split-adjusted series without dividend adjustment therefore measures price changes rather than the full return including cash dividends.

That convention does not, by itself, establish the adjustment history of this CSV. The yfinance software can replace Close with Adj Close when automatic adjustment is enabled, retaining the name Close (yfinance contributors, n.d.). The repository does not document the download settings needed to resolve that distinction. This study adopts the supplied Close values and applies no additional split or dividend adjustment. The results are therefore defined as returns on those price series; they are not claimed to be independently verified dividend-excluded or total shareholder returns.

Table 1. Variables used in the empirical analysis.

Variable	Definition	Use in analysis
Date	Trading date	Indexes each daily observation
Close	Selected Close field in the supplied CSV	Used to calculate the daily log return
Daily log return	Natural logarithm of consecutive Close ratios	Primary variable tested for normality
Mean	Average daily return	Measures central tendency
Standard deviation	Sample standard deviation of daily returns	Measures dispersion
Skewness	Bias-corrected standardized third moment	Measures asymmetry
Kurtosis	Bias-corrected Pearson kurtosis	Summarizes extreme standardized observations

4 Methodology

Calculating daily log returns

For each company, the daily logarithmic return is calculated from consecutive available trading-day prices:

$$r_t = ln(\frac{P_t}{P_{t-1}})$$

Here, P denotes the selected Close series, and t indexes trading observations. The previous observation can be separated by a weekend or holiday. The calculations use decimal log returns; multiplying by 100 expresses them in percent. These are logarithmic changes, so an extreme value such as −20.40% is not identical to a simple percentage price change.

Descriptive statistics

The empirical distribution is summarized by its mean, sample standard deviation, skewness, Pearson kurtosis, minimum, and maximum. Standard deviation uses a denominator of n − 1. Skewness and kurtosis use SciPy's bias-corrected estimators, as shown in Appendix A. Normal-distribution benchmarks are zero skewness and Pearson kurtosis of 3. Excess kurtosis equals Pearson kurtosis minus 3. Kurtosis emphasizes extreme standardized observations; it is not, by itself, a measurement of the height of a distribution's central peak.

Graphical diagnostics

Histograms display the empirical distributions with normal densities fitted using each sample's mean and standard deviation. The normal Q-Q plots compare ordered returns with theoretical standard-normal quantiles and a fitted reference line. Box plots show the median and interquartile range. Their whiskers extend to the most extreme observations within 1.5 interquartile ranges of the box, and observations beyond the whiskers appear separately. Such points are retained in all calculations.

Formal normality tests

The Shapiro–Wilk test compares ordered observations with the pattern expected under normality (Shapiro & Wilk, 1965). The Jarque–Bera test evaluates the joint departure of skewness and kurtosis from their normal values (Jarque & Bera, 1980). For each stock, the null hypothesis is that the daily return distribution is normal; the alternative is that it is not normal. The significance level is 5%. A reported p-value below 0.05 leads to rejection under the test's assumptions.

Shapiro–Wilk and the graphical diagnostics provide the main assessment; Jarque–Bera is supplementary. SciPy cautions that the Jarque–Bera chi-square approximation requires a large sample and gives more than 2,000 observations as a guideline. The present sample has 1,459 observations. SciPy's separate warning about Shapiro–Wilk p-values applies above 5,000 observations (SciPy community, n.d.-a, n.d.-b). Both tests use conventional independent-sample reference distributions. Their reported p-values are therefore nominal when applied to potentially dependent financial returns. The tests are assessed alongside the observed distributional departures.

5 Results

Descriptive statistics

Table 2 reports statistics for 1,459 daily log returns per stock. Mean returns are 0.0871% for AAPL, 0.0840% for MSFT, and 0.2336% for NVDA. NVDA has the largest standard deviation, 3.3602%, and MSFT the smallest, 1.8769%.

Table 2. Descriptive statistics of daily log returns.

Stock	N	Mean	Std. dev.	Skewness	Kurtosis	Min	Max
AAPL	1,459	0.0871%	2.0259%	0.0284	9.2732	-13.77%	14.26%
MSFT	1,459	0.0840%	1.8769%	-0.1663	10.8366	-15.95%	13.29%
NVDA	1,459	0.2336%	3.3602%	0.0409	7.0510	-20.40%	21.81%

Note. N = 1,459 returns per stock, dated January 3, 2020 to October 22, 2025. Mean, standard deviation, minimum, and maximum are in percent. Skewness and Pearson kurtosis are dimensionless and bias-corrected. Source: calculations from Sakibul786 (n.d.).

