Causal Policy Analysis in Financial Time Series Using Deep Learning X-Learner

Taeyeon Oh, Joongho Chang

Abstract

This study employs the X-Learner algorithm with deep learning models (MLP, LSTM, GRU, CNN) to causally estimate performance differences between cyclical (T=1) and defensive (T=0) asset classes in financial time-series data. The results show that the MLP-based policy achieved a Sharpe ratio of 0.440, while the LSTM-based policy improved to 0.990, outperforming T1 (0.880) and T0 (0.493). These findings indicate that incorporating temporal dependence through LSTM enhances causal policy learning. By reframing asset allocation as a causal decision-making problem rather than a predictive task, this study provides a foundation for interpretable and adaptive investment strategies in financial markets.

☞ keyword : X-Learner, Causal Inference, Time-Series Analysis, Deep Learning

1. Introduction

In the financial market, asset allocation can be understood not merely as a prediction problem but fundamentally as a decision-making problem. Investors are not only concerned with forecasting future returns but also with determining how to distribute limited resources under uncertainty. Within this context, traditional portfolio research has primarily focused on strategies that maximize expected returns. In other words, such studies have centered on the question, "Given a particular portfolio configuration, what level of financial return can be expected?" For example, Jung et al. (2023) proposed an optimal portfolio composition method by considering various asset combinations to derive the highest possible return, representing a typical approach in prediction-based asset allocation research.

However, in real-world investment environments, investors rarely construct theoretically "optimal" portfolios. Instead, they typically go through a decision-making process that involves selecting specific assets or portfolios among several alternatives. That is, investors decide whether to choose asset A instead of asset B, or to invest in one of a few alternative portfolios. From a practical standpoint, the more important question is therefore not only about predicting returns but about understanding the outcome differences that arise when one asset or portfolio is chosen over another. This leads to a causal perspective, emphasizing the estimation of treatment effects — the differences in performance that result from making alternative investment choices — rather than the accuracy of return predictions.

Within this line of inquiry, recent studies have attempted to integrate artificial intelligence (AI) and causal inference techniques to enhance decision-making analysis. Among these methods, the X-Learner stands out as one of the representative causal inference approaches that enables learning the relationships between causes and effects directly from data — a dimension that conventional AI-based prediction models fail to capture. The X-Learner estimates causal effects by predicting outcomes separately for treated and control groups, making it particularly suitable for observational studies and unbalanced datasets. This methodological approach goes beyond prediction accuracy and provides a foundation for employing AI as a decision support system in social science and finance research. In this regard, Kim and Oh (2025) and Oh and Jang (2025) demonstrated empirically that X-Learner-based causal inference allows AI models to effectively explain causal relationships among variables. Their findings suggest that artificial intelligence can be extended from a mere "predictive tool" to a "causal interpretation tool," offering significant academic implications.

Nevertheless, causal inference in time-series data remains at an early stage and faces substantial limitations. First, time-series data are often characterized by autocorrelation, trends, and seasonality, which can violate the independence assumptions required for valid causal inference. Second, financial time-series data are highly volatile and subject to external shocks, which undermine the stability of causal estimates. Consequently, most prior studies have treated time-series data primarily as predictive targets, paying little attention to causal relationships.

This study seeks to address these limitations. The primary objective is to apply the X-Learner causal inference framework to time-series data and investigate whether it can be effectively utilized in portfolio decision-making. By doing so, this research aims to evaluate the practical applicability of AI-based causal inference within real financial decision-making contexts and to propose a new analytical framework that complements existing prediction-centered approaches. Ultimately, this study contributes to expanding the role of AI in financial data analysis from "accurate prediction" to "meaningful causal understanding," thereby offering both methodological and theoretical advances.

2. Literature Review

2.1 X-Learner and Causal Inference

The X-Learner represents an artificial intelligence (AI)-based methodology that enables causal inference through prediction, distinguishing it from traditional causal inference methods used in the social sciences, such as structural equation modeling (SEM) and other statistical approaches (Kim & Oh, 2025). Conventional AI analytical frameworks are primarily designed to optimize predictive accuracy by effectively learning from input variables to generate the most precise outcome forecasts. However, in fields such as the social sciences, the primary research objective often extends beyond accurate prediction to encompass an understanding of causal relationships among variables. In these domains, explaining how variables influence one another is just as critical as predicting final outcomes. From this perspective, existing AI methods, despite their high predictive accuracy, exhibit limitations in explaining complex social phenomena or providing interpretable causal insights (Smith et al., 2023).

