Resilience of Nepal's stock market to systemic shocks: A case study of COVID-19

Avinash Maskey[1]

Abstract

COVID-19 created a major upheaval and shock for Nepalese stock market but the direction and nature of its effect on NEPSE remain unclear because the pandemic coincided with change interest rates, inflation, policy restrictions, trading activity and investor behavior. This study examines the causal effect of the COVID-19 period on NEPSE daily returns using a Double Machine Learning Framework and daily data from 2018 to 2023. NEPSE daily returns were considered as an outcome variable while the treatment is measured by a binary COVID-period indicator and alternatively by daily changes in the Oxford stringency index. The model controls stationarised macroeconomic and pandemic related covariates including base rate, inflation, COVID-19 cases, policy stringency and turnover. The results show that the binary COVID-period indicator is associated with average daily returns of about 21 basis points above the macro-implied counterfactual with stable estimates across adjustment sets and estimators. However, daily changes in policy stringency show no significant dose-response effect. Robustness checks using sensitivity analysis and placebo permutation test supports the main estimate. The findings reveal that NEPSE's pandemic-period performance was driven less by daily policy changes and more by liquidity conditions, retail participation and speculative market behavior. These results highlight the need for stronger monitoring of retail-driven trading activity during crisis periods

Keyword: COVID-19, Nepal Stock Exchange, Double Machine Learning, Causal Inference, Frontier Markets, Stock Market Resilience, Behavioral Finance

JEL classification: C14, C32, G12, G14, G15

[1]Kathmandu University, School of Education, Lalitpur, Nepal.
Corresponding Email: avinash.maskey@kusoed.edu.np

INTRODUCTION

Systemic shocks and financial market resilience have long been of great academic and policy interest. Stock markets, especially those in emerging economies, play an important role in the mobilization of investment and savings, and are a good indicator of economic vulnerabilities. Their capacity to withstand big shocks, however, is very variable, especially in the less-developed financial markets of frontier markets. From this point of view, the COVID-19 pandemic is a historical once-in-a-lifetime phenomenon that put the structural integrity and behavioral mechanisms of financial systems under global stress.

While the impact of the pandemic on developed countries and some larger emerging markets has been the subject of numerous studies, much remains to be understood about the pandemic's causal impact on smaller frontier economies like Nepal. Despite being a small exchange by global standards and the relative isolation of the country's capital markets, the Nepal Stock Exchange (NEPSE) was highly volatile and recovered in the COVID-19 era. But whether this was a real structural resilience or a pure consequence of the temporary behavioral dynamics and rapid liquidity fluctuations is still an empirical question. This is a challenge that needs methodological approaches that can distinguish causal relations while accounting for a large number of confounding relationships, especially in the case of financial and macroeconomic data, in crises.

One of the first aspects of systemic shocks that is usually analyzed is the near-term fragility of financial markets to external shocks, and their ability to recover after the immediate shock. According to classical literature, institutional quality, liquidity depth, investor profile, and macroeconomic fundamentals are the key factors that determine the resilience of markets (Brunnermeier et al., 2009). In times of unprecedented events globally, such as the COVID-19 pandemic, behavioral responses such as panic selling, herd behavior, and speculative bubbles can overwhelm traditional fundamentals, particularly in retail-dominated markets (Corbet et al., 2020; Ramelli & Wagner, 2020).

The COVID-19 pandemic led to synchronized stock market declines, a fast evolution in the volatility patterns, and unprecedented policy measures on a global scale. In the first months of 2020, the major indices tumbled but then rose at an unexpected pace following the implementation of broad monetary and fiscal policies by the governments and central banks (Ashraf, 2020; Zaremba et al., 2020). However, these reactions varied in terms of their intensity and extent in different regions due to pandemic dynamics, policy responses, structural and behavioral market characteristics, etc. (Basuony et al., 2021). While Nepal is a frontier market that is relatively less studied in terms of systemic resilience, the large developed and emerging markets are relatively well studied. The initial evidence suggests that NEPSE was very volatile in the initial shock but recovered relatively quickly (Ghising, 2025; Karki, 2022). It is helpful to examine and differentiate between these trends and between resiliency, investor sentiment, and short-term liquidity impacts on the markets.

Additionally, Nepal's distinctive economic and institutional environment makes it challenging to directly compare with other economies. The NEPSE is a relatively new stock exchange that features a high presence of retail investors, a lack of foreign participation, low trading volume, and high political and regulatory risk. These are all factors that can exacerbate the potential for behavior distortions when a crisis arises, and the importance of credible causal identification grows even more urgent.

Typical financial crises involve data with many variables, noise, and non-linearity, which is a problem for traditional econometric methods. Simple OLS event studies/regressions of time series are not appropriate for separating out multiple treatment effects from contemporaneous shocks. To overcome these shortcomings, this study uses a Double Machine Learning (DML) framework (Chernozhukov et al., 2018), which combines flexible machine learning models with orthogonalized estimation techniques to yield unbiased causal estimates with weaker assumptions than traditional methods. In recent years, DML has been used in the financial sector for financial risk assessment (Li et al., 2025), vaccine effectiveness (Jiang et al., 2025), as well as corporate financial resilience (Song et al., 2025). Additionally, Cao and Ren (2022) and Sen (1968) discussed that the estimation of counterfactual trajectory can be enhanced by the application of a machine learning algorithm-based synthetic control methodology, including the Ridge regression model and Theil–Sen estimator in the stress situation.

Hence, the goal of this study is to address the important gap in the literature by employing causal machine learning techniques to assess the impact of the COVID-19 systemic shock on the Nepalese stock market. Using a unique data set across daily NEPSE index levels, turnover, macro-economic indicators, and measures related to the stringency of policy responses for the period 2018–2023, the research compares the actual path of the market during the pandemic period with the counterfactual path that would have occurred in the absence of the pandemic. The study uses a unique data set covering daily NEPSE index levels, turnover, macro-economic indicators, and policy stringency measures between 2018 and 2023, to generate counterfactual market trajectories and compute the Average Treatment Effect (ATE) of the pandemic period on daily returns.

This study adds both empirically and methodologically. Empirically, this research is the first, and to the best of our knowledge, the only, causal estimation of the impact of COVID-19 on NEPSE, which can give an idea of the extent and resilience of the market. In terms of methodology, it uses cutting-edge causal inference techniques when investigating frontier markets, thus establishing a precedent for future research in similar settings. It contributes to the general knowledge of how systemic shock spreads through thin and sensitive financial systems. Hence, the following were the major research questions for this research:

- RQ1: To what extent did the COVID-19 period affect NEPSE daily returns after controlling for macroeconomic conditions, policy response, pandemic intensity and market behavior?
- RQ: Does the estimated COVID-19 effect remain consistent across different model specifications including alternative treatment definitions, adjustment sets and OLS and DML estimators?
- RQ3: Do changes in volatility, skewness, kurtosis and event window abnormal returns indicate that NEPSE's pandemic period performance was influenced by speculative and sentiment–driven behavior rather than fundamental based repricing?

REVIEW OFLITERATURE

According to Brunnermeier et al. (2009), one can define financial resilience as a market's ability to withstand systemic shocks and avoid any long-term dislocation or outright collapse. That is an especially pertinent notion in frontier markets like Nepal, where you have structural weaknesses in the form of political risk, institutional inefficiencies, and a dearth of liquidity. The pandemic offered a kind of natural experiment to put the resilience of such economies to the test. After all, COVID-19 was an unprecedented shock that sent synchronized tremors through asset classes worldwide, causing volatility to spike and markets to sell off in a hurry amid thin liquidity (Khan et al., 2023; Ramelli & Wagner, 2020). Yet when it comes to how a place like Nepal has fared, there has been little in the way of rigorous causal analysis to date.

The empirical findings from the research revealed that COVID-19's impact on global financial markets was great, but it was heterogeneous. To support this, Ashraf (2020) found that the stock markets were negatively affected and reacted to the increase in both COVID-19 cases and fatalities. This also included investors' sentiment that drove early-stage volatility. In a similar context, government responses, fiscal stimuli, and lockdowns significantly moderate market volatility across different countries in the world (Zaremba et al., 2020). In contrast to this, Corbet et al. (2020) investigated and found that traditional safe havens like gold and emerging assets such as cryptocurrencies showed contagion patterns during the COVID-19 pandemic. Kumar et al. (2023) examined asset market interconnectedness in greater detail by employing copula models. The research found an increased correlation between oil prices and stock indices, highlighting the systemic characteristics of the shock. In another study, the depth of market reactions highly correlated with pre-existing macro-financial vulnerabilities by suggesting that institutional strength buffered some economies against deeper downturns (Basuony et al., 2021).

Emerging and frontier markets experienced sharper and more prolonged disruptions compared to developed markets (Topcu & Gulal, 2020). Cao and Ren (2022) provided evidence that in the U.S., COVID-19 reshaped corporate risk structures, while in emerging markets, market stress was amplified by weaker financial infrastructures.

In Nepal's context, Ghising (2025) documented significant NEPSE volatility during early pandemic phases, followed by a speculative surge as lockdowns eased. Karki (2022) noted that both monetary easing and retail investor optimism played a role in an unexpected market rebound, though it lacked strong macroeconomic backing.

However, these descriptive studies do not provide strong causal attribution, highlighting the need for more sophisticated methods such as causal machine learning. Traditional econometric approaches often falter under systemic shocks, where high-dimensional confounding and dynamic treatment effects complicate inference. Chernozhukov et al. (2018) introduced the

Double Machine Learning (DML) framework to address this, allowing robust causal estimation even under complex observational settings.

