

the journal of
financial data science

Frank J. Fabozzi
 Editor

**Marcos López
 de Prado**
 Editor

Joseph Simonian
 Editor

Francesco A. Fabozzi
 Managing Editor

Mitchell Gang
 Production Editor

Deborah Brouwer
 Production and Design
 Manager

Mark Adelson
 Content Director

Rosie Instance
 Marketing Manager

William Law
 Account Manager—
 Asia/Middle East

Ryan C. Meyers
 Subscription Sales
 Director—Global

David Rowe
 Reprints Manager and
 Advertising Director

Cathy Scott
 General Manager
 and Publisher

Several articles, two of which were published in this journal, have shown how reinforcement learning can be used to take trading costs into account in hedging decisions. In the lead article of this issue, “Deep Hedging of Derivatives Using Reinforcement Learning,” Jay Cao, Jacky Chen, John Hull, and Zissis Poulos extend the standard reinforcement learning approach by utilizing multiple Q-functions for the purpose of increasing the range of objective functions that can be used and by using algorithms that allow the state space and action space to be continuous. The authors suggest an approach where a relatively simple valuation model is used in conjunction with more complex models for the evolution of the asset price. This allows good hedges to be developed for asset price processes that are not associated with analytic pricing models.

Deep sequence models have been applied to predicting asset returns. These models are flexible enough to capture the high-dimensionality, nonlinear, interactive, low signal-to-noise, and dynamic nature of financial data. More specifically, these models can outperform the conventionally used models because of their ability to detect path-dependence patterns. Lin William Cong, Ke Tang, Jingyuan Wang, and Yang Zhang in their article “Deep Sequence Modeling: Development and Applications in Asset Pricing,” show how to predict asset returns and measure risk premiums by applying deep sequence modeling. They begin by providing an overview of the development of deep sequence models, introducing their applications in asset pricing, and discussing the advantages and limitations of deep sequence models. A comparative analysis of these methods using data on US equities is then provided in the second part of the article where the authors demonstrate how sequence modeling benefits investors in general by incorporating complex historical path dependence. They report that long short-term memory has the best performance in terms of out-of-sample predictive R-squared, and long short-term memory with an attention mechanism has the best portfolio performance when excluding microcap stocks.

In the formulation of an investment process, it is critical to build a view of causal relations among economic entities. Because of the complex and opaque nature of many market interactions, this can be challenging. Various models of economic causality have been proposed to both explain the past and aide investors in the investment process such as causal networks. Such networks provide an efficient framework for assisting with investment decisions that are supported by both quantitative and qualitative evidence. When building causal networks, the addition of more causes adds to the issue of computational complexity because of the necessity to calculate the combined impact of larger and larger sets of causes. In “Causal Uncertainty in Capital Markets: A Robust Noisy-Or Framework for Portfolio Management,” Joseph Simonian argues that among the various approaches to causal networks, the “noisy-or model” offers the means to calculate the aggregate effect of causes in a linear manner assuming that the causal probability values used by model builders are completely reliable. To address the question of uncertainty, Simonian provides a robust, uncertainty-adjusted noisy-or framework that draws on evidence-based subjective logic (i.e., a many-valued logic explicitly designed to assess the reliability

of the evidence supporting an investor's beliefs). The framework provides an investor with the ability to apply the noisy-or model such that it is consistent with the strength of both an investor's evidence and overall research methodology.

The sizing of investment positions often depends on a trader's conviction in a trade idea. However, there is relatively little discussion in the investment literature about how to size an investment position in a data-driven way. Inspired by recent innovations in financial machine learning, Trent Spears, Stefan Zohren, and Stephen Roberts in "Investment Sizing with Deep Learning Prediction Uncertainties for High-Frequency Eurodollar Futures Trading," argue that uncertainty estimates obtained from deep learning models can be useful inputs for influencing the relative allocation of risk capital across trades, and demonstrate how such models can improve relative investment performance. In particular, the authors apply deep learning models to predict changes in the Eurodollar futures curve, as a function of a recent history of asset price data, within a high-frequency domain.

Another application of deep learning techniques is suggested by Jonathan Iworiso and Spyridon Vrontos in "On the Predictability of the Equity Premium Using Deep Learning Techniques." Here several deep learning techniques—H2O deep neural networks, a stacked autoencoder, long short-term memory, and the fusion of some of these techniques—are studied for estimating the equity premium (response variable) using financial variables (such as bond yields, exchange rates, microeconomic and macroeconomic variables) as the predictors. Their empirical analysis suggests that the stacked autoencoder-with-H2O deep learning produced the best predictive performance and economically significant results. Furthermore, their findings were robust across all out-of-sample periods.