Pearson kurtosis is 9.2732 for AAPL, 10.8366 for MSFT, and 7.0510 for NVDA. Each value exceeds the normal benchmark of 3. The corresponding excess kurtosis values are 6.2732, 7.8366, and 4.0510. Skewness is 0.0284, −0.1663, and 0.0409, respectively.

The minimum and maximum log returns are −13.77% and 14.26% for AAPL, −15.95% and 13.29% for MSFT, and −20.40% and 21.81% for NVDA. All three samples contain both positive and negative observations.

Distributional plots

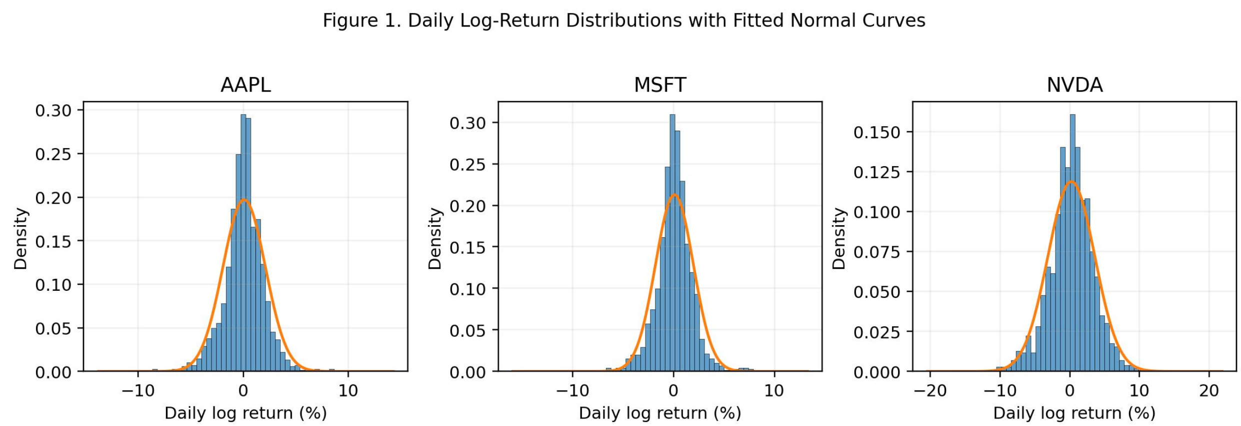

Figure 1. Daily log-return distributions with fitted normal curves.

Note. Returns are in percent; curves use each stock's sample mean and standard deviation. N = 1,459 per panel.

In Figure 1, the tallest central histogram bars exceed the fitted normal-density peaks for all three stocks. Observations extend to the positive and negative extremes reported in Table 2.

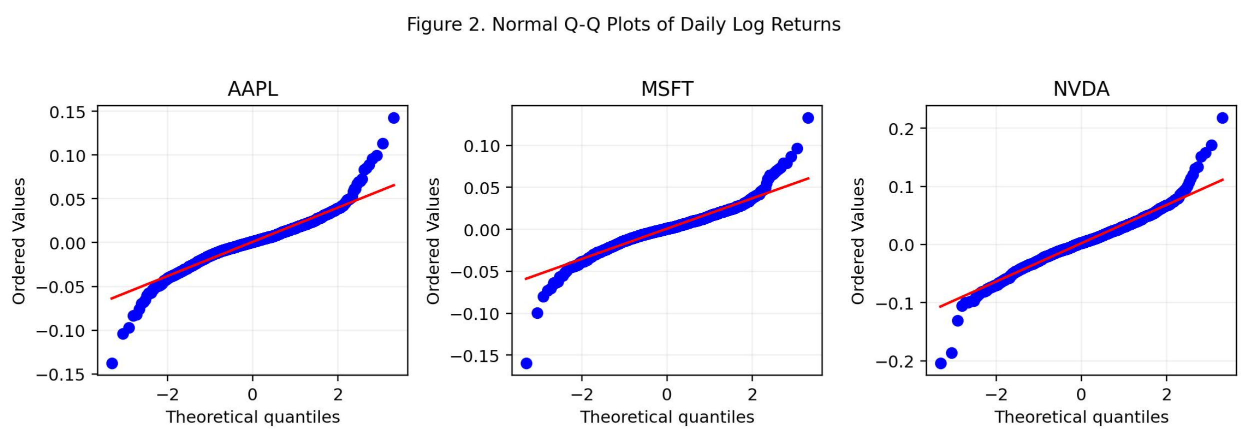

Figure 2. Normal Q-Q plots of daily log returns.

