Assessing the Effects of Expanded Input Elicitation and Machine Learning-Based Priming on Crowd Stock Prediction

Harika Bhogaraju[1], Arushi Jain[2], Jyotika Jaiswal[2], and Adolfo R. Escobedo[3](✉)

[1] Arizona State University, Tempe, AZ 85281, USA
[2] Vellore Institute of Technology, Chennai, India
[3] North Carolina State University, Raleigh, NC 27606, USA
arescobedo@ncsu.edu

Abstract. The stock market is affected by a seemingly infinite number of factors, making it highly uncertain yet impactful. A large determinant of stock performance is public sentiment, which can often be volatile. To integrate human inputs in a more structured and effective manner, this study explores a combination of the wisdom of crowds concept and machine learning (ML) for stock price prediction. A crowdsourcing study is developed to test three ways to elicit stock predictions from the crowd. The study also assesses the impact of priming participants with estimates provided by an Long Short Term Model (LSTM) model herein developed for this context.

Keywords: Wisdom of Crowds · Machine Learning · Priming · Stock Prediction

1 Introduction

Human beings have always had an inherent interest in predicting the unpredictable. The journey may have started with fortune tellers, oracles, and prophets, using sun signs, planetary positions, and star alignments to predict the future. But today, artificial intelligence takes their place using statistics, big data analytics, and at the heart of it all, machine learning (ML). ML has permeated every modern industry to provide predictive insights. A non-comprehensive list is as follows. In healthcare, ML models have been used to predict COVID-19 surges, death numbers, and recovery rates during the pandemic [17]. In manufacturing, they have been applied to improve the efficiency of smart factories [19]. In surveillance, Karpathy et al. [12] evaluate the performance of Convolutional Neural Networks(CNNs) on broadly 487 categories of videos, totaling 1 million YouTube videos to further surveillance automation. In the mental health space, Fathi et al. [7] train an ANFIS model on a large dataset with seven input features to detect social anxiety disorder in healthcare.

N. T. Nguyen et al. (Eds.): ICCCI 2023, CCIS 1864, pp. 3–16, 2023.
https://doi.org/10.1007/978-3-031-41774-0_1

There are seemingly more applications of ML emerging for forecasting highly uncertain yet impactful events. One representative use case is the stock market, whose irregular nature puzzles day traders, investors, and scientists alike. While stock trading is often thought of as a 'zero-sum game', there are tangible factors that play a role in measuring the performance of a stock. For instance, historical prices and other financial data about a company are often good indicators of the potential of a stock to be profitable for an investor. With the advances in ML, processing large quantities of financial data has become a routine task. Using neural networks, millions of data points can be processed to produce a prediction of the stock price of a company. For example, Shen and Shafiq [23] compare different ML models to determine which model has more accurate predictions; a comprehensive Long Short Term Memory (LSTM) network achieves the highest accuracy.

LSTMs have certainly brought us a step closer to unraveling the stock market, but what makes stock trading so mercurial is public sentiment due to its very subjective and inconsistent nature. Moreover, the stock market is a victim of second-order effects; trading activity leads to chain reactions of consequent trading activity. This makes public sentiment difficult to track, and yet it remains an essential determinant in stock performance. Many have tried to find concrete ways to measure public sentiment. For example, Das et al. [4] use tweets, social media comments, news headlines, and other news sources to predict stock prices during the COVID-19 pandemic. The paper concludes that using public sentiment from social media and news in combination with stock data improves stock prediction accuracy. However, social media users are often easily influenced by others' opinions resulting in biased and polarized public sentiment that does not accurately represent personal views [20]. Similarly, Checkley et al [3] find that bloggers can add uncertainty to the market rather than bring clarity. The study concludes that using sentiment collected from blogs improves forecasts of volatility and trading volume but not of price direction.