Recent studies (e.g. (Jiang et al., 2025; Song et al., 2025)) demonstrate the application of DML to fields such as vaccine effectiveness and corporate resilience, highlighting its versatility. In finance, DML enables researchers to unravel the true effect of systemic events from spurious correlations arising from endogenous market behaviors.

Counterfactual modeling techniques such as Ridge regression and Theil–Sen estimators (Sen, 1968) further strengthen causal inference by creating synthetic control groups against which real-world market movements can be benchmarked.

Behavioral finance provides crucial lenses for understanding market dynamics during crises. Marks (2018) argued that collective psychology fundamentally shapes market cycles such as greed, fear, and herd behavior often dominates rational calculations during extreme events.

Ramelli and Wagner (2020) illustrated how panic-driven selloffs characterized early COVID-19 market reactions. Meanwhile, Corbet et al. (2020) noted the irrational movement of funds between traditional and alternative assets, driven largely by investor sentiment rather than fundamentals.

Given NEPSE's dominance by retail investors, behavioral distortions likely played a pivotal role in shaping market trajectories during COVID-19, an aspect that causal machine learning can help quantify and control. The Nepalese stock market (NEPSE) operates under distinct constraints: shallow depth, limited foreign participation, and retail-driven volatility.

While Nepal Rastra Bank's monetary policies and government interventions provided some stability (Nepal Rastra Bank, 2024), their exact causal impacts on market behavior remain insufficiently explored. Existing studies (Ghising, 2025; Karki, 2022) rely heavily on event studies or OLS regressions, which are prone to endogeneity biases, thus limiting their policy relevance.

World Context: COVID-19 and Global Equity Markets

The initial COVID-19 financial research confirmed that the financial markets did not react to the pandemic at a uniform pace or slowly. Baker et al. (2020)

demonstrate that, compared with previous pandemics, the stock-market response to COVID-19 has been historically idiosyncratic; Al-Awadhi et al. (2020) and Ashraf (2020) find that during the initial stages of the pandemic, confirmed cases and deaths were correlated with reduced market returns. A similar finding is also reached by Zhang et al. (2020) from the perspective of global financial instability: Uncertainty surged as the pandemic spread to the global level. Ramelli and Wagner (2020) extend the picture to the firm level, demonstrating real-time market re-pricing of sectoral exposure, leverage, and international linkages.

A second stream of work concerns transmission mechanisms, where headline losses are followed up. Zaremba et al. (2020) find that government interventions were highly correlated with volatility across countries, and Haroon and Rizvi (2020) find that investor attention and media intensity enhanced market reactions. More deaths, panic, and lockdowns, as described by Baig et al. (2021), are related to the deterioration of liquidity and the rise in volatility of the U.S. equity markets. Basuony et al. (2021) build on this by demonstrating that contraction of markets was not just the result of the pandemic but was also driven by prior macro-financial vulnerabilities. Overall, these studies indicate that COVID-19 should not be viewed as a stand-alone event impacting a stable market. It worked simultaneously via policy, liquidity, sentiment, and balance sheet.

Emerging and Frontier Markets

The emerging market literature adds layers to the developed-market narrative that are relevant to Nepal. Topcu and Gulal (2020) demonstrate that the negative effect of COVID-19 on emerging stock markets was significant, with the magnitude of the effect and its duration differing from one region to another. Similarly, Salisu et al. (2020) concludes that the pandemic uncertainty had an impact on emerging market stocks, further supporting the idea that markets that are thin are more affected by fear, information shocks, and the quick changes in investor positioning. As pointed out by Szczygielski et al. (2021), the level of uncertainty experienced by COVID had an uneven impact across regions, with different markets showing increased levels of resilience.

Comparative evidence is important not because NEPSE is not a smaller version of a developed market, but because NEPSE should not be seen in this light. Frontier markets are less liquid, more retail-driven, and less institutionalized. Some studies during the pandemic have demonstrated that restrictions and policy measures during that period had not only an impact on the returns but on the functioning of the market as well, such as in the work of Alaoui Mdaghri et al. (2021) and Zaremba et al. (2020). Evidence from the South Asian region is sparse, but existing evidence suggests that COVID-19 impacted market behavior and reduced market dynamism. This helps to ensure that a design is developed and not just extrapolated from bigger markets.

Nepalese Context: COVID-19 and NEPSE

There is less but valuable evidence specific to Nepal. To directly look at the impact of the pandemic period and relate the performance of NEPSE with the variables related to COVID, the lockdown period, and the rate of interest, Karki (2022) has done a study. Also, Malla et al. (2022) have reported that COVID-19 had a significant impact on the behavior of investors in the stock market in Nepal, whereas Pokhrel (2023) has provided a descriptive overview of the performance of NEPSE during the lockdown period. Dangol et al. (2023) are interested in volatility during catastrophic events, while Rana (2022) documents the persistence of return volatility during COVID-19 using the GARCH-family model's approach. Hamal and Gautam (2021) pool other existing evidence on stock volatility, returns, and government responses in the pandemic.

These studies, in the Nepalese and global context, are useful documents of the episode, but they also highlight the gap. The vast majority are descriptive, event-based, or volatility-based. They demonstrate that NEPSE was disturbed, volatility increased, the market recovered, and the nature of investors' activity evolved. The first question they don't ask themselves is how much of the observed return behavior is due to the pandemic regime, that is, how much is due to money easing and inflation, policy restrictions, cases, and trading conditions, all held constant. The limited number of Nepal studies is not the only reason for the gap. The gap is the lack of a credible causal estimate of the difference in the NEPSE returns due to the pandemic after accounting for contemporaneous macro-financial movements.

DATA SOURCES AND ESTIMATION DIAGNOSTICS

The data in this study comprises Daily NEPSE data and macroeconomic variables from 2 January 2018 to 28 December 2023. NEPSE index and turnover are sourced from publicly available data in Sharesansar (2024). The base interest rate and year on year inflation rate are sourced from Nepal Rastra Bank publications (Nepal Rastra Bank, 2024). Daily COVID-19 cases and Oxford COVID-19 government response tracker stringency index are obtained from Our World in Data (Mathieu et al., 2020). Since macroeconomic variables are commonly reported monthly, the variables were forward filled to align them with daily NEPSE return.

The COVID-19 treatment variable is considered as binary indicator. The trading from 24 March 2020 to 31 December 2021 is referred as one and zero for all trading days. This timeline is considered because 24 March 2020 marks the starting of Nepal National lockdown and the beginning of pandemic disruptions in regard to market activity, liquidity conditions and investor sentiments. The trading halt period in lockdown from 22 March to 12 May 2020 and a short interruption in late June 2020 are excluded from the study. These exclusions were necessary because authentic daily returns cannot be observed when the market is dormant.

The first COVID-19 case in Nepal was officially reported on 23 January 2020. However, the Our world in Data case recorded the first observation on 26 January 2020 indicating a reporting lag of three days. So, for the event study analysis the study has used the official confirmation date as the event anchor and the causal treatment window starts on 24 March 2020. After considering the trading halt exclusions, the final working sample data contains 1,371 trading days of which 362 trading days fall within COVID-19 treatment window.

Descriptive Statistics

Table 1 reports the descriptive behavior of NEPSE daily returns. In the pre-COVID treatment period, the mean daily returns were almost zero at ($M = -0.0002$, $SD = 0.0115$). The excess kurtosis was reported at 5.26, indicating that returns were already fat tailed before the start of pandemic which seems common in frontier equity markets. During the COVID shock window, the mean increases to ($M = 0.0019$, $SD = 0.0165$) daily, while the excess kurtosis declined to 1.82, suggesting that return movements became less vibrant than in the pre-COVID period because trading activity became larger and more frequent as new retail investors entered the market. In the post-COVID recovery period, the mean moves back to zero at ($M = -0.0003$, $SD = 0.0141$, $S_K = 0.73$) suggesting a stronger tendency toward positive returns movements in the recovery phase.

Table 1: Descriptive Statistics

Table 1 reports the descriptive statistics of NEPSE daily returns for the period pre-COVID, COVID-19 shock and post-COVID recovery. The table entails the number of observations (N), mean (M), standard deviation (SD), skewness (SK) and excess kurtosis. These statistics provide a brief overview of the return distribution, volatility, asymmetry and tail behavior of NEPSE across different phase of the pandemic.

Period	N	M	SD	S_K	Excess Kurtosis
Pre-COVID (2018/01–2020/03)	539	-0.0002	0.0115	-0.085	5.259
COVID shock (2020/05–2021/12)	362	+0.0019	0.0165	+0.065	1.818
Post-COVID (2022/01–2023/12)	470	-0.0003	0.0141	+0.725	1.198

Stationary, Heteroskedasticity and the Design Choice

Prior to estimating any treatment effect, the time series characteristics of every variable must be examined. This is very crucial because financial and macroeconomic data contain trends and changing volatility. The existence of unit roots or conditional heteroskedasticity would undermine inference based on traditional standard errors and would alter the orthogonalization step at the DML procedure. Augmented Dickey-Fuller tests (Dickey & Fuller, 1979) and KPSS tests (Kwiatkowski et al., 1992) are reported in Table 2 for each series in levels.

