European Research Studies Journal
Volume XXIX, Issue 3, 2026

pp. 1038-1060

From Oil Prices to Stock Trading: A Machine Learning Approach

Submitted 23/07/26, 1st revision 21/08/26, 2nd revision 31/08/26, accepted 30/09/26

Paul Handro[1]

Abstract:

Purpose: *Our study investigates to what extent the predictive information contained in the relationship between the price of West Texas Intermediate (WTI) crude oil and the evolution of the share price of the integrated energy company Chevron (CVX) can be exploited, through machine learning (ML) such as Gaussian Process (GP) and Deep Gaussian Process (DGP) models, to generate trading signals with economic utility. Specifically, the study examines whether the predictions of these models can be transformed into positive economic outcomes by applying long/short rules, including contrarian interpretation of the signal, and whether increasing the architectural depth from a single-layer GP to two and three layers DGPs generates incremental improvements in economic performance.*
Design/Methodology/Approach: *The methodology is based on a ML framework using GP and DGP models to investigate the predictive relationship between WTI crude oil prices and CVX stock prices. Three model architectures, comprising one, two, and three layers, are estimated and compared in terms of their statistical and economic performance. The model results are subsequently translated into LONG/SHORT trading signals and evaluated through an out-of-sample backtesting framework. Both the direct and contrarian signals generated by the model are examined to determine whether the statistical predictive information can be converted into economically exploitable trading signals. Furthermore, in our analysis we investigate the sensitivity of trading performance to market volatility regimes.*
Findings: *The results highlight the contrarian nature of predictive information and show that economic performance depends on the rule of transforming predictions into signals, while increasing the model depth does not guarantee superior economic outcomes.*
Practical implications*: The findings of our research have implications for both financial modeling and the design of an algorithmic trading system. Our evidence underlines that the contrarian strategies based on DGP predictions can outperform a simple buy-and-hold strategy, suggesting that the market inefficiencies may persist even in the case of a liquid asset such as Chevron.*
Originality/Value: *Our research contributes to the vast literature on the application of ML models in investments by demonstrating that the model is not evaluated solely by its predictive capacity, but rather by considering the entire process by which predictive information is transformed into investment decisions and, ultimately, into economic value.*

Keywords: *Gaussian process regression; Deep Gaussian processes; Oil–stock price nexus; Economic value of forecasts; Algorithmic trading.*

[1]*Ph.D., Candidate, West University of Timisoara, Romania, ORCID: 0009-0003-7951-1005*
e-mail: paul.handro76@e-uvt.ro;

JEL Classification: *C53, G17, Q47.*

Paper type: *Research article.*

1. Introduction

The modern financial market is characterized by a dynamic state of high complexity, periods of increased volatility and a lot of uncertainty. These features can be amplified by exogenous pressure, such as the economic environment, geopolitical factors, even climate change or endogenous pressure within the market, such as investors' sentiment, transactional volumes or technical analysis and many others. These forces create an environment that is difficult to predict and which is governed by a large uncertainty.

The recent history, starting with the pandemic breakout March 2020, the inflationary pressure from 2022, the geopolitical wars, the speculative behaviour on the crypto market, the election cycles and isolationism's new agenda in some politics, together with the US aggressive tariffs negotiation, raises the market's volatility and risk tail in the markets.

In this context, the traditional widely use autoregressive linear models, as ARMA or GARCH models, have shown their limitations because they assume linear relations between variables and normal distributions. These assumptions limit their capacity to capture the stylized facts from financial data, such as leverage effect, volatility clustering or other nonlinearities in the latent processes. The rise of ML techniques, especially deep learning, which have a better ability to handle big data and acknowledge the nonlinearity of the input features and the output, can provide better predictions than linear models (Jiang, 2021).

Moreover, there is an ongoing academic debate regarding the financial time series feature as a deterministic one or a probabilistic one (Harvey, 1997). However, practitioners and academia are accepting that we cannot predict exactly (deterministic), conversely in terms of probabilities we can capture the uncertainty of our prediction.

Additionally, Rao (2010) concluded that between two major treatments of the financial time series (deterministic and stochastic) there should be also considered as important the suggestion of checking for Perron’s structural breaks as per the deterministic trend. However, due to the particularity of the stochastic approach that

has the capabilities to quantify the uncertainty (risk), the Bayesian approach in modelling is receiving more and more attention (Zhu *et al.,* 2018).

In the literature we found different types of Bayesian approaches to ML models. One of the new approaches that is gaining popularity lately is the Gaussian Process Regression (GPR) a non-parametric model which can generalize a Gaussian probability distribution (Williams and Rasmussen, 2006) and capture the uncertainty (Bisht *et al.,* 2022; Suphawan *et al.,* 2022). The GPR models were initially used in Geostatistics (Bisht *et al.,* 2022) and after that, it received recognition in other fields such as electricity usage (Zeng *et al.,* 2020), in chemistry (Deringer *et al.,* 2021), in robotics (Deisenroth *et al.,* 2015) and many other domains.

One of the advantages of using GP is that it can handle diverse functions with a few parameters, even when we do not have enough data. On the other hand, the single-layer GP models are limited by the flexibility of the kernels to capture the complex relationships. To overcome this disadvantage and to retain the advantages of a single-layer GP, a DGP with multiple layers can be used (Salimbeni and Deisenroth, 2017).

However, in the stock market prediction and other financial time series, we identified a scarcity regarding the use of DGP. Although DGP are recognized as powerful tool for uncertainty quantification in deep learning, it remains largely unexplored in financial applications, as underlined by recent surveys (Blasco *et al.,* 2024).

Although ML models have demonstrated a high capacity to capture nonlinear relationships in financial series, a significant part of the literature continues to evaluate them predominantly through predictive performance indicators. However, statistical accuracy does not automatically translate into economic value, because the investment outcome also depends on the rule by which the prediction is transformed into a long or short position, the transaction horizon and the implementation costs.

Buczynski *et al.* (2021) show that conventional predictive metrics used in isolation provide insufficient information about the investment utility of a model. Therefore, the first gap lies not in the absence of economic assessments, but in the still limited analysis of the relationship between predictive performance, the mechanism for transforming predictions into trading signals, and the economic outcomes. Starting from this, the study formulates the first research question:

RQ1: To what extent does the predictive performance of GP/DGP models translate into economic value when predictions are transformed into trading signals?

A second gap concerns the economic incremental effect of model depth. Salimbeni and Deisenroth (2017) show that adding layers to DGP architectures can expand the capacity to represent nonlinear relationships and improve predictive performance

without necessarily generating overfitting. However, the available evidence on financial series suggests that additional complexity does not automatically guarantee superior economic performance (see e.g., Kynigakis and Panopoulou, 2022), moreover, systematic comparisons between GP and DGP with successive depths, evaluated simultaneously predictively and economically, remain limited.