The *wisdom of crowds* can be a promising approach to obtaining a more stable estimate of public sentiment since it gathers people's inputs in a more structured and controlled setting. This concept builds on the notion that "the sum is greater than the sum of the parts", to derive aggregate judgments that are often more accurate and reliable than those of any individual, including subject matter experts [25]. To achieve these effects, each crowd member's opinion must be elicited independently from others. In addition, this private information must be obtained from a diverse set of individuals. The wisdom of crowds has many commonplace applications including in the healthcare industry, where it is used for diagnosis, nutrition, surveillance, and public health [27]. It is also utilized in more narrow applications such as recommendation systems [22]. In a context closely related to this work, prediction markets rely on the wisdom of crowds to derive accurate estimates by engaging large participant pools in simple prediction tasks like binary classification (up-down movement of stocks) and simple price estimations. However, more complex forms of input elicitation that may require higher cognitive power from the participant pool but potentially enhance the wisdom of crowd effects have received little attention.

To explore this potential, this paper develops a wisdom of crowds approach that combines different input elicitation techniques and machine learning-based priming to augment stock performance prediction. The approach is deployed via a custom user interface on a crowdsourcing platform (Amazon MTurk) in order to engage a diverse pool of participants who can draw on their own independent knowledge to provide the requested estimates.

2 Literature Review

2.1 Machine Learning in Finance

It is important to understand specific financial indicators that have been integral to analyzing stocks. An elementary analysis of a stock can be performed by analyzing five key variables: Opening price, high of the day, low of the day, volume, and adjusted closing price [21]. The indicators are primarily used in making short-term predictions. On the other hand, long-term predictions tend to use 52-week high, price-earnings ratio (P/E ratio), and moving averages (MAs). Performance metrics such as these use varying ratios of revenue, profits, debt, and assets to quantify the health of a company.

While these indicators are good at summarizing the current health of a stock, ML can be a great source for looking into the future. Many different neural networks including LSTMs, MDWDNNs (Multi-Dimensional Wide Neural Networks), and MDRNNs (Multi-Dimensional Recurrent Neural Networks) have been developed to predict time-stamped data such as temperature fluctuations and annual retail sales. For example, Khodabakhsh et al. [14] use crude oil and byproduct flow rates to predict a petrochemical plant's crude oil purchase amount. Although various neural network models can be employed to predict stock prices [1,18], LSTMs are simple, easily accessible, and more popular than their more complex counterparts. Therefore, we adopt LSTMs in the present study to shift attention to the crowd rather than on model performance, since the latter has been the primary focus of the majority of other works in this space. In the featured context, LSTMs have been used to read historical stock prices and predict future stock price values [16]. Nevertheless, LSTMs have evolved to utilize more than historical data alone [10,15].

2.2 Public Sentiment and Stock Performance

In addition to financial factors, public sentiment can be used as a feature to support ML model predictions. For example, Siah and Myers [24] measured public sentiment through Google Trends and Bloomberg Businessweek. "Positive" and "Negative" words were first categorized and then data was mined to obtain a public sentiment score. The study found that Google Trends were unhelpful in improving model accuracy. Such findings highlight the volatility of social media and news as sources of public sentiment and the need to gather data in a more structured manner.

2.3 Wisdom of Crowds

The wisdom of crowds can be applied to obtain a more stable form of public sentiment. This concept can be leveraged by eliciting and aggregating the judgments or predictions from multiple participants to reduce errors and thereby better approximate the ground truth. Traditional wisdom of crowds techniques can be applied to gauge the potential value of a stock by eliciting price estimates (numerical input) or up/down estimates (binary choice) [6]. However, the former can lead to marked under/overestimation due to cognitive biases [9], while the latter may fail to capture nuanced information. In fact, a recommended practice in prediction markets is to normalize the market price and to interpret the elicited crowd inputs as the probability that the market believes in the stock [2]. The latter observation motivates experimenting with a wider variety of input elicitation techniques for this context. This is further supported by recent studies showing that richer crowd information can be extracted on human computation tasks by eliciting and combining multiple modalities of estimates [13,28].