Note. Ordered returns are in decimal units (0.10 = 10%); theoretical quantiles are standard normal. Red lines are fitted references. N = 1,459 per panel.

In Figure 2, observations in the lower tails fall below the reference lines, while observations in the upper tails lie above them. The departures are greater at the extreme quantiles than near the center.

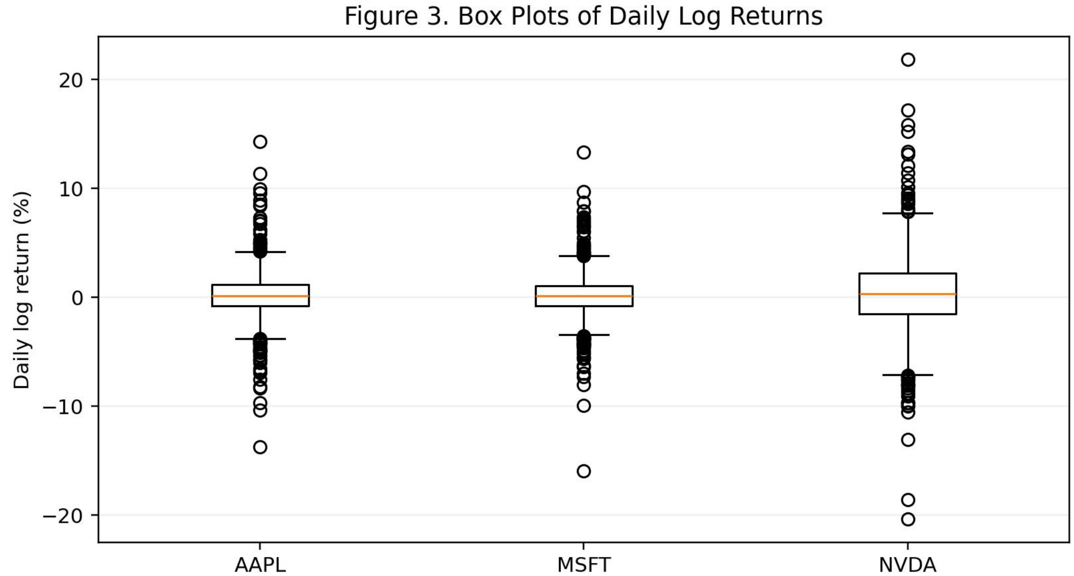

Figure 3. Box plots of daily log returns.

Note. Returns are in percent. Circles mark observations beyond the 1.5-interquartile-range whiskers. N = 1,459 per stock. Figures 1–3 use the same sample as Table 2.

Figure 3 shows observations beyond both whiskers for each stock. NVDA has the widest box and the largest overall range of plotted returns.

Formal normality tests

Table 3 reports Jarque–Bera statistics of 2373.04, 3710.63, and 989.18 for AAPL, MSFT, and NVDA. The corresponding Shapiro–Wilk statistics are 0.9336, 0.9289, and 0.9635. Both tests report p-values below 0.001 for each company. At the nominal 5% level, each test rejects normality for all three stocks.

Table 3. Formal tests of the normality of daily log returns.

Stock	Jarque–Bera statistic	JB p-value	Shapiro–Wilk statistic	SW p-value	Decision at 5%
AAPL	2373.04	<0.001	0.9336	<0.001	Reject H_0
MSFT	3710.63	<0.001	0.9289	<0.001	Reject H_0
NVDA	989.18	<0.001	0.9635	<0.001	Reject H_0

Note. Decisions use the nominal 5% significance level. Jarque–Bera p-values use an asymptotic approximation; neither test adjusts for serial dependence. Very small p-values are reported as <0.001, rather than zero. Source: calculations from Sakibul786 (n.d.).

6 Discussion

The combined evidence indicates that a single normal distribution does not adequately describe the pooled daily returns of AAPL, MSFT, or NVDA in this sample. Elevated kurtosis and the Q-Q tail departures show why rejection is more than a response to a small, visually negligible discrepancy.

Both tests reach the same decision, although this agreement does not remove their shared dependence assumptions or the finite-sample limitation of Jarque–Bera.