Consequently, it is not clear whether the predictive advantages of a deeper architecture also translate into incremental increases in the profitability of a trading model. For this reason, we can say that this gap underpins the second question:

RQ2: How does increasing model depth – from a single layer GP to two – and three layers DGP architectures - affects model predictive and economic performance of the associated trading strategies.

In this context, to provide empirical evidence, our starting point was to analyze the relationship between two asset classes for constructing a prediction model based on a GP and DGP framework.

The oil price, expressed through WTI, represents a key indicator for the global energy market. It directly influences both the revenue side and the cost of the integrated oil and gas company. On one hand, an increase in oil prices enhances the profitability in the upstream segment (exploration and production); on the other hand, it creates pressure on the refining and distribution margins within the downstream segment.

Regarding the selected company for our research, we focus on CVX, one of the top companies in the energy industry that operates as an integrated energy firm, given its exposure to all the segments: upstream, midstream, and downstream. Further, the choice of CVX is motivated by its integratedposition across the oil and gas value chain, the highliquidity of its shares, and strong dependence on WTI price dynamics.

Moreover, the statistical correlations between the two above mentioned variables base on Pearson (p= 0.72) and Spearman (p= 0.66) confirm a strong long-term relationship which consistent with previous studies such as Degiannakis *et al.* (2018) and Broadstock and Filis (2014).

Next, the paper presents the relevant literature, describes the data and methodology used, analyzes the empirical data, and synthesizes the main conclusions and implications of the research.

2. Literature Review

Regression based on GP is a Bayesian non-parametric method that has gained increased attention in finance applications and research. The advantage of the predictions generated by GP model in comparison with the deterministic models is

that of a full predictive distribution (Williams and Rasmussen, 2006), which means both the predicted value and the uncertainty. The GPR models are both for classification and regression applications; the first is predicting continuous quantities, and the second is assigning discrete class labels (Williams and Rasmussen, 2006).

In the literature, we find both simple applications of the GP model, a hybrid approach, and an extension of the model. For example, Suphawan *et al.* (2022) is comparing the results obtained by the GPR in predicting the stock price of Thailand (SET) with the predictions of deterministic parametric models such as artificial neural network (ANN) and recurrent neural networks (RNN). The dataset is divided in two, gravitation around the pandemic period (COVID-19): one includes the whole period (2015-june 2020) and the second is a subset that is from January 2020 to June 2020.

The results are compared using metrics such as MAE, RMSE, MAPE, and NSE and found that GPR outperforms the ANN and RNN, is less sensitive to the number of lags, but is more sensitive to the outliers. Kim et al. (2022) forecast the price of WTI and Brent oil using as data input the selection of the most influential stocks out of 74 companies that are part of the SandP 500.

Based on that, the authors used three prediction models: deep learning, vine copula regression with Nonlinear Principal Component Analysis, and GPR. The results revealed that the GPR model is outperforming the deep learning model, particularly a back propagation neural network, and is highly competitive with vine copula regression, even outperforming it in terms of MAD for the Brent forecast.

Abunasser *et al.* (2023) comparing the efficiency of 14 ML algorithms in stock prediction for the period 2015-2020, found that GPR produced superior predictions in terms of R^2, MSE, and MAE. The input feature were enhanced beside the raw data with technical analysis indicators such as MA, RSI and oscillators.

Next, we will explore the hybrid usage of GP models in the literature. Ou and Wang (2011) proposed a hybrid and Bayesian approach to forecast the volatility, by combining GARCH type with GPR capturing the non-linearity and the asymmetry from Nasdaq indexes. The results showed that the hybrid and Bayesian approach offers better predictions than classical models, especially when linear and polynomial kernels are used.

Tu *et al.* (2024) using the closing price of three stocks listed one of the Chinese stock market (SSE) is comparing the predictions obtained by a GPR model, with the prediction of a GPR with combined kernels. Moreover, it is adding the features of an ARIMA model to a GPRC to enhance the capacity to capture both the linear and non-linear relationships between variables. The results prove that the hybrid Gaussian Process Regression with a Combined Covariance Function (GPRC) model

is superior to GPR and the combination ARIMA+GPRC outperforms the standalone models ARIMA, ANN and GPRC. Wang et al.(2021a) is using GPR to predict the uncertainty at the final of a chain compound of decomposition algorithm (SSA), a SVM with enhance weight to predict the decomposed elements, adding a RNN to integrate the previous results from the EWSVM in a single-valued forecast.

The model was used to forecast two indices (Nikkei 225 Index, and Hang Seng Index), and the data comprises more than 4600 days with five input features of raw data. Finally, the authors are comparing the model with the other eight and conclude that the hybrid model SSA-EWSVM-RNN-GPR shows good forecast accuracy.

Additionally, Wang *et al.* (2021b) using this time a variational mode decomposition and an autoencoder algorithm is used to capture the relevant signals and to reduce the dimension and the noise. This output is used in two layers of deep learning (RNN and LSTM), and in the final stage, the GPR is used to quantify the uncertainty. The authors conclude that the framework provides a reliable forecast.

In the same way, He and Wang (2024) is the coupling of the GPR model with a selection of algorithms and ML prediction techniques for forecasting three Chinese indices, one US, and one European. In this framework, the GPR model had an important role in probabilistic inference and extracting the relevant data by isolating the noise and adding robustness to the final results.

Göçken *et al.* (2019) using Harmony Search, a meta-heuristic algorithm, configured the optimal structure of 6 deterministic ML models and one Bayesian model, namely GPR. Based on the HS algorithm, six technical indicators were selected from the 22 available in the study as input features for the GPR model.

The authors outline two important aspects in modelling: first, the importance of the transfer function, and, in the case of GPR, the importance of the covariance function; and second, input feature selection. The models were tested on three stocks part of the BIST-100, and the result showed that HS-GPR did not perform the best; instead HS-JRNN outperformed.

Due to some obvious limitations of using a simple GP (single-layer), such as outlier sensitivity, low scalability, kernel's expressivity (Salimbeni and Deisenroth, 2017) some hybrid approaches or extended version of GP models are necessary to be used for improving the accuracy of the prediction.

Winarni *et al.* (2022) using as training data of 2200 closing prices of one US stock, addresses the scalability limitation of GP models. The authors overcome this limitation by using an extension, namely the partition resampling technique combined with hyperparameter averaging. The method preserves the GPR structure while reducing computation time, making it valuable for other applications in stock predictions.