2.4 ML and Crowdsourcing

ML and crowdsourcing have complementary strengths, meaning that they can work better together than alone. This has been shown across various applications. For example, Yasmin et al. [29] use crowdsourcing inputs as features within several standard ML algorithms for image classification; they show that these hybrid algorithms perform better than a state-of-the-art fully automated classification algorithm, when the training sets are small to medium in size. In addition, Demartini et al. [5] combine inputs from machine learning algorithms, crowdsourcing workers, and experts to combat the spread of online misinformation. Similar approaches can be extrapolated to many different use cases including the stock market. Indeed, Hill and Ready-Campbell [8] use over 2 million users' stock picks on the investment platform CAPS to train an ML model that ranks stocks and builds stock portfolios. These works provide a glimpse of the potential advantages of combining ML and the wisdom of crowds. Following in these footsteps, this paper employs a combination of an ML model and crowdsourcing inputs to predict the performance of stocks.

3 Experiment Design

Three experiments were deployed on Amazon MTurk, a crowdsourcing platform, on a total of 308 participants. Each experiment was deployed on the first weekend of the months of September 2022 (Experiment A), October 2022 (Experiment B), and November 2022 (Experiment C). Weekends were chosen so as to avoid frequent and momentary fluctuations in the current stock price. Experiments A and B use a common set of 22 stocks, and Experiment C uses 10 stocks selected from the larger set of the two other experiments.

The user interface asked participants to provide estimates on several stocks through a combination of different input elicitation techniques and representations of ML predictions (i.e., priming), the latter of which were obtained from the outputs of an LSTM model. Depending on the complexity of the question, a time limit of either 90 s or 120 s was imposed. It is important to remark that Experiments A and B were conducted at a time when the Dow Jones Index (DJI) was trending downward, while Experiment C was conducted at a time when the DJI was trending upward, as is shown in Fig. 1.

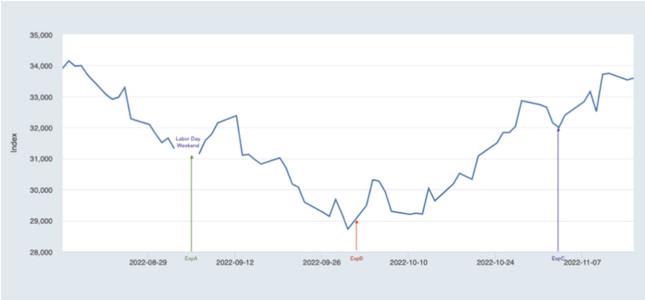

Fig. 1. The Dow Jones Index from Aug 15, 2022 to Nov 15, 2022.

3.1 The ML Model

An LSTM model is trained on the historical stock prices of a select set of stocks to obtain short-term (1 month) and long-term predictions (3 months and 1 year). It is optimized using the Adam optimization algorithm [26]. This univariate model uses the closing price of a stock recorded over the past 3 years from Aug 31, 2019 to Aug 31, 2022, to train and the last 60 days to make a prediction for Day 61; the sliding window is altered to incorporate a new prediction for Day 62. For Experiment A, this process is repeated until a prediction for Sep 30, 2022 and a prediction for Nov 30, 2022 are acquired. For Experiment B, a similar process is repeated using data from Sep 30, 2019 to Sep 30, 2022 to retrieve a prediction for Oct 31, 2022 and for December 31, 2022. For Experiment C, a similar process is repeated using data from Oct 31, 2019 to Oct 31, 2022 to retrieve a prediction for Nov 30, 2022 and for January 31, 2023. Due to time constraints, the accuracy of the three-month predictions could not be verified.

3.2 Input Elicitation Techniques

Each study consists of a total of 10 questions that deploy a selection of three types of input elicitation techniques: percentage growth via a slider, price movement via a graph, and investment allocation via a pie chart. The percentage growth questions measure the projected percentage growth of a stock on a scale of -25%

to +25%; −25% corresponds to the leftmost point on the scale (in the darkest red), and +25% corresponds to the rightmost point on the scale (in the darkest green). The slider increments by 5% for each change in color/shade allowing users to choose a percentage of growth within the given range. Each related question asks users to provide predictions over different time windows. Based on the data obtained in Experiments A and B, the granularity of the slider was reduced for Experiment C. For questions that use ML-based priming, a note with the percentage growth predicted by the LSTM model for the stock is provided above the slider. Figure 2 shows the user interface for this type of input elicitation.

Fig. 2. Interface Example of a Percentage Growth Question.