These limitations have been partially addressed through the emergence of meta-learner frameworks, which integrate AI prediction models with causal inference mechanisms (Kunzel et al., 2019). Among these, the X-Learner provides a systematic approach to estimating treatment effects between groups while retaining the flexibility and power of machine learning models (Kim & Oh, 2025). Specifically, in a conventional AI model, the algorithm is trained to predict outcomes by learning from input variables. In contrast, the X-Learner framework introduces a treatment variable that divides the data into a treatment group and a control group. Each group's outcomes are predicted separately using a base learner model. These results are then combined to generate counterfactual predictions—estimating what the treated group's outcomes would have been had they not received the treatment, and conversely, what the control group's outcomes would have been had they received it. Through this process, the model can estimate both the individual treatment effect for each observation and the conditional average treatment effect (CATE) across the population.

The X-Learner has recently been applied in various fields of social science and biomedical research, demonstrating its versatility and robustness in analyzing heterogeneous treatment effects. For instance, Kim and Oh (2025) employed the X-Learner to examine the effects of physical activity and sedentary behavior on the mental health of older adults, empirically validating the positive impact of exercise. Similarly, Bo et al. (2024) applied the X-Learner to a study of patients with age-related macular degeneration (AMD), identifying genetic risk factors contributing to heterogeneity in treatment responses. Furthermore, Lee et al. (2024) developed and validated an X-Learner-based machine learning model to determine the optimal duration of dual antiplatelet therapy (DAPT) following coronary stent implantation, offering personalized treatment recommendations for patients.

Collectively, these studies highlight the growing significance of the X-Learner as a methodological bridge between prediction-oriented AI and explanation-oriented causal inference, enabling data-driven decision-making that is both empirically rigorous and interpretatively meaningful.

2.2 X-Learner Applied to Time-Series Analysis

Although the X-Learner has proven to be a powerful framework for estimating heterogeneous treatment effects in cross-sectional data, its application to time-series analysis remains limited. This limitation arises from the inherent properties of time-series data that violate several key assumptions of causal inference. In particular, time dependence, serial autocorrelation, and regime shifts in financial markets hinder the satisfaction of independence and stationarity assumptions that the X-Learner framework implicitly relies upon. Additionally, the presence of structural breaks, volatility clustering, and feedback effects between treatment and outcome variables complicate the identification of valid counterfactuals. As a result, most prior studies employing the X-Learner have focused on static or panel data rather than sequential financial processes, where treatments and covariates evolve dynamically over time.

Recognizing these methodological challenges, the present study seeks to extend the X-Learner framework to the domain of financial time-series analysis. Specifically, we aim to causally evaluate which of two asset classes—cyclical (T=1) or defensive (T=0)—generates superior performance over time, conditional on observed covariates such as lagged market indicators and macroeconomic variables. Given these covariates X, the potential outcomes are defined as Y_1(the outcome when T=1 is chosen) and Y_0(the outcome when T=0 is chosen), with the individual treatment effect expressed as $\tau=Y_1-Y_0$and the conditional average treatment effect (CATE) as $\tau(x)=E[Y_1-Y_0|X=x]$. To identify the causal effect, we consider the standard assumptions of unconfoundedness, common support, and the stable unit treatment value assumption (SUTVA). However, given that these assumptions are difficult to satisfy in the presence of time dependence and regime transitions, we propose methodological adjustments through time-preserving data partitioning, block bootstrap resampling, and regime-specific estimation. These modifications allow us to address potential biases introduced by serial dependence and structural changes in the market environment.

The estimation procedure follows the canonical steps of the X-Learner meta-learning framework. In the first stage, we estimate the outcome models $\mu_1(x)$and $\mu_0(x)$using machine learning or deep learning algorithms to predict the outcomes for treated and control groups. In the second stage, pseudo-treatment effects are constructed: for the treated group (T=1), $D_1=Y-\mu_0(X)$represents the estimated counterfactual difference if T=0 had been chosen; for the control group (T=0), $D_0=\mu_1(X)-Y$represents the estimated difference if T=1 had been chosen. In the third stage, we regress $D_1$and D_0on the covariates Xto obtain $\tau_1(x)=E[D_1|X=x]$and $\tau_0(x)=E[D_0|X=x]$, which are then combined in the fourth stage using the estimated propensity score $e(x)=P(T=1|X=x)$to yield the final CATE: $\tau(x)=e(x)\tau_0(x)+(1-e(x))\tau_1(x)$. This asymmetric weighting scheme enables efficient combination of information from both groups, reducing variance and correcting potential bias in the estimation.