Wolff *et al.* (2020) employed a multi-output Gaussian process to model the co-movement of three distinct asset classes represented by four financial series, such as commodity (gold and oil), stocks (Nasdaq) and currency (USD Index). Moreover, the authors tested a set of expressive kernels that fit better and found that Multi-Output Spectral Mixture obtained the best results but has the highest complexity.

DGP are considered a new class of algorithms that deal with the limitation of other ML algorithms, mainly governed by the black box method, that do not offer information regarding the confidence level of the output (Biggio *et al.,* 2021; Casolaro *et al.,* 2023). The precursor of the DGP is the initial Lawrence and Moore's (2007), DG- latent variable model (LVM), which formalized the GP with multiple layers over successive latent spaces.

On the other hand, the model had the limitation of using only MAP (Maximum a Posteriori) as an inference method, which generates the risk of overfitting. Further, Damianou and Lawrence (2013) formalized the standard DGP, where each layer governs the latent inputs to the next. Unlike the previous models, this model captures hierarchical structures and uncertainty through fully Bayesian inference. Salimbeni and Deisenroth (2017) showed that using doubly stochastic inference for a DGP, many regression and classification applications can be used without intervention. Moreover, their model is overperforming a single layer GP significantly.

Dwarakanath *et al.* (2022) developed two new algorithms for buy and sell decisions and they extend these to a DGP. The advantage of the new model is that it includes the uncertainty in the final decision, which gives more information to the decision-maker. The model was tested with success on real data such as intraday day stock price, closing price, and treasury bills yield curve, and performed better than Deep optimal stopping.

Ming *et al.* (2023) propose a DGP model that uses stochastic imputation, which treats latent layers as missing data, enabling an efficient inference and analytically tractable predictions of means and variances. The results of the DGP model with stochastic imputation inference are comparable to the fully Bayesian inference, but with lower computational cost.

Although numerous studies document the predictive superiority of ML methods over linear models and, in certain configurations, relevant investment gains (Gu *et al.,* 2020), statistical performance does not automatically translate into economic value. Leung *et al.* (2021) show that statistical advantages may generate more limited economic benefits, and Avramov *et al.* (2021) highlight the deterioration of profitability in the presence of transaction costs and high turnover.

At the same time, the economic outcome depends on how predictions are transformed into positions, not just on their accuracy (Chalvatzis and Hristu-

Varsakelis, 2020). Moreover, Buczynski *et al.* (2021), argue that predictive metrics provide limited information on the economic utility of a model and that investment performance should be evaluated separately from the statistical accuracy of the predictions.

Thus, a first gap that remains in our analysis is that the understanding of the relationship between predictive performance, the signal transformation rule and the resulting economic performance is still incomplete.

The second identified gap concerns the incremental effect of model depth on economic performance. Although increasing depth can improve the predictive ability of DGP models, it is not sufficiently clarified whether moving from one GP to two or three layer DGPs also generates successive improvements in trading outcomes.

Salimbeni and Deisenroth (2017) show that adding layers to DGP architectures can improve predictive performance in general regression and classification problems, without necessarily inducing overfitting.

However, available financial evidence suggests that the relationship between complexity and performance is not necessarily monotonic: Kynigakis and Panopoulou (2022) find that shrinkage methods and shallow neural networks achieve the best individual results, and Cakici *et al.* (2024) identify limited marginal benefits of complexity. Although Dwarakanath *et al.* (2022) demonstrate the utility of GP/DGP in financial decisions, it remains insufficiently clarified whether moving from GP to DGP with successive depths produces incremental economic improvements in trading strategies.

Regarding the correlation between oil and stocks there is an emerging consensus in the literature that oil prices have major implications for stock market performance, but the nature and intensity of these effects remain hotly debated. The bibliometric review by Bashir (2022) shows that research on oil price shocks and spillover effects on stock market returns has expanded rapidly in the last two decades, with hundreds of studies published investigating this nexus from various perspectives.

While some works (Ciner, 2001; Huang *et al.,* 1996; Jones and Kaul, 1996) have identified direct, mostly negative, relationships between oil shocks and stock market returns, others (Faff and Brailsford, 1999) show that the effects of oil shocks differ by sector: extractive and resource industries benefit from price increases, while sectors such as transport and packaging are negatively affected.

Recent literature also makes increasingly clear distinctions between developed and emerging markets, between oil importers and exporters, and between aggregate, sectoral, and firm-level analyses (Degiannakis et al., 2018; Smyth and Narayan, 2018).

Additionally, Crawford *et al.* (2021) show that firm fixed effects explain much more variation in earnings and investment than industry or time effects, and investors react particularly to oil-related earnings.

Much more recent studies (e.g., Logun, 2025) analyze the relationship between oil prices and international stock markets from a nonlinear perspective. The author's findings suggest that the interdependence between stock markets and oil prices is complex, asymmetric, and dependent on the level of market development. These findigs supports why nonlinearity should be taken into account in portfolio diversification, risk management, and economic policymaking strategies.

Overall, this framework confirms that the oil–stock market relationship is complex, context-dependent, and actor-structure-dependent, which justifies the application of flexible modeling methods capable of capturing these heterogeneities.

3. Data and Methodology

In this section, we describe the data used for the empirical application. We collected daily price series for the WTI (CL=F) and CVX pairs from Yahoo Finance, using the R yfR package (Perlin, 2023). From these series, we calculated adjusted logarithmic returns, based on daily closing prices.

The analyzed interval covers the period January 2020 to June 2025 , a period characterized by high volatility in commodity market and stocks, including the initial shock of the COVID-19 pandemic and subsequent developments generated by the shocks in oil demand, followed by the energy crises due to the March 2022 the Ukraine's invasion.

The data series were synchronized and missing values were removed using the na.omit() procedure. We divided the data set into two subsamples: (i) the training set, consisting of the first 1200 observations, used for model estimation, and (ii) the test set, consisting of the remaining observations, used for evaluating predictive performance.

The results of the Pearson (0.72) and Spearman (0.66) correlations indicate a positive and significant relationship between the price of Brent oil and the price of Chevron shares, suggesting that oil market developments are directly reflected in the stock market dynamics of this company. This result is consistent with the existing literature, which highlights that the effects of oil prices on stock markets depend on the level of data aggregation and the specifics of the analyzed firms, especially in the case of integrated energy companies (Degiannakis *et al.,* 2018).

Moreover, Smyth and Narayan (2018) highlight that the more recent literature not only distinguishes between aggregate, sectoral or firm-level analyses, but also considers the location of companies, i.e. whether they operate in an oil-exporting or

oil-importing country, as well as their role as producers or consumers. The choice of the WTI– Chevron pair aligns with this approach, as WTI represents a US oil quote and Chevron is a company integrated into the US energy sector.