An algorithm that has been used in time series analysis for measuring the optimal alignment between two temporal sequences is dynamic time warping. This time series tool has been successfully applied to data mining and speech recognition. In "Dynamic Time Warping: S&P 500 Sector ETF Pattern Matching Trading Strategy," Alexander Fleiss, Che Liu, Gihyen Eom, Serena Yu, and Wo Zhang explore how dynamic time warping can be applied to a quantitative equity trading strategy. Specifically, the authors apply dynamic time warping to the S&P 500 sector exchange-traded funds (ETFs) by first implementing a pattern-matching trading system to identify trends and then estimating a decision-making dictionary from the windows of ETF prices that can be used to determine the entry points for trading. Using the validation set to construct a portfolio that minimizes the expected volatility by estimating the optimal weights for each component, they identify profitable quantitative strategies. They demonstrate the flexibility of the pattern-matching trading strategy by varying the strategies to adapt to both the pre-COVID-19 period and the post-COVID-19 period.

Tristan Lim and Chin Sin Ong provide another application of dynamic time series warping to portfolio management in their article "Portfolio Diversification Using Shape-Based Clustering." It is well-known that portfolio diversification depends on the correlation between portfolio assets—the lower the correlation the greater the improvement of a portfolio's risk-return profile. However, it is reasonable to expect that relying on descriptive statistics, and specifically, correlation, to achieve an improved risk-return profile via diversification may not yield the most optimal multiperiod portfolio because stocks included in a portfolio can exhibit different price trends over time,

even when using the same computed pairwise correlation. Using dynamic time-series warping, Lim and Ong use the distance measure to aggregate stocks into like-trending clusters across time as a portfolio diversification tool. Their findings suggest that the shape-based clustering technique can be used for three investment decisions: (1) portfolio allocation and rebalancing, (2) dynamic predictive portfolio construction, and (3) individual stock selection through outlier identification.

When constructing a predictive model using machine learning, feature selection involves eliminating redundant (non-informative) input variables in order to reduce computational costs and/or improve the performance of the model. The feature importance score measures how much information a feature contributes when building a supervised learning model. For each feature in the dataset, an importance score or ranking is computed, thereby allowing the features to be ranked. Several algorithms have been used for feature selection, the more popular ones being MDA, LIME, and SHAP. However, most feature selection algorithms expose the user to intrinsic randomness. That is, there is no guarantee that the importance score or rank of each feature will remain the same and that the same features will be selected for each sample. For MDA, this randomness is attributable to the random permutations of the values of one feature at a time. For LIME, it is due to the random perturbations of all features at the same time for each sample and for SHAP the random permutations of the feature sequence in both forward and reverse directions. In “The Best Way to Select Features? Comparing MDA, LIME, and SHAP,” Xin Man and Ernest P. Chan propose a novel ranked-based instability index to measure the stability of the three feature selection algorithms. They refer to this measure as the “instability index” and apply it to evaluate the stability of the three feature importance scoring algorithms. The instability of these three algorithms is compared using two public datasets, a public S&P 500 dataset, and a proprietary financial trading dataset. They find that LIME is more stable than MDA and SHAP on the features with high importance scores and improves the trading strategy that they examine in terms of a higher Sharpe ratio and higher cumulative return.

In economics, covered interest rate parity is the principle that the yield on domestic and foreign debt instruments should be the same after adjusting for the currency exchange rate between the two countries when calculating the payoff in the domestic currency. Yet, empirically it has been reported that covered interest rate parity violation is persistent and varies systematically, providing traders with an arbitrage opportunity. In “Deviations from Covered Interest Rate Parity: The Case of British Pound Sterling versus Euro,” Frank Lehrbass and Thamara Sandra Schuster argue that due to the presence of nonlinearities, an appropriate way to identify and capitalize on the arbitrage opportunity is to use deep learning methods. Taking into account event-driven factors such as Brexit-related politics, the authors investigate the deviations from covered interest rate parity in the case of the British pound and the euro. They find that the foreign exchange derivatives market for these two currencies do in fact deviate from the covered interest rate parity daily which has importance for arbitrage desks. Specifically, they demonstrate that when arbitrage opportunities arise, when to look for better alternatives than hedging with forwards.

Francesco A. Fabozzi
Managing Editor