The price movement questions use an interactive graph to display projected stock prices and elicit the participant's estimate. The historical stock price for the past year is provided to the user, who is then asked to move from a fixed data point indicating the stock price for the upcoming month to the user's estimated price. For questions that use ML-based priming, the upcoming month's price is set to the predicted value; for questions that do not, the upcoming month's price is set to the same value as the current month. Figure 3 shows the user interface for this type of input elicitation.

Each investment allocation question asks users to allocate a budget of virtual funds across a basket of stocks. The total amount of virtual money available to the user is $100,000. A user can invest the full amount in any single company or distribute it among multiple companies. An interactive pie chart is used to illustrate the fund proportions allocated by the user. Users move the slider of each stock to specify the amount of virtual money they would like to allocate to it. The users are presented with the current stock price of all the companies in the basket but, for the question involving ML-based priming, a projected

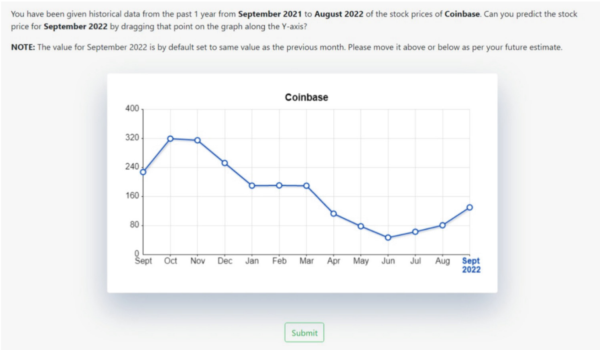

Fig. 3. Interface Example of a Price Movement Question.

percentage growth for each of the stocks in the basket is also shown. Figure 4 shows the user interface for this type of input elicitation.

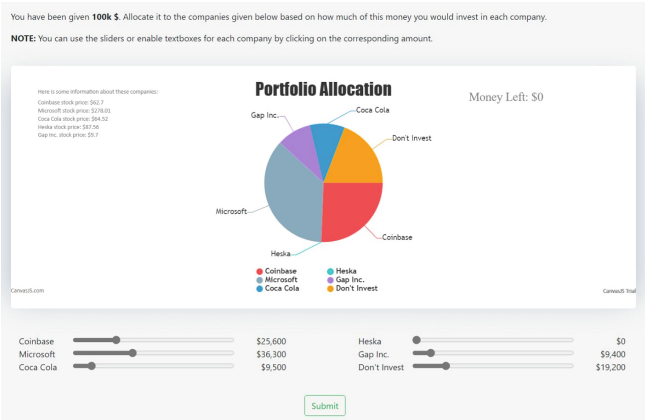

Fig. 4. Interface Example of an Investment Allocation Question.

4 Analysis

This section assesses the performance of the tested input elicitation techniques based on the resulting differences between the actual growth and the median of the crowd's predicted growth for the individual stocks. In addition, it analyzes the effect of priming on participant performance in three ways: (1) by measuring the correlation between the predictions by ML vs predictions by the crowd and between the actual price vs the crowd's predictions; (2) by comparing the differences between the actual growth and the median of the crowd's predicted growth; and (3) by comparing the rankings of stocks that achieved the highest actual growth with those that received the highest investment from the crowd and with those that were deemed most profitable by the LSTM model.

4.1 Input Elicitation Techniques

First, the percentage growth technique is compared with the price movement technique. The results are summarized in Table 1, which compares the difference between the actual growth and the median of the crowd's predicted growth across all three experiments.

Table 1. Comparison of the accuracy achieved via percentage growth questions vs via price movement questions.

	Crowd's Predicted Growth - Actual Growth (%)		
Question Type	Experiment A	Experiment B	Experiment C
Percentage Growth Questions	15.62	16.12	4.91
Price Movement Questions	5.3	6.66	75.55

In Experiments A and B, the responses to the percentage-growth questions had a more considerable difference from the actual percentage growth than the responses to the graph-based price movement questions. In Experiment A, across all graph-based price movement questions, the crowd's predicted percentage growth is 5.30% over the actual percentage growth of the stocks; on the other hand, the percentage slider questions are 15.62% over the actual percentage growth of the stocks. A similar trend is observed in Experiment B, where the crowd's predicted percentage growth is 6.66% over the actual percentage growth of the stocks for graph-based price movement questions, while it is 16.26% over the actual percentage growth of the stocks for percentage growth questions.