Table 1. *Descriptive statistics*

Ticker	*Sample*	*Mean*	*SD*	*Min*	*Max*	*Skewness*	*Kurtosis*
CL=F	*All (2020-01-01 to 2025-07-01)*	*0.0005*	*0.0329*	*-0.2822*	*0.3196*	*0.0240*	*26.9220*
CVX	*All (2020-01-01 to 2025-07-01)*	*0.0002*	*0.0223*	*-0.2501*	*0.2049*	*-1.0990*	*26.4320*
CL=F	*Training (first 1200 obs.)*	*0.0007*	*0.0343*	*-0.2822*	*0.3196*	*0.0398*	*26.2470*
CVX	*Training (first 1200 obs.)*	*0.0002*	*0.0230*	*-0.2501*	*0.2049*	*-1.0733*	*26.5160*
CL=F	*Test (179 obs.)*	*-0.0009*	*0.0219*	*-0.0896*	*0.0701*	*-0.6706*	*5.4000*
CVX	*Test (179 obs.)*	*-0.0003*	*0.0171*	*-0.0858*	*0.0643*	*-1.4000*	*9.5500*

Source: *Author's own calculation.*

The high values of kurtosis reported in Table 1 confirm that the distributions of returns have fat tails, a phenomenon well documented in the financial literature. Cont (2001) identifies the presence of fat tails as one of the fundamental stylized facts of financial series, suggesting that the probability of extreme events occurring is significantly higher than that predicted by a normal distribution. In the same vein, Tsay (2005) shows that for a normal distribution, kurtosis has the value 3, and a positive excess of kurtosis ($K - 3 > 0$) indicates leptokurticity, reflecting a higher concentration of the mass of the distribution in the tails.

The contrast between the extreme values of kurtosis obtained on the entire sample (26.9 for WTI and 26.4 for Chevron) and the lower, but still significantly positive, values in the test sample (5.4 for WTI and 9.5 for Chevron) highlights the fact that the intensity of the phenomenon varies over time. Although the distributions appear closer to normal in the recent period, they remain leptokurtic, which highlights the persistence of the risk of extreme events.

These structures can have important methodological implications: the use of models that assume strictly normal innovations can lead to an underestimation of the probability of severe shocks.

Therefore, the structure of the data suggests the opportunity for modeling in a GP or a DGP framework, which can capture nonlinear dependencies and complex relationships between variables.

As documented in the literature (Crawford *et al.,* 2021; Degiannakis *et al.,* 2018; Smyth and Narayan, 2018), we used WTI crude oil as an explanatory variable and Chevron (CVX) as a dependent variable.

To forecast CVX price direction, we apply the GP and DGP with one, two, and three layers, based on the Matérn kernel and Markov Chain Monte Carlo (MCMC) methods. Prior to model estimation as Table 2 shows we test the performance of the model for different lags of WTI on Chevron returns.

The root mean square error (RMSE) and log predictive score or log-score (SCORE) values indicate small differences between specifications, but the 4-day lag yields the smallest RMSE (0.7853) and the least negative SCORE (−0.356). This result suggests that the influence of oil movements on Chevron returns manifests itself with a lag of several days, consistent with the idea that markets gradually integrate oil price information into the valuation of energy companies.

Table 2. *The model performance*

Metric	Lag 1	Lag 2	lag 3	Lag 4
RMSE	0.7927	0.7863	0.7925	0.7853
SCORE	-0.3862	-0.3654	-0.3933	-0.356

Source: *Author's own calculation.*

Further, we split the data between the training dataset, which includes the first 1.200 observations, and the testing dataset, which includes the last 179 observations.

Based on the literature metrics, we assess the performance of each layer using RMSE, Continuous Ranked Probability Score (CRPS), and log-score, which evaluate the absolute level of accuracy. We apply the Diebold-Mariano test to see if the differences between the models are statistically significant.

We started the methodology using a GPR model where the target variable is estimated as:

$$y(x) = f(x) + \varepsilon, f(x) \sim GP\big(m(x), k(x, x')\big), \varepsilon \sim N(0, \sigma 2). \quad (1)$$

We used a Matérn kernel defined as follows:

$$k_{\nu}(r) = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}r}{1}\right)^{\nu} K_{\nu}\left(\frac{\sqrt{2\nu}r}{l}\right), \quad (2)$$

where $r = |x - x'|$, ℓ is the scale length and ν the degree of smoothness.

Furthermore, we add a second and third layer, developing the model into DGPwith the latent function that results from the composition of several GP:

$$\mathrm{y} = f^{(l)} \circ f^{(l-1)} \circ \cdots \circ f^{(1)}(\mathrm{x}) + \varepsilon\,, \quad (3)$$

with each layer $\mathrm{f}^{(l)} \sim \mathrm{GP}(\mathrm{m}_l, \mathrm{k}_l), l = 1, \ldots, L$.

The hyperparameters have been estimated using a fully Bayesian model, implying a MCMC with 20.000 iterations. To foster and efficiently perform the calculations, we applied the Vecchia approximation documented by Katzfuss and Guinness (2021) and Vecchia (1988).

4. Research Results and Discussion

Table 3 presents the forecasting performance of the GP and DGP models. The results show that the deep architectures (two- and three-layer DGP) achieve lower RMSE (point accuracy) and CRPS (distributional accuracy) values compared to the simple GP, suggesting a better marginal accuracy in both point and distributed forecasting. In particular, the three-layer DGP model records the best SCORE (–0.068), confirming the improvement in forecast quality with model depth, although the marginal gains beyond two layers are small.

Table 3. *The forecasting performance of the GP and DGP*

Model	RMSE	CRPS	SCORE
One-layer GP	0.788876	0.500031	-0.3622
Two-layer DGP	0.780742	0.491333	-0.12826
Three-layer DGP	0.780667	0.493554	-0.0678

Source: *Author's own calculation.*

The Diebold–Mariano (DM) as can be seen in Table 4 tests do not reject the null hypothesis of equal predictive accuracy between the models, regardless of the forecast horizon considered (p-values being consistently greater than 0.6). This indicates that although the DGP models provide slightly better performances in terms of RMSE and CRPS indicators, the differences compared to the single-layer GP model are not statistically significant.

Table 4. *The Diebold–Mariano tests*

Horizon	Comparison	Statistic	p-value
1	f1 vs f2	0.87778	0.81
1	f1 vs f3	0.46257	0.6782
5	f1 vs f2	0.43712	0.669
5	f1 vs f3	0.28242	0.6112
15	f1 vs f2	0.29559	0.6162
15	f1 vs f3	0.2176	0.5861
25	f1 vs f2	0.29716	0.6168
25	f1 vs f3	0.26747	0.6054
50	f1 vs f2	0.35623	0.6392
50	f1 vs f3	0.62506	0.734
100	f1 vs f2	0.4342	0.6679
100	f1 vs f3	0.79815	0.7876

Source: *Author's own calculation.*

Overall, the results suggest that DGP models are better at capturing nonlinear dependencies and fat-tailed behaviors of data, but the additional benefits of deeper architectures are not statistically significant.

Next, based on these forecasting results, we investigated whether the GP and DGP models can be used for trading decisions based on the directional price signal. For

this purpose, the variation of the WTI oil price (symbol CL=F) was used as a predictor for the evolution of theCVX stock, and the signals generated by the models (positive or negative) were interpreted as LONG (buy) or SHORT (sell) recommendations for the next four days.

The results obtained in the test period (179 observations, lag = 4 days) highlight that the raw signals generated by the models had a contrarian behavior with respect to the actual market movement, indicating a negative correlation between the estimated and the observed direction. Consequently, reversing the model signal (–sign) led to a significant improvement in the directional accuracy and profitability of the trading simulation.

Although the forecasting results showed a marginal performance of the DGP models compared to the classic GP model, the practical application of these models in the context of trading decisions revealed an interesting, contrarian behavior.

In the simulation, the variation in the price of WTI oil (symbol CL=F) was used as an explanatory variable for the prediction of the price of CVX, with a four-day lag. This choice reflects the assumption that recent changes in the price of oil influence, with a short delay, the stock performance of integrated companies in the energy sector.

However, the directional analysis of the signals generated by the models indicated a negative correlation between the forecasted movement and the actual movement of the CVX price. Specifically, in more than 55% of the cases, the signal estimated by the model had the opposite sign to the actual variation of the share price, which suggests that the GP and DGP models learned an inverted pattern in the relationship between WTI and CVX.

Table 5. *Empirical results*

Model	Accuracy (original)	Accuracy (inverted)	Cumulative profit (inverted)
GP 1 Layer	40.60%	59.40%	6.15
DGP 2 Layers	46.90%	53.10%	4.29
DGP 3 layers	45.70%	54.30%	3.87

Source: *Author's own calculation.*

The empirical results, as can be seen in Table 5, highlight a contrarian behavior of the DGP models. In their original form, the signals generated by the models showed a directional accuracy below the 50% threshold, indicating a higher probability of wrong predictions than correct ones.

However, after reversing the sign of the prediction (i.e., interpreting the model signal in the opposite direction), the directional accuracy increased considerably, ranging from 54% to 59% depending on the architecture used. In addition, the cumulative

return expressed in standardized units (z-score) went from negative values in the case of the original signal to positive values after the reversal, empirically confirming the contrarian nature of the strategies based on the DGP models. This standardized positive performance, of up to 6.15 standard deviations above the mean, suggests that the models consistently capture opposite market reactions, providing a statistical basis for interpreting the opposite price behavior.

From an economic perspective, this contrarian behavior is consistent with the fundamental dynamics of the relationship between the price of crude oil and the shares of integrated oil companies, such as Chevron. In the short term, a rapid increase in the WTI price often signals supply tensions in the raw materials market and an increase in refining costs, which temporarily erode downstream margins.

Thus, even if in the long term a high oil price is favorable to overall profitability, in short horizons (1–5 days) the immediate effect on the stock price can be negative, due to speculative adjustments and pressures on operating costs. This short-term reverse reaction explains why models based solely on the variation of oil tend to generate signals that mispredict the immediate direction of CVX shares — a rise in WTI is initially interpreted as a positive signal, but is empirically followed by a decline in the stock, confirming the contrarian nature identified statistically.

The evaluation of predictive performance was extended by simulating trading strategies based on signals generated by the GP and DGP models, using a four-days window, corresponding to the time lag identified between the price of WTI oil and CVX.

The logic of the strategy is based on the assumption that the GP model with one, two, or three layers provides a predictive signal that tends to follow the trend movement of the training period. Reversing the signal turns the prediction into a contrarian strategy, designed to exploit the directional overcompensation of the model and generate economic profit.

The results summarized in the Table 6 show that, over the entire test period (179 observations), 175 transactions were made, almost evenly distributed between long and short positions.

Table 6. *Trading summary results*

Trades based on GP 1 Layer signals						
Trade type	Number of trades	Profit (Value)	Number of trades	Profit (Value)	Total number of trades	Profit /Loss
LONG	32	-144	54	175	86	30.82
SHORT	44	-140	45	178	89	37.95
Grand Total	76	-284	99	353	175	68.76
Trades based on DGP 2 Layers signals						

Trade type	Number of trades	Profit (Value)	Number of trades	Profit (Value)	Total number of trades	Profit /Loss
LONG	43	-99	46	107	89	8.23
SHORT	39	-79	47	95	86	15.36
Grand Total	82	-178	93	202	175	23.59
Trades based on DGP 3 Layers signals						
Trade type	Number of trades	Profit (Value)	Number of trades	Profit (Value)	Total number of trades	Profit /Loss
LONG	34	-159	53	154	87	-4.37
SHORT	45	-161	43	163	88	2.76
Grand Total	79	-319	96	318	175	-1.61

Source: *Author's own calculation.*

The GP model (single-layer) performs the best, with a cumulative profit of +68.76 USD, coming from 99 profitable trades out of 175. The distribution of the profitability is distributed comparatively evenly between long positions (+30.82 USD) and short positions (+37.95 USD), showing the model's capability to catch short-term trend reversals.

Next, the DGP model with two layers remains in positive performance, obtaining +23.59 USD, but with a lower level of profitable signals. The total number of profitable trades decreases to 93, and the profit distribution indicates a slight dominance of short positions (+15.36 USD), showing the same ability to better exploit the trend reversals as the one-layer model.

Conversely, the three-layer DGP model generates a negative result (-1.61 USD), with 96 profitable trades, but a fragile balance between losses from long positions (-4.37 USD) with marginal gains from short positions (+2.76 USD). We can argue that the model remains directionally correct in about half of the trades; the complexity introduced by the third layer leads to excessive smoothing of the predictions.

Finally, we can argue that the model depth does not automatically translate into higher profitability. Moreover, simpler models (one-layer GP) seem to capture local price movements better, while deeper versions (two- or three-layer DGP) tend to provide more stable predictions, but are less exploitable in trading.

The results also indicate that, for short trading horizons, the model's ability to correctly capture the direction of price movements and sensitivity to local variations are more operationally relevant than the overall prediction accuracy, measured by mean errors such as RMSE or CRPS.

The selection of the two variables was based on both economic and statistical arguments. From an economic perspective, the price of oil is a fundamental determinant of the financial performance of integrated oil companies, especially for

the upstream segment, which justifies a positive relationship between the dynamics of the WTI price and that of Chevron shares.

From a statistical perspective, we evaluated the relationship between the two variables through the Pearson (p=0.72) and Spearman (p=0.66) correlation coefficients, and the results indicate a strong positive relationship in the long run.

We report these results in the context of the existing literature to validate the robustness of our findings. In this line, we find that prior research such as, Degiannakis et al., (2018) and Broadstock and Filis (2014), confirms the existence of a positive correlation between the price of oil and the stock market performance of the integrated firms, more accentuated in the period of oil price expansion.

The existence of other evidence in the literature adds relevance to the choice of the pair (WTI, CVX) for our analysis, providing a potentially valid framework for applying DGP models in investigating nonlinear relationships and dynamic dependencies in the energy sector.

However, before estimating the model, we tested its performance with different lags. The performance results were analyzed using metrics such as RMSE and SCORE, and we found that the most efficient lag is four days. At first glance, the results showed poor performance, but with a sign reversal —namely, applying a contrarian behavior in the short term —the performance of the model was significantly improved.

The finding regarding the improved results based on an opposite sign does not contradict the positive structural relationship observed in the long run, but rather reflects at least two important aspects: (i) the delayed and asymmetric response of the stock market to oil market shocks and (ii) the model overestimating the trend from the training period.

Specifically, a sudden increase in oil prices can generate short-term pressures on the downstream segments (refining and distribution), which causes a temporary reverse adjustment of Chevron's stock price. The observed reversal process is consistent with Narayan and Sharma (2011), from two standpoints: firstly, the underreaction hypothesis, which postulates that the investors process the public information gradually, reacting with delay to price shocks. Secondly, the results confirm the threshold effect hypothesis, which posits that the magnitude of the oil variation determines different reactions in the capital market.

Moreover, the ML-based contrarian strategies have been previously addressed in the literature as solutions to correct systematic errors of “overreaction” or trend overestimation. Huang *et al.* (2011) propose a dual-classifier framework that combines behavioral finance principles with advanced ML algorithms, such as AdaBoost and Support Vector Machines. This system learns to identify stocks with a

high probability of return reversal, transforming the classic contrarian strategy into an intelligent mechanism for adaptive portfolio adjustment. Similarly, Zheng (2024) formulates a functional contrarian strategy in a two-asset framework, where portfolio exposure is dynamically generated according to deviations from a reference price, providing a formal tool for rebalancing investment decisions.

Our results obtained during the backtesting period based on a contrarian interpretation of the model output, provide promising economic results for two out of a total of three layers considered. Practically, when applied to the CVX price series, a simple SHORT/LONG strategy with a short time horizon of four days based on the inverted DGP signal, the strategy yielded positive results, especially for the one and two layer configurations.

Similar to the conclusions of Salimbeni and Deisenroth (2017), who demonstrated that DGP architectures tend to smooth the latent function and reduce local variations, our results confirm that additional model depth causes a decrease in sensitivity to rapid price movements.

This smoothness effect explains why, within contrarian strategies on short horizons, DGP models with two or three layers lose local accuracy in favor of an overall stability of the prediction.

In the same sense, the study by Göçken *et al.* (2019) emphasized the importance of the covariance function (kernel) in determining the performance of the GPR model, showing that its selection influences the ability to capture local variations.

In addition, the studies of Biggio *et al.* (2021) and Casolaro *et al.* (2023) argue that the fundamental advantage of DGPs lies in their ability to provide a robust probabilistic estimate of uncertainty and to avoid the "black box" nature of deep learning algorithms. Our results confirm this feature of stability and robustness.

Also, Dwarakanath *et al.* (2022) have shown that DGPs can provide superior results in automated trading decisions, due to the integration of uncertainty into the decision-making process.

Moreover, recent studies indicate that embedding GPR within hybrid forecasting pipelines or combining it with decomposition, feature-selection, and feature-extraction techniques can improve point or interval prediction relative to standalone or less elaborate benchmarks. Tu *et al.* (2024) report that ARIMA–GPRC outperforms GPRC, while Wang et al. (2021a, 2021b) and He and Wang (2024) document predictive gains from SSA–EWSVM–RNN–GPR, VMD–AE–RNN–LSTM–GPR, and RFECV–SAE–GPR frameworks across several stock markets.

However, these studies neither isolate the effect of GP/DGP depth nor establish whether greater model complexity generates superior economic trading performance.

5. Conclusions

In this study, we modeled the complex relationship between the price of crude oil, expressed by the WTI index, used as an explanatory variable (predictor), and the share price of the integrated energy company CVX, used as a dependent variable.

In relation to the first research question, the results show that the predictive performance of GP/DGP models can be translated into economic value, but this transformation is not direct or automatic, but depends on the rule used to generate trading signals. The initial signals recorded a directional accuracy below 50% and negative economic results, while their reversal increased the directional accuracy to 54–59% and allowed for positive cumulative profits for two of the three architectures.

The results thus suggest that the models capture contrarian information, which becomes economically exploitable only after appropriate signal transformation. Regarding the second question, increasing the model depth did not lead to successive improvements in economic performance.

For the 175 simulated trades, with an exposure equivalent to one share and a holding horizon of four days, the single-layer GP achieved the highest cumulative gross profit (+68.76 USD), followed by the two-layer DGP (+23.59 USD), while the three-layer DGP generated a loss (-1.61 USD). Therefore, in the empirical framework analyzed, the simpler architecture proved to be the most economically efficient, and adding layers did not guarantee a superior exploitation of predictive information.

The study contributes to the literature in several ways. First, it integrates predictive and economic evaluation into a common framework, highlighting that the statistical performance of a model is not sufficient for the stability of its investment utility.

Second, by comparing the initial signal with the inverted signal, the analysis demonstrates the essential role of the rules by which predictions are transformed into long/short positions and shows that the information provided by the model can become economically exploitable through a contrarian interpretation.

Third, the comparison between a one-layer GP and two- and three-layer DGPs allows the assessment of the incremental effect of depth on predictive and economic performance, results indicating that deeper architecture does not automatically generate superior trading results.

Finally, the study complements a literature predominantly focused on the relationship between oil prices and aggregate stock indices, moving the analysis to the level of an individual stock and applying GP/DGP models to the relationship between WTI and CVX, to highlight energy company-specific reactions that stock index dynamics may mask.

In this context, although the results are promising, there are some limitations that should be mentioned.

First, the analysis was limited to a single stock (CVX) and a single explanatory factor (WTI price), which restricts the generalizability of the conclusions to the entire energy sector or to other asset classes and sectors.

Second, the four-day holding horizon was chosen empirically and may not be optimal in other volatility regimes.

Third, the DGP models are computationally complex and their training can be affected by local minima, which influences the stability and convergence of the parameters.

Finally, trading costs and liquidity constraints were not included in the backtest simulation, which may slightly overestimate the net profitability of the proposed strategy.

Future research can extend this work in three main directions: methodological, empirical and practical.

Based on this work, some future research avenues can be exploited. First perspective, from a methodological standpoint, future research should take into account testing the applicability of other kernel functions within GP and the DGP architecture, such as Rational Quadratic, Spectral Mixture, or Periodic kernels, to estimate their capacity to capture regime-dependent or cyclical behaviour.

A comparative framework that includes other kernel architectures, applied to one or more pairs of assets can provide additional evidence regarding the performance of each of these kernels in different contexts and asset classes. Taking into consideration the empirical perspective, our analysis could be extended to different stocks linked to commodities.

Also, including different time horizons and volatility regimes can evaluate the model's robustness under different market conditions. Additionally, incorporating liquidity risk, trading costs, and risk-adjusted performance indicators provided by Sortino and Sharpe ratios, the results would provide a more realistic assessment of the economic and practical applicability of the described strategy.

The findings of our research have implications for both financial modeling and the design of an algorithmic trading system. Our evidence underlines that the contrarian strategies based on GDP predictions can outperform a simple buy-and-hold strategy, suggesting that the market inefficiencies may persist even in the case of a liquid asset such as Chevron.

Moreover, extending this idea, the adaptive ML models can capture nonlinear corrections and behavioral asymmetries that traditional econometric models tend to neglect.

References:

Abunasser, B.S., Daud, S.M., Abu-Naser, S.S. 2023. Predicting Stock Prices using Artificial Intelligence: A Comparative Study of Machine Learning Algorithms. International Journal of Advances in Soft Computing and Its Applications, 15(3).

Avramov, D., Cheng, S., Metzker, L. 2021. Machine Learning vs. Economic Restrictions: Evidence from Stock Return Predictability (SSRN Scholarly Paper No. 3450322). Social Science Research Network. https://doi.org/10.2139/ssrn.3450322.

Bashir, M.F. 2022. Oil price shocks, stock market returns, and volatility spillovers: A bibliometric analysis and its implications. Environmental Science and Pollution Research, 29(16), 22809-22828. https://doi.org/10.1007/s11356-021-18314-4.

Biggio, L., Wieland, A., Chao, M.A., Kastanis, I., Fink, O. 2021. Uncertainty-Aware Prognosis via Deep Gaussian Process. IEEE Access, 9, 123517-123527. https://doi.org/10.1109/ACCESS.2021.3110049.

Bisht, A., Chahar, A., Kabthiyal, A., Goel, A. 2022. Stock Prediction using Gaussian Process Regression. 6th International Conference on Computing Methodologies and Communication (ICCMC), 693-699. https://doi.org/10.1109/ICCMC53470.2022.9754114.

Blasco, T., Sánchez, J.S., García, V. 2024. A survey on uncertainty quantification in deep learning for financial time series prediction. Neurocomputing, 576, 127339.

Broadstock, D.C., Filis, G. 2014. Oil price shocks and stock market returns: New evidence from the United States and China. Journal of International Financial Markets, Institutions and Money, 33, 417-433. https://doi.org/10.1016/j.intfin.2014.09.007.

Buczynski, W., Cuzzolin, F., Sahakian, B. 2021. A review of machine learning experiments in equity investment decision-making: Why most published research findings do not live up to their promise in real life. International Journal of Data Science and Analytics, 11(3), 221-242.

Cakici, N., Shahzad, S.J.H., Będowska-Sójka, B., Zaremba, A. 2024. Machine learning and the cross-section of cryptocurrency returns. International Review of Financial Analysis, 94, 103244. https://doi.org/10.1016/j.irfa.2024.103244.

Casolaro, A., Capone, V., Iannuzzo, G., Camastra, F. 2023. Deep Learning for Time Series Forecasting: Advances and Open Problems. Information, 14(11), Article 11. https://doi.org/10.3390/info14110598.

Chalvatzis, C., Hristu-Varsakelis, D. 2020. High-performance stock index trading via neural networks and trees. Applied Soft Computing, 96, 106567. https://doi.org/10.1016/j.asoc.2020.106567.

Ciner, C. 2001. Energy shocks and financial markets: Nonlinear linkages. Studies in Nonlinear Dynamics and Econometrics, 5(3).

Cont, R. 2001. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative finance, 1(2), 223.

Crawford, S., Markarian, G., Muslu, V., Price, R. 2021. Oil prices, earnings, and stock returns. Review of Accounting Studies, 26(1), 218-257. https://doi.org/10.1007/s11142-020-09556-7.

Damianou, A., Lawrence, N.D. 2013. Deep gaussian processes. Artificial Intelligence and Statistics, 207-215.

Degiannakis, S., Filis, G., Arora, V. 2018. Oil Prices and Stock Markets: A Review of the Theory and Empirical Evidence. The Energy Journal, 39(5), 85–-130. https://doi.org/10.5547/01956574.39.5.sdeg.

Deisenroth, M.P., Fox, D., Rasmussen, C.E. 2015. Gaussian Processes for Data-Efficient Learning in Robotics and Control. IEEE Transactions on Pattern Analysis and Machine Intelligence, 37(2), 408-423. https://doi.org/10.1109/TPAMI.2013.218.

Deringer, V.L., Bartók, A.P., Bernstein, N., Wilkins, D.M., Ceriotti, M., Csányi, G. 2021. Gaussian Process Regression for Materials and Molecules. Chemical Reviews, 121(16), 10073-10141. https://doi.org/10.1021/acs.chemrev.1c00022.

Dwarakanath, K., Dervovic, D., Tavallali, P., Vyetrenko, S., Balch, T. 2022. Optimal Stopping with Gaussian Processes. Proceedings of the Third ACM International Conference on AI in Finance, 497-505.

Faff, R.W., Brailsford, T.J. 1999. Oil price risk and the Australian stock market. Journal of Energy Finance and Development, 4(1), 69-87. https://doi.org/10.1016/S1085-7443(99)00005-8.

Göçken, M., Özçalıcı, M., Boru, A., Dosdoğru, A.T. 2019. Stock price prediction using hybrid soft computing models incorporating parameter tuning and input variable selection. Neural Computing and Applications, 31(2), 577-592.

Gu, S., Kelly, B., Xiu, D. 2020. Empirical asset pricing via machine learning. The Review of Financial Studies, 33(5), 2223-2273.

Harvey, A. 1997. Trends, Cycles and Autoregressions. The Economic Journal, 107(440), 192-201. https://doi.org/10.1111/1468-0297.00152.

