Artificial Swarm Intelligence Amplifies Accuracy when Predicting Financial Markets

Louis Rosenberg
Unanimous A.I.
San Francisco, CA, USA
Louis@Unanimous.ai

Niccolo Pescetelli
University of Oxford
Clarendon, UK
niccolo.pescetelli@chch.ox.ac.uk

Gregg Willcox
Unanimous A.I.
San Francisco, CA, USA
gregg@Unanimous.ai

***Abstract*— Across the natural world, many species have evolved methods for amplifying their decision-making accuracy by thinking together in real-time closed-loop systems. Known as Swarm Intelligence (SI) in the field of biology, the process has been deeply studied in schools of fish, flocks of bird, and swarms of bees. The present research looks at human groups and tests their ability to make financial predictions by forming online systems modeled after natural swarms. Specifically, groups of financial traders were tasked with predicting the weekly trends of four common market indices (SPX, GLD, GDX, and Crude Oil) over a period of 14 consecutive weeks. Results showed that individual participants, who averaged 61% accuracy when predicting weekly trends on their own, amplified their accuracy to 77% when predicting together as real-time swarms. These results reflect a 26% increase in financial prediction accuracy and show high statistical significance (p=0.001). This suggests that enabling groups of traders to form real-time systems online, governed by swarm intelligence algorithms, has the potential to significantly increase the accuracy of financial forecasts.**

Keywords*— Swarm Intelligence, Artificial Swarm Intelligence, Collective Intelligence, Human Swarming, Artificial Intelligence.

I. Introduction

Artificial Swarm Intelligence (ASI) has been shown to amplify the intelligence of human groups by connecting users into real-time systems modeled after biological swarms [1, 2]. Prior studies have shown that "human swarms" can produce more accurate predictions than traditional "Wisdom of Crowd" methods such as votes, polls, and surveys [3]. For example, a 2015 study tested the ability of human swarms to forecast the outcome of college football games. A swarm comprised of 75 amateur sports fans was tasked with predicting 10 college bowl games. As individuals, the participants averaged 50% accuracy when predicting outcomes against the spread. When thinking together in real-time swarms, those same participants achieved 70% accuracy against the spread [2]. Similar increases have been found in other studies, including a long-term test that required participants to predict a set of 50 soccer matches in the English Premier League and showed a 31% increase in accuracy when participants were connected in swarms [4].

While prior studies have documented the ability of artificial swarms to amplify the predictive ability of human groups when forecasting sporting events, political races, and media awards such as the Oscars and Grammys, no formal study has been performed to assess whether swarm-based predictions of financial markets can achieve similar improvements. To address this need, a fourteen-week study was conducted that tasked human swarms of financial traders with making weekly predictions regarding the change in four financial indices (SPX, GLD, GDX, and Crude Oil). The objective was to assess whether a statistically significant improvement would be measured when comparing individual predictions to swarm predictions. In addition, swarm performance was compared with the traditional "Wisdom of Crowd" method of using the most popular prediction across the participant pool as the group forecast. Thus, the present study compared the predictive abilities of three cases – Individuals, Crowds, and Swarms.

II. SWARMS AS INTELLIGENT SYSTEMS

The primary difference between "crowds" and "swarms" is that in crowd-based methods, individual participants provide their input in isolation (for statistical aggregation after the fact), while in swarm-based methods, groups "think together" as real-time systems governed by intelligence algorithms and converge on solutions in synchrony. The swarming process is generally modeled after biological systems such as schools of fish and swarms of bees. The present research uses Swarm AI technology from Unanimous A.I. Inc, which is modeled largely on honeybee swarms. This model was chosen for the current study because honeybee swarms are known to significantly amplify the accuracy of critical decisions by enabling members to form real-time systems – i.e. *"hive minds"* – that can solve problems as a unified and amplified intelligence.