One plausible explanation for this difference in performance can be that 5% increments in the slider of the percentage-growth questions are not sufficiently granular, leading to large differences between the crowd's responses. To explore this issue, the increments for Experiment C were altered to reduce the increment gap. The modified slider contains the values -20%, -10%, -5%, -2%, -1%, 0%, 1%, 2%, 5%, 10%, 20%.

After modifying the slider scale, participant performance improved considerably. In Experiment C, the crowd's predicted percentage growth is 4.91% over the actual percentage growth of the stocks. On the other hand, the price movement questions performed much worse in this experiment than in the previous two. Specifically, the crowd's predicted percentage growth is 75.55% over the actual percentage growth of the stocks. However, an interesting observation can be made here. Participants overestimated a few stocks (META, TSLA, AAL, and X) by nearly 3 times their actual price. A reason for this could be that at the time of this study, these stocks were very popular, as measured by high Google Trends scores (>60) causing people to be overly optimistic about their performance. In addition, unlike Experiments A and B, the market was trending upward at the time Experiment C was deployed, which can explain why the participants overestimated the prices on popular stocks. Indeed, when META, TSLA, AAL, and X are removed, the price movement questions outperformed the percentage growth questions, with a difference of only 0.12% between the actual growth and the crowd's predicted growth.

In summary, with a more granular calibration of the slider, the percentage movement questions yield closer estimations than the price movement questions but, after removing outliers, the price movement questions elicit more accurate predictions. These results motivate the need to further research these two techniques before a more definitive statement on their comparative performance can be made. Due to the nature of the input elicitation techniques, only the percentage growth and price movement questions can be directly compared to each other. Next, the investment allocation questions reveal some interesting findings. This information can be found in Table 2, which shows the performance of investment allocation questions across all three studies.

Table 2. Rankings of stocks based on their actual profitability vs based on the proportion of funds allocated to them by the crowd.

Rankings of Stocks					
Experiment A		Experiment B		Experiment C	
Actual	Crowd	Actual	Crowd	Actual	Crowd
COIN	KO	GPS	KO	META	X
MSFT	GPS	KO	COIN	X	META
KO	COIN	COIN	GPS	KO	AAPL
GPS	HSKA	MSFT	MSFT	AAPL	KO
HSKA	MSFT	HSKA	HSKA	VISA	VISA

In Column 1, stocks are ordered from most to least profitable in Experiment A. In Column 2, the stocks are ordered based on the proportion of funds invested by participants in each stock from highest to lowest. This is repeated for all the experiments across the rest of the columns. Based on the data in the table, it appears that the crowd tended to invest most of their funds in the most profitable stocks. In Experiment A, KO and COIN were heavily invested stocks; they are also the among the most profitable companies among the set of stocks given to users. In Experiment B, most participants invested the largest proportion of money in KO, which has the second-highest actual growth. Lastly in Study C, most people invested the most on X and the second-most on META. The actual growth of META and X were the first and second highest, respectively, among the bucket of stocks for this question.

4.2 Priming

The correlations between the crowd's predicted price, ML's predicted price, and the actual price of a stock were measured using Pearson's correlation coefficient. This measure calculates the linear correlation between two vectors, with −1 representing a total negative linear relationship, +1 representing a total positive linear relationship, and 0 representing no linear correlation.

Table 3 contains the correlation coefficients between ML's predicted prices, the crowd's predicted prices, and the actual price. For questions with priming, the correlation between the crowd's predictions and ML along with the correlation between the crowd's predictions and actual price are represented. For questions without priming, only the correlation between the crowd's prediction and the actual price is shown. The experimental results show that the crowd's predictions were influenced by ML. Indeed, the table shows that there is a strong positive linear correlation between ML's predicted prices and the median of the crowd's predicted prices. In addition, for questions with ML-based priming, there was a higher correlation between the median prices predicted by participants and those predicted by ML.

Table 3. Correlation between the predictions by ML vs predictions by the crowd and between the actual price vs the crowd's predictions, with and without priming.

	Pearson's Correlation					
	Experiment A		Experiment B		Experiment C	
Question Type	ML & Crowd	Actual & Crowd	ML & Crowd	Actual & Crowd	ML & Crowd	Actual & Crowd
Priming	0.9244	0.8628	0.9091	0.8918	0.925	0.9485
No priming	-	0.7381	-	0.8216	-	0.6982

Next, the effect of priming on each input elicitation technique is analyzed. However, due to previously discussed differences in the nature of the input elicitation techniques (e.g., mismatches in data types) and the time frame of certain questions, certain direct comparisons could not be made.

First, the effect of priming on the percentage growth questions is observed. The data pertaining to this can be found in Table 4, which contains the difference between the actual growth of a stock and the crowd's predicted growth across all three experiments.

Experiments A and B did not feature questions with priming within the 1-month time period and, thus, could not be used. However, the data from Experiment C suggests that priming did not help participants to make more accurate predictions for the percentage growth questions. In particular, the absolute value of the difference between the actual growth and the median of the crowd's predicted growth is higher for questions where an ML suggestion was provided. However, there is an interesting observation to be made here. Priming appears

Table 4. Assessment of the effects of priming on percentage growth questions.

Question Type	Crowd's Predicted Growth - Actual Growth (%)		
	Experiment A	Experiment B	Experiment C
Priming	-	-	4.91
No priming	−15.62	16.26	−2.67

to have made the crowd's predictions more optimistic, specifically, the difference between the actual growth and the crowd's predicted growth is 4.91% with priming, while it is −2.67% without priming. That said, the difference between the two numbers is not large enough to close the debate on priming.

When it comes to the price movement questions, ML-based priming boosted participant performance (up to 3 times better). This information can be found in Table 5 which contains the difference between actual growth and the median of the crowd's predicted growth across all three experiments.

Table 5. Assessment of the effects of priming on price movement questions.

Question Type	Crowd's Predicted Growth - Actual Growth (%)		
	Experiment A	Experiment B	Experiment C
Priming	17.93	3.75	31.96
No priming	18.41	12.45	114.33

Based on the data in the table, priming seems to help participants make more accurate predictions. One explanation for this could be that participants were able to refine their estimations against the ML prediction. Just as the percentage growth questions give participants a set of values to choose from, the predicted ML price serves as a suggestion for the price movement question. Lastly, when looking at the investment allocation questions, priming seems to have influenced the crowd's investment decision-making processes. This information can be found in Table 6, which contains the data from the investment allocation questions that are primed across all three experiments.

The first three table columns apply to Experiment A: Column 1 orders the stocks from most to least profitable, column 2 from highest to lowest proportion of funds allocated by the crowd, and column 3 from most to least profitable as predicted by the ML model. This is repeated for Experiments B and C in the remaining table columns. In Experiments A and C, most people invest in stocks that are deemed most profitable by the ML model (BAC and PYPL in Exp A and NFLX and TSLA in Exp C). In Experiment B, there is an equal split between people who invested the largest proportion of their money in the least profitable and the most profitable stocks (XOM, PYPL). This could speak to participants' risk-taking appetite. Low-risk investors tend to invest in stocks

Table 6. Rankings of stocks based on actual profitability, proportion of funds allocated by the crowd, and ML's profitability projections.

Rankings of Stocks								
Experiment A			Experiment B			Experiment C		
Actual	Crowd	ML	Actual	Crowd	ML	Actual	Crowd	ML
PFE	GOOG	XOM	XOM	GOOG	GOOG	GPS	AAL	SRPT
PYPL	BAC	BAC	BAC	(PYPL, XOM)	PFE	SRPT	NFLX	NFLX
XOM	PYPL	PYPL	PFE	(BAC, PFE)	XOM	NFLX	TSLA	TSLA
BAC	PFE	GOOG	GOOG		BAC	AAL	GPS	AAL
GOOG	XOM	PFE	PYPL		PYPL	TSLA	SRPT	GPS

that are already profitable with expectations of slow growth, whereas high-risk investors tend to invest in stocks that are currently not profitable with the hopes of rapid growth in a short span of time.

5 Conclusion

The relationship between the wisdom of crowds and ML is complex but certainly worth exploring in the pursuit of building more effective hybrid AI systems. The stock market is one case for hybrid AI systems to improve forecasts. In order to utilize crowdsourcing, it is important to determine which input elicitation techniques tend to yield higher-quality estimates. In the case of this study, while the percentage growth questions perform better with a more granular calibration of the slider, the price movement questions often yield more accurate results on certain types of stocks. Further testing of these techniques is needed on a more consistent selection of stocks to advance this research. Investment allocation (via a pie chart) is also an interesting approach to analyzing the crowd's investment tendencies. Most people invested in companies that had the highest actual growth. A wider variety of stocks with more context on the previous performance of the stocks is a promising direction of further study.

In addition to the right crowdsourcing technique, at the heart of hybrid AI systems is establishing a symbiotic relationship between people and ML models. Based on participant data from this study, there was a strong correlation between the crowd's predictions and those made by the ML model. In fact, there is a stronger correlation between the crowd's predictions and the actual stock price when the crowd is exposed to suggestions from ML than when they are not.

While ML was not helpful in improving participant performance for percentage growth questions, there is a considerable improvement in participant performance for price movement questions. This could have broader implications on other ML crowdsourcing tasks. For questions that can be ambiguous or rather open-ended, ML might be a great source of support and direction for participants. This could be tested further with more types of activities across different contexts beyond the stock market. In addition, participants were aligned with ML's projected growth of stocks for investment allocation questions in all three studies. To reach a more definitive conclusion, additional studies with a larger participant pool need to be conducted. Based on the participant data, 101 out of 308 participants claim to engage in stock trading on a monthly basis, 73 participants traded on a weekly basis, and 55 never traded. Using these different expertise levels, more nuanced aggregation techniques can be employed to better engage the crowd in the prediction process. Another important area to explore is the different prediction windows presented to users. Considering the accessibility and speed at which information disseminates, prediction windows within hours or minutes might be more relevant and could lead to more accurate predictions [11]. Lastly, ML model performance can be tested by experimenting with different types of models and input features.

Acknowledgements. The authors thank all participants in this study, which received institutional IRB approval prior to deployment. The lead PI of the project (the fourth author) and one of the students (the first author) also gratefully acknowledge support from the National Science Foundation under Award Number 1850355.

References

1. Machine learning approaches in stock price prediction: A systematic review. https://iopscience.iop.org/article/10.1088/1742-6596/2161/1/012065
2. Bassamboo, A., Cui, R., Moreno, A.: Wisdom of crowds in operations: Forecasting using prediction markets (2015). Available at SSRN 2679663
3. Checkley, M.S., Higón, D.A., Alles, H.: The hasty wisdom of the mob: How market sentiment predicts stock market behavior. Expert Syst. Appl. **77**, 256-263 (2017). https://www.sciencedirect.com/science/article/abs/pii/S0957417417300398
4. Das, N., Sadhukhan, B., Chatterjee, T., Chakrabarti, S.: Effect of public sentiment on stock market movement prediction during the COVID-19 outbreak. Soc. Netw. Anal. Min. **12**(1), 92 (2022)
5. Demartini, G., Mizzaro, S., Spina, D.: Human-in-the-loop artificial intelligence for fighting online misinformation: Challenges and opportunities. IEEE Data Eng. Bull. **43**(3), 65–74 (2020)
6. Endress, T., et al.: deliberated intuition in stock price forecasting. Econ. Sociol. **11**(3), 11–27 (2018)
7. Fathi, S., Ahmadi, M., Birashk, B., Dehnad, A.: Development and use of a clinical decision support system for the diagnosis of social anxiety disorder. Comput. Methods Programs Biomed. **190**, 105354 (2020)
8. Hill, S., Ready-Campbell, N.: Expert stock picker: the wisdom of (experts in) crowds. Int. J. Electron. Commer. **15**(3), 73–102 (2011)
9. Honda, H., Kagawa, R., Shirasuna, M.: On the round number bias and wisdom of crowds in different response formats for numerical estimation. Sci. Rep. **12**(1), 1–18 (2022)
10. Huang, Y., Capretz, L.F., Ho, D.: Machine learning for stock prediction based on fundamental analysis. In: 2021 IEEE Symposium Series on Computational Intelligence (SSCI), pp. 01–10. IEEE (2021)
11. Martins, C.J.L., et al.: Information diffusion, trading speed and their potential impact on price efficiency-Literature review. Borsa Istanbul Rev. **22**(1), 122-132 (2021). https://www.sciencedirect.com/science/article/pii/S2214845021000193

12. Karpathy, A., Toderici, G., Shetty, S., Leung, T., Sukthankar, R., Fei-Fei, L.: Large-scale video classification with convolutional neural networks. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, pp. 1725–1732 (2014)
13. Kemmer, R., Yoo, Y., Escobedo, A., Maciejewski, R.: Enhancing collective estimates by aggregating cardinal and ordinal inputs. https://ojs.aaai.org/index.php/HCOMP/article/view/7465
14. Khodabakhsh, A., Ari, I., Bakır, M., Alagoz, S.M.: Forecasting multivariate time-series data using LSTM and mini-batches. In: Bohlouli, M., Sadeghi Bigham, B., Narimani, Z., Vasighi, M., Ansari, E. (eds.) CiDaS 2019. LNDECT, vol. 45, pp. 121–129. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-37309-2_10
15. links open overlay panelAdil Moghar a, A., a, b, has never been easy to invest in a set of assets, A.: Stock market prediction using lstm recurrent neural network (2020). https://www.sciencedirect.com/science/article/pii/S1877050920304865
16. Mohanty, S., Vijay, A., Gopakumar, N.: Stockbot: Using lstms to predict stock prices. arXiv preprint arXiv:2207.06605 (2022)
17. Mojjada, R.K., Yadav, A., Prabhu, A., Natarajan, Y.: Machine learning models for COVID-19 future forecasting. Materials Today: Proceedings (2020)
18. Wang, Q., Xu, W., Zheng, H.: Combining the wisdom of crowds and technical analysis for financial market prediction using deep random subspace ensembles. Neurocomputing, **299**, 51-61 (2018). https://www.sciencedirect.com/science/article/abs/pii/S0925231218303540
19. Rai, R., Tiwari, M.K., Ivanov, D., Dolgui, A.: Machine learning in manufacturing and industry 4.0 applications. Int. J. Prod. Res. **59**, 4773–4778 (2021)
20. Remias, R.: President Trump's Tweets and their Effect on the Stock Market: The Relationship Between Social Media, Politics, and Emotional Economic Decision-Making. Ph.D. thesis, Wittenberg University (2021)
21. Seethalakshmi, R.: Analysis of stock market predictor variables using linear regression. Int. J. Pure and Appl. Math. **119**(15), 369–378 (2018)
22. Shang, S., Hui, P., Kulkarni, S.R., Cuff, P.W.: Wisdom of the crowd: incorporating social influence in recommendation models. In: 2011 IEEE 17th International Conference on Parallel and Distributed Systems, pp. 835–840. IEEE (2011)
23. Shen, J., Shafiq, M.O.: Short-term stock market price trend prediction using a comprehensive deep learning system. J. big Data **7**(1), 1–33 (2020)
24. Siah, K.W., Myers, P.: Stock market prediction through technical and public sentiment analysis (2016)
25. Surowiecki, J.: The wisdom of crowds. Anchor (2005)
26. Team, K.: Keras documentation: Adam. https://keras.io/api/optimizers/adam/#::text=Adam
27. Wazny, K.: Applications of crowdsourcing in health: an overview. J. Global Health **8**(1), 010502 (2018)
28. Yasmin, R., Grassel, J.T., Hassan, M.M., Fuentes, O., Escobedo, A.R.: Enhancing image classification capabilities of crowdsourcing-based methods through expanded input elicitation. https://ojs.aaai.org/index.php/HCOMP/article/view/18949
29. Yasmin, R., Hassan, M.M., Grassel, J.T., Bhogaraju, H., Escobedo, A.R., Fuentes, O.: Improving crowdsourcing-based image classification through expanded input elicitation and machine learning. Front. Artif. Intell. **5**, 848056 (2022)