Table 2: Stationarity Test for Variables

The Augmented Dickey-Fuller (ADF) test calculates the null hypothesis of a unit root, while the KPSS test calculates the null hypothesis of stationarity. The result is based on the combined interpretation of both the tests at the 5% significance level.

Series	ADF stat	ADF p	KPSS stat	Verdict
NEPSE Index (level)	-1.25	0.651	3.44	Non-stationary
NEPSE Daily Return	-9.94	<0.001	0.18	Stationary
Turnover	-2.58	0.097	1.34	Non-stationary
log Turnover	-2.07	0.255	3.04	Non-stationary
Base Interest Rate	-1.36	0.601	1.3	Non-stationary
YoY Inflation	-1.95	0.311	2.53	Non-stationary
Stringency Index	-2.19	0.21	2.42	Non-stationary
COVID new cases	-5.76	<0.001	0.98	Inconclusive

The diagnostic results are clear. NEPSE daily returns are stationary as expected for equity returns. The other covariates fail to reject the unit-root null hypothesis. The forward filling of monthly series created a long flat segment that mimic random walk behavior and the stringency index and COVID cases count are zero before 2020 and therefore shows a structural break rather than true unit root. Aggregating such variables with a stationarity result can create spurious regression identified by Granger and Newbold (1974) where standard errors are unreliable even when the point estimation is consistent.

Hence, the design is stationarised before estimation. The base rate, year-on-year inflation and stringency index are first differenced. Turnover and COVID new cases are first difference of their natural logarithm with log1p are applied to handle the zero observations. After the series has been transformed re-testing, it has confirmed stationarity at levels reported in Table 3.

Table 3: Transformed Stationarity Test Results

Table 3 presents ADF and KPSS test results after applying the transformations in estimation. Δ denotes the first-difference operator, turnover and COVID-19 cases are log-transformed before differencing with $\log(1+cases)$ *used to account for zero case observations. The ADF test calculates the null hypothesis of unit root, while the KPSS test evaluates the null hypothesis of stationarity. The results obtained indicate that all variables used in DML estimation satisfy stationarity conditions at the 5% significance level.*

Transformed Series	ADF Stat	ADF p	KPSS Stat	Verdict
Δ Base Rate	-5.26	<0.001	0.52	Stationary (borderline)
Δ YoY Inflation	-36.97	<0.001	0.10	Stationary
Δ Stringency	-23.26	<0.001	0.06	Stationary
Δ log Turnover	-14.30	<0.001	0.04	Stationary
Δ log(1+Cases)	-13.87	<0.001	0.02	Stationary
Daily Return (Y)	-9.94	<0.001	0.18	Stationary

Conditional heteroskedasticity was examined as a second diagnostic because equity returns often exhibit periods of clustered volatility. Engle's ARCH-LM test (Engle, 1982) on daily returns yielded a statistic of 159.06 at lag ten ($p < 0.001$) and Breusch-Pagan test on a constant mean residual reported an LM statistic of 32.96 ($p < 0.001$). Both tests strongly reject homoskedasticity confirming the volatility-clustering pattern. Accordingly, HC3 heteroskedasticity-consistent standard errors were reported.

To explore the volatility pattern more clearly, GARCH (1,1) model was fitted with Student-t innovations NEPSE daily return. The use of Student-t innovations seems credible in this study because the return distributions comprise of heavy tails and extreme observations. Figure 1 exhibits the estimated conditional volatility path, the shaded are denotes the COVID-19 treatment window.

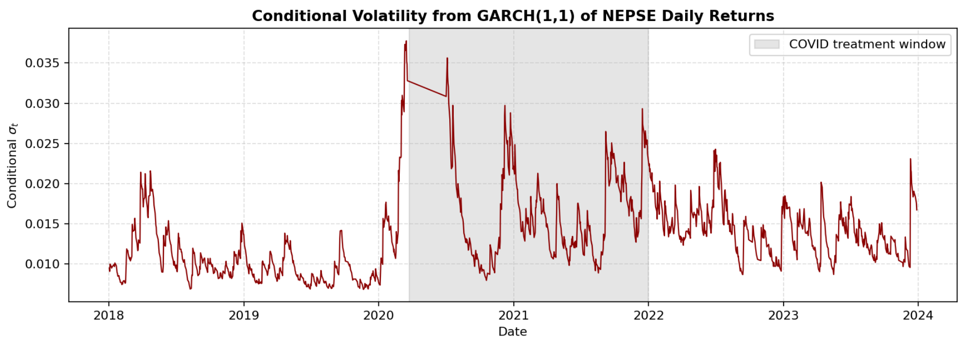

Figure 1: Conditional Volatility of NEPSE Daily Returns from a GARCH (1,1) Model

The figure exhibits conditional volatility estimated from GARCH (1,1) model with Student-t innovations. The shaded area represents the COVID-19 treatment window. Conditional volatility sharply increases around the beginning of the pandemic and increases through the treatment window before settling at higher post-pandemic levels.

Further, the study also examines how higher moments of the return distribution shift around the pandemic. Figure 2 plots the 60-day rolling skewness and excess kurtosis of daily NEPSE returns. These two measure are very useful because they capture the key aspects of market behavior which are not visible from the average returns alone. Skewness exhibits whether returns are unusually tilted towards large gains or losses, while excess kurtosis indicates whether huge returns movements became more frequent. It is noted that rolling skewness turned sharply negative in early 2020 indicating stronger downside pressure during the initial phase of the pandemic shock. Then it gradually recovered as the market adjusted to the new environment. Also, excess kurtosis increased repeatedly in late 2020 and mid-2021, suggesting that return movements were common during parts of the pandemic period. These outlines suggest that COVID-19 affected not only the average level of returns but also the distributional behavior of the market.

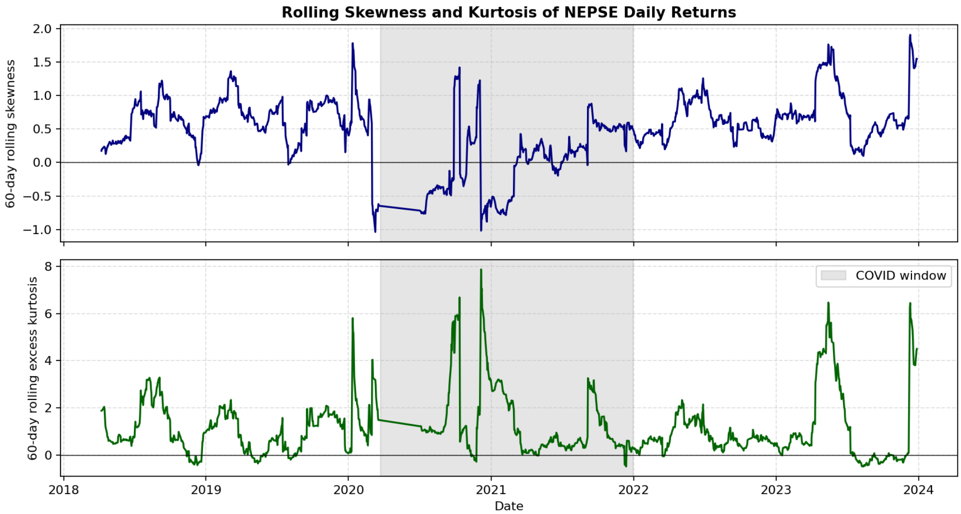

Figure 2: Sixty-Day Rolling Skewness and Excess Kurtosis of NEPSE Daily Returns

The upper graph presents 60-day rolling skewness while the lower graph reports 60-day rolling excess kurtosis. The shaded region represents COVID-19 treatment window. The figure shows a sharp negative shift in skewness during the early pandemic period and sustains increase in excess kurtosis during late 2020 and mid-2021, suggesting changes in asymmetry and tail risk during the COVID-19 pandemic.

Figure 3 exhibits the NEPSE index, COVID-19 case counts, policy stringency index, turnover, base rate and inflation rate in a common timeline. This helps to identify the market movement within the pandemic and macroeconomic environment. The figure shows three important shifts during COVID-19 treatment, a sharp rise in policy stringency at the beginning of pandemic, a sustained decline in the base rate through 2020 and 2021 and shows a clear increase in market turnover after the trading has resumed. This pattern shows that NEPSE's pandemic period fluctuations occurred alongside major shifts in public health restrictions, monetary conditions and investor activity which altogether account for the empirical design.