The decision-making processes that govern the behavior of honeybee swarms have been studied since the 1950s and have been shown to be remarkably similar to the decision-making processes in neurological brains [5,6]. Both employ large populations of simple excitable units (i.e., bees and neurons) that work in parallel to integrate noisy evidence, weigh competing alternatives, and converge on decisions in synchrony. In both, outcomes are arrived at through a real-time competition among sub-populations of excitable units. When one sub-population exceeds a threshold level of support, the

978-1-5386-1104-3/17/$31.00

corresponding alternative is chosen. In honeybees, this enables the group to converge on optimal decisions, picking the best solution to complex problems (i.e. selecting a new home location) over 80% of the time [7,8,9].

The similarity between "brains" and "swarms" becomes even more apparent when comparing decision-making models that represent each. For example, the decision process in primate brains is often modeled as mutually inhibitory leaky integrators that aggregate incoming evidence from competing neural populations [10]. A common framework for primate decision is the Usher-McClelland model in Figure 1 below.

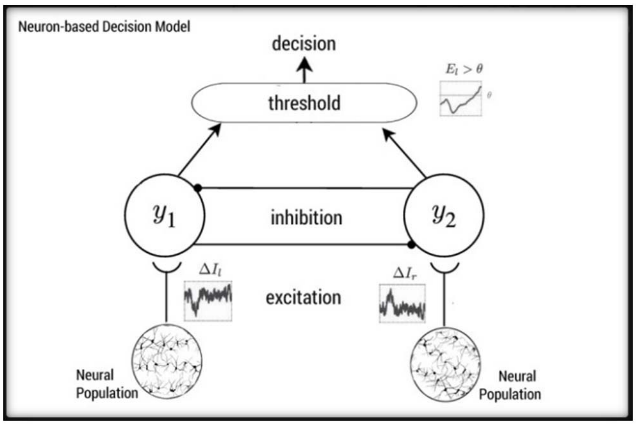

Fig. 1. Usher-McClelland model of neurological decision-making

This neurological decision model can be directly compared to swarm-based decision models, for example the honey-bee model represented in Figure 2 below. As shown, swarm-based decisions follow a very similar process, aggregating input from sub-populations of swarm members through mutual excitation and inhibition, until a threshold is exceeded.

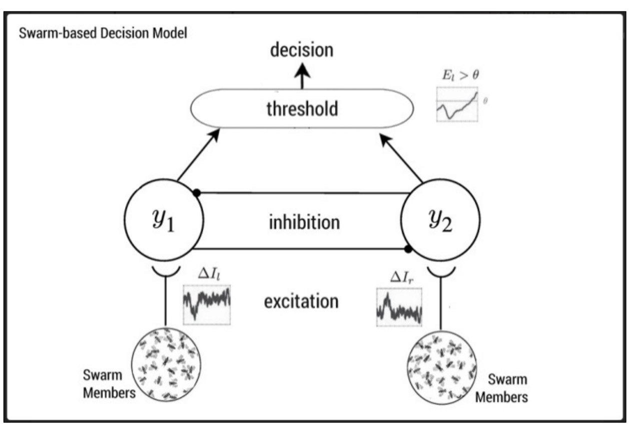

Fig. 2. Mutually inhibitory decision-making model in bee swarms

When viewed in this context, it becomes apparent that, while brains are systems of neurons structured so intelligence emerges, swarms are systems of brains structured so amplified intelligence emerges. Thus, the objective of the current study is to connect human financial traders into synchronous systems that are structured so that an amplified intelligence emerges.

III. ENABLING "HUMAN SWARMS"

Unlike many other social species, humans have not evolved the natural ability to form closed-loop systems that enable real-time swarming. That's because we lack the subtle connections that other organisms use to establish high speed feedback-loops among members. Schooling fish detect vibrations in the water around them. Flocking birds detect subtle motions propagating through the population. Swarming bees use complex body vibrations called a "waggle dance." To enable real-time swarming among groups of networked humans, specialized user interfaces, intelligence algorithms, and networking paradigms are required to close the loop among all members.

To address this need, a technology called Swarm AI was developed to enable human groups to congregate online as real-time swarms, connecting synchronously from anywhere in the world. It was first deployed in 2015 in an online platform called UNU that allow distributed groups of users to form closed-loop swarms using standard web-browsers [1]. Modeled after the decision-making process of honeybee swarms, the online system allows groups of distributed users to work in parallel to (a) integrate noisy evidence, (b) weigh competing alternatives, and (c) converge on decisions in synchrony, while also allowing all participants to perceive and react to the changing system in real-time, thereby closing a feedback loop around the full population of participants.

