Can AI Replace Stock Analysts? Evidence from Deep Learning Financial Statements

G. Nathan Dong *

`

March 31, 2024

ABSTRACT

Equity research analysts have long been regarded as experts in predicting the performance of stocks. With the rapid advancements in artificial intelligence (AI), the question arises as to whether neural-network based deep-learning technology can surpass human analysts in this task. In this research, we apply a model-free deep-learning technique to equity research by training a simple neural-network (without pre-specified financial models or pre-defined accounting ratios) to predict 12-month-ahead stock price. Our AI model that runs on an inexpensive computer equipment is capable of processing vast amounts of financial statement raw data over a 14-year period in a fraction of the time and cost it would take a human analyst ($1,450 in equipment, 15 hours in time, and $2.5 in electricity, to be exact).† It outperforms human analysts in 12-month-ahead target price forecasts by a large margin. Analysts and AI models are more likely to make similar predictions of the target price for firms that show signs of potential investment and future uncertainty, such as higher levels of cash, capital expenditures, and return volatility on the day of forecast. Overall, the results of this research support the notion that AI has the potential to replace human analysts in certain aspects of predicting financial performance.

Keywords: Analyst forecast, Target price, Artificial intelligence, Deep learning, Fintech
JEL Codes: C10, C45, G30, G23

* Department of Finance, Carroll School of Management, Boston College, Chestnut Hill, MA 02467. Tel: (617) 552-6426. E-mail: nathan.dong@bc.edu.. All errors remain our responsibility.
† Detailed description of software and hardware setups used in this research can be found on the author's website: http://sites.google.com/view/gang-nathan-dong/ai-research-in-finance.



— The New York Times ‡

INTRODUCTION

Equity research analysts play an important role in the financial market by interpreting public information, discovering new information, and providing recommendations.§ Part of their job is to utilize their knowledge and analysis of the market to forecast price targets for stocks, which are then used by investors to make informed decisions (Brav and Lehavy, 2003). It involves analyzing financial statement data, studying market trends, and offering educated speculations about the future performance of a particular stock (Beyer et al., 2010).** However, with the advancements in artificial intelligence (AI) technology, it naturally begs the question of whether AI has the capability to replace this otherwise repetitive and labor-intensive task of forecasting target prices.

The use of AI in predicting performance has attracted considerable attention as an alternative method to traditional equity research. Utilizing advanced algorithms and machine learning techniques, AI systems can analyze vast amounts of data to identify patterns and trends that could be unnoticed by even the most experienced professionals. Specifically, the deep-learning technique based on neural-network has been applied to various fields and industries such as entertainment (Igami, 2020), monetary policy (Gorodnichenko et al., 2023), macroeconomics (Maliar et al., 2021), microeconomics (Khachiyan et al., 2022) and corporate governance (Erel et al., 2021). Previous studies have shown that AI is certainly capable of analyzing complex datasets and finding non-linear relationships between different financial variables in a very small fraction of time compared to that of a human analyst. This could potentially lead to more accurate predictions with reduced margin of error. In addition, AI is

‡ This is the title of a NYT article contributed by Nathaniel Popper and published in the New York Times' Sunday Magazine with the headline "Stocks & Bots" on February 28, 2016.

§ The relative importance of analysts' role in information interpretation versus information discovery has been studied in the accounting literature (Livnat and Zhang, 2012).

** The focus of this study is on sell-side analysts, rather than buy-side analysts, but they are closely related because sell-side analysts often provide buy-side analysts with in-depth industry knowledge and access to company management (Brown et al., 2016).

not influenced by emotions, which sometimes can cloud the judgment of human analysts, thereby eliminating possible biases (Green et al., 2024).

The debate over whether artificial intelligence can outperform human stock analysts in predicting performance remains ongoing. While AI has demonstrated remarkable capabilities in analyzing huge amounts of data with a high level of speed and accuracy and making predictions based on patterns and probabilities, human analysts bring in intuition, experience, and an understanding of market dynamics that cannot be replicated by a machine (Joos et al., 2016). Undoubtedly, AI lacks the intuition and critical thinking skills that human analysts possess because equity research is not just about crunching numbers; it also involves understanding the underlying factors that may be affecting the performance of a stock. This kind of qualitative analysis can never be perfectly replicated by machines (Iskhakov et al., 2020). Indeed, a human-defined variable, like a competition measure such as HHI, or an accounting ratio such as ROA, that has been used by the analysts, are often included in the AI deep-learning (Cao et al., Forthcoming). The use of pre-specified models or pre-defined ratios in AI deep-learning can certainly make the predictions more reliable. Unfortunately, no analysis has attempted to compare the performance of analyst forecasts with that of model-free deep-learning, which involves training a raw neutral-network that has no prior knowledge of financial models and accounting ratios.

The purpose of this paper is to fill this gap by offering an empirical comparison of the accuracy in forecasting 12-month-ahead stock prices by both human analysts and AI model-free predictions.†† We developed a simple neural-network model and utilized it to deep-learn raw data of financial statements, historical stock prices and interest rates. As Birru, Gokkaya, Liu and Stulz (2022) points out, analyst target prices are typically derived from an analyst's long-term evaluations of a firm's prospects over an extended period. These evaluations are based on a combination of fundamental information and financial models (Birru et al., 2022). Why can't an AI model, which has access to the same financial data, develop its own model to accurately forecast stock prices 12 months in advance? The first challenge is determining the type of data that should be used to train the AI model to ensure reliable price prediction and fair comparison with human analysts. Training a neural-network using a comprehensive dataset over an extended period of time makes the comparison between humans and AI unfair. In this research,

†† The price forecasted by the analysts, also known as the target prices are predominantly one-year-ahead prices (Brav and Lehavy, 2003).

we utilize three sets of historical data: quarterly financial statements, stock prices, and interest rates.

A second challenge is deciding which model to build. If we simply use a pre-built model like the discounted cash flow (DCF) and let the computer to predict stock price, the comparison is no longer between human and AI, but rather a replacement of humans by machine. We prefer a model-free algorithm, which is the estimation of an optimal model that is free-form, meaning it is not constrained by predetermined structures or assumptions. In this paper, we rely on the deep-learning capabilities of neural-network models to explore and interpret data in ways that cannot be achieved through conventional methods. The unpredictable nature of the stock market and external factors make it challenging for both humans and AI to consistently predict future performance accurately. Free-form modeling using neural-networks offers a more holistic understanding of historical data. This approach allows for more nuanced and dynamic representations of a firm's past performance, potentially leading to deeper insights and more accurate predictions.

This paper makes two contributions to the literature. First, it provides empirical evidence of the deep-learning capability of neural-network models to accurately forecast target prices. It identifies the scope of the data needed to train the model, specifically historical financial statements, stock prices, and interest rate data. After a 5,000-epoch training that takes 15 hours and costs less than three dollars in electricity, a basic AI model running on $1,500 computer equipment outperforms human analysts in predicting stock prices 12 months ahead. Secondly, this research examines the underlying economic factors that drive the deviation in predicted target prices between analysts and the AI model. It shows that the forecast difference is smaller in firms that have a lower market-to-book ratio and a higher level of cash, capital expenditures, stock price, and volatility on the day of the forecast. Overall, the results of this research support the notion that AI has the potential to replace human analysts in certain aspects of predicting financial performance.

The remainder of the paper is organized as follows. Section II introduces the neutral-network model and the technical details of model training and price prediction. Section III presents the sample data and measurement choices. Section IV analyzes the model performance by comparing the prices between analyst forecasts and AI predictions. Section V conducts robustness tests by constructing various "deeper-learning" and "shallow-learning" models.

Section VI provides a summary and concluding marks.

II. METHOD OF MODEL-FREE DEEP LEARNING

Deep learning, sometimes referred to as artificial neural networks, is a machine learning technique that can train computer models to do what comes naturally to humans. It is able to handle a large number of input variables and is not affected by multicolinearity. This is because neural networks use multiple hidden layers to process the input variables and capture the relationships between them. Each hidden layer in a neural network learns and extracts different features from the input data, reducing the impact of multicolinearity (De Veaux and Ungar, 1994). In contrast, linear regression models like OLS (Ordinary Least Squares) rely on the assumption of independent predictor variables, and when multicolinearity is present, it becomes difficult to determine which variable is actually contributing to the outcome. This can result in unreliable and inflated parameter estimates, rendering the model less effective in making accurate predictions (Dormann et al., 2013). Additionally, neural networks are able to handle missing data and outliers effectively, which can further contribute to the accuracy of their predictions.

Therefore, we believe that using a neural network to deep-learn a panel dataset of financial statements and other market data for predicting risk and outcomes can be a promising approach that can offer valuable insights for investors and financial analysts. For example, a deep-learning model has been developed to study the non-linear relationship between mortgage risk factors and loan performance. The nonlinearities revealed by the deep-learning model have profound implications for risk management, investment management, and mortgage-backed securities (Sadhwani et al., 2021).[⌗⌗]

2.1. Neural network model and in-sample data training

In this study, we trained a multi-layer neural-network model using a panel data of financial statements, stock prices and interest rates. Training this model is a gradual and repetitive process aimed at enhancing the model's capacity to analyze historical data and generate predictions. At its core, a neural network is composed of layers of interconnected nodes that collaborate to process input data and produce output (as shown in Figure 1). These

⌗⌗ A similar work can be found in using machine learning to predict loan default using European data (Barbaglia et al., 2023).

layers are essential for the functioning of a neural network, as each layer performs a specific task that contributes to the overall prediction made by the network. The first layer, known as the input layer, is responsible for receiving the financial statement and other market data and passing it on to the next layer. This layer plays a crucial role in the prediction process by establishing the initial values for the network's weights and biases. The next layer is the hidden layer, where most of the complex computations take place. This layer is responsible for extracting features from the input data, which are then passed on to the next layer. The number of hidden layers in a neural network can vary depending on the complexity of the problem being solved. These layers are where the network's learning takes place, as the weights and biases are continuously adjusted based on the input data. Finally, the output layer is where the prediction is made. This layer takes the information from the hidden layers and produces the final output, which is the predicted 12-month stock price (to be compared with the actual price). It is important to note that we use five layers in the neural-network. The choice of this number is completely arbitrary because we want to keep the model as simple as possible while maintaining the level of sophistication needed to make the predictions as accurate as possible. In the robustness section, we will evaluate the performance of a model with only three layers.