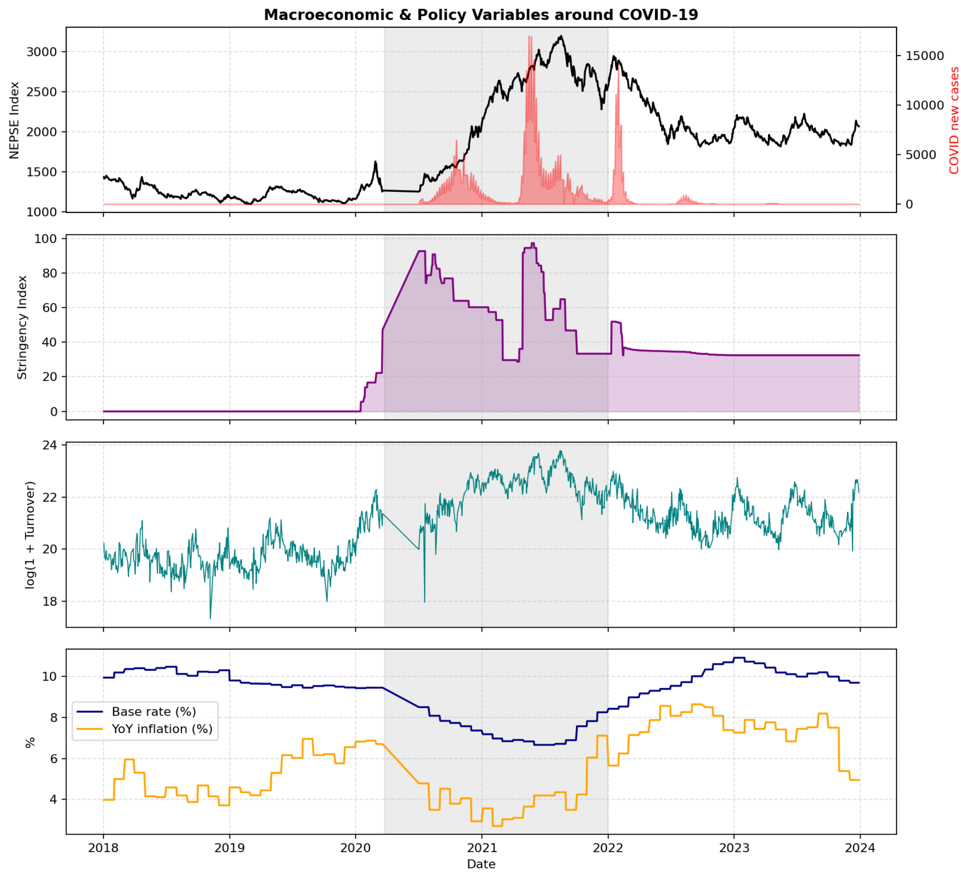

Figure 3: Macroeconomic and Policy Variables around COVID-19

The figure shows fluctuations of macroeconomic and policy variables, from top NEPSE Index with daily covid new cases, the OXFORD stringency policy index, log (1 + turnover) and base rate alongside year-on-year inflation. The shaded region in all graphs indicates COVID-19 treatment window.

METHODOLOGY

The study has adopted a causal inference design to estimate the effect of COVID-19 period on NEPSE daily returns. The analysis uses daily market data in conjunction with macroeconomic, policy, pandemic and trading activity. Since financial and macroeconomic time-series data contains trends and changing volatility, the study first conducts stationarity tests and heteroskedasticity checks before estimating the treatment effect. The main treatment effect is estimated using Double Machine Learning (DML) with Ridge regression used in the first stage nuisance models and HC3 robust standard errors used for inference. Also, additional checks such as continuous stringency treatment, sensitivity analysis, placebo permutation test and event-study cumulative abnormal returns are used to assess the stability and credibility of findings.

The Double Machine Learning Framework

Let Y_t denote the daily NEPSE return on day t, T_t the binary COVID-19 treatment indicator, and X_t a vector of stationarised macro-financial covariates. In the continuous-treatment specification, T_t is replaced by the daily change in the Oxford stringency index. For the binary specification, the parameter of interest is the average treatment effect (ATE):

$$\tau = E[Y_t(1) - Y_t(0)] \tag{1}$$

Only one potential outcome is observed for each trading day, so the ATE cannot be estimated through a direct comparison of observed outcomes. The study therefore uses the partialling-out approach of Double Machine Learning proposed by Chernozhukov et al. (2018). The procedure begins by estimating two nuisance functions:

$$m(x) = E[Y_t|X_t = x] \tag{2}$$

$$p(x) = E[T_t|X_t = x] \tag{3}$$

The first function estimates the expected daily return conditional on the covariates, while the second estimates the expected treatment exposure conditional on the same covariates. Both nuisance functions are estimated using Ridge regression on a standardized design matrix. The regularization parameter λ is selected through five-fold cross-validation over a logarithmic grid, $\lambda \in [10^{-4}, 10^{4}]$. For the outcome nuisance model, Ridge minimises the penalised residual sum of squares:

$$\hat{\beta}_{\text{Ridge}} = \arg\min_{\beta} \left\{ \sum_t (Y_t - X_t\beta)^2 + \lambda\|\beta\|^2 \right\} \tag{4}$$

The same regularized estimation logic is applied to the treatment nuisance model. To reduce overfitting, the sample is divided into five folds. In each fold, the nuisance models are trained on four folds and used to generate predictions for the held-out fold. This cross-fitting procedure is repeated across ten random seeds to reduce sensitivity to any single fold assignment. The revisualized outcome and treatment are then constructed as:

$$\tilde{Y}_t = Y_t - \hat{m}(X_t), \tilde{T}_t = T_t - \hat{p}(X_t) \tag{5}$$

The treatment effect is estimated in the final stage by regressing the revisualized outcome on the revisualized treatment:

$$\hat{\tau} = \left(\sum_t \tilde{T}_t^2 \right)^{-1} \sum_t \tilde{T}_t \tilde{Y}_t \tag{6}$$

Standard errors are computed using the HC3 heteroskedasticity-consistent estimator (MacKinnon & White, 1985). This correction is appropriate because the diagnostic tests show volatility clustering and unequal error variance in the daily return series. The reported point estimate is the average of the second-stage coefficients across the ten random seeds. The reported standard error combines the average within-seed HC3 standard error with the dispersion of the seed-level estimates:

$$SE(\hat{\tau}) = \sqrt{SE_{avg}^2 + SD_{seed}^2} \tag{7}$$

where SE_{avg} is the mean within-seed HC3 standard error and SD_{seed} is the standard deviation of the treatment-effect estimates across seeds. Under the maintained conditional ignorability assumption, where treatment assignment is independent of potential outcomes after conditioning on X_t, Equation (6) provides a consistent estimate of the ATE defined in Equation (1).

Why DML and Ridge?

DML is used because the empirical background requires more than a simple before and after comparison. The COVID-19 period coincided with changes in interest rates, inflation, policy restrictions, case intensity and trading activity. These variables may affect returns at the same time, and their relationships with the market are not likely to be strictly linear. DML is well suited to this setting because it first removes the part of both the outcome and the treatment that can be explained by the covariates and then estimates the remaining treatment effect. Its cross-fitting procedure also limits overfitting and reduces the regularization bias that can arise when machine-learning predictions are used directly for causal inference (Chernozhukov et al., 2018). This makes the final stage estimate more appropriate for hypothesis testing than a purely predictive model.

The choice of Ridge regression as the first-stage learner is based on empirical validation rather than assumption. Five candidate estimators were compared on the pre-treatment sample using rolling-origin cross-validation: Ridge, Lasso, Gradient Boosting, Random Forest, and the Theil-Sen estimator (Sen, 1968). The validation design uses a 360-day initial training window, a 20-day forecast horizon, and a 15-day forward step, producing 11 rolling folds. Table 4 reports the resulting out-of-sample RMSE values.

Table 4: Out-of-Sample RMSE from Rolling-Origin Cross-Validation on the Pre-Treatment Sample

Table 4. presents lower RMSE which indicates better predictive performance. The comparison uses 11 rolling-origin folds. Ridge is retained as the primary first-stage estimator based on its stable out-of-sample performance and the Diebold-Mariano predictive-accuracy test.

Estimator	Mean RMSE	Median RMSE	SD
Lasso ($\alpha = 0.1$)	60.35	66.34	21.68
Ridge ($\alpha = 1.0$)	60.49	66.67	21.68
Gradient Boosting	62.35	55.26	35.46
Random Forest	65.26	65.59	35.21
Theil-Sen	86.17	56.83	79.65

The results show that Ridge and Lasso perform almost identically. Their obtained mean RMSE values are 60.49 and 60.35 respectively with the same standard deviation of 21.68 across folds. This indicates that both regularized linear estimators adjust well in the pre-treatment data. Gradient Boosting and Random Forest produce a slightly higher mean RMSE values which stood at 62.35 and 65.26 and shows a greater variation across folds. Theil-Sen performs well in a typical fold as reflected in its median RMSE of 56.83, but its mean RMSE rises sharply to 86.17 because of a few poorly predicted folds. Its large standard deviation of 79.65 confirms that its performance is less stable across the validation.

A formal Diebold-Mariano test (Diebold & Mariano, 1995) comparing Ridge and Theil-Sen further supports the choice of Ridge. With the Harvey, Harvey et al. (1997) small-sample correction, the corrected statistic is 3.673 with a p-value of 0.0003, rejecting equal predictive accuracy in favor of Ridge. Ridge is therefore used as the primary nuisance estimator in the DML procedure. Theil-Sen is retained only as a robustness and transparency benchmark in the counterfactual analysis, allowing readers to see whether the visual counterfactual is sensitive to the choice of estimator.