As shown in Figure 3, swarms answer questions by moving a graphical puck to select among a set of alternatives. Each participant provides input by manipulating a graphical magnet with a mouse or touchscreen. By positioning their magnet with respect to the moving puck, real-time participants express and impart their personal intent on the swarm as a whole. The input from each user is not a discrete vote, but a stream of vectors that varies freely over time. Because the full population of users can adjust their intent continuously in real-time, the swarm moves, not based on the input of any individual, but based on the dynamics of the full system. This enables a complex negotiation among all members at once, empowering the group to collectively explore the decision-space and converge on the most agreeable solution in synchrony.

Fig. 3. A human swarm answering a question in real-time

It is important to note that participants do not only vary the direction of their intent, but also modulate the magnitude of their intent by adjusting the distance between their magnet and the puck. Because the puck is in continuous motion across the decision-space, users need to continually move their magnet so that it stays close to the puck's outer rim. This is significant, for it requires participants to be engaged continuously throughout the decision process, evaluating and re-evaluating their intent as they convey their contribution. If they stop adjusting their magnet with respect to the changing position of the puck, the distance grows and their applied sentiment wanes.

Thus, like bees vibrating their bodies to express sentiment in a biological swarm, or neurons firing activation signals to express conviction levels within a biological neural-network, the participants in an artificial swarm must continuously update and express their changing preferences during the decision process, or lose their influence over the collective outcome. In addition, intelligence algorithms monitor the behaviors of all swarm members in real-time, inferring their implied conviction based upon their relative motions over time. This reveals a range of behavioral characteristics within the swarm population and weights their contributions accordingly, from entrenched participants to flexible participants to fickle participants.

IV. Financial Prediction Study

To assess the ability of human swarms to amplify their accuracy in financial predictions, a study was conducted over a fourteen week period using groups of volunteers who were unaffiliated with the research team. The participants were all self-identified as "active traders" who follow the financial markets daily and make financial trades regularly. Each weekly group consisted of between 7 to 36 participants. To establish a baseline, all participants provided their weekly forecasts as individuals using a standard online survey. The group then congregated online as a real-time swarm using the UNU swarming platform to make synchronous forecasts.

Across the fourteen week period, predictions were made for the following financial indices: (a) the S&P 500 (**SPX**), (b) the gold shares index fund (**GLD**), (c) the gold miners index fund (**GDX**), and (d) the crude oil index (**CRUDE**). The forecasts were generated every Tuesday at market close. The participants were asked to predict if each index would be higher or lower from the current price at market close on Friday (i.e. 72 hours later). Predictions were recorded first from individuals on private surveys, then from swarms working together as a system. In addition, participants were asked to qualify the expected change in price by indicating if the predicted move would be "by a little" or "by a lot." This was included as a means for evoking participant confidence in their directional forecast rather than as a true predictor of magnitude.

Figure 4 shows a snapshot of a human swarm comprised of 24 participants in the process of predicting a weekly change in GDX price. As shown in the figure, the swarm is given four options to choose from, enabling the swarm to identify which direction the index will move, as well as express a general sense of magnitude. The magnitude indicator is helpful as it causes the swarm to split into multiple different factions and then converge over time on a solution that maximizes their collective confidence and conviction. It's important to note that Figure 4 shows an instantaneous snapshot of the swarm as it moves over time towards a final answer. The full process of converging upon a solution generally required between 10 and 30 seconds of real-time interactions within the swarm.