The large observed movements relative to the fitted normal curves are consistent with the earlier evidence of Mandelbrot (1963) and Fama (1965). Here, “heavier tails” describes the empirical departures from the normal benchmark. The study does not estimate a tail exponent or establish a particular mathematical class of heavy-tailed distributions.

The comparison across companies separates return dispersion from distributional shape. NVDA has the greatest standard deviation, but MSFT has the highest kurtosis. Kurtosis scales observations by their dispersion, so a stock can have lower overall volatility while showing more pronounced extreme observations relative to that volatility. The statistics therefore describe different features of the data.

All three skewness estimates are fairly close to zero, despite the rejection of normality. The main visible departures concern the extremes on both sides of the distributions. This is consistent with Peiró’s (1999) distinction between symmetry and normality. It does not establish exact symmetry, because the study has not performed a separate symmetry test.

The volatility literature offers a possible explanation for the pooled shape. If quiet and volatile days are combined, a single normal curve can miss both the concentration of small movements and the extremes. Volatility clustering makes this explanation relevant to financial data (Cont, 2001; Bollerslev, 1986). However, the present plots do not show temporal ordering, and no test of squared-return dependence or GARCH model is estimated. Clustering is therefore a literature-based interpretation, not an additional empirical finding of this paper.

The inclusion of 2020 also matters when interpreting the period as a whole. Baker et al. (2020) document exceptional daily market movements during the early pandemic, while Wątorek et al. (2021) describe changes in return-distribution behavior during the disruption. These studies make COVID-19 a plausible contributor to the sample’s extremes. They do not establish that it caused the kurtosis observed here, and the full-period statistics cannot isolate its contribution.

A Student’s t or asymmetric Student-t distribution could accommodate the observed departures more flexibly than a normal curve (Zhu & Galbraith, 2010). That is a candidate extension rather than a fitted result. The evidence in this paper establishes poor agreement with normality without selecting a replacement model.

Rejecting normality does not establish that future return directions can be predicted from past prices. It also does not show that every financial model using a normal component is inappropriate. The direct implication is narrower: estimates that use one fitted normal distribution to describe the full sample may give too little weight to extreme daily movements.

7 Limitations

The sample covers only three large U.S. technology companies. It cannot represent other industries, smaller firms, international equities, or other asset classes. It also pools nearly six years into one distribution. Research on the COVID-19 period documents unusual market movements and changes in return distributions (Baker et al., 2020; Wątorek et al., 2021). Because no subperiod comparison is conducted here, the analysis cannot quantify how much the 2020 episode contributes to the full-sample kurtosis or determine whether the same conclusion holds in every year.

The analysis evaluates normality without fitting and comparing alternatives. Student's t, asymmetric Student-t, and stable distributions are supported as candidates by the literature (Mandelbrot, 1963; Zhu & Galbraith, 2010), but rejecting normality does not establish that any one of them fits these observations well. A comparison would need parameter estimation, goodness-of-fit diagnostics, and an assessment of performance on observations not used for fitting.

Financial returns may also be dependent in magnitude. The conventional normality-test p-values do not account for volatility clustering, and the full-sample results cannot distinguish changing conditional variance from a non-normal conditional innovation distribution. The GARCH framework provides one way to make that distinction (Bollerslev, 1986). Given these limits and the Jarque–Bera sample-size caveat, the graphical and descriptive evidence is essential to the conclusion.

Finally, the adjustment provenance of the supplied Close fields limits the economic interpretation of the returns. The distinction between split adjustment and dividend adjustment is explained in Section 3, but the CSV does not preserve a separate comparison series. No sensitivity test using independently verified dividend-adjusted and dividend-excluded prices is performed. It would therefore be unsupported to claim that changing the adjustment basis must leave every statistic or test decision unchanged.

8 Conclusion

This study examined the question "Do daily stock returns follow a normal distribution?" using AAPL, MSFT, and NVDA over a sample window from January 1, 2020 to October 22, 2025. Consecutive Close observations yield 1,459 daily logarithmic returns per company.

For all three companies, the pooled returns depart substantially from a normal distribution. Pearson kurtosis ranges from 7.05 to 10.84, the normal Q-Q plots show deviations in both tails, and both formal tests report p-values below 0.001. These findings support rejecting a single normal distribution as an adequate description of this sample, subject to the stated test assumptions.