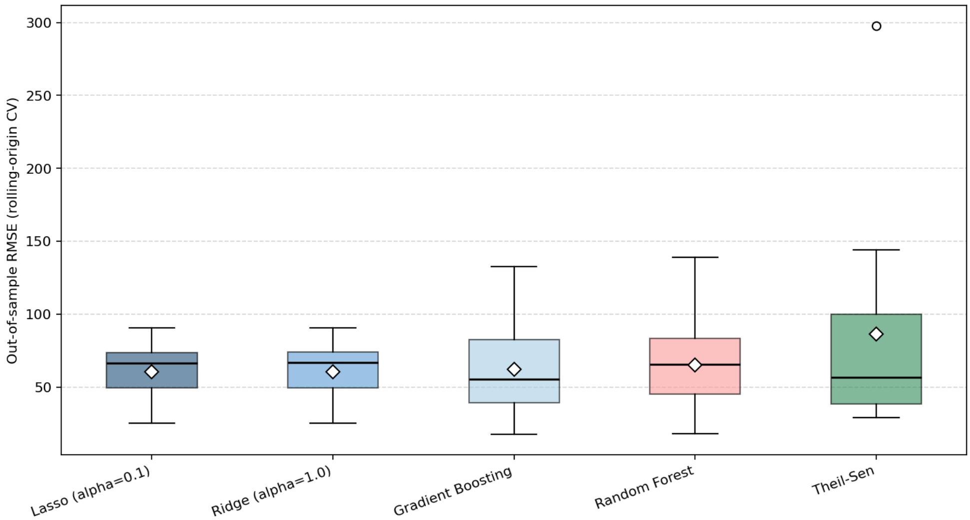

Figure 4: Predictive Accuracy across Counterfactual Estimators

The figure exhibits the distribution of out-of-sample RMSE values for five counterfactual estimators such as Lasso, Ridge, Gradient Boosting, Random Forest and Theil-sen. The diamond markers indicate mean RMSE while horizontal lines indicate median RMSE. Ridge and Lasso show the most stable predictive performance across folds. Thiel-sen achieves a competitive median RSME but has highest mean RMSE indicating it sensitivity to small number of unusual folds.

Q-Q plots of the pre-treatment residuals are presented in Figure 5. The plots show clear variations from normality for both the Ridge and Theil-Sen specifications particularly in the tails of the distribution. The Jarque-Bera tests also reject the null hypothesis of normally distributed residuals in both cases. This pattern is consistent with the heavy-tailed behavior which can be easily observed in financial return data. Accordingly, the empirical analysis reports HC3 heteroskedasticity-consistent standard errors rather than relying on classical OLS standard errors.

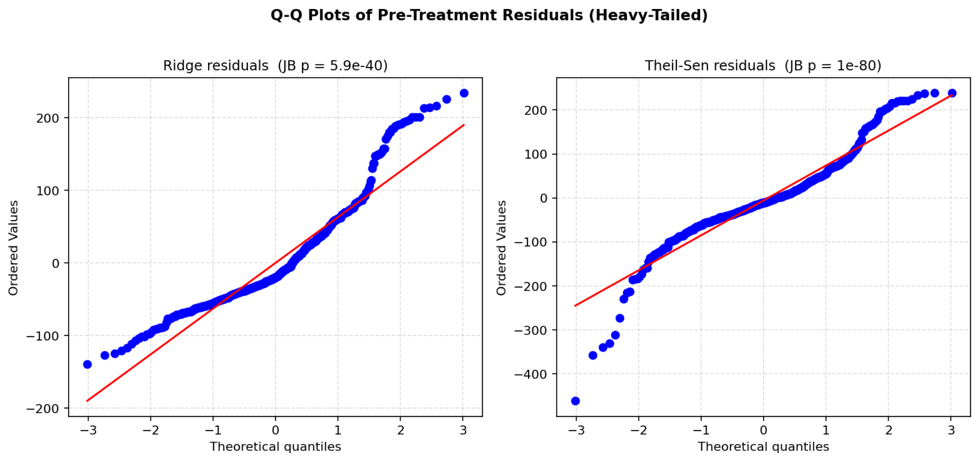

Figure 5: Q-Q Plots of Pre-Treatment Residuals from Ridge and Theil-Sen Models

In Figure 5, the left panel reports residuals from the Ridge specification while the right panel reports residuals from the Theil-Sen specification. Both residual distributions show heavy-tailed behavior and Jarque-Bera tests reject normality in both cases. This supports the use of heteroskedasticity-consistent inference in the main estimations.

Treatment Operationalization and Adjustment Sets The empirical analysis reports estimates using three nested adjustment sets and two treatment specifications. The first adjustment set, M1, includes the differenced base rate and differenced year-on-year inflation. These variables are treated as core macroeconomic confounders because they are determined outside the daily stock-market return process but may still influence market performance. The second adjustment set, M2 adds differenced log COVID-19 cases and differenced stringency, allowing the model to account for observable pandemic intensity and policy response. The third adjustment set, M3, includes differenced log turnover. Turnover is treated with caution because it may lie on the pathway between the pandemic shock and market returns. Hence, the M3 estimate is interpreted as a controlled direct effect, representing the part of the pandemic-period effect that does not operate through trading activity (Pearl, 2009).

The main treatment specification is the binary COVID-19 treatment indicator which is described in the data section. As a robustness check, the analysis also uses the daily change in the Oxford stringency index as a continuous treatment measure. This continuous specification is restricted to the

2020 to 2022 sub-period, during which the stringency index shows meaningful variation.

Counterfactual Construction

Together with the DML estimation, a synthetic counterfactual path is constructed for the NEPSE index. Ridge regression is fitted on the pre-treatment sample using smoothed macro-financial covariates and the fitted relationship is then extended into the treatment window to approximate the market path that would have been expected in the absence of the COVID-19 shock period. The Theil-Sen estimator is fitted on the same pre-treatment data as a robustness benchmark.

This counterfactual exercise is descriptive rather than causal. Its purpose is to show the size and persistence of the gap between the observed NEPSE index and the macro-implied benchmark. The actual treatment-effect evidence is provided by the DML estimates while the counterfactual path helps in explaining the economic meaning of those estimates.

Sensitivity to Unmeasured Confounding

The causal interpretation of the DML estimate depends on the conditional ignorability assumption. This means that after conditioning on the observed covariates X_t, treatment assignment T_t is assumed to be mean-independent of the potential outcomes. This assumption is important in the study because unobserved factors can influence NEPSE during the pandemic period such as global investor sentiment, vaccine news, cross-border financial spillovers and domestic expectations.

To assess the sensitivity of the estimated effect to such omitted factors, the analysis applies the partial$-R^2$ sensitivity framework developed by Cinelli and Hazlett (2020). For a hypothetical unmeasured confounder Z, the sensitivity analysis depends on two strength parameters:

$$R^2(Y, Z|T) \text{ and } R^2(T, Z|X) \tag{8}$$

The first parameter measures the explanatory strength of the unmeasured confounder for the outcome, after accounting for treatment and observed controls. The second measures the explanatory strength of the same confounder for treatment assignment, after accounting for observed controls. The robustness value indicates how strongly an omitted confounder would need to be related to both treatment and outcome to reduce the estimated coefficient to zero. The robustness value at the 5% significance level gives the corresponding threshold required to make the estimate statistically insignificant. The benchmark comparisons based on observed covariates are used to make these thresholds substantively interpretable.

Placebo Permutation Test

A placebo permutation test is conducted as an additional check on the DML procedure. In this test, a fake treatment window with the same length as the actual COVID-19 treatment window is assigned randomly within the pre-treatment sample. The full stationarised DML procedure is then repeated for each placebo assignment. The analysis uses 1,000 placebo placements, which allows the permutation $p - value$ to be estimated with reasonable precision (Good, 2005).

The purpose of this test is to examine whether the estimation procedure produces large treatment effects even when the treatment window is artificially assigned. If the procedure is well behaved, the placebo coefficients should cluster around zero, while the actual COVID-19 estimate should appear in the tail of the placebo distribution. The placebo exercise is therefore used as supportive evidence for the specificity of the estimated treatment effect.

Auxiliary Diagnostics

Two additional diagnostic exercises are used to interpret the market behavior surrounding the pandemic period. First, cumulative abnormal returns are computed over an eleven-day event window around three key dates, the first confirmed COVID-19 case in Nepal on 23 January 2020, the second-wave peak on 11 May 2021 and the full reopening on 1 September 2021. Expected returns are estimated from a 220-day pre-event window.

Second, the rolling skewness and excess kurtosis reported earlier are used to examine whether the pandemic period changed the shape of the return distribution. These diagnostics are not treated as separate causal tests but instead they provide additional evidence on volatility, asymmetry, tail behavior and the broader behavioral footprint of the COVID-19 shock period.

RESULTS

Primary Causal Estimate

Table 5 explains the estimated average treatment effect of the binary COVID-19 indicator on NEPSE daily index returns. The estimates are stated across three adjustments sets and two estimators, OLS and DML with HC3 robust standards errors. The result is stable across all six specifications. The COVID-19 period is associated with average daily returns between 0.20 and 0.22 percentage points above the macro-implied counterfactual with $p-values$ ranging from 0.023 to 0.049.

The proximity of the OLS and DML estimates is essential. Across all specifications, the difference between the two estimators is less than three basis points. This indicates that, once the covariates are properly stationarised the remaining non-linear confounding captured by the DML first stage seems limited. The null hypothesis of no average COVID-period effect is rejected in all combinators of estimator and adjustment set.

Table 5: Binary Average Treatment Effect of the COVID-19 Period on NEPSE Daily Returns

All covariates in the estimation are in stationary form. Standard errors are HC3 heteroskedasticity-consistent. DML estimates use Ridge as first-stage models with five-fold cross-fitting and ten random seeds. M3 includes differenced log turnover and is interpreted as a Controlled Direct Effect (CDE).