[Insert Figure 1]

During the training process, the neural-network layers undergo a series of transformations known as forward propagation, which ultimately enhance the model's capability to accurately predict the stock price. At the beginning of the training process, the weights and biases of the neural network layers are randomly initialized. These weights and biases determine the strength of the connections between neurons and the threshold at which they fire. As the training progresses, the model is presented with a set of in-sample training data, and the input data is fed forward through the layers of the neural network. As the data passes through each layer, the outputs of neurons are calculated using the current weights and biases. The output of the final layer (predicted 12-month ahead price) is then compared to the desired output (realized stock price), and the difference between the two is measured using a loss function.

2.2. Loss function and optimizer

The Mean Squared Error (MSE) loss function is used to compute the error or loss of the model, representing the extent to which the predictions deviate from the desired output. MSE is the most popular loss function. It directs the optimizer to minimize the squared differences between the predicted and the actual 12-month-ahead stock prices. More specifically, the optimizer, a mathematical function based on the concept of gradient descent (i.e., the iterative reduction of the loss function by following the gradient in a greedy manner), adjusts the weights of each neurons using the gradients. Its function can be as simple as subtracting the gradients from the weights. More complex optimizers can be faster and more efficient, dramatically improving model performance. We chose the Adaptive Moment Estimation (Adam) optimizer, which utilizes adaptive optimization methods, for our model training because it consistently converges faster and has become the default optimization algorithm in many social science research fields. The learning rate of the Adam optimizer in our models is set to 0.001, which is the default value in Pytorch, an open-source machine learning (ML) library using the Python programming language.

The goal of training a neural network is to minimize the loss function by adjusting the weights and biases of neurons in each layers. This is done through backpropagation, where the error (i.e., the value from the loss function using MSE) is propagated backward from the output layer to the hidden layers, and the weights and biases are updated accordingly. Backpropagation allows the network to learn from its mistakes in the current epoch and improve its predictions in the subsequent epoch. This process is repeated multiple times (epochs), with the weights and biases being adjusted after each iteration, until the network reaches a point where the error is minimized, indicating that the predictions of price are accurate. As the weights and biases are updated, the neural network layers undergo a process called learning. Neurons that receive similar inputs most of the time will have their connections strengthened. Neurons that receive rare inputs will have their connections weakened. This process of learning helps the network to become more specialized in recognizing patterns and making accurate predictions.

As training progresses, the layers of neurons become more efficient at extracting important characteristics from the data. Additional fine-tuning is made by modifying the weights and biases to give more importance to specific features that are more valuable for

making accurate prediction of price. This process is commonly referred to as feature learning. In addition to the tuning of the weights and biases during training, we also apply regularization techniques including dropout and L1 by adding a penalty term to the loss function to prevent over-fitting, where the model is too specialized in the training data and fails generalize to new data. Regularization helps avoid over-fitting by inducing the model to learn more general patterns from new data. We first use "dropout" as the regularization method in our model, with a dropout rate of 0.25. During training, the model will take the output from its previous layer, randomly select 25% of the neurons, zero them out, and then pass the output to the next layer, effectively ignoring them. Essentially, this involves randomly dropping one-fourth of the neurons in the network, requiring the other neurons to take over and represent the information needed to make the predictions for the remaining neurons. This may thus force the network to become less sensitive to the exact weights of neurons, potentially reducing the chances of over-fitting the training data and ultimately improving the model's accuracy. Secondly, we use L2 regularization embedded in the Adam (Adaptive Moment Estimation) optimizer, with a weight decay of 0.001, to penalize the squared weights in the loss function. This penalty term is proportional to the sum of the squared weights in the neural network. By adding this term, the loss function becomes:

$$\text{Loss} = \text{Original Loss} + \lambda * \Sigma(\text{Weight}^2)$$

Here, λ is the regularization parameter, and the summation is over all the weights in the network. The regularization parameter controls the strength of the penalty term, with larger values leading to stronger regularization. The effect of L2 regularization is to discourage the model from assigning large weights to individual features. This is because large weights can lead to overfitting, as they allow the model to memorize specific patterns in the training data. By penalizing large weights, L2 regularization encourages the model to distribute the importance of features more evenly, which results in simpler, more generalizable patterns.

2.3. Out-of-sample prediction and total costs of training

Once the neural network has been trained, it is then used to make predictions on the out-of-sample panel dataset. The new dataset (financial statement and other market data that are one calendar-quarter later than the in-sample data) is fed into the network, and the forward propagation process begins. The data is processed through the layers of neurons until it reaches

the final layer, where the prediction is produced. Of course, the accuracy of the prediction depends on the quality of the training data and the complexity of the neural network. In essence, a trained neural network makes predictions by leveraging its capacity to identify patterns and relationships in the data it receives, along with the knowledge it has acquired during the training process.

Table 1 lists the equipment and software used for this research. Because the operating system (Ubuntu) and programming tool (Python) are open-source, which means they are free and available for everybody to use for any purpose, the total costs of this research mainly come from the hardware equipment ($1,450), with half of it allocated to an Nvidia GPU ($700). On the other hand, while each training process includes 5,000 epochs and takes about 15 hours to finish, the electricity consumption ($2.52) is minuscule compared to the vast amount of energy used in mining cryptocurrencies (Dong, 2024) and training large language model (LLM) models like the ChatGPT (Yang et al., 2024).

[Insert Table 1 Here]

III. SAMPLE DATA

The primary source of data for training the neural-network model is a combination of three sources: Compustat, CRSP and IBES. The quarterly and year-end financial statement data from the Compustat 10Q and 10K reports from 2010 to 2023 are linked to the CRSP monthly stock file and the IBES-CRSP Link (Beta) from the WRDS to calculate the month-end market value and ensure that each firm-quarter observation includes at least one analyst's forecast from the IBES database. This generates a panel dataset with 175,718 firm-quarter observations. It is then merged with the IBES Detail History's Price Target, and the forecast period is restricted to 12 months. Both the stock price on the forecast day and the actual realized price after the 12-month period are collected only if the price level is lower than $1,000 to prevent potential distortions caused by outliers. In addition, four historical stock prices (3-month, 6-month, 9-month, and 12-month) prior to the forecast day are also included. Because there are multiple analysts covering the same firm in the same quarter, the sample size increases to 266,092.

One of the tools that equity research analysts often utilize to predict future stock prices is the DCF model. The cost of capital is a key input in DCF analysis, as it directly impacts the present value of the future cash flows. We can simplify the estimation of a firm's cost of capital

by using a risk-free rate and a risky rate as proxies for the rate that a company pays on its funds. We collect two types of interest rate data from the Federal Reserve Economic Database maintained by the Research Division of the Federal Reserve Bank of St. Louis. The market yield on U.S. Treasury securities at a 10-year constant maturity (TCMNOM-Y10) is used to represent the risk-free rate, while the Moody's Seasoned Baa Corporate Bond Yield (DBAA) serves as a proxy for the cost of capital associated with higher investment risk.

To construct the training data for the neural-network in deep-learning, each row (firm-quarter) of the panel dataset includes the historical stock prices (on the prediction day, 3-month prior, 6-month prior, 9-month prior, 12-month prior, and 12-month later), interest rates (10-year risk-free and 1-year Baa-grade), and financial statements with data from the four prior quarters. In the Compustat 10Q/10K dataset, there are 598 unique accounting variables, many of which have missing values in almost every reporting period. After dropping the mostly blank-value accounting variables, there are 154 variables that have non-missing values at least half of the time. The detailed list of accounting variables included in the training data and their definitions can be found in Table 2. The training data consists of a total of 167 columns, which include firm identification (2), time (3), stock price (6), interest rate (2), and financial statement data (154).

[Insert Table 2 Here]

During the training, the Python code will also locate three prior quarters of financial statement data (154 × 3 = 462) for each firm-year observation and append them to the end of each row, increasing the number of variables to 629 inside the Python and Nvidia GPU's memory. The entire dataset is dynamically split into the in-sample dataset for training and the out-of-sample dataset for prediction. For example, to make a 12-month prediction in the first quarter of 2022 (Q1), the neural network is trained for 5,000 epochs using the four consecutive quarterly financial statements from 2020 (Q1, Q2, Q3, and Q4). After the model is trained, it uses the quarterly financial statement data from 2022 Q1 to predict stock prices. Because it is a 12-month prediction, it is forecasting the price for the first quarter of 2023 (Q1). This predicted stock price will be compared with the analyst-forecasted price and the realized stock price in Q1 of 2023.

IV. EMPIRICAL ANALYSIS

4.1. Evaluating AI predictions of stock price

Once the neural network has been trained, it can be used to make predictions on new data that it has not seen before. The input data is fed into the network, and the forward propagation process begins. The data is processed through the layers of neurons until it reaches the final layer, where the prediction is produced. The model compares the predicted price with the actual price and calculates the MSE of the prediction. This helps improve the calibration of the model in the next epoch of the training process. After 5,000 epochs of training using the in-sample dataset, the model makes predictions based on the out-of-sample data. Of course, the accuracy of the prediction depends on the quality of the training data and the complexity of the neural network. The main neutral model used in this research has a total of five layers, and the input size of the first layer is set to the number of variables in the in-sample dataset.

4.1.1. Deep-learning financial statement data

In the first model, we only use data from the financial statement, so the input size is four times the number of accounting variables (154 × 4 = 616). In the Python code, this is specified as:

```
self.linear1 = nn.Linear(616, 4096)
self.linear2 = nn.Linear(4096, 8192)
self.linear3 = nn.Linear(8192, 4096)
self.linear4 = nn.Linear(4096, 2048)
self.linear5 = nn.Linear(2048, 1)
```

It is noted that the choice of the number of epochs (5,000), the number of layers (5) and the size of each layer (2048, 4096, and 8192) in the neutral-network model is entirely arbitrary. In the robustness check section, we will change the number of epochs to 2,500 and 10,000, and reduce the number of layers to three.