The answer applies to the observed daily returns and their chosen price basis. It does not establish an optimal alternative model or imply predictability. Comparisons across market periods and fitted distributions would address different questions from the normality assessment completed here.

References

Baker, S. R., Bloom, N., Davis, S. J., Kost, K., Sammon, M., & Viratyosin, T. (2020). The unprecedented stock market reaction to COVID-19. *The Review of Asset Pricing Studies, 10*(4), 742–758. https://doi.org/10.1093/rapstu/raaa008

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics, 31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

Cont, R. (2001). Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance, 1*(2), 223–236. https://doi.org/10.1080/713665670

Fama, E. F. (1965). The behavior of stock-market prices. *The Journal of Business, 38*(1), 34–105. https://doi.org/10.1086/294743

Jarque, C. M., & Bera, A. K. (1980). Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters, 6*(3), 255–259. https://doi.org/10.1016/0165-1765(80)90024-5

Mandelbrot, B. (1963). The variation of certain speculative prices. *The Journal of Business, 36*(4), 394–419. https://doi.org/10.1086/294632

Peiró, A. (1999). Skewness in financial returns. *Journal of Banking & Finance, 23*(6), 847–862. https://doi.org/10.1016/S0378-4266(98)00119-8

Sakibul786. (n.d.). *stock-ohclv-dataset* [Data set]. GitHub. https://github.com/Sakibul786/stock-ohclv-dataset

SciPy community. (n.d.-a). *scipy.stats.jarque_bera*. SciPy documentation. https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.jarque_bera.html

SciPy community. (n.d.-b). *scipy.stats.shapiro*. SciPy documentation. https://docs.scipy.org/doc/scipy/reference/generated/scipy.stats.shapiro.html

Shapiro, S. S., & Wilk, M. B. (1965). An analysis of variance test for normality (complete samples). *Biometrika, 52*(3–4), 591–611. https://doi.org/10.1093/biomet/52.3-4.591

Wątorek, M., Kwapień, J., & Drożdż, S. (2021). Financial return distributions: Past, present, and COVID-19. *Entropy, 23*(7), Article 884. https://doi.org/10.3390/e23070884

Yahoo Finance. (n.d.). *Apple Inc. (AAPL) historical data*. https://ca.finance.yahoo.com/quote/AAPL/history/

yfinance contributors. (n.d.). *utils.py* [Source code]. GitHub. https://github.com/ranaroussi/yfinance/blob/main/yfinance/utils.py

Zhu, D., & Galbraith, J. W. (2010). A generalized asymmetric Student-t distribution with application to financial econometrics. *Journal of Econometrics, 157*(2), 297–305. https://doi.org/10.1016/j.jeconom.2010.01.013

Appendix A Reproducibility Code

The following code reproduces Tables 2 and 3 from 3_Companies_OHLCV.csv. It checks the selected observations and computes returns after date filtering. No additional corporate-action adjustment or outlier removal is applied.

```
import numpy as np
import pandas as pd
from scipy import stats

df = pd.read_csv("3_Companies_OHLCV.csv", parse_dates=["Date"])
df = df.sort_values("Date")
df = df[df["Date"].between("2020-01-01", "2025-10-22")].copy()
cols = {
    "AAPL": "AAPL_('Close', 'AAPL')",
    "MSFT": "MSFT_('Close', 'MSFT')",
    "NVDA": "NVDA_('Close', 'NVDA')"
}
assert len(df) == 1460
assert not df["Date"].duplicated().any()

for stock, col in cols.items():
    assert df[col].notna().all() and (df[col] > 0).all()
    r = np.log(df[col] / df[col].shift(1)).dropna()
    assert len(r) == 1459
    summary = {
        "N": len(r),
        "Mean (%)": r.mean() * 100,
        "Std. dev. (%)": r.std(ddof=1) * 100,
        "Skewness": stats.skew(r, bias=False),
        "Pearson kurtosis": stats.kurtosis(
            r, fisher=False, bias=False),
        "Min (%)": r.min() * 100,
        "Max (%)": r.max() * 100
    }
    print(stock, summary)
    jb = stats.jarque_bera(r)
    sw = stats.shapiro(r)
    for name, test in [("Jarque-Bera", jb), ("Shapiro-Wilk", sw)]:
        p = "<0.001" if test.pvalue < 0.001 else f"{test.pvalue:.4f}"
        decision = "Reject H0" if test.pvalue < 0.05 else "Do not reject H0"
        print(name, test.statistic, p, decision)
```