Specification	Est.	Coef.	SE	95% CI Low	95% CI High	p
M1 · ΔBase, ΔInfl	OLS	+0.00218	0.00096	+0.00031	+0.00406	0.023
M1 · ΔBase, ΔInfl	DML	+0.00211	0.00096	+0.00022	+0.00399	0.028
M2 · + ΔCases, ΔString	OLS	+0.00219	0.00097	+0.00029	+0.00410	0.024
M2 · + ΔCases, ΔString	DML	+0.00201	0.00102	+0.00002	+0.00400	0.048
M3 · + Δlog Turnover (CDE)	OLS	+0.00217	0.00097	+0.00027	+0.00408	0.025
M3 · + Δlog Turnover (CDE)	DML	+0.00200	0.00102	+0.00001	+0.00399	0.049

The estimated effects are moderate in daily terms but seem very meaningful over the full treatment period. It is noted that an average return increase of approximately 21 basis points per trading day, which is accumulated over the 362 trading days in the treatment window corresponds to a cumulative gap of roughly 75 to 85 log-points above the macro-implied path. This interpretation resonates well with Figure 6, where the actual NEPSE index remains visibly above the Ridge-based counterfactual for much of the COVID-19 treatment period. Overall, with the descriptive results in Table 1, the evidence indicates that the pandemic period was not only volatile than the surrounding periods but also associated with higher average returns after accounting for interest rates, inflation, policy stringency and COVID cases.

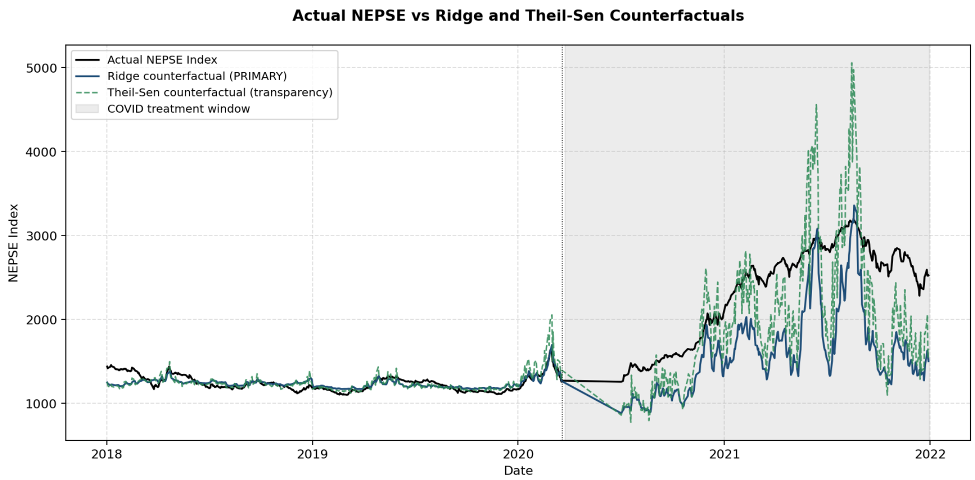

Figure 6: Actual NEPSE and Synthetic Counterfactual Paths

Compares the observed NEPSE index with Ridge and Theil-Sen counterfactual paths estimated from pre-treatment data. The dotted vertical line denotes the beginning of the treatment window, and the shaded area represents the COVID-19 treatment period. The observed NEPSE index remains above both counterfactual paths in the treatment period, which is consistent with the positive average treatment effect reported in Table 5.

Figure 7 shown below provides a compact visual summary of the estimates. Panel A reports the binary treatment across the three adjustment sets and two estimators. Panel B reports the corresponding continuous stringency estimates. The binary treatment estimates remain positive across all specifications while the continuous-stringency confidence intervals consistently cross zero. This visual pattern states that the estimated effect is connected to the broader COVID-19 period rather than day-to-day changes in policy stringency.

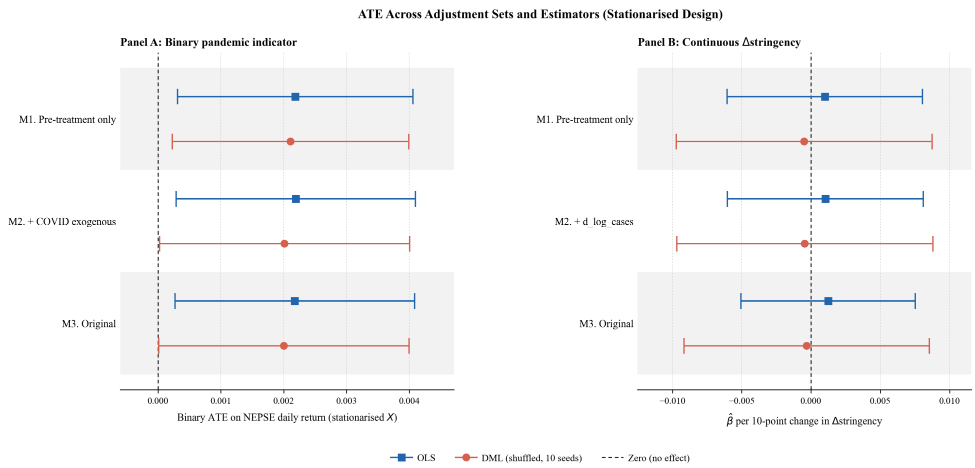

Figure 7: Average Treatment Effect Across Adjustment Sets and Estimators

Panel A reports the binary average treatment effect of the COVID-19 period. Panel B reports the continuous dose-response estimate based on 10-point change in differenced stringency. Bars here represent 95% confidence intervals computed using HC3 robust standard errors.

Continuous Treatment: Dose Response in Policy Stringency

Table 6 below replaces the binary COVID-19 indicator with the daily change in Oxford stringency index for the 2020 to 2022 sub-period. This model tests whether NEPSE daily returns moved when the government tightened or relaxed restrictions on a given day rather than in response to the COVID-19 pandemic as a whole.

The estimated effects are close to zero across all three adjustment sets and both estimators. All p-values remain above 0.70 indicating there is no statistically significant dose-response relationship. Also, when the coefficient is expressed as the effect of a 10-point change in stringency the implied daily return effect remains below 0.13 basis points in absolute value. The result hence suggests that day-to-day changes in restriction intensity did not independently move NEPSE returns once the pandemic-period effect was taken into account.

The positive effect which is reported in Table 5 appears reflect the wider pandemic environment rather than specific daily changes in government stringency. In simpler terms, the market response was linked to broader shift in liquidity conditions, investor behavior and macroeconomic adjustment during the COVID-19 period but aligned to each marginal tightening or loosening of restrictions.

Table 6: Continuous Dose-Response Effect of Differenced Stringency on NEPSE Daily Returns

Table 6 presents the estimated effect of a one-point and ten-point change in the differenced Oxford stringency index over the 2020 to 2022 sub-period.

Specification	Est.	β per 1 pt	β per 10 pts	p
M1 · ΔBase, ΔInfl	OLS	$+9.87 \times 10^{-5}$	$+9.87 \times 10^{-4}$	0.784
M1 · ΔBase, ΔInfl	DML	-4.97×10^{-5}	-4.97×10^{-4}	0.916
M2 · + Δlog Cases	OLS	$+1.03 \times 10^{-4}$	$+1.03 \times 10^{-3}$	0.775
M2 · + Δlog Cases	DML	-4.52×10^{-5}	-4.52×10^{-4}	0.924
M3 · + Δlog Turnover	OLS	$+1.23 \times 10^{-4}$	$+1.23 \times 10^{-3}$	0.702
M3 · + Δlog Turnover	DML	-3.19×10^{-5}	-3.19×10^{-4}	0.944

Why Stationarising the Design Matters?

Table 7 exhibits why pre-estimation diagnostics matter for the main findings. The M1 specification is estimated in two ways. The first version uses the base rate and year-on-year inflation in levels and the second version uses the first differences.

The differences between the two estimates are informative as when the covariates are used in levels the estimated effect of the COVID-19 period remains positive, but it is not statistically significant. The OLS coefficient is +0.00160 with a $p - value$ of 0.499 while the DML coefficient is +0.00155 with a $p - value$ of 0.513. When the same covariates are used in first differences the estimated effect becomes larger and statistically significant. The OLS coefficient increases to +0.00218 with a $p - value$ of 0.023 and the DML coefficient increases to +0.00211 with a $p - value$ of 0.028.

The contrast shows that the estimated treatment effect is very sensitive to the time-series features of the control variables. Since the level variables show non-stationary behavior while daily returns are stationary the level specification is vulnerable to the spurious-regression problem discussed by

Granger and Newbold (1974).

Table 7: Binary Average Treatment Effect Under Level and First Differenced Covariate Designs

Table 7 compares the M1 model using macroeconomic control levels and in first differences. The first-differenced is retained as the preferred estimation design because the stationarity diagnostics indicate that the macroeconomic controls should not be used in levels.

Design	OLS coef.	OLS p	DML coef.	DML p
Levels (invalid asymptotic inference)	0.0016	0.499	0.00155	0.513
Differences (valid)	0.00218	0.023	0.00211	0.028

Robustness to Unmeasured Confounding

Table 8 illustrates the Cinelli and Hazlett (2020) sensitivity analysis for the M1 binary treatment estimate. The estimated ATE is reported at +0.00218 with an HC3 $t - statistic$ of 2.57. Partial R2 of the treatment with the outcome after conditioning on the observed covariates is 0.48%. The robustness value is reported at 6.71% which states that an unobserved confounder would need to explain 6.71% of the residual in both the treatment and the outcome to reduce the estimated effect to zero. The robustness value at the 5% significance threshold is 1.62%.