In order to evaluate the quality of this AI model's performance in predicting prices, we can compare the 12-month predicted price (P_{AI}) with the realized stock price after 12 months (P_{Realized}) by estimating an accuracy function in the following equation:

$$Accuracy_{\text{AI}} = \left| \frac{P_{\text{Realized}} - P_{\text{AI}}}{P_{\text{Realized}}} \right|$$

The measure is the absolute difference between the realized stock price and the analyst's target price, scaled by the realized price. A lower value of this measure indicates a higher ability to

predict price. Similarly, we can define the accuracy of analyst forecasts by comparing the target price (P_{Analyst}) with the actual stock price (P_{Realized}) as:

$$Accuracy_{\text{Analyst}} = \left| \frac{P_{\text{Realized}} - P_{\text{Analyst}}}{P_{\text{Realized}}} \right|$$

The average accuracy of analyst forecasts in the entire sample from 2010 to 2023 is reported in column (1) of Table 3 Panel A. The statistically significant t-statistic, based on the accuracy measure (2.396) and the standard error (0.025), suggests that, on average, the analysts set inaccurate target prices. This is generally consistent with some earlier findings indicating that analysts have limited abilities to provide accurate target price forecasts (Bradshaw et al., 2013), but it does not necessarily agree with other studies (Bilinski et al., 2013). However, the focus of this research is on the out-of-sample prediction ability of the AI model after being trained on the in-sample data. Column (2) of Table 3 Panel A reports a slightly lower accuracy measure (3.837). In fact, the difference between the accuracy of analyst forecasts and that of AI predictions is significantly negative (–1.441) at the one-percent level. In other words, analysts have a better ability to forecast the 12-month-ahead stock price.

[Insert Table 3 Here]

4.1.2. Deep-learning financial statement and interest rate data

Interest rates are a crucial element of a country's monetary policy and are closely monitored by analysts because they can impact the cost of borrowing and spending. This, in turn, can affect stock prices as companies may adjust their financing and investment decisions based on interest rates. For instance, when interest rates are low, companies may be more inclined to borrow and invest in expanding their business. This can lead to potential growth and an increase in stock prices. Moreover, equity research analysts also pay attention to central bank decisions on interest rates, as any changes in these rates can significantly impact the economy and the stock market. If the central bank raises interest rates, it can result in higher costs for businesses, potentially leading to slower economic growth and a decrease in stock prices. On the other hand, a decrease in interest rates may signal a strong economy and result in increased consumer spending, which could have a positive impact on stock prices.

Equity research analysts utilize interest rate data not only to monitor interest rates on a

macro level but also to analyze specific companies. If a company has a high level of debt, an increase in interest rates may raise their cost of borrowing, ultimately affecting their earnings and stock price. Conversely, a company with a strong balance sheet and low debt may be less affected by changes in interest rates.

Given the importance of interest rates as both a macroeconomic indicator and a cost factor in a DCF-valuation model, we include two interest rate measures in the training data for the second model. The second row in Table 3 Panel A reports the difference in accuracy measures between analyst forecasts and AI predictions. The smaller difference (-0.992) suggests that while the AI model still performs worse than analysts, additional information about the interest rates in each quarter certainly helps improve the model's prediction ability.

4.1.3. Deep-learning financial statement, interest rate and historical stock price

The market efficiency hypothesis suggests that all relevant information is already incorporated into the stock price. Therefore, the predictions made by equity research analysts for future stock prices would largely be based on the current stock price. Current stock prices reflect the collective expectations and beliefs of market participants about the company's future performance. This includes factors such as the company's financials, management team, industry trends, and economic conditions. This information is constantly being updated and absorbed by the market; hence, the current stock price is considered the most accurate reflection of the company's value at any given point in time. Moreover, equity research analysts primarily use fundamental analysis to predict stock prices. This involves analyzing a company's financial data, management performance, and economic factors to determine its intrinsic value. The current stock price provides a benchmark for comparing the intrinsic value of a company with its market value. If the current stock price is lower than the intrinsic value, it may indicate that the stock is undervalued and has the potential to increase in the future. Conversely, if the current stock price exceeds the intrinsic value, it may suggest an overvalued stock that could experience a decline in the future.

Since analysts heavily rely on the current stock price to make predictions about future stock prices, it is crucial to include this factor in the deep-learning model. Not surprisingly, the accuracy of AI predictions improves dramatically after the model is trained on both financial statement and historical stock price data. The sign of the difference in prediction accuracy

between the analyst forecast and AI prediction has changed from negative to positive (0.600), and it is statistically significant at the one-percent level (third row in Table 3 Panel A).

The poor accuracy of analyst forecasts deserves attention because both human analysts and AI models are exposed to the same historical accounting and price data. The significant deviation in accuracy suggests that either human analysts were incompetent or biased when studying companies with these characteristics. However, it is also likely that the announcement of the analyst's target price has caused market over-reaction, leading to a change in the subsequent market price. This is consistent with the evidence that analysts create value for firms under their coverage by enhancing their investor recognition rather than by monitoring or reducing information asymmetry (Li and You, 2015). In other words, the analyst may have made an accurate forecasts, but his or her target price has caused a significant drift in the realized price over the next 12 months (Dechow and You, 2020).

Finally, when we include all three types of data (financial statements, stock prices, and interest rates) in the training dataset, the improvement in prediction accuracy actually reduces to 0.355 in the fourth row in Table 3 Panel A. However, the smaller average value does not necessarily indicate that the model with two sets of data is better than the one with three. The answer depends on the model's goodness of fit.

4.2 Evaluating AI model's goodneet of fit

We first calculate the mean squared error (MSE) which is the average squared difference between the realized price and the predicted price:

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(P_{\text{Realized}} - P_{\text{Forecast}})^2$$

The first table in Table 3 Panel B reports the MSE for analyst forecast, the MSE for AI predictions, and the difference between the two. Clearly, the model with all three types of data (financial statements, interest rates and stock prices) in the fourth row has the lowest MSE (1982.8) than other models, and it is significantly smaller than the MSE of analyst forecast (4223.9).

We also estimate the Predicted R-squared, which is the proportion of variance in the realized stock prices that can be explained by the variance in the out-of-sample predicted prices:

$$Predicted\ R^2 = 1 - \frac{\sum_{i=1}^{N}(P_{\text{Realized}} - P_{\text{Forecast}})^2}{\sum_{i=1}^{N}(P_{\text{Realized}} - \overline{P}_{\text{Realized}})^2}$$

Again, the model with all historical data has the best prediction ability, as suggested by the highest value of the Predicted R-squared (0.64) in the fourth row in the second table in Table 3 Panel B. On the other hand, the model deep-learning only financial statement data has a Predicted R-squared of 0.07, which is much smaller than the Predicted R-squared of the analyst forecast (0.23) as shown in the first row. This suggests that while the model deep-learning accounting data does not explain much of the price variability (7%), supplementing with historical stock price and interest rate data dramatically improves the fit of the prediction model, being able to explain 64% of the variation in the realized stock price.

4.3. Breakdown of prediction accuracy comparison

We then break down the comparison into four quartiles of the stock price on the day when both the analyst and the AI model make a prediction and report the differences in prediction accuracy for each price segment separately in Table 4 Panel A. The most significant improvement in AI prediction accuracy over human analysts can be found when the stock price is low on the prediction day. The difference is 1.439 when the stock price is below $18.36, which is the cutoff price between the first and the second quartiles. This improvement in prediction is more than ten times larger than the other price ranges when the stock price is higher than $18.36 and smaller than $39.94, as in the second quartile. For stock prices that are higher than the medium value, the accuracy of analyst forecasts is slightly better than that of AI predictions, particularly when the stock price is above $74.74, which is the cutoff price between the third and forth quartiles (0.408 for analysts vs. 0.416 for AI).

[Insert Table 4 Here]

When we examine the differences in accuracy between analyst forecasts and AI predictions at the year level, it is interesting to find that AI predictions retain most of their edge, in ten out of 13 years, as shown in Table 4 Panel B. Analyst forecasts perform better than AI predictions only in 2017, 2018, and 2022.

4.4. Determinants of prediction deviation between analyst and AI model

The results presented thus far have shown the ability of AI to analyze large amounts of financial statement data and generate reliable stock price predictions; however, empirically it is equally of interest to examine how and why the properties of the predictions by deep-learning models, differ from those produced by traditional human analysis. While both methods aim to provide insights into the stock market and accounting information, they differ in their approach, methodology, and the nature of the output they produce (Coleman et al., 2022). In this section, we construct a set of measures of firm characteristics to conduct a series of fixed-effects regressions to investigate the factors that could potentially drive the deviation between analyst forecasts and AI predictions. Table 5 Panel A provides a list of variables and their definitions.

[Insert Table 5 Here]

The dependent variable is the difference in the predicted prices of the analyst and AI model, scaled by the average price of the two, and we call it "Deviation":

$$Deviation_{\text{Analyst,AI}} = \frac{\left|P_{\text{Analyst}} - P_{\text{AI}}\right|}{\frac{P_{\text{Analyst}} - P_{\text{AI}}}{2}}$$

Panel B reports the summary statistics of all variables used in the subsequent regression analysis. Overall, the firms that receive analyst coverage for their 12-month target price have quite a high level of institutional ownership. Institutional investors hold, on average, 73% of the outstanding shares, and the median value is even higher (82%). In addition, the Pearson's correlation coefficients among all independent variables to be used in the regression analysis are shown in the lower triangle of Table 6, and the Spearman's rank correlations are shown above the diagonal in the same table. There is some degree of correlation between the level of institutional ownership and the stock price (0.33).

[Insert Table 6 Here]

In the first set of regression tests using quarterly time fixed-effects in Table 7, there is some evidence that analysts and the AI model make similar predictions for firms that have

larger total assets, lower leverage, and higher ROA, cash, and capital expenditures (Column 1 and Column 3 with stock return volatility). When we add the level of stock price (the logarithm of price on the prediction day), the effects of leverage and ROA become insignificant (Column 2 with price and Column 4 with both price and volatility). It should be noted that both human analysts and the AI model have access to all accounting information and stock prices on the prediction date (all included in the training dataset). Therefore, the significant deviation effects of cash, capital expenditures, and stock price must be caused by analysts' incompetence and bias when they were studying companies with these characteristics. It is also important to emphasize the changing role of the market-to-book variable in Columns (2) and (4). It is interesting to note that the stock price variable captures a significant amount of effects from the market-to-book variable, turning the sign of the market-to-book ratio from negative to positive (Column 2). Subsequently, the variable of return volatility captures the entire effect of the market-to-book variable in Column (4).

Extensive research in finance suggests that the ownership of institutional investors is closely related to analysts' forecasts of the target price. Analysts play a more important function for smaller firms that have less institutional ownership (Asquith et al., 2005), and analyst forecasts are more valuable (i.e., stronger market reaction) for more opaque firms that have lower institutional ownership (Loh and Stulz, 2018). Because the training data does not contain institutional ownership data, the AI model does not have any knowledge about this information. This suggests that for firms with lower institutional ownership, the deviation in predicted price between the analysts and the AI model may increase. In other words, a higher degree of institutional ownership can reduce the deviation. While the negative sign of institutional ownership reported in Column 5 does seem to support the hypothesis, it is statistically insignificant.