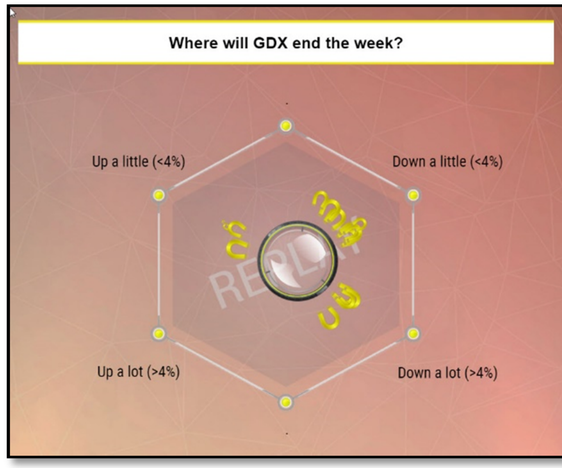

Fig. 4. Snapshot of a human swarm predicting GDX in real-time

A process called Faction Analysis has been developed to help researchers visualize how the participants in a human swarm adjust their support over time and converge upon the answer they can best agree upon. Figure 5 shows a faction analysis plot for the swarm above, the colors indicating how the swarm was initially split (at t=0.0) and then converged over the time upon the "down a little" option (at t=12.5), which in this case was an accurate prediction.

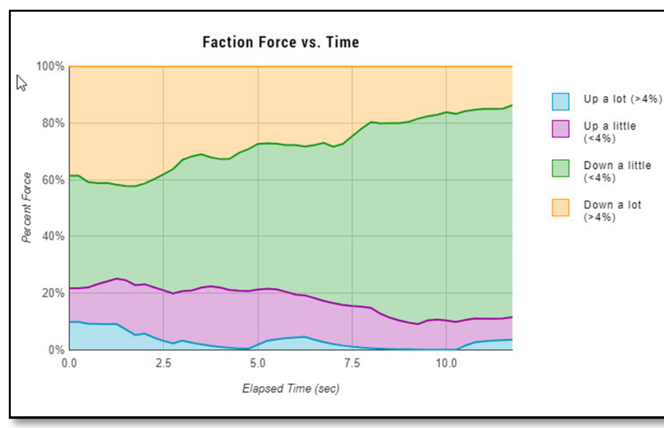

Fig. 5. Faction Analysis for human swarm predicting GDX

V. Analysis and Results

For each of the fourteen weeks in the testing period, a set of predictions were made for each of the four market indices (SPX, GLD, GDX, CRUDE), providing 56 sets of predictions. Results were generated indicating: (a) Individual Accuracy – computed as the average performance across the pool of human subjects, (b) Crowd Accuracy – computed by taking the most popular prediction from the participant pool and using that to compute accuracy over time, and (c) Swarm Accuracy – computed by assessing the accuracy of the predictions made by the swarms each week.

To assess whether the human swarms predicted the directional change in market indices more accurately than individuals, we compared swarm performance with individual performance using a bootstrapping procedure. For each of the four investment categories (SPX, GLD, GDX, CRUDE) and each prediction week, we selected the answer provided by an individual sampled at random among the individuals who provided a response for that particular week and investment type. Answers were averaged across the four investment types and the 14 weeks to obtain a percentage accuracy measure. The procedure was repeated 10,000 times in order to obtain a distribution of probabilities for making a correct prediction.

The distribution, shown in Figure 6 as a probability density function, represents the probability of an individual making a correct prediction when responses are randomly sampled from the individual answers provided. With a median accuracy of **61%,** the individuals were moderately better than random guessing when predicting the directional change in these market indicators. The red line in Figure 6 shows the empirical performance of the swarms, which at **77%** accuracy was significantly higher performing as compared to individuals. The probability of individuals scoring better than the swarm, using a random sampling procedure, was extremely low (p=0.002) indicating a highly significant result.

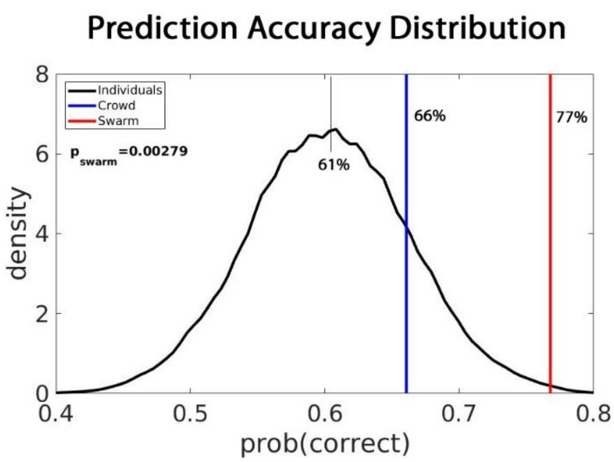

Fig. 6. Individual Accuracy vs Swarm Accuracy when predicting the directional change in all four indices in the subsequent 72-hour period.