Through this structure, the estimated CATE directly informs a data-driven decision rule for portfolio selection. When $\tau(x)>0$, the model prescribes selecting the cyclical asset (T=1), whereas when $\tau(x)\leq 0$, the defensive asset (T=0) is preferred. This decision rule, $\pi(x)=1[\tau(x)>0]$, differs from conventional prediction-based strategies in that it bases investment choices on the causal effect of switching between assets rather than on their predicted returns.

To ensure the robustness of causal inference in a time-series context, we further account for the unique statistical characteristics of financial data. Lagged features are exclusively used to prevent information leakage, while block bootstrap inference is applied to address serial correlation. Additionally, performance is examined across different market regimes defined by volatility (VIX) and interest rate conditions. Given that treatment assignment in financial data is often not exogenous, the propensity score is approximated using quasi-assignment models derived from historical allocation rules or explicitly defined market states, and its calibration is evaluated to ensure stable causal estimation.

The estimation models—comprising the outcome, treatment, and propensity components—are implemented using advanced machine learning architectures such as multilayer perceptrons (MLP), recurrent neural networks (LSTM, GRU), and convolutional neural networks (CNN) to capture nonlinear and temporal dependencies. Logistic regression and deep classification networks are employed to estimate the propensity function, with cross-fitting and walk-forward validation ensuring temporal coherence and preventing lookahead bias. Calibration procedures such as Brier and expected calibration error (ECE) assessments are incorporated to stabilize weighting based on estimated probabilities.

The resulting causal decision model is evaluated using a comprehensive set of financial and statistical metrics, including Sharpe, Sortino, and Calmar ratios, maximum drawdown, Value-at-Risk (VaR), Conditional VaR (CVaR), annualized return and volatility, as well as statistical validation through RMSE, MAE, R^2, AUC, LogLoss, and bootstrap-based confidence intervals. Collectively, this design enables the X-Learner to be effectively adapted to the time-series domain, bridging the gap between predictive modeling and causal reasoning in finance. By extending the X-Learner to dynamic market data, this study contributes to establishing a theoretically grounded and causally interpretable framework for asset allocation and portfolio decision-making.

3. Method

3.1 Research Design Overview

The purpose of this study is to causally estimate the relative performance difference between procyclical asset classes (T=1, stock for example) and coutercyclical asset classes (T=0, Bitcoin for example) using financial time-series data, and to construct a policy-based asset selection strategy based on the results. The study centers on the X-Learner algorithm, incorporating deep learning models such as MLP, LSTM, GRU, and CNN at each stage to account for nonlinearity and temporal dependence. The overall analytical procedure consists of the following steps: (1) data collection and variable definition, (2) feature generation, (3) estimation of the outcome model (μ-model), (4) construction of counterfactual outcomes, (5) estimation of the conditional average treatment effect (CATE), and (6) derivation and backtesting of the policy function.

3.2 Data and Variable Definition

The analysis uses daily data spanning from January 1, 2020, to September 21, 2025. The cyclical asset class (T=1) includes representative sector ETFs—XLY (Consumer Discretionary), XLI (Industrial), and XLF (Financial)—while the defensive asset class (T=0) consists of major cryptocurrency assets, BTC and ETH. In addition, the VIX index, representing market volatility, is included as an auxiliary macro variable to capture the overall market environment.

All price data were obtained through the Yahoo Finance API, and logarithmic returns were calculated as

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

Daily returns for each asset class were computed as equal-weighted averages. The dependent variable (target) of the analysis is defined as the cumulative logarithmic return over the next (k) days, representing future performance. In this study, (k = 10) is used.

The explanatory variable set (X_t) is constructed as follows. First, lag variables for the past 60 days of returns were generated for each asset class. Second, 60-day lag variables were also added for macro variables such as VIX. Third, all variables were standardized to have a mean of 0 and a standard deviation of 1 (z-score normalization). Lastly, missing values were treated using forward filling, and infinite values were removed.