Overall, in this set of benchmark regression results, the degree of similarity in stock price forecasts is mainly driven by lower market-to-book and higher cash, capital expenditures, stock price, and volatility. This suggests that the analysts and the AI model are more likely to agree with other in terms of their future performance about the firms that exhibit signs of investment and uncertainty on the day of forecast.

[Insert Table 7 Here]

We repeat the same set of regression analyses with additional industry fixed-effects based on the firm's 2-digit SIC code to capture unit-invariant heterogeneity due to both time and industry, and the results remain largely unchanged as shown in Table 8.

[Insert Table 8 Here]

However, when we replace the industry fixed-effects by the firm fixed-effects to examine the "within variation" of the same firm over time, the effects of market-to-book, ROA, and capital expenditures become insignificant (Table 9). Analysts and the AI model made similar predictions when the firm's stock price and return volatility were higher on the day of the forecast. Again, institutional ownership does not play an important role in driving the similarity of their performance in the target price forecast.

[Insert Table 9 Here]

Finally, when we include the analyst fixed-effects, rather than the firm or industry fixed-effects, in our regression specifications, the effects of market-to-book, cash, and capital expenditure become significant again, along with R&D, stock price, and return volatility (Table 10). This is similar to the benchmark regression results reported earlier, showing that the degree of similarity in stock price forecasts is mainly driven by lower market-to-book and higher cash, capital expenditures, stock price, and volatility. After the analyst-level fixed-effects remove the unmeasured forecast ability of each individual analyst, both the R&D expenditures and the ownership of institutional investors show slightly negative effects, although only significant at the 5-percent and 10-percent levels, respectively, reducing the difference in 12-month-ahead price forecasts between the analysts and the AI model. Because the AI model is not trained on the information of institutional ownership, the negative effect of institutional ownership does support the hypothesis that analyst forecasts are more valuable for more opaque firms that have lower institutional ownership (Loh and Stulz, 2018).

[Insert Table 10 Here]

It is interesting to note that if we assume that AI-generated predictions are based on impartial observation and objective analysis (Coleman et al., 2022), then the deviation of analysts' forecasts from AI predictions reflects the degree of optimism bias or the need to curry favor with companies' management.

V. ROBUSTNESS CHECKS

The number of epochs, or the number of times the entire dataset is fed into the neural network during training, is a critical factor in building a fit and efficient AI model. The impact of epoch numbers is particularly significant when shifting from shallow-learning to deep-learning. In shallow-learning, the neural network consists of only one or two hidden layers and fewer number of epochs, making it relatively easy and fast to train. Certainly, in the early stages of training, the neural-network may have a high error rate and produce inaccurate predictions. The training process continues with each epoch to further adjust weights and biases, reduce the error rate, and improve accuracy. On the other hand, as neural networks become "deeper," with multiple hidden layers, the number of epochs needed for successful training significantly increases.

5.1. Prediction performance of a "Deeper"-learning model

In this robustness section, we use the same in-sample dataset to train three neural-networks in one "deeper-learning" model and two "shallow-learning" models. The "deeper-learning" neural-network requires more repetition, 10,000 epochs rather than 5,000 epochs as used in the previous training efforts, to learn the non-linear relationship between the financial statement variables (along with historical stock prices and interest rates) and the 12-month realized stock price. The accuracy measures of both analyst forecast and AI prediction using the "deeper-learning" model are reported in Columns (1) and (2), respectively, in the first row of Table 11. The difference of 0.573, which is significant at the one-percent level, is much larger than the difference of 0.355 reported in the model using 5,000 epochs. However, the measure of model fit does not improve compared with the previous results. The MSE of the model with 10,000-epoch training is 2080.7, whereas the MSE of the model with 5,000-epoch training is 1982.8. Similarly, the Predicted R-squared of the 10,000-epoch model is 0.62 compared to that of

the 5,000-epoch model, which is 0.64. This comparison suggests that while the neural-network is able to recognize a broader range of patterns and make more precise out-of-sample predictions after being trained with more epochs using the same in-sample dataset, the improvement in efficiency is still limited.

[Insert Table 11 Here]

5.2. Prediction performance of two "Shallow"-learning models

Another crucial factor to consider is the potential for over-fitting in a deep-learning model when a network becomes too specialized in recognizing patterns in the training data, resulting in poor performance on out-of-sample prediction. A "shallow" neural-network can often achieve good results with a small number of epochs as the network can quickly learn and adapt to the training data. In the second robustness check, we reduce the number of epochs from 5,000 to 2,500, and the accuracy comparison is reported in the second row of Table 11. Certainly, the improvement in accuracy (0.125) in this model with only 2,500 epochs is smaller than that of the model with 5,000 epochs (0.355). In addition, both the MSE (2673.3) and the Predicted R-squared (0.51) are worse than that of the 5,000-epoch model (1982.8 and 0.64, respectively). This result is not very surprising because a model with fewer iterations simply can not learn the underlying patterns and relationships within the data adequately.

Another challenge with deep-learning models is the risk of overfitting. With more layers, the network becomes more complex and has the potential to memorize the training data rather than learning the underlying patterns. This could be due to the problem of vanishing gradients: as the number of layers increases, the gradients that are used to update the weights of each neuron in the neural-network model become very small. It results in slower training and can even cause the network to stop learning altogether, leading to poor performance in making predictions using the out-of-sample data.

In the last robustness check, we build a "shallow-learning" model that has only three layers (one input, one hidden, and one output layer), rather than the five layers used in the previous models. This three-layer neural-network is relatively simple and easy to train as it has a limited number of parameters, helping stabilize the gradients and facilitating faster and more stable training. The forecast accuracy of this 3-layer model with 5,000-epoch training is reported

in the third row of Table 11. It performs well in predicting 12-month stock price using the out-of-sample data with an accuracy score of 2.112, which is slightly better than the analyst forecast's 2.396 and very similar to the benchmark 5-layer 5,000-epoch model's 2,041. However, its efficiency, as measured by its goodness of fit, is worse. The MSE is 5366.8 compared to the analysts' 4225.9 and the benchmark model's 2983.8, and the Predicted R-squared is only 0.02 compared to the analysts' 0.23 and the benchmark model's 0.64. This suggests that its limited depth in the neural network severely limits its ability to learn complex functions and patterns from the data.

5.2. Possibility of using linear regressions to deep-learn financial statements

When using linear regression models such as the OLS to analyze financial statements, the presence of multicollinearity can lead to inaccurate and unreliable results. In accounting data, it is common for different variables to be interdependent but have a strong correlation with each other. For example, in the income statement, the revenue and cost of goods sold are highly correlated, as they are both affected by the level of sales. Similarly, on the balance sheet, assets and liabilities are interrelated, as the purchase of assets often involves taking on additional liabilities. Without creating meaningful financial ratios (as the purpose of this research is to let artificial intelligence learn by itself), the high correlation between accounting variables can lead to problems in generating regression estimates, often causing the coefficients to become unstable or have large standard errors. In our experiments of in-sample OLS regressions and out-of-sample predictions, the data constantly cause estimation of errors "matrix not positive definite" in Stata.

One of the key advantages of neural networks over linear regression models, such as the OLS, is that they are less vulnerable to collinearity-type problems because they tend to be overparameterized, as in our research (using 629 variables; see the section of sample data for more detailed discussion). A large number of parameters are fitted using backward propagation, and the learned weights of neurons from such a repetitive training process create redundancies that make things that affect any small subset of features (such as multicollinearity) less important (Fang et al., 2022). Therefore, the ability of a network of interconnected nodes to learn and extract patterns from the predictors helps generate accurate predictions, regardless of whether the predictors are correlated or not. In some cases, data multicollinearity can even

improve the performance of neural networks, as it provides additional information for the network to learn from.

VI. CONCLUSION

In recent years, advancements in artificial intelligence have revolutionized the field of stock market analysis. AI systems have the ability to analyze vast amounts of data at a speed and accuracy that far surpasses that of human analysts. By using machine-learning algorithms, AI can identify patterns and trends in the market that may go unnoticed by human analysts. In this research, we are able to train a simple neural-network model to deep-learn financial statement data without using pre-specified financial models or pre-defined accounting ratios. It slightly underperforms human analysts in 12-month-ahead target price forecasts; however, with additional data on historical stock prices and interest rates in the last three, six, nine, and twelve months, AI models can make more accurate predictions than analysts.

The deviation in predicted target prices between the analysts and the AI model is smaller, in other words, they are more likely to agree with each other, in firms that exhibit signs of investment and uncertainty, such as having a higher level of cash, capital expenditures, and return volatility on the day of forecast. The empirical findings in this paper demonstrate that AI does have the potential to outperform stock analysts in predicting the performance of stocks. While human intuition and experience are invaluable in the world of finance, the sheer processing power and data analysis capabilities of AI can provide a significant advantage in the prediction of stock performance. As AI technology continues to evolve, it is likely that we will see an increasing number of investors turning to AI systems for their market predictions.

REFERENCE

Asquith, P., Mikhail, M.B., Au, A.S., 2005. Information content of equity analyst reports. Journal of Financial Economics 75, 245–282.

Barbaglia, L., Manzan, S., Tosetti, E., 2023. Forecasting Loan Default in Europe with Machine Learning*. Journal of Financial Econometrics 21, 569–596.

Beyer, A., Cohen, D.A., Lys, T.Z., Walther, B.R., 2010. The financial reporting environment: Review of the recent literature. Journal of Accounting and Economics 50, 296–343.

Bilinski, P., Lyssimachou, D., Walker, M., 2013. Target Price Accuracy: International Evidence. The Accounting Review 88, 825–851.

Birru, J., Gokkaya, S., Liu, X., Stulz, R.M., 2022. Are Analyst Short-Term Trade Ideas Valuable? The Journal of Finance 77, 1829–1875.

Bradshaw, M.T., Brown, L.D., Huang, K., 2013. Do sell-side analysts exhibit differential target price forecasting ability? Review of Accounting Studies 18, 930–955.

Brav, A., Lehavy, R., 2003. An Empirical Analysis of Analysts' Target Prices: Short-term Informativeness and Long-term Dynamics. The Journal of Finance 58, 1933–1967.