Table 8: Cinelli-Hazlett Sensitivity Bounds for the M1 Binary Average Treatment Effect

RV = Robustness Value. The robustness value is 6.71% and RV at $\alpha = 0.05$ is 1.62%. Each benchmark reports the adjusted estimate and adjusted t-statistic under a hypothetical unmeasured confounder with strength equal to k times the named observed covariate.

Benchmark	$R^2(T, Z \mid X)$	$R^2(Y, Z \mid T, X)$	Adjusted estimate	Adjusted t
Unadjusted estimate	—	—	+0.00218	2.567
$1\times$ Δ base rate	0.0016	0.0016	+0.00213	2.508
$2\times$ Δ base rate	0.0032	0.0031	+0.00208	2.450
$3\times$ Δ base rate	0.0047	0.0047	+0.00203	2.392
$1\times$ Δ inflation	0.0003	0.0000	+0.00218	2.564
$2\times$ Δ inflation	0.0005	0.0000	+0.00218	2.561
$3\times$ Δ inflation	0.0008	0.0000	+0.00218	2.559

The benchmark results in this study suggest that the estimate is not easily overturned by the omitted-variable. Neither the differenced base rate nor differenced inflation comes near the threshold even when their observed explanatory power is multiplied by three. For instance, a hypothetical confounder three times as strong as the differenced base rate still leaves the estimated positive at +0.00203 with an adjusted $t-statistic$ of 2.392. Also, it is noted that inflation-based benchmarks produce smaller changes.

For the estimated effect to become statistically insignificant an omitted factor would need to be approximately 14 times stronger than the base rate benchmark. To reduce the estimated effect completely to zero it would need to be about 42 times stronger. Also, this does not mean hidden confounding is impossible because observational studies can never rule it out completely.

In-Time Placebo Permutation Test

In-time placebo permutation test was conducted to check whether the estimated COVID-19 period effect could be reproduced by randomly assigning a treatment window within the pre-treatment sample. A placebo window of the same length as the actual COVID-19 treatment window was randomly placed within the pre-treatment period and the full stationarised DML procedure was repeated 1,000 times.

The observed COVID-period estimate was reported at +0.002348 with $p = 0.006$. In terms of comparison, the mean placebo estimate was reported at -0.000836. The observed estimate falls approximately six standard deviations above the placebo mean and the permutation p-value is below 0.001. It was observed that none of the randomly assigned placebo windows produced an estimate comparable to the observed COVID-period effect. These results support the interpretation that the positive return effect is specific to the actual COVID-19 period rather than a mechanical outcome of the estimation procedure.

Also, it should be noted that placebo results should be read with caution because the placebo distribution is slightly negative on average and the nominal rejection rate is 41.1% which is higher than the conventional 5% benchmark. These results reflect the structure of the pre-treatment NEPSE timeline, where placebo windows may capture existing market trends and volatility rather than purely random variation. So, in this study placebo test can be treated as evidence rather than formal test of correct size. Hence, the

Cinelli-Hazlett sensitivity analysis remains the primary assessment of robustness to unmeasured confounding.

Event Study Cumulative Abnormal Returns

Table 9 presents cumulative abnormal results around three key COVID-19 dates. Figure 8 plots the corresponding Cumulative Abnormal Returns (CAR) over the eleven-day event window. The first confirmed COVID-19 case in Nepal was dated on 23 January 2020 which was followed by an abnormal return of -1.13% on the same-day. This initial decline was short lived and eleven-day CAR turned positive at +2.50%, indicating that first health shock appeared to have a brief negative reaction rather than a sustained market decline.

The second wave peak on 11 May 2021 showed a different behavior. On the same day abnormal returns stood positive at +0.64%, the five-day CAR reached +2.21% and the eleven-day CAR remained positive at +1.86% suggesting that by the second wave the market had already adjusted to pandemic situations and continued to rise even after the COVID cases count increased.

The full reopening on 1 September 2021 was followed by the sharpest negative adjustment. On the same day CAR negatively stood at -0.25%. But later the decline started to deepen reaching -5.81% over the eleven-day window. This pattern explains that part of pandemic-period rally was supported by conditions which are specific to that period, including liquidity, restricted movement and increase in retail trading.

Table 9: Cumulative Abnormal Returns Around Key COVID-19 Events

Table 9 presents an expected return estimated over a 220-day pre-event window. CAR [0] reports the abnormal return on the event day. CAR [0,4] and CAR [0,10] report cumulative abnormal returns over the five-day and eleven-day event windows respectively.

Event	CAR [0] (%)	CAR [0, 4] (%)	CAR [0, 10] (%)
First Nepal COVID case (2020-01-23)	−1.133	−0.130	+2.496
Second wave peak (2021-05-11)	+0.635	+2.206	+1.858
Full reopening (2021-09-01)	−0.248	−2.510	−5.806

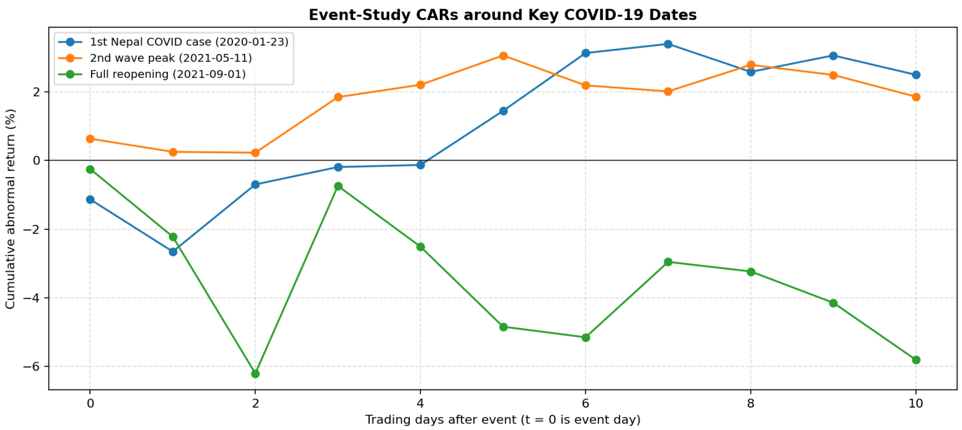

Figure 8: Cumulative Abnormal Returns Around Key COVID-19 Events

Figure 8 exhibits a cumulative abnormal return over the eleven-day window following each event. The first confirmed COVID-19 case and the second wave peak are followed by positive cumulative movement over the event window while the full reopening was followed by a sustained negative adjustment.

CONCLUSION

The study applied a Double Machine Learning framework to daily NEPSE data from 2018 to 2023 to estimate the causal impact of the COVID-19 pandemic on Nepali equity returns. After stationarising the design and using HC3-robust inference, the COVID period is associated with an average daily return of roughly twenty-one basis points above the macro-implied counterfactual. The estimate is stable across adjustment sets and across OLS and DML estimators also the estimate is robust to unmeasured confounders of plausible empirical strength as quantified by the Cinelli–Hazlett framework and not driven by within-period variation in policy stringency. An in-time placebo permutation test across 1,000 random reassignments confirms that the real estimate is specific to the actual COVID period, with a permutation p-value below 0.001 and the real estimate sitting six standard deviations above the placebo mean. In addition, auxiliary diagnostics documents a parallel shift in conditional volatility, higher moments, and event-study cumulative abnormal returns which are all consistent with a behavioral-rebound channel. NEPSE's resilience under the pandemic should not be understood as a case where COVID-19 had no effect. Its performance appears to have improved because lower interest rates, easier liquidity and rising retail trading activity supported the market during that period.

Policy Implications

The findings have four practical implications for Nepali regulators. First, the behavioral footprint of the pandemic period (in turnover, in higher moments, and in the divergence between returns and fundamentals) suggests that surveillance capacity should be expanded specifically for periods of monetary easing and increase in retail-inflow. In operational terms, this would mean publishing a regular dashboard that tracks turnover concentration, the share of total trading volume in the top decile of stocks, and circuit-breaker activations, with automatic regulatory review triggered when thresholds are crossed. Second, the surge in retail entry that reasonably drove the rebound highlights the value of mandatory short-form investor education at the point of new demat account opening, with a focus on volatility, leverage, and concentration risk. Third, the role of liquidity provision during the trading-halt period argues for formalizing the crisis-window liquidity facility that central bank used informally during 2020 into a standing rule-based mechanism with pre-announced eligibility and pricing. Fourth, the results show that market returns were influenced by overall crisis sentiment than by daily changes in policy restrictions. This suggests the need for clear and coordinated communication during periods of market stress. NEPSE, Nepal Rastra Bank and the Ministry of Finance should communicate frequently to reduce uncertainty and limit rumor-driven volatility.

Limitations and Future Research Directions

This study has a few drawbacks. First, some macroeconomic variables are monthly while NEPSE returns are daily. Forward filling was used to match the frequencies, but this may only smooth short-run movements in interest rates and inflation. Second, the binary COVID-19 indicators treat a complex pandemic period as one broad treatment window, so it may not capture differences across lockdown, reopening, second wave and recovery phases. Third, though Cinelli-Hazlett robustness value of 6.71% suggests reasonable steadiness against omitted confounding, unobserved factors such as investor sentiment, informal liquidity and cross-broader market influence cannot be fully ruled out.