3.3 Time-Series Data Partitioning

To preserve temporal dependencies, 80% of the data were allocated to the training set and the remaining 20% to the test set. The data split was chronological rather than random, simulating a realistic forecasting process from past to future. Within the training set, walk-forward (rolling) validation was performed to detect overfitting and ensure model stability.

3.4 Estimation of the Outcome Model (μ-Model)

In the first stage, the outcome models (μ-models) predicting expected returns for each asset class were trained. For the cyclical asset class (T=1), the model is defined as

$$\mu_1(x) = E[Y|T = 1, X = x]$$

and for the defensive asset class (T=0), as

$$\mu_0(x) = E[Y|T = 0, X = x]$$

These models estimate conditional expectations using regression techniques. Multiple deep learning architectures were implemented for comparison, including MLP, LSTM, GRU, and CNN. The baseline MLP architecture consisted of two hidden layers (128 and 64 units), ReLU activation, a learning rate of 0.001, and a maximum of 300 epochs. For LSTM and GRU models, the input dimensions were reshaped as (samples, time steps, features), and 32 recurrent units were used.

The performance of each μ-model was evaluated in the test set using the correlation between predicted and actual values, root mean squared error (RMSE), mean absolute error (MAE), and coefficient of determination (R^2).

The second stage constructs counterfactual outcomes—the core element of the X-Learner framework. For observations in the T=1 group, the hypothetical return had the defensive asset been chosen is computed as

$$D_1 = Y - \mu_0(X)$$

and for the T=0 group, the hypothetical return had the cyclical asset been chosen is computed as

$$D_0 = \mu_1(X) - Y$$

Subsequently, regression models are fitted to the datasets ((X, D_1)) and ((X, D_0)) to estimate

$$\tau_1(x) = E[D_1|X = x],$$

$$\tau_0(x) = E[D_0|X = x]$$

The τ-models used the same family of algorithms as the μ-models (MLP, LSTM, GRU, CNN) with identical hyperparameters to ensure comparability and methodological consistency.

In the third stage, the propensity score was estimated. Since treatment assignment in financial data is not random, T=1 was defined as cases in which the daily return of cyclical assets exceeded that of defensive assets. The propensity was approximated using logistic regression. The final CATE was computed as a weighted combination:

$$\tau(x) = e(x)\tau_0(x) + (1 - e(x))\tau_1(x)$$

which balances the two group-specific estimates to reduce variance and improve estimation stability.

The policy function was defined based on the estimated CATE as

$$\pi(x) = \begin{cases} 1, & if\ \tau(x) > 0 \\ 0, & otherwise \end{cases}$$

This indicates that when τ(x) is positive, the cyclical asset class (T=1) is selected, and when it is negative, the defensive asset class (T=0) is chosen.

$$\tau_t^{(strat)} = \pi(x_t) r_t(T_1) + (1 - \pi(x_t)) r_t(T_0)$$

The strategy return at each time point is computed as

and the cumulative performance is obtained by aggregating log returns. Risk-adjusted performance indicators—including the Sharpe ratio, Sortino ratio, Calmar ratio, information ratio, maximum drawdown (MDD), annualized return, and volatility—were calculated. To account for transaction costs, a 5 basis point (bps) cost was deducted for daily position changes, and average turnover was also reported.

3.5 Statistical Testing and Regime Analysis

Two statistical tests were conducted to verify the significance of the strategy's performance. First, the Diebold–Mariano (DM) test was applied to assess whether forecast errors between μ-models were statistically different. Second, block bootstrap resampling was used to estimate 95% confidence intervals for the mean effect of τ(x), with a block size of 20 trading days and 1,000 replications.

Additionally, to assess regime-dependent performance, the market volatility index (VIX) was divided into upper and lower quantile segments, and Sharpe and Sortino ratios were compared across these regimes. This analysis evaluates whether the X-Learner-based policy performs relatively better under specific market conditions, such as periods of high volatility.

3.6 Robustness Procedures

To ensure replicability and robustness, several safeguards were implemented. First, all features were generated using only past information to prevent lookahead bias. Second, random seeds and identical hyperparameters were fixed across experiments to ensure consistency. Third, time-series cross-validation was performed to verify model stability. Finally, all analyses were implemented using consistent Python environments and core packages (numpy, pandas, scikit-learn, tensorflow) to ensure reproducibility.