Brown, L.D., Call, A.C., Clement, M.B., Sharp, N.Y., 2016. The activities of buy-side analysts and the determinants of their stock recommendations. Journal of Accounting and Economics 62, 139–156.

Cao, S., Jiang, W., Wang, J.L., Yang, B., Forthcoming. From Man vs. Machine to Man Machine: The Art and AI of Stock Analyses. Journal of Financial Economics.

Coleman, B., Merkley, K., Pacelli, J., 2022. Human Versus Machine: A Comparison of Robo-Analyst and Traditional Research Analyst Investment Recommendations. The Accounting Review 97, 221–244.

De Veaux, R.D., Ungar, L.H., 1994. Multicollinearity: A tale of two nonparametric regressions, in: Cheeseman, P., Oldford, R.W. (Eds.), Selecting Models from Data. Springer, New York, NY, pp. 393–402.

Dechow, P.M., You, H., 2020. Understanding the Determinants of Analyst Target Price Implied Returns. The Accounting Review 95, 125–149.

Dong, G.N., 2024. Carbon Footprint and Market Value of Cryptocurrencies and the Real Economy, in: FinTech and Green Investment: Transforming Challenges into Opportunities. World Scientific Publishing, Singapore.

Dormann, C.F., Elith, J., Bacher, S., Buchmann, C., Carl, G., Carré, G., Marquéz, J.R.G., Gruber, B., Lafourcade, B., Leitão, P.J., Münkemüller, T., McClean, C., Osborne, P.E., Reineking,

B., Schröder, B., Skidmore, A.K., Zurell, D., Lautenbach, S., 2013. Collinearity: a review of methods to deal with it and a simulation study evaluating their performance. Ecography 36, 27–46.

Erel, I., Stern, L.H., Tan, C., Weisbach, M.S., 2021. Selecting Directors Using Machine Learning. The Review of Financial Studies 34, 3226–3264.

Fang, C., Gu, Y., Zhang, W., Zhang, T., 2022. Convex Formulation of Overparameterized Deep Neural Networks. IEEE Transactions on Information Theory 68, 5340–5352.

Gorodnichenko, Y., Pham, T., Talavera, O., 2023. The Voice of Monetary Policy. American Economic Review 113, 548–584.

Green, J., Hand, J.R.M., Sikochi, A., 2024. The asymmetric mispricing information in analysts' target prices. Review of Accounting Studies 29, 889–915.

Igami, M., 2020. Artificial intelligence as structural estimation: Deep Blue, Bonanza, and AlphaGo. The Econometrics Journal 23, S1–S24.

Iskhakov, F., Rust, J., Schjerning, B., 2020. Machine learning and structural econometrics: contrasts and synergies. The Econometrics Journal 23, S81–S124.

Joos, P., Piotroski, J.D., Srinivasan, S., 2016. Can analysts assess fundamental risk and valuation uncertainty? An empirical analysis of scenario-based value estimates. Journal of Financial Economics 121, 645–663.

Khachiyan, A., Thomas, A., Zhou, H., Hanson, G., Cloninger, A., Rosing, T., Khandelwal, A.K., 2022. Using Neural Networks to Predict Microspatial Economic Growth. American Economic Review: Insights 4, 491–506.

Li, K.K., You, H., 2015. What is the value of sell-side analysts? Evidence from coverage initiations and terminations. Journal of Accounting and Economics 60, 141–160.

Livnat, J., Zhang, Y., 2012. Information interpretation or information discovery: which role of analysts do investors value more? Review of Accounting Studies 17, 612–641.

Loh, R.K., Stulz, R.M., 2018. Is Sell-Side Research More Valuable in Bad Times? The Journal of Finance 73, 959–1013.

Maliar, L., Maliar, S., Winant, P., 2021. Deep learning for solving dynamic economic models. Journal of Monetary Economics 122, 76–101.

Sadhwani, A., Giesecke, K., Sirignano, J., 2021. Deep Learning for Mortgage Risk. Journal of Financial Econometrics 19, 313–368.

Yang, J., Jin, H., Tang, R., Han, X., Feng, Q., Jiang, H., Zhong, S., Yin, B., Hu, X., 2024. Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond. ACM Transactions on Knowledge Discovery from Data.

Figure 1. Neural network model

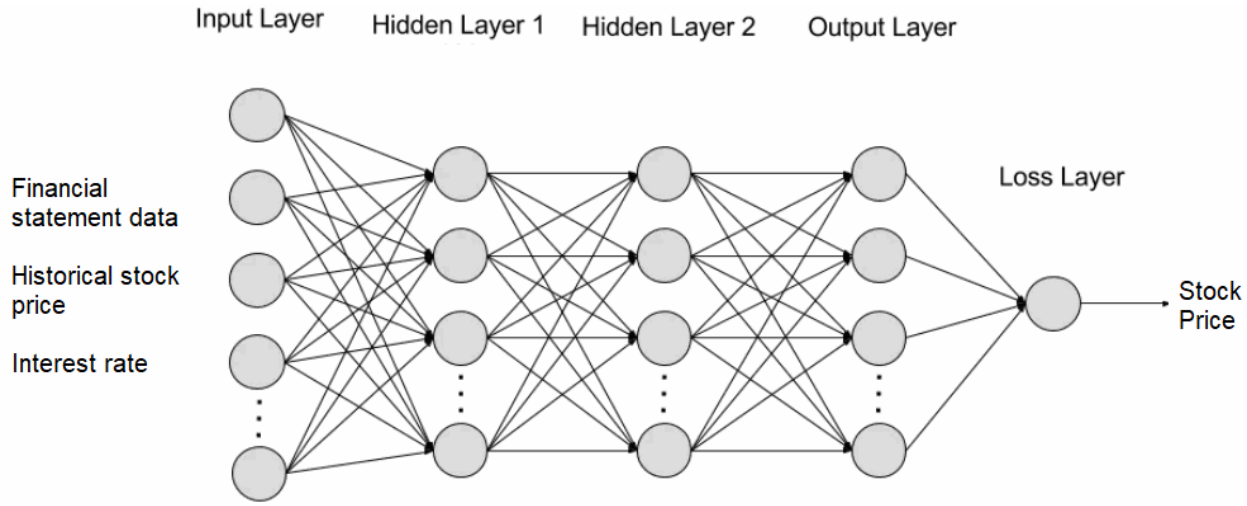


Table 1. List of software, equipment, training time, and power consumption

Software		TDP (Watts)	MSRP ($)
Operating system	Ubuntu 22.04 Linux distribution based on Debian	0	Free
Coding tool	Python 3.10	0	Free
Neural network library	Pytorch 2.2	0	Free
Statistical model library	Scikit-learn 1.4	0	Free
Hardware			
Motherboard	Gigabyte Z590 Aorus Ultra LGA1200	50W	$400
CPU	Intel Core i5-11400T 3.70 GHz 6-core 12-thread	35W	$200
DIMM memory	Samsung 64GB 2666MHz 4 × 16GB	5W	$100
Nvidia GPU	Zotac Gaming GeForce RTX 3080 Trinity OC	320W	$700
SSD Storage	Western Digital 1TB SN580 NVMe	10W	$50
Total cost of hardware		420W	$1,450
Energy consumption			
Total running time	Each training includes 5,000 epochs	15 hours	
Unit cost of electricity	Average price per kWh in Boston area (after tax)	$0.40	
Total cost of electricity		$2.52	

Table 2. List of financial statement variables used in deep-learning

Variable Name	Description
ACCHGQ	Accounting Changes - Cumulative Effect
ACOMINCQ	Accumulated Other Comprehensive Income
ACOQ	Current Assets - Other - Total
ACTQ	Current Assets - Total
ALTOQ	Other Long-term Assets
ANCQ	Non-Current Assets - Total
ANOQ	Assets Netting & Other Adjustments
AOCIDERGLQ	Accum Other Comp Inc - Derivatives Unrealized Gain/Loss
AOCIOTHERQ	Accum Other Comp Inc - Other Adjustments
AOCIPENQ	Accum Other Comp Inc - Min Pension Liab Adj
AOCISECGLQ	Accum Other Comp Inc - Unreal G/L Ret Int in Sec Assets
AOL2Q	Assets Level2
AOQ	Assets - Other - Total
APQ	Account Payable/Creditors - Trade
AQPL1Q	Assets Level1
ATQ	Assets - Total
AUL3Q	Assets Level3
CAPSQ	Capital Surplus/Share Premium Reserve
CEQQ	Common/Ordinary Equity - Total
CHEQ	Cash and Short-Term Investments
CHQ	Cash
CIBEGNIQ	Comp Inc - Beginning Net Income
CICURRQ	Comp Inc - Currency Trans Adj
CIDERGLQ	Comp Inc - Derivative Gains/Losses
CIMIIQ	Comprehensive Income - Noncontrolling Interest
CIOTHERQ	Comp Inc - Other Adj
CIPENQ	Comp Inc - Minimum Pension Adj
CIQ	Comprehensive Income - Total
CISECGLQ	Comp Inc - Securities Gains/Losses
CITOTALQ	Comprehensive Income - Parent
COGSQ	Cost of Goods Sold
CSH12Q	Common Shares Used to Calculate Earnings Per Share - 12M Moving
CSHFD12	Common Shares Used to Calc Earnings Per Share - Fully Diluted - 12M Moving
CSHFDQ	Com Shares for Diluted EPS
CSHIQ	Common Shares Issued
CSHOPQ	Total Shares Repurchased - Quarter
CSHOQ	Common Shares Outstanding
CSHPRQ	Common Shares Used to Calculate Earnings Per Share - Basic
CSTKCVQ	Carrying Value
CSTKEQ	Common Stock Equivalents - Dollar Savings
CSTKQ	Common/Ordinary Stock
DCOMQ	Deferred Compensation
DD1Q	Long-Term Debt Due in One Year
DILADQ	Dilution Adjustment
DILAVQ	Dilution Available - Excluding Extraordinary Items
DLCQ	Debt in Current Liabilities
DLTTQ	Long-Term Debt - Total
DOQ	Discontinued Operations
DPACTQ	Depreciation, Depletion and Amortization

(Table 2 continued)