A similar analysis was done using the more traditional "Wisdom of Crowd" method of taking the most popular predictions across the pool of individuals as the forecast. This is shown as the blue line in the figure above. At 66% accurate, the crowd was significantly lower performing than the swarm. In addition, the probability of an individual scoring better than the crowd was less conclusive (p=0.164). For these reasons, the results suggest that real-time swarming is a significantly more accurate and more reliable method for amplifying the intelligence of a human population when making financial forecasts. Looking at the results as a percentage increase, the swarms, on average, were **26%** more accurate when predicting the directional movement in the financial indices than the individual financial traders who comprised those swarms.

In addition to analyzing the predictive accuracy across all four indices in aggregate (as shown in Figure 6 above), it is also instructive to assess performance with respect to each of the four financial categories in isolation. This is shown in Figure 7 below. Across 14 weeks, the swarm outperformed the individual traders and the crowd-based forecasts in all four instances. Of particular interest is the performance of the swarm when predicting GDX, as we see the crowd underperformed the individual traders while the swarm significantly over-performed the individual traders. This suggests that while crowds can amplify errors, swarms are more resistant, converging on correct results in instances when the majority of participants selected an incorrect prediction.

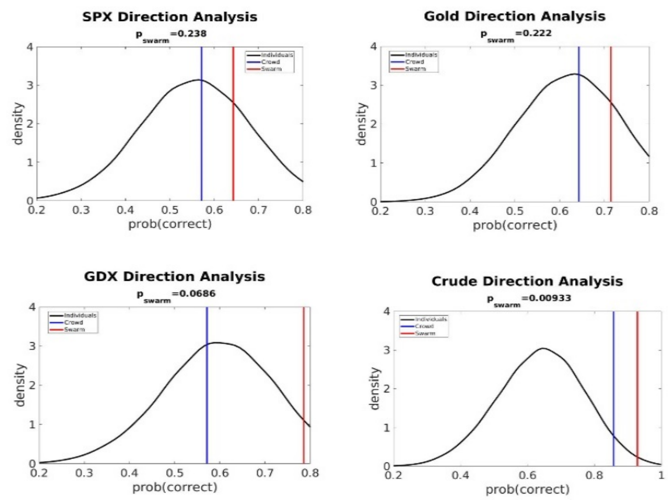

Fig. 7. Individual Accuracy vs Swarm Accuracy when predicting the directional change in each individual index in the subsequent 72-hour period.

Focusing specifically on the ability of swarms to amplify the accuracy of human forecasters and thereby enable more accurate financial predictions, the accuracy improvements for each of the four indices above are summarized in Table 1 below. As shown, the swarm amplified the accuracy of the participants by differing amounts across the four financial categories. The largest accuracy increase was recorded in crude oil predictions, which registered an impressive **28%** point net gain, corresponding to a **43%** amplification in total accuracy.

Financial Index	Individual Accuracy	Swarm Accuracy	Net Increase	Accuracy Amplication
SPX	55%	64%	9%	16%
GLD	62%	71%	9%	15%
CRUDE	65%	93%	28%	43%
GDX	62%	79%	17%	27%

Table 1. Individual Accuracy vs Swarm Accuracy across each index

In addition to assessing the ability to predict the directional trend of each financial index, it is also instructive to assess the ability of individuals, crowds, and swarms to qualify their predictions further by indicating if the weekly change would be "by a little" or "by a lot." While these are loose metrics, each was tied to a specific threshold. For example, when predicting GDX, the threshold was defined for the participants as a 4% change, meaning if the index changed by less than 4% it was classified as "by a little" and if the index changed by more than 4% it was classified as "by a lot." Because all participants were required to provide this additional qualifier when predicting on the survey and in the swarm, the same analysis can be performed for the primary results above.