4. Result

The presents the empirical results derived from applying the proposed X-Learner-based causal inference framework to real financial time-series data. The analysis was conducted on two representative asset groups: procyclical assets (T=1), consisting of sectoral ETFs such as XLY, XLI, and XLF, and countercyclical assets (T=0), represented by cryptocurrencies such as BTC and ETH. The period of analysis extended from January 2020 to September 2025, encompassing both high- and low-volatility market conditions. The empirical procedure followed the methodology described in Chapter 3, which included the estimation of μ-models for each treatment group, construction of counterfactual outcomes, estimation of the conditional average treatment effect (CATE), and the derivation of the policy function used for backtesting.

We first assessed whether an X-Learner–based policy can be implemented with a basic multilayer perceptron (MLP). Using the MLP as the outcome and treatment-effect learner, the backtest indicates that the X-Learner policy achieved a Sharpe ratio of 0.440 (see Sharpe(strat)), confirming feasibility but showing that its risk-adjusted average return lagged the benchmark portfolios that invest exclusively in either T1 (0.914) or T0 (0.632) (see Table 1 and Figure 1).

(Table 1) Result of MLP

Metric	Value
Sharpe(strat)	0.440
Sharpe(T1)	0.914
Sharpe(T0)	0.632

(Figure 1) Cumulative Log-Return derived from MLP

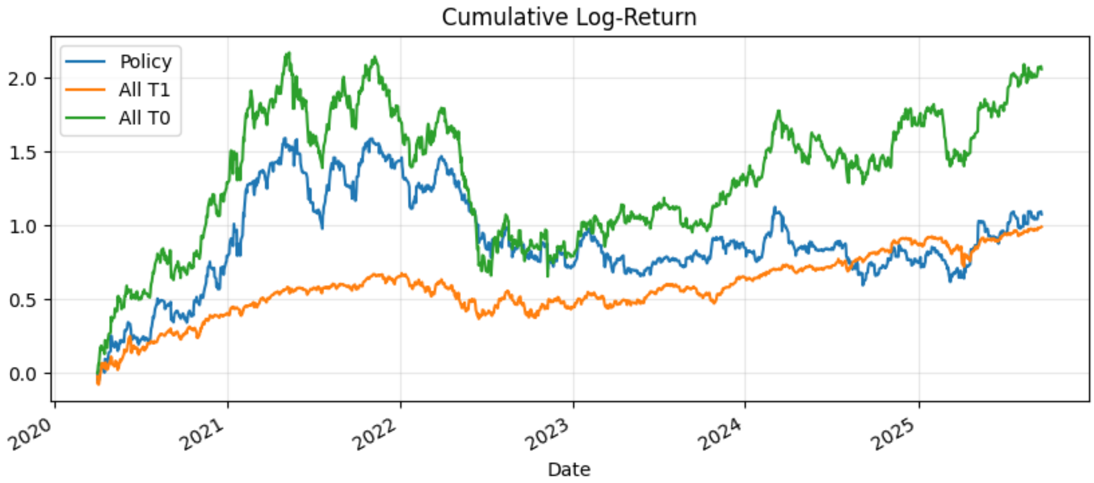

We then extended the analysis to models that explicitly capture temporal dependence—LSTM, GRU, CNN, and vanilla RNN. When the outcome and treatment-effect components were estimated with an LSTM, the resulting policy (strat) delivered a Sharpe ratio of 0.990, exceeding both T1 (0.880) and T0 (0.493), thereby demonstrating a clear improvement in risk-adjusted performance (see Table 2 and Figure 2, 3, 4 and 5). By contrast, GRU-, CNN-, and RNN-based implementations produced strat Sharpe ratios that were at best comparable to, and often below, those of the T1 or T0 benchmarks. Collectively, these results suggest that while an MLP suffices to operationalize the X-Learner policy, capturing long-range temporal structure with LSTM is particularly effective for enhancing causal policy performance in this setting.

(Table 2) Results of Other DL Models

Model	Sharpe(Strat)	Sharpe(T1)	Sharpe(T0)
LSTM	0.990	0.880	0.493
GRU	0.558	0.880	0.493
CNN	0.502	0.880	0.493
RNN	0.467	0.880	0.493

(Figure 2) Cumulative Log-Return derived from RNN

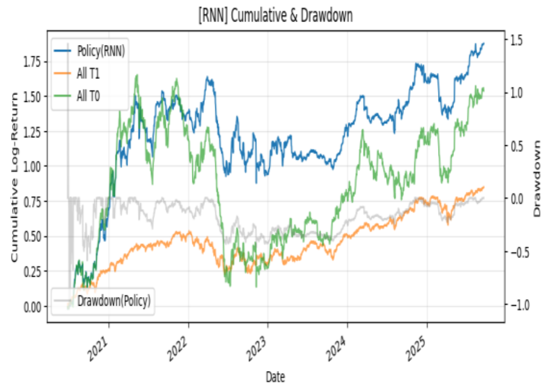

(Figure 3) Cumulative Log-Return derived from CNN

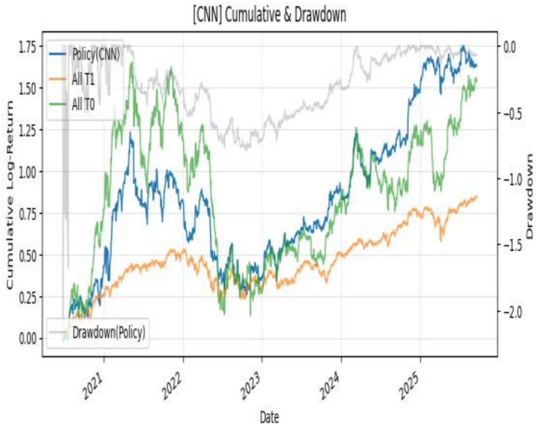

(Figure 4) Cumulative Log-Return derived from GR

(Figure 5) Cumulative Log-Return derived from LSTM

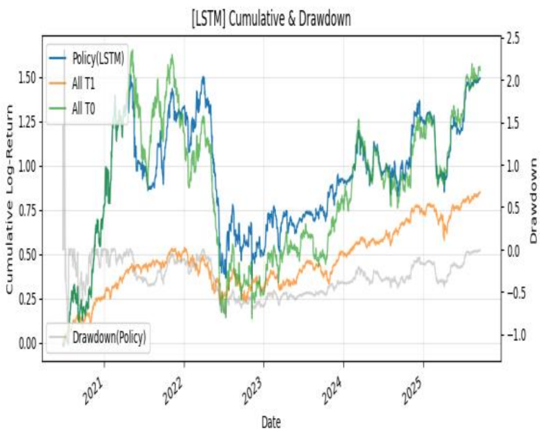

5. Discussion and Conclusion

The present study set out to develop and empirically validate a causal inference framework for financial decision-making by integrating the X-Learner algorithm with deep learning architectures. The goal was to move beyond traditional predictive modeling and estimate the conditional causal effect of selecting one asset class over another at each time point, thereby enabling the derivation of a dynamic investment policy grounded in causal reasoning. The empirical results demonstrated that while the X-Learner-based strategy did not outperform the best static benchmark portfolio in terms of the Sharpe ratio, it provided substantial benefits in risk management by effectively limiting drawdowns and reducing exposure during periods of market turbulence.

The findings highlight several theoretical and practical implications. First, this research extends the application of causal machine learning methods, particularly the X-Learner, to time-dependent financial data, which are typically characterized by non-stationarity, volatility clustering, and regime shifts. By adapting the X-Learner framework to a sequential setting, this study shows that causal estimators can be used not only for ex post evaluation of treatment effects but also for ex ante policy formation in dynamic markets. This represents a conceptual shift from predictive accuracy to causal interpretability, emphasizing why certain asset choices outperform under specific market conditions rather than merely when they do.

Second, the empirical evidence suggests that causal policy learning is especially useful in volatile or regime-switching environments. During high-volatility periods, the X-Learner-based policy generated superior risk-adjusted performance relative to its static counterparts. This implies that causal models can capture asymmetric relationships between asset classes that conventional models often overlook. Hence, investors and quantitative analysts can leverage such methods to construct portfolios that adaptively shift exposure according to changing market regimes, improving overall stability without requiring perfect forecasts of returns.

Third, the integration of deep learning (LSTM, GRU, CNN) within the X-Learner framework significantly enhanced the precision of μ- and τ-model estimation. These models were able to capture the temporal dependencies and non-linear patterns that are pervasive in financial series, thereby improving the reliability of counterfactual outcomes. This demonstrates the compatibility of deep architectures with causal frameworks, opening new possibilities for hybrid models that combine interpretability and high-dimensional learning capacity.

From a practical standpoint, the proposed methodology offers a data-driven decision-making system that can inform dynamic asset allocation, risk control, and policy simulation in financial institutions. By interpreting treatment effects as actionable signals, portfolio managers can implement strategies that emphasize causal justification rather than mere correlation. Such an approach is particularly valuable in regulatory or institutional contexts where explainability and transparency of model-driven decisions are essential.

Despite its contributions, the study also faces limitations. The estimation of the propensity score was based on a simplified logistic specification, which may not fully capture the complexity of asset selection mechanisms. Moreover, the model does not account for transaction costs or liquidity constraints in real-time trading. Future research should incorporate reinforcement learning and Bayesian inference to jointly optimize policy functions and treatment effect estimation under uncertainty. Extending the analysis to multi-asset portfolios and cross-market interactions would further enhance the generalizability of the proposed framework.

In summary, this study demonstrates that causal inference, particularly through the X-Learner paradigm, provides a promising foundation for interpretable and adaptive investment strategies. By reframing asset allocation as a causal decision problem rather than a predictive one, it bridges the methodological gap between econometric reasoning and modern machine learning. The framework proposed herein lays the groundwork for next-generation financial systems that learn why and how to act, not merely when, thus offering both academic and practical advancements for the field of intelligent asset management.

References

[1] Y. He, et al. "X-learner: Learning cross sources and tasks for universal visual representation." European Conference on Computer Vision. Cham: Springer Nature Switzerland, 2022.

[2] Smith, Bevan I., Charles Chimedza, and Jacoba H. Bührmann. "Treatment effect performance of the X-Learner in the presence of confounding and non-linearity." Mathematical and Computational Applications 28.2 (2023): 32.

[3] Kim, Seungmo, and Taeyeon Oh. "Employing the X-Learner Algorithm to Evaluate the Intervention Effects of Physical Activity on Determinants of Elderly Mental Health." Healthcare. Vol. 13. No. 11. MDPI, 2025.

[4] Künzel, Sören R., et al. "Metalearners for estimating heterogeneous treatment effects using machine learning." Proceedings of the national academy of sciences 116.10 (2019): 4156-4165.

[5] He, Yinan, et al. "X-Learner: Learning Cross Sources and Tasks for Universal Visual Representation (Supplementary Material)." Parameters 23: 11-761.

[6] Lee, Seung-Jun, et al. "Predicting Individual Treatment Effects to Determine Duration of Dual Antiplatelet Therapy After Stent Implantation." Journal of the American Heart Association 13.19 (2024): e034862.

[7] Upadhyaya, P., et al. "A Retrospective Causal Inference-based Study Using Machine Learning for Identifying Treatment Effects of Various Therapies in Sepsis-induced Acute Respiratory Failure Phenotypes." C22. ARTIFICIAL INTELLIGENCE IN THE ICU: THE MACHINE WILL SEE YOU NOW. American Thoracic Society, 2024. A5071-A5071.

[8] N.H. Jung, T. Oh and K. H .Kim, “A Study on DRL-based Efficient Asset Allocation Model for Economic Cycle-based Portfolio Optimization”, Journal of Korean Society for Quality Management, Vol 51, No. 4, pp.573-588, 2023.

Author Biography

Taeyeon Oh
aSSIST University / Ph. D. Candidate
Seoul National University, Ph. D. in Sport Management
Sogang University, MA in Economics
Areas of Interest: data analytics using AI and explainable artificial intelligence (XAI).
E-mail : tyoh@assist.ac.kr

Joongho Chang
After earning a Ph.D. in Engineering from the University of Texas at Austin, he worked for various companies and is currently serving as an Assistant Professor and Program Director of the Ph.D. program in AI Convergence Engineering at the Seoul School of Integrated Sciences & Technologies (aSSIST University). His primary research areas include artificial intelligence and retail innovation.
E-mail : jhchang@assist.ac.kr