Variable Name	Description
DPQ	Depreciation and Amortization - Total
DRCQ	Deferred Revenue - Current
DRLTQ	Deferred Revenue - Long-term
DVPQ	Dividends - Preferred/Preference
EPSF12	Earnings Per Share
EPSFI12	Earnings Per Share
EPSFIQ	Earnings Per Share
EPSFXQ	Earnings Per Share
EPSPI12	Earnings Per Share
EPSPIQ	Earnings Per Share
EPSPXQ	Earnings Per Share
EPSX12	Earnings Per Share
ESOPCTQ	Common ESOP Obligation - Total
ESOPNRQ	Preferred ESOP Obligation - Non-Redeemable
ESOPRQ	Preferred ESOP Obligation - Redeemable
ESOPTQ	Preferred ESOP Obligation - Total
GDWLQ	Goodwill
IBADJ12	Income Before Extra Items - Adj for Common Stock Equivalents - 12MM
IBADJQ	Income Before Extraordinary Items - Adjusted for Common Stock Equivalents
IBCOMQ	Income Before Extraordinary Items - Available for Common
IBMIIQ	Income before Extraordinary Items and Noncontrolling Interests
IBQ	Income Before Extraordinary Items
ICAPTQ	Invested Capital - Total - Quarterly
INTANOQ	Other Intangibles
INTANQ	Intangible Assets - Total
INVFGQ	Inventory - Finished Goods
INVOQ	Inventory - Other
INVRMQ	Inventory - Raw Materials
INVTQ	Inventories - Total
INVWIPQ	Inventory - Work in Process
IVLTQ	Total Long-term Investments
IVSTQ	Short-Term Investments- Total
LCOQ	Current Liabilities - Other - Total
LCTQ	Current Liabilities - Total
LLTQ	Long-Term Liabilities
LNOQ	Liabilities Netting & Other Adjustments
LOL2Q	Liabilities Level2
LOQ	Liabilities - Other
LOXDRQ	Liabilities - Other - Excluding Deferred Revenue
LQPL1Q	Liabilities Level1
LSEQ	Liabilities and Stockholders Equity - Total
LTMIBQ	Liabilities - Total and Noncontrolling Interest
LTQ	Liabilities - Total
LUL3Q	Liabilities Level3
MIBNQ	Noncontrolling Interests - Nonredeemable - Balance Sheet
MIBQ	Noncontrolling Interest - Redeemable - Balance Sheet
MIBTQ	Noncontrolling Interests - Total - Balance Sheet
MIIQ	Noncontrolling Interest - Income Account
MSAQ	Accum Other Comp Inc - Marketable Security Adjustments

(Table 2 continued)

Variable Name	Description
NIQ	Net Income
NOPIQ	Non-Operating Income
NPQ	Notes Payable
OEPF12	Earnings Per Share - Diluted - from Operations - 12MM
OEPS12	Earnings Per Share from Operations - 12 Months Moving
OEPSXQ	Earnings Per Share - Diluted - from Operations
OIADPQ	Operating Income After Depreciation - Quarterly
OIBDPQ	Operating Income Before Depreciation - Quarterly
OPEPSQ	Earnings Per Share from Operations
PIQ	Pretax Income
PNRSHOQ	Nonred Pfd Shares Outs
PPEGTQ	Property, Plant and Equipment - Total
PPENTQ	Property Plant and Equipment - Total
PRCRAQ	Repurchase Price - Average per share Quarter
PRSHOQ	Redeem Pfd Shares Outs
PSTKNQ	Preferred/Preference Stock - Nonredeemable
PSTKQ	Preferred/Preference Stock
PSTKRQ	Preferred/Preference Stock - Redeemable
RDIPAQ	In Process R&D Expense After-tax
RDIPDQ	In Process R&D Expense Diluted EPS Effect
RDIPEPSQ	In Process R&D Expense Basic EPS Effect
RDIPQ	In Process R&D
RECDQ	Receivables - Estimated Doubtful
RECTAQ	Accum Other Comp Inc - Cumulative Translation Adjustments
RECTOQ	Receivables - Current Other incl Tax Refunds
RECTQ	Receivables - Total
RECTRQ	Receivables - Trade
REQ	Retained Earnings
REUNAQ	Unadjusted Retained Earnings
REVTQ	Revenue - Total
SALEQ	Sales/Turnover
SEQOQ	Other Stockholders- Equity Adjustments
SEQQ	Stockholders Equity > Parent > Index Fundamental > Quarterly
SPIQ	Special Items
STKCOQ	Stock Compensation Expense
TEQQ	Stockholders Equity - Total
TFVAQ	Total Fair Value Assets
TFVLQ	Total Fair Value Liabilities
TSTKNQ	Treasury Stock - Number of Common Shares
TSTKQ	Treasury Stock - Total
TXDBAQ	Deferred Tax Asset - Long Term
TXDBCAQ	Current Deferred Tax Asset
TXDBCLQ	Current Deferred Tax Liability
TXDBQ	Deferred Taxes - Balance Sheet
TXDITCQ	Deferred Taxes and Investment Tax Credit
TXPQ	Income Taxes Payable
TXTQ	Income Taxes - Total
TXWQ	Excise Taxes
WCAPQ	Working Capital

(Table 2 continued)

Variable Name	Description
XACCQ	Accrued Expenses
XIDOQ	Extraordinary Items and Discontinued Operations
XINTQ	Interest and Related Expense- Total
XIQ	Extraordinary Items
XOPRQ	Operating Expense- Total
XRDQ	Research and Development Expense
XSGAQ	Selling, General and Administrative Expenses
TFVLQ	Total Fair Value Liabilities
TSTKNQ	Treasury Stock - Number of Common Shares
TSTKQ	Treasury Stock - Total
TXDBAQ	Deferred Tax Asset - Long Term
TXDBCAQ	Current Deferred Tax Asset
TXDBCLQ	Current Deferred Tax Liability
TXDBQ	Deferred Taxes - Balance Sheet
TXDITCQ	Deferred Taxes and Investment Tax Credit
TXPQ	Income Taxes Payable
TXTQ	Income Taxes - Total
TXWQ	Excise Taxes
WCAPQ	Working Capital
XACCQ	Accrued Expenses
XIDOQ	Extraordinary Items and Discontinued Operations
XINTQ	Interest and Related Expense- Total
XIQ	Extraordinary Items
XOPRQ	Operating Expense- Total
XRDQ	Research and Development Expense
XSGAQ	Selling, General and Administrative Expenses

Table 3. Univariate comparison of price prediction accuracy

Panel A. Difference-in-difference test

This table reports the accuracy of analyst forecasts and the associated paired t-test in column (1). The accuracy is defined as the average value of the absolute differences between the realized stock price and analyst's target price scaled by the realized price:

$$Accuracy_{\text{Analyst}} = \left| \frac{P_{\text{Realized}} - P_{\text{Analyst}}}{P_{\text{Realized}}} \right| \quad (1)$$

Column (2) reports the accuracy of AI predictions and the associated paired t-test. The accuracy is defined as the average value of the absolute differences between the realized stock price and AI-predicted price scaled by the realized price:

$$Accuracy_{\text{AI}} = \left| \frac{P_{\text{Realized}} - P_{\text{AI}}}{P_{\text{Realized}}} \right| \quad (2)$$

The difference between the above two values in (1) and (2) and the associated paired t-test result are reported in the third column in the table. Standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively.

AI Prediction Uses Data in			N	Accuracy of Analyst Forecast	Accuracy of AI Prediction	Difference
Financial Statement	Interest Rate	Stock Price		(1)	(2)	(1) – (2)
Yes			251,347	2.396*** (0.025)	3.837*** (0.025)	–1.441*** (0.025)
Yes	Yes		251,347	2.396*** (0.025)	3.388*** (0.023)	–0.992*** (0.026)
Yes		Yes	251,347	2.396*** (0.025)	1.796*** (0.011)	+0.600*** (0.021)
Yes	Yes	Yes	251,347	2.396*** (0.025)	2.041*** (0.013)	+0.355*** (0.021)

Panel B. Goodness of fit

This reports two measures of goodness of fit in the prediction of stock price. The first measure is the mean squared error which is the average squared difference between the realized price and the predicted price:

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(P_{\text{Realized}} - P_{\text{Forecast}})^2$$

The difference between the MSE of analyst forecast and MSE of AI out-of-sample prediction are reported in the third column in the table.

AI Prediction Uses Data in			N	$MSE_{\text{Analyst Forecast}}$	$MSE_{\text{AI out-of-sample Prediction}}$	Difference
Financial Statement	Interest Rate	Stock Price				
Yes			251,347	4223.9	5076.3	–852.4
Yes	Yes		251,347	4223.9	6382.8	–2158.9
Yes		Yes	251,347	4223.9	2490.1	1733.8
Yes	Yes	Yes	251,347	4223.9	1982.8	2241.1

The second measure is the Predicted R-squared which is the proportion of variance in the realized stock prices that can be explained by the variance in the out-of-sample predicted prices:

$$Predicted\ R^2 = 1 - \frac{\sum_{i=1}^{N}(P_{\text{Realized}} - P_{\text{Forecast}})^2}{\sum_{i=1}^{N}(P_{\text{Realized}} - \overline{P}_{\text{Realized}})^2}$$

The difference between the predicted R^2 of analyst forecast and the predicted R^2 of AI out-of-sample prediction are reported in the third column in the table.

AI Prediction Uses Data in			N	$R^2_{\text{Analyst Forecast}}$	$R^2_{\text{AI out-of-sample Prediction}}$	Difference
Financial Statement	Interest Rate	Stock Price				
Yes			251,347	0.23	0.07	0.16
Yes	Yes		251,347	0.23	–0.17	0.40
Yes		Yes	251,347	0.23	0.55	–0.32
Yes	Yes	Yes	251,347	0.23	0.64	–0.41

Table 4. Breakdown of prediction accuracy comparison

Panel A. Breakdown by stock price quartiles

This table reports the accuracy of analyst forecasts and the associated paired t-test in column (1). The accuracy is defined as the average value of the absolute differences between the realized stock price and analyst's target price scaled by the realized price:

$$Accuracy_{\text{Analyst}} = \left| \frac{P_{\text{Realized}} - P_{\text{Analyst}}}{P_{\text{Realized}}} \right| \quad (1)$$

AI prediction uses stock prices, interest rates and financial statements of the last four quarters. Column (2) reports the accuracy of AI predictions and the associated paired t-test. The accuracy is defined as the average value of the absolute differences between the realized stock price and AI-predicted price scaled by the realized price:

$$Accuracy_{\text{AI}} = \left| \frac{P_{\text{Realized}} - P_{\text{AI}}}{P_{\text{Realized}}} \right| \quad (2)$$

The difference between the above two values in (1) and (2) and the associated paired t-test result are reported in the third column in the table. Standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively.