Figure 8 below shows a probability density function that represents the probability of an individual making a correct prediction (in both direction and value) across all fourteen weeks and all four market indices. With a median accuracy of **44%,** the individuals were less accurate with this additional qualifier added. The red line in Figure 8 shows the empirical accuracy of the swarms, which at **57%** accuracy, was still significantly higher performing as compared to individuals. The probability of individuals scoring better than the swarm, using a random sampling procedure, was extremely low (p=0.01) indicating a highly significant result.

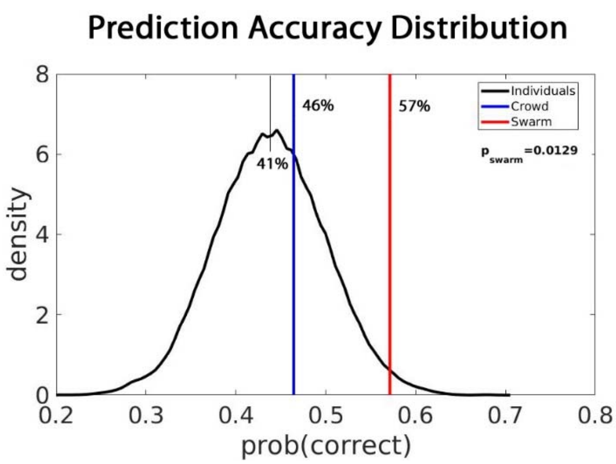

Fig. 8. Individual Accuracy vs Swarm Accuracy when predicting the direction and value change in all four indices in the subsequent 72-hour period.

As also shown in Figure 8 above, the swarm predictions were more accurate than the crowd predictions, which at 46% accurate was not statistically better than the individual forecasts (p=0.32). Thus, even with the added constraint of predicting if the weekly move would be "by a little" or "by a lot" (which increased difficulty), the swarm demonstrated significant benefits over both individual and the crowd-based forecasts. This suggests that swarm-based forecasting is not only a benefit when the individuals are most often correct (as in Figure 6), but also a benefit when the individuals are most often incorrect (as in Figure 8), further supporting robustness.

VI. CONCLUSIONS

Can real-time swarms of financial traders outperform the predictive accuracy of individual traders? The results of the current study suggest that swarms can significantly increase prediction accuracy when forecasting the directional movement of certain financial metrics. The results also show that the swarming process is more accurate and more repeatable than traditional crowd-based forecasting. Additional research is warranted to further validate the benefits of Artificial Swarm Intelligence for financial forecasting applications. Of particular interest is the ability of swarms to amplify prediction accuracy in longer term predictions, as the current study used a relatively short 72-hour forecasting window. Other topics recommended for ongoing research include exploring swarms of larger sizes, comparing participants of varying expertise levels, and testing improved swarming algorithms.

ACKNOWLEDGMENT

Thanks to Chris Hornbostel and David Baltaxe for their efforts in coordinating the weekly swarms of financial traders that generated predictions. Also, thanks to Unanimous A.I. for the use of the unu.ai platform for this ongoing work.

REFERENCES

[1] Rosenberg, L.B., "Human Swarms, a real-time method for collective intelligence." Proceedings of the European Conference on Artificial Life 2015, pp. 658-659

[2] Rosenberg, Louis. "Artificial Swarm Intelligence vs Human Experts," Neural Networks (IJCNN), 2016 International Joint Conference on. IEEE. J. Clerk Maxwell, A Treatise on Electricity and Magnetism, 3rd ed., vol. 2. Oxford: Clarendon, 1892, pp.68–73.

[3] Rosenberg, Louis. Baltaxe, David and Pescetelli, Nicollo. "Crowds vs Swarms, a Comparison of Intelligence," IEEE 2016 Swarm/Human Blended Intelligence (SHBI), Cleveland, OH, 2016, pp. 1-4.

[4] Baltaxe, David, Rosenberg, Louis and N. Pescetelli, "Amplifying Prediction Accuracy using Human Swarms", Collective Intelligence 2017. New York, NY ; 2017.

[5] Seeley T.D, Buhrman S.C 2001 "Nest-site selection in honey bees: how well do swarms implement the 'best-of-N' decision rule?" Behav. Ecol. