Is it "Priced In"? Classifying Efficient Markets

Team Members: Neophytos Charalambides, Anuvetha Govindarajan, Tanuj Mathur,
Ram Purandhar Reddy Sudha, Pinky Thomas.

1. Problem Statement: In this project, we attempt to answer the question *"Is the market aware of the situation which has taken place regarding a certain event or trend relating to a company, and has it reacted to it?*". This is commonly referred to as "Priced In", i.e. "are all available information regarding a particular event, news, or potential outcome, already reflected in the current price of a financial asset, stock market index, or ticker symbol?''. In this summary, event refers to the situation where earnings were reported, ticker to a publicly traded company, and index to a measure of stock market performances, e.g. S&P 500 and NASDAQ.

The above question is crucial for investors, traders, and financial analysts, as it helps them understand how much of the market's expectations about future events are already reflected in current prices. Some instances where knowing the answer to the above question is crucial are risk management, and coming up with investment strategies and decisions. A central issue with answering the above question, is that there is no one universal definition for being "Priced In". Additionally, the complexity of financial markets and the behavior of market participants makes it even more challenging. Nonetheless, we were able to come up with an efficient method for answering the question at stake.

2. Methodology: Throughout the project, we consider the pipeline depicted below.

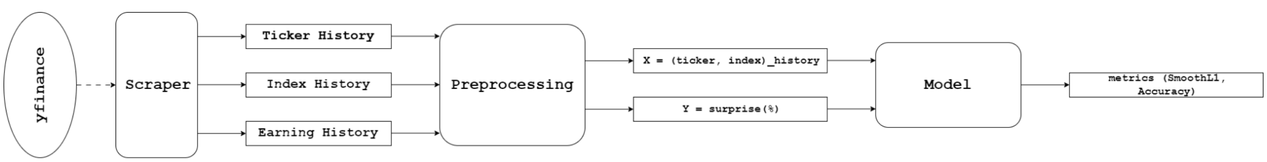

Fig.1: Data Processing and Modeling Pipeline

Data: We gathered the ticker and index history which contain pricings at the opening and closing times, as well as the highest and lowest pricings recorded for each trading day for the past five years from Yahoo Finance and considered NASDAQ top 100 tickers and data from a total of 5 market indexes. The data also includes the volume of the tickers and indexes, which is the number of shares or contracts traded in a security or market over that day. We also obtain the earning history, which is what we use for determining events. For each ticker at each time stamp, the earning history contains the EPS (Earnings Per Share) Estimate, the Reported EPS, and the Surprise percentage. Before proceeding to our data analysis, we cleaned the data in order to make it easy to manipulate in the different models we attempted. Any duplicate events were removed, and only data which had valid entries was considered (i.e removed all NaN values).

Model: The main model considered is a Recurrent Neural Network (RNN), to analyze whether or not stocks from certain tech companies can be predicted to have been considered "priced in'' or not. RNN is a type of neural network where the output from one time step is used as input for the next time step. This structure allows RNNs to maintain information across a sequence of data points, capturing temporal correlations. The input features for our model include the ticker history and the index history for 60 days prior to the event as well as 60 days after the event. Specifically, they include the `closing price`, `volume` and `volatility` of the ticker, as well as the `closing price` of the market. We have normalized each of these variables with respect to their values on the day of the event. Surprise (%) is used as the output label. The Mean Absolute Error loss function (L1 loss function) and SmoothL1 loss function are used to measure the average absolute difference between the predicted and actual values.

3. Results: The Model results in an **80% accuracy** on the testing dataset for classifying the events based on the Surprise (positive vs negative).

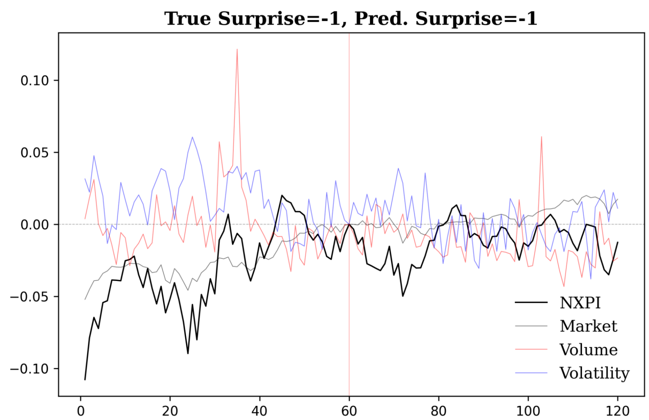

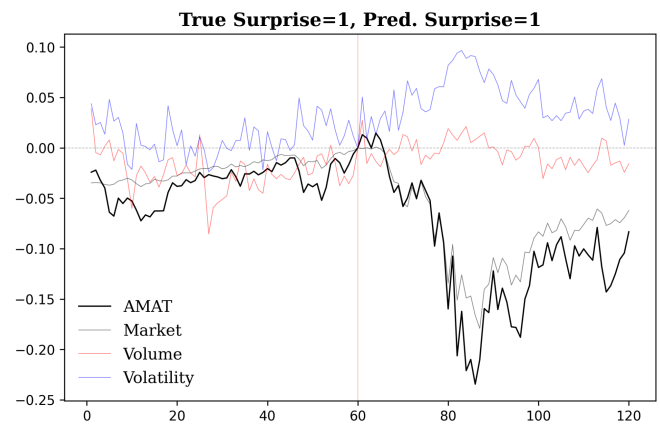

Fig.2 Fig.3

RNN Prediction of Market Reactions to Earnings Announcements

The graphs above represent the prediction results of the RNN model for stock price analysis. True surprise indicates the sign of actual surprise value for the event. The “Pred. Surprise” is the sign of the RNN-predicted surprise value.

For Figure 2, a True surprise value of -1 and predicted surprise value of -1 imply that the RNN model has correctly predicted that the stock has underperformed relative to the market expectations. Similarly, in Figure 3 we see that the model correctly predicted that the stock would outperform market expectations.

The model effectively captures the temporal correlations in the stock prices, market indexes, volume, and volatility data. The alignment between the predicted and actual surprises indicates the RNN model's potential in forecasting stock price reactions to earnings announcements.

4. Conclusions: Our results demonstrate that for events such as Earnings Reports, the market efficiently captures the effects through either anticipation before the event or re-adjustments after the event. The RNN model's 80% accuracy in predicting the overall effect of events (positive or negative) based on market & stock history highlights its robustness and practicality in applications. Understanding and leveraging the model’s predictions can aid investors in making informed decisions by anticipating market reactions based on historical data.

Future directions for this research could include extending the analysis to other types of market events to evaluate the model’s versatility and general applicability. Additionally, incorporating larger datasets could improve accuracy, and exploring sentiment analysis may further enhance the model's predictive capabilities.