Sample	N	Accuracy of Analyst Forecast (1)	Accuracy of AI Prediction (2)	Difference (1) – (2)
1st quartile of stock price on the forecast day: (0, $18.36]	62,844	7.703*** (0.094)	6.264*** (0.048)	+1.439*** (0.083)
2nd quartile of stock price on the forecast day: ($18.36, $39.94]	62,857	1.030*** (0.020)	0.982*** (0.007)	+0.048*** (0.019)
3rd quartile of stock price on the forecast day: ($39.94, $74.41]	62,818	0.441*** (0.004)	0.499*** (0.003)	–0.058*** (0.003)
4th quartile of stock price on the forecast day: ($74.41, $466.22]	62,828	0.408*** (0.004)	0.416*** (0.003)	–0.008** (0.003)

Panel B. Breakdown by year

This table reports the accuracy of analyst forecasts and the associated paired t-test in column (1). AI prediction uses stock prices, interest rates and financial statements of the last four quarters. Column (2) reports the accuracy of AI predictions and the associated paired t-test.The difference between the above two values in (1) and (2) and the associated paired t-test result are reported in the third column in the table. Standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively.

Sample	N	Accuracy of Analyst Forecast (1)	Accuracy of AI Prediction (2)	Difference (1) – (2)
2010	6,622	2.053*** (0.117)	0.605*** (0.024)	+1.448*** (0.108)
2011	19,930	1.688*** (0.062)	1.652*** (0.029)	+0,.036 (0.054)
2012	18,863	1.851*** (0.073)	1.098*** (0.027)	+0.753*** (0.065)
2013	20,331	1.904*** (0.068)	0.885*** (0.019)	+1.019*** (0.063)
2014	20,524	3.829*** (0.108)	1.917*** (0.048)	+1.912*** (0.098)
2015	20,360	2.908*** (0.098)	1.740*** (0.033)	+1.168*** (0.087)
2016	21,054	2.208*** (0.082)	1.827*** (0.037)	+0.381*** (0.074)
2017	20,427	2.031*** (0.076)	2.229*** (0.045)	–0.198*** (0.065)
2018	20,215	3.478*** (0.113)	4.420*** (0.080)	–0.942*** (0.099)
2019	24,114	1.631*** (0.064)	1.514*** (0.035)	+0.117** (0.058)
2020	24,439	1.541*** (0.068)	1.367*** (0.027)	+0.174*** (0.055)
2021	20,456	3.659*** (0.121)	2.788*** (0.058)	+0.871*** (0.087)
2022	14,012	2.389*** (0.111)	4.441*** (0.094)	–2.052*** (0.097)

Table 5. Variable definition and summary statistics

Panel A. Definition of variables used in the regression analysis

Variable	Definition and data source	Source
Market Capitalization (million $)	Average stock price × Average shares outstanding in a quarter	CRSP Compustat
Market to book	Market Capitalization / Stockholders Equity	
Total Assets (million $)	Total assets at the end of each quarter	Compustat 10Q, 10K
Debt to asset	(Short-term debt + Long-term debt) / Total assets	
Return on assets	Net income (loss) / Total assets	
Cash to assets	Cash and short-term investments / Total assets	
R&D expenses to assets	Research and development expense / Total assets	
Capital Expenditures to assets	(Capital expenditures – Sale of Property) /Total assets	
Stock price on prediction day ($)	Average stock price in the quarter when an analyst sets the target price	CRSP IBES Detail
Stock return daily volatility (90-day prior)	Stock return volatility estimated during the 90-day period prior to the date when an analyst sets the 12-month target price	
Realized stock price ($)	Average stock price in the quarter when an analyst sets the 12-month target price	
Analyst's target price ($)	Average stock price in the quarter of the analyst's target date (12 months after the forecast date)	
Institutional investor ownership	Total shares owned by institutional investors/Total shares outstanding	Thomson Reuters S34
AI predicted stock price ($)	AI predicted stock price on the analyst's target date (12 months after the forecast date)	Deep Learning
Deviation between analyst and AI prediction (using financial statements, stock price and interest rate) (using financial statements and stock price) (using financial statements and interest rate) (using financial statements only)	$Deviation_{\text{Analyst,AI}} = \frac{\left\lvert P_{\text{Analyst}} - P_{\text{AI}} \right\rvert}{\frac{P_{\text{Analyst}} - P_{\text{AI}}}{2}}$	

Panel B. Summary statistics of variables used in the regression analysis

Variable	N	Mean	Std. Dev.	Med	Min	Max
Total Assets (million $)	251,347	24445.8	77765.9	4611.9	28.46	813891
Market Capitalization (million $)	251,347	18661.2	43474.7	4468.3	35.07	360109.7
Market to book	251,347	4.939	9.853	2.508	0.258	99.54
Debt to asset	251,347	0.273	0.186	0.267	0	0.796
Profit margin	251,347	-1.112	9.734	0.0585	-119.3	1.150
Return on assets	251,347	0.00224	0.0460	0.00932	-0.269	0.126
Cash to assets	251,347	0.173	0.206	0.0931	0.0000382	0.943
R&D expenses to assets	251,347	0.0101	0.0226	0	0	0.152
Capital Expenditures to assets	251,347	0.0371	0.0531	0.0185	0	0.351
Stock price on prediction day ($)	251,347	59.61	67.09	39.99	0.969	465.8
Stock return daily volatility (90-day prior)	251,347	0.0277	0.0173	0.0230	0.00718	0.118
Institutional ownership	251,347	0.730	0.275	0.823	0	0.996
Realized stock price ($)	251,347	63.73	73.98	41.76	0.555	500.9
Analyst's target price ($)	251,347	72.27	82.12	48	2.500	570
AI predicted stock price ($)	251,347	71.66	63.07	54.20	4.104	433.5
Deviation between analyst and AI prediction (using financial statements, stock price and interest rate)	251,347	0.956	0.583	0.927	0.00906	1.989
Deviation between analyst and AI prediction (using financial statements and stock price)	251,347	0.960	0.578	0.936	0.00933	1.989
Deviation between analyst and AI prediction (using financial statements and interest rate)	251,347	1.075	0.594	1.109	0.0110	1.994
Deviation between analyst and AI prediction (using financial statements only)	251,347	1.084	0.592	1.129	0.0119	1.993

Table 6. Correlation matrix

The Pearson's correlation coefficients are shown in the lower triangle, and the Spearman's rank correlations are shown above the diagonal.

	Log(Total assets)	Market to book	Debt to asset	Return on assets	Cash to assets	R&D to assets	CapEx to assets	Log(Stock price)	Return volatility	Institutional ownership
Log(Total assets)		-0.154	0.246	0.146	-0.344	-0.305	-0.002	0.423	-0.405	-0.045
Market to book	-0.087		0.014	0.235	0.3495	0.360	-0.028	0.419	-0.083	0.091
Debt to asset	0.219	0.173		-0.122	-0.382	-0.253	0.071	-0.034	-0.002	0.037
Return on assets	0.273	-0.020	-0.030		-0.013	-0.121	0.071	-0.034	-0.002	0.059
Cash to assets	-0.430	0.189	-0.337	-0.315		0.546	-0.244	-0.008	0.197	-0.030
R&D to assets	-0.389	0.191	-0.232	-0.469	0.676		-0.176	0.002	0.188	-0.047
CapEx to assets	-0.039	-0.051	0.057	-0.004	-0.204	-0.148		-0.027	0.061	0.057
Log(Stock price)	0.465	0.184	-0.042	0.382	-0.085	-0.144	-0.094		-0.439	0.220
Return volatility	-0.340	0.022	0.049	-0.379	0.246	0.257	0.096	-0.463		-0.023
Institutional ownership	0.077	0.040	0.011	0.133	-0.036	-0.066	-0.006	0.330	-0.115	

Table 7. Determinants of price prediction deviation between analyst and AI prediction

The dependent variable is the deviation between analyst and AI prediction (using financial statements, stock price and interest rate):

$$Deviation_{\text{Analyst,AI}} = \frac{\left|P_{\text{Analyst}} - P_{\text{AI}}\right|}{\frac{P_{\text{Analyst}} - P_{\text{AI}}}{2}}$$

The independent variables include the natural logarithm of total assets, market capitalization to book value of equity, debt to total assets, return on assets (ROA), cash to total assets, R&D expenses to total assets, capital expenditures to total assets, the natural logarithm of stock price on the prediction day, and the daily stock return volatility 90-day prior to the prediction day. All specifications use OLS regressions with time (quarterly) fixed-effects. Robust standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively.

	(1)	(2)	(3)	(4)	(5)
Log(Total assets)	-0.0324***	0.00111	-0.0309***	-0.00154	-0.00194
	(0.00510)	(0.00407)	(0.00499)	(0.00409)	(0.00392)
Market to book	-0.00230***	0.000684***	-0.00224***	0.000773***	0.000748***
	(0.000377)	(0.000244)	(0.000376)	(0.000236)	(0.000239)
Debt to asset	0.0944***	-0.0174	0.0896***	-0.00928	-0.00787
	(0.0212)	(0.0163)	(0.0199)	(0.0163)	(0.0161)
Return on assets	-0.754***	0.130	-0.699***	-0.00227	-0.00129
	(0.133)	(0.147)	(0.145)	(0.142)	(0.142)
Cash to assets	-0.189***	-0.133***	-0.194***	-0.102***	-0.103***
	(0.0299)	(0.0261)	(0.0302)	(0.0247)	(0.0247)
R&D to assets	-0.0912	0.173	-0.0919	0.206	0.198
	(0.198)	(0.176)	(0.197)	(0.178)	(0.176)
CapEx to assets	-0.310***	-0.349***	-0.337***	-0.240**	-0.239**
	(0.100)	(0.106)	(0.105)	(0.101)	(0.102)
Log(Stock price)		-0.119***		-0.133***	-0.131***
		(0.0114)		(0.0123)	(0.0128)
Return volatility			0.612	-2.646***	-2.639***
			(0.418)	(0.380)	(0.378)
Institutional ownership					-0.0187
					(0.0149)
Constant	1.260***	1.407***	1.233***	1.539***	1.549***
	(0.0453)	(0.0548)	(0.0461)	(0.0645)	(0.0611)
Time FE	Yes	Yes	Yes	Yes	Yes
N	251,347	251,347	251,347	251,347	251,347
R-squared (within)	0.0163	0.0487	0.0164	0.0517	0.0517
R-squared (between)	0.00451	0.0763	0.000894	0.119	0.119
R-squared (overall)	0.0161	0.0486	0.0149	0.0558	0.0559
F-statistic	24.99	31.00	22.63	26.61	32.33

Table 8. Determinants of price prediction deviation with time and industry fixed effects

The dependent variable is the deviation between analyst and AI prediction (using financial statements, stock price and interest rate):

$$Deviation_{\text{Analyst,AI}} = \frac{\left|P_{\text{Analyst}} - P_{\text{AI}}\right|}{\frac{P_{\text{Analyst}} - P_{\text{AI}}}{2}}$$

The independent variables include the natural logarithm of total assets, market capitalization to book value of equity, debt to total assets, return on assets (ROA), cash to total assets, R&D expenses to total assets, capital expenditures to total assets, the natural logarithm of stock price on the prediction day, and the daily stock return volatility 90-day prior to the prediction day. All specifications use OLS regressions with time (quarterly) and industry (2-digit SIC code) fixed-effects. Robust standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively.