He, X., Wang, J. 2024. A Hybrid Forecasting System Based on Comprehensive Feature Selection and Intelligent Optimization for Stock Price Index Forecasting. Mathematics (2227-7390), 12(23). https://search.ebscohost.com/login.aspx?direct=trueandprofile=ehostandscope=siteandauthtype=crawlerandjrnl=22277390andAN=181656407andh=hIqG4TQRQH1kQdMaKn1D8iTWn%2F9oBrYwZSOVzbHn7XRj24zlkgwR8ARXl4yuz18KJ%2FFaw8yteWR4kcMNPtOMhQ%3D%3Dandcrl=c.

Huang, R.D., Masulis, R.W., Stoll, H.R. 1996. Energy shocks and financial markets. Journal of Futures Markets, 16(1), 1-27.

Huang, S.H., Lai, S.H., Tai, S.H. 2011. A learning-based contrarian trading strategy via a dual-classifier model. ACM Trans. Intell. Syst. Technol., 2(3), 20, 1-20. https://doi.org/10.1145/1961189.1961192.

Jiang, W. 2021. Applications of deep learning in stock market prediction: Recent progress. Expert Systems with Applications, 184, 115537. https://doi.org/10.1016/j.eswa.2021.115537.

Jones, C.M., Kaul, G. 1996. Oil and the stock markets. The Journal of Finance, 51(2), 463–491.

Katzfuss, M., Guinness, J. 2021. A general framework for Vecchia approximations of Gaussian processes. Statistical Science, 36(1), 124-141.

Kim, J.M., Han, H.H., Kim, S. 2022. Forecasting crude oil prices with major SandP 500 stock prices: Deep learning, Gaussian process, and vine copula. Axioms, 11(8), 375.

Kynigakis, I., Panopoulou, E. 2022. Does model complexity add value to asset allocation? Evidence from machine learning forecasting models. Journal of Applied Econometrics, 37(3), 603-639. https://doi.org/10.1002/jae.2885.

Lawrence, N.D., Moore, A.J. 2007. Hierarchical Gaussian process latent variable models. Proceedings of the 24th International Conference on Machine Learning, 481-488.

Leung, E., Lohre, H., Mischlich, D., Shea, Y., Stroh, M. 2021. The promises and pitfalls of machine learning for predicting stock returns. Available at SSRN 3546725. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3546725.

Logun, A. 2025. The Relationship between World Stock Markets and Oil Prices: A Nonlinear Analysis. İşletme Araştırmaları Dergisi, 17(1), 31-44.

Ming, D., Williamson, D., Guillas, S. 2023. Deep Gaussian Process Emulation using Stochastic Imputation. Technometrics, 65(2), 150-161. https://doi.org/10.1080/00401706.2022.2124311.

Narayan, P.K., Sharma, S.S. 2011. New evidence on oil price and firm returns. Journal of Banking and Finance, 35(12), 3253-3262. https://doi.org/10.1016/j.jbankfin.2011.05.010.

Ou, P., Wang, H. 2011. Volatility Prediction by Treed Gaussian Process With Limiting Linear Model. International Journal of Modelling and Simulation. https://www.tandfonline.com/doi/abs/10.2316/Journal.205.2011.2.205-5498.

Perlin, M. 2023. yfR: Downloads and Organizes Financial Data from Yahoo Finance [Computer software]. https://github.com/ropensci/yfR.

Rao, B.B. 2010. Deterministic and stochastic trends in the time series models: A guide for the applied economist. Applied Economics, 42(17), 2193-2202. https://doi.org/10.1080/00036840701765494.

Salimbeni, H., Deisenroth, M. 2017. Doubly Stochastic Variational Inference for Deep Gaussian Processes. Advances in Neural Information Processing Systems, 30. https://proceedings.neurips.cc/paper/2017/hash/8208974663db80265e9bfe7b222dcb18-Abstract.html.

Smyth, R., Narayan, P.K. 2018. What do we know about oil prices and stock returns? International Review of Financial Analysis, 57, 148-156. https://doi.org/10.1016/j.irfa.2018.03.010.

Suphawan, K., Kardkasem, R., Chaisee, K. 2022. A Gaussian Process Regression Model for Forecasting Stock Exchange of Thailand. Trends in Sciences, 19(6), Article 6. https://doi.org/10.48048/tis.2022.3045.

Tsay, R.S. 2005. Analysis of financial time series. John wiley and sons.

Tu, S., Huang, J., Mu, H., Lu, J., Li, Y. 2024. Combining Autoregressive Integrated Moving Average Model and Gaussian Process Regression to Improve Stock Price Forecast. Mathematics, 12(8), Article 8. https://doi.org/10.3390/math12081187.

Vecchia, A.V. 1988. Estimation and model identification for continuous spatial processes. Journal of the Royal Statistical Society Series B: Statistical Methodology, 50(2), 297-312.

Wang, J., Feng, L., Li, Y., He, J., Feng, C. 2021. Deep Nonlinear Ensemble Framework for Stock Index Forecasting and Uncertainty Analysis. Cognitive Computation, 13(6), 1574-1592. https://doi.org/10.1007/s12559-021-09961-3.

Wang, J., He, J., Feng, C., Feng, L., Li, Y. 2021. Stock index prediction and uncertainty analysis using multi-scale nonlinear ensemble paradigm of optimal feature extraction, two-stage deep learning and Gaussian process regression. Applied Soft Computing, 113, 107898.

Williams, C.K., Rasmussen, C.E. 2006. Gaussian processes for machine learning (Vol. 2, Issue 3). MIT press Cambridge, MA.

Winarni, S., Indratno, S.W., Sari, K.N. 2022. Gaussian process regression using partition resampling model. Int. J. Math. Computer Sci., 17, 853-864.

Wolff, T. de, Cuevas, A., Tobar, F. 2020. Gaussian Process Imputation of Multiple Financial Series. ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech

and Signal Processing (ICASSP), 8444-8448. https://doi.org/10.1109/ICASSP40776.2020.9054102.

Zeng, A., Ho, H., Yu, Y. 2020. Prediction of building electricity usage using Gaussian Process Regression. Journal of Building Engineering, 28, 101054. https://doi.org/10.1016/j.jobe.2019.101054.

Zheng, B. 2024. An Introduction to Contrarian Functional Trading (SSRN Scholarly Paper No. 4700707). Social Science Research Network. https://doi.org/10.2139/ssrn.4700707.

Zhu, F., Quan, W., Zheng, Z., Wan, S. 2018. A Bayesian Learning Method for Financial Time-Series Analysis. IEEE Access, 6, 38959-38966. https://doi.org/10.1109/ACCESS.2018.2853998.