Future researchers can extend this work by using higher-frequency data, firm level returns, sectoral indices, order-book records and direct measures of investor sentiment. They can also apply dynamic treatment effect models to scrutinize how the COVID-19 effect changed across different pandemic phases. A regional comparison with India, Sri Lanka, Bangladesh and other South

Asian markets would also help determine whether the NEPSE response was specific to Nepal or part of a wider frontier market pattern.

References

Alaoui Mdaghri, A., Raghibi, M., Thanh, T. L., & Oubdi, L. (2021). Stock market liquidity, the great lockdown and COVID-19 global pandemic nexus in MENA countries. *Review of Behavioral Finance*, *13*(1), 51–68. https://doi.org/10.1108/RBF-06-2020-0132

Al-Awadhi, A. M., Alsaifi, K., Al-Awadhi, A., & Alhammadi, S. (2020). Death and contagious infectious diseases: Impact of the COVID-19 virus on stock market returns. *Journal of Behavioral and Experimental Finance*, *27*, 100326. https://doi.org/10.1016/j.jbef.2020.100326

Ashraf, B. N. (2020). Stock markets' reaction to COVID-19: Cases or fatalities? *Research in International Business and Finance*, *54*, 101249. https://doi.org/10.1016/j.ribaf.2020.101249

Baig, A. S., Butt, H. A., Haroon, O., & Rizvi, S. A. R. (2021). Deaths, panic, lockdowns and US equity markets: The case of COVID-19 pandemic. *Finance Research Letters*, *38*, 101701. https://doi.org/10.1016/j.frl.2020.101701

Baker, S. R., Bloom, N., Davis, J. 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

Basuony, M. A. K., Bouaddi, M., Ali, H., & EmadEldeen, R. (2021). The effect of COVID-19 pandemic on global stock markets: Return, volatility, and bad state probability dynamics. *Journal of Public Affairs*, *21*(S1), e2761. https://doi.org/10.1002/pa.2761

Brunnermeier, M. K., Crockett, A., Goodhart, C. A., Persaud, A. D., & Shin, H. S. (2009). *The fundamental principles of financial regulation*. Geneva Reports on the World Economy.

Cao, L., & Ren, J. (2022). Machine learning shows that the COVID-19 pandemic is impacting U.S. public companies unequally by changing risk structures. *PLOS ONE*, *17*(6), e0269582. https://doi.org/10.1371/journal.pone.0269582

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, *21*(1), C1–C68. https://doi.org/10.1111/ectj.12097

Cinelli, C., & Hazlett, C. (2020). Making sense of sensitivity: Extending omitted variable bias. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *82*(1), 39–67. https://doi.org/10.1111/rssb.12348

Corbet, S., Larkin, C., & Lucey, B. (2020). The contagion effects of the COVID-19 pandemic: Evidence from gold and cryptocurrencies. *Finance Research Letters*, *35*, 101554. https://doi.org/10.1016/j.frl.2020.101554

Dangol, P., Rai, K. C., Sharma, K., Gc, A., Shrestha, N., & Khanal, A. (2023). Nepalese stock market volatility during catastrophic events (COVID-19). *New Perspective: Journal of Business and Economics*, *6*(1), 91–98. https://doi.org/10.3126/npjbe.v6i1.58922

Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, *74*(366), 427–431. https://doi.org/10.1080/01621459.1979.10482531

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics*, *13*(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, *50*(4), 987–1007. https://doi.org/10.2307/1912773

Ghising, S. (2025). *The influence of macroeconomic indicators on the Nepalese stock market* [Doctoral dissertation, Tribhuvan University] [Unpublished doctoral thesis].

Good, P. I. (2005). *Permutation, parametric, and bootstrap tests of hypotheses* (3rd). Springer.

Granger, C. W. J., & Newbold, P. (1974). Spurious regressions in econometrics. *Journal of Econometrics*, *2*(2), 111–120. https://doi.org/10.1016/0304-4076(74)90034-7

Hamal, J. B., & Gautam, R. R. (2021). Capital market response to COVID-19 pandemic: A systematic review on stock volatility and performance. *Marsyangdi Journal*, *2*(1), 27–49. https://doi.org/10.3126/mj.v2i1.39963

Haroon, O., & Rizvi, S. A. R. (2020). COVID-19: Media coverage and financial markets behavior—A sectoral inquiry. *Journal of Behavioral and Experimental Finance*, *27*, 100343. https://doi.org/10.1016/j.jbef.2020.100343

Harvey, D. I., Leybourne, S. J., & Newbold, P. (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting*, *13*(2), 281–291. https://doi.org/10.1016/S0169-2070(96)00719-4

Jiang, C., Talbot, D., Carazo, S., & Schnitzer, M. E. (2025). A double machine learning approach for the evaluation of COVID-19 vaccine effectiveness under the test-negative design: Analysis of Quebec administrative data. *Statistics in Medicine*, *44*(5), e70025. https://doi.org/10.1002/sim.70025

Karki, D. (2022). Navigating the new normal: Performance of stock market during pandemic. *Pravaha*, *28*(1), 119–132.

Khan, M., Kayani, U. N., Khan, M., Mughal, K. S., & Haseeb, M. (2023). COVID-19 pandemic & financial market volatility: Evidence from GARCH models. *Journal of Risk and Financial Management*, *16*(1), 50. https://doi.org/10.3390/jrfm16010050

Kumar, S., Kumar, A., & Singh, G. (2023). Gold, crude oil, Bitcoin and Indian stock market: Recent confirmation from nonlinear ARDL analysis. *Journal of Economic Studies*, *50*(4), 734–751. https://doi.org/10.1108/JES-01-2022-0007

Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, *54*(1–3), 159–178. https://doi.org/10.1016/0304-4076(92)90104-Y

Li, X., Li, W., Yu, X., Han, Z., & Jin, Q. (2025). Financial risk assessment of imbalanced data based on nonlinear causal time-series network. *Information Processing & Management*, *62*(3), 104025. https://doi.org/10.1016/j.ipm.2024.104025

MacKinnon, J. G., & White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of Econometrics*, *29*(3), 305–325. https://doi.org/10.1016/0304-4076(85)90158-7

Malla, K., Chhetri, R., & Bhandari, G. (2022). Impact of COVID-19 on Nepalese stock market. *The Lumbini Journal of Business and Economics*, *10*(1–2), 78–87. https://doi.org/10.3126/ljbe.v10i1-2.54207

Marks, H. (2018). *Mastering the market cycle: Getting the odds on your side*. Houghton Mifflin Harcourt.

Mathieu, E., Ritchie, H., Rodés-Guirao, L., Appel, C., Gavrilov, D., Giattino, C., Hasell, J., Macdonald, B., Dattani, S., Beltekian, D., Ortiz-Ospina, E., & Roser, M. (2020). COVID-19 pandemic. https://ourworldindata.org/coronavirus

Nepal Rastra Bank. (2024). *Current macroeconomic situation reports* (tech. rep.). NRB.

Pearl, J. (2009). Causal inference in statistics: An overview. *Statistics Surveys*, *3*, 96–146. https://doi.org/10.1214/09-SS057

Pokhrel, S. (2023). Stock market performance during COVID-19 lockdown: A descriptive analysis of Nepal Stock Exchange. *Adhyayan Journal*, *10*(10), 24–33. https://doi.org/10.3126/aj.v10i10.57346

Ramelli, S., & Wagner, A. F. (2020). Feverish stock price reactions to COVID-19. *The Review of Corporate Finance Studies*, *9*(3), 622–655. https://doi.org/10.1093/rcfs/cfaa012

Rana, S. B. (2022). Stock returns volatility in Nepal: Evidence during COVID-19. *Butwal Campus Journal*, *5*(1), 12–25. https://doi.org/10.3126/bcj.v5i1.50179

Salisu, A. A., Sikiru, A. A., & Vo, X. V. (2020). Pandemics and the emerging stock markets. *Borsa Istanbul Review*, *20*(S1), S40–S48. https://doi.org/10.1016/j.bir.2020.11.004

Sen, P. K. (1968). Estimates of the regression coefficient based on Kendall's tau. *Journal of the American Statistical Association*, *63*(324), 1379–1389. https://doi.org/10.1080/01621459.1968.10480934

Sharesansar. (2024). Nepal Stock Exchange indices and turnover data. https://www.sharesansar.com/

Song, F., Huang, Y., & Liu, C. (2025). Extended application of double machine learning in corporate financial resilience research: Based on data factor marketization. *Systems*, *13*(4), 292. https://doi.org/10.3390/systems13040292

Szczygielski, J. J., Bwanya, P. R., Charteris, A., & Brzeszczyński, J. (2021). The only certainty is uncertainty: An analysis of the impact of COVID-19 uncertainty on regional stock markets. *Finance Research Letters*, *43*, 101945. https://doi.org/10.1016/j.frl.2021.101945

Topcu, M., & Gulal, O. S. (2020). The impact of COVID-19 on emerging stock markets. *Finance Research Letters*, *36*, 101691. https://doi.org/10.1016/j.frl.2020.101691

Zaremba, A., Kizys, R., Aharon, D. Y., & Demir, E. (2020). Infected markets: Novel coronavirus, government interventions, and stock return volatility around the globe. *Finance Research Letters*, *35*, 101597. https://doi.org/10.1016/j.frl.2020.101597

Zhang, D., Hu, M., & Ji, Q. (2020). Financial markets under the global pandemic of COVID-19. *Finance Research Letters*, *36*, 101528. https://doi.org/10.1016/j.frl.2020.101528