	(1)	(2)	(3)	(4)	(5)
Log(Total assets)	-0.0374***	0.00126	-0.0356***	-0.000901	-0.00122
	(0.00324)	(0.00383)	(0.00329)	(0.00390)	(0.00402)
Market to book	-0.00204***	0.000797**	-0.00198***	0.000864***	0.000845***
	(0.000320)	(0.000318)	(0.000297)	(0.000313)	(0.000317)
Debt to asset	0.107***	-0.0262	0.100***	-0.0172	-0.0157
	(0.0264)	(0.0245)	(0.0259)	(0.0231)	(0.0238)
Return on assets	-0.706***	0.0792	-0.652***	-0.0265	-0.0252
	(0.0640)	(0.0846)	(0.0760)	(0.0748)	(0.0760)
Cash to assets	-0.191***	-0.132***	-0.197***	-0.108***	-0.108***
	(0.0282)	(0.0168)	(0.0264)	(0.0163)	(0.0162)
R&D to assets	0.0534	0.0530	0.0402	0.0962	0.0948
	(0.284)	(0.201)	(0.280)	(0.198)	(0.199)
CapEx to assets	-0.380***	-0.210***	-0.391***	-0.157***	-0.159***
	(0.0471)	(0.0426)	(0.0467)	(0.0376)	(0.0378)
Log(Stock price)		-0.123***		-0.134***	-0.133***
		(0.00585)		(0.00603)	(0.00663)
Return volatility			0.712	-2.330***	-2.331***
			(0.564)	(0.327)	(0.328)
Institutional ownership					-0.0142
					(0.0154)
Constant	1.264***	1.366***	1.237***	1.465***	1.473***
	(0.0369)	(0.0356)	(0.0345)	(0.0349)	(0.0340)
Time FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes	Yes
N	251,347	251,347	251,347	251,347	251,347
R-squared (within)	0.0576	0.0873	0.0578	0.0894	0.0894
R-squared (between)	0.662	0.580	0.657	0.628	0.634
R-squared (overall)	0.0574	0.0887	0.0576	0.0916	0.0916
F-statistic	1190.6	937.7	1184.6	1224.4	1300.5

Table 9. Determinants of price prediction deviation with time and firm fixed effects

The dependent variable is the deviation between analyst and AI prediction (using financial statements, stock price and interest rate):

$$Deviation_{\text{Analyst,AI}} = \frac{\left|P_{\text{Analyst}} - P_{\text{AI}}\right|}{\frac{P_{\text{Analyst}} - P_{\text{AI}}}{2}}$$

The independent variables include the natural logarithm of total assets, market capitalization to book value of equity, debt to total assets, return on assets (ROA), cash to total assets, R&D expenses to total assets, capital expenditures to total assets, the natural logarithm of stock price on the prediction day, and the daily stock return volatility 90-day prior to the prediction day. All specifications use OLS regressions with time (quarterly) and firm (gvkey) fixed-effects. Robust standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively.

	(1)	(2)	(3)	(4)	(5)
Log(Total assets)	-0.0368***	0.00506	-0.0361***	0.00684	0.00532
	(0.0128)	(0.0126)	(0.0128)	(0.0125)	(0.0126)
Market to book	-0.000978**	0.000397	-0.000954**	0.000452	0.000454
	(0.000426)	(0.000400)	(0.000424)	(0.000400)	(0.000400)
Debt to asset	0.0891**	-0.0154	0.0829**	-0.00924	-0.00764
	(0.0389)	(0.0368)	(0.0385)	(0.0367)	(0.0367)
Return on assets	-0.427***	-0.141	-0.408***	-0.162*	-0.161*
	(0.0910)	(0.0891)	(0.0907)	(0.0890)	(0.0889)
Cash to assets	-0.119***	-0.0680*	-0.118***	-0.0651*	-0.0670*
	(0.0371)	(0.0369)	(0.0371)	(0.0368)	(0.0369)
R&D to assets	-0.534*	-0.553*	-0.545*	-0.529*	-0.527*
	(0.318)	(0.309)	(0.318)	(0.309)	(0.310)
CapEx to assets	-0.104	0.0413	-0.103	0.0504	0.0515
	(0.0970)	(0.0961)	(0.0970)	(0.0957)	(0.0955)
Log(Stock price)		-0.0964***		-0.104***	-0.105***
		(0.00897)		(0.00896)	(0.00890)
Return volatility			0.437	-1.027***	-1.032***
			(0.300)	(0.275)	(0.276)
Institutional ownership					0.0351
					(0.0236)
Constant	1.246***	1.244***	1.232***	1.277***	1.267***
	(0.101)	(0.0991)	(0.101)	(0.0987)	(0.0980)
Time FE	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
N	251,347	251,347	251,347	251,347	251,347
R-squared (within)	0.0434	0.0492	0.0435	0.0495	0.0496
R-squared (between)	0.167	0.281	0.168	0.289	0.288
R-squared (overall)	0.0558	0.0860	0.0558	0.0877	0.0872
F-statistic	36.91	38.35	36.31	38.01	37.64

Table 10. Determinants of price prediction deviation with time and analyst fixed effects

The dependent variable is the deviation between analyst and AI prediction (using financial statements, stock price and interest rate):

$$Deviation_{\text{Analyst,AI}} = \frac{\left|P_{\text{Analyst}} - P_{\text{AI}}\right|}{\frac{P_{\text{Analyst}} - P_{\text{AI}}}{2}}$$

The independent variables include the natural logarithm of total assets, market capitalization to book value of equity, debt to total assets, return on assets (ROA), cash to total assets, R&D expenses to total assets, capital expenditures to total assets, the natural logarithm of stock price on the prediction day, and the daily stock return volatility 90-day prior to the prediction day. All specifications use OLS regressions with time (quarterly) and analyst (amaskcd) fixed-effects. Robust standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively.

	(1)	(2)	(3)	(4)	(5)
Log(Total assets)	-0.0319***	0.00475***	-0.0297***	0.00239*	0.00213
	(0.00143)	(0.00136)	(0.00143)	(0.00136)	(0.00136)
Market to book	-0.00198***	0.000628***	-0.00191***	0.000676***	0.000662***
	(0.000157)	(0.000148)	(0.000156)	(0.000147)	(0.000147)
Debt to asset	0.0984***	-0.0418***	0.0898***	-0.0323***	-0.0314***
	(0.0104)	(0.00970)	(0.0104)	(0.00969)	(0.00971)
Return on assets	-0.691***	0.0438	-0.631***	-0.0405	-0.0398
	(0.0390)	(0.0386)	(0.0400)	(0.0386)	(0.0386)
Cash to assets	-0.157***	-0.107***	-0.163***	-0.0914***	-0.0916***
	(0.0119)	(0.0111)	(0.0120)	(0.0112)	(0.0112)
R&D to assets	0.187*	0.192*	0.183*	0.202**	0.199**
	(0.109)	(0.102)	(0.109)	(0.101)	(0.101)
CapEx to assets	-0.293***	-0.133***	-0.300***	-0.105***	-0.107***
	(0.0385)	(0.0379)	(0.0386)	(0.0377)	(0.0378)
Log(Stock price)		-0.124***		-0.134***	-0.133***
		(0.00214)		(0.00225)	(0.00234)
Return volatility			0.840***	-1.938***	-1.940***
			(0.135)	(0.126)	(0.126)
Institutional ownership					-0.0112*
					(0.00590)
Constant	1.217***	1.346***	1.184***	1.432***	1.439***
	(0.0149)	(0.0148)	(0.0152)	(0.0157)	(0.0161)
Time FE	Yes	Yes	Yes	Yes	Yes
Analyst FE	Yes	Yes	Yes	Yes	Yes
N	251347	251347	251347	251347	251347
R-squared (within)	0.0494	0.0768	0.0497	0.0782	0.0782
R-squared (between)	0.100	0.151	0.0979	0.162	0.162
R-squared (overall)	0.0574	0.0883	0.0576	0.0912	0.0912
F-statistic	120.9	168.4	121.0	167.3	164.9

Table 11. Robustness checks using "deeper" learning model and "shallow" learning model

This table reports the accuracy of analyst forecasts and the associated paired t-test in column (1). Column (2) reports the accuracy of AI predictions and the associated paired t-test. The difference between the above two values in (1) and (2) and the associated paired t-test result are reported in the third column in the table. Standard errors are shown in the parentheses with ***, ** and * indicating its statistical significant level of 1%, 5% and 10% respectively. Two measures of goodness of fit (MSE and predicted R^2 are reported underneath the accuracy statistics in the table.

AI Prediction Method	N	Accuracy of Analyst Forecast	Accuracy of AI Prediction	Difference
		(1)	(2)	(1) – (2)
"Deeper" learning model	251,347	2.396***	1.823***	+0.573***
		(0.025)	(0.011)	(0.021)
5-layer neural network	MSE	4223.9	2080.7	
10,000 epochs	Predicted R^2	0.23	0.62	
"Shallow" learning model	251,347		2.271***	+0.125***
			(0.015)	(0.022)
5-layer neural network	MSE		2673.3	
2,500 epochs	Predicted R^2		0.51	
"Shallow" learning model	251,347		2.112***	+0.284***
			(0.014)	(0.023)
3-layer neural network	MSE		5366.8	
5,000 epochs	Predicted R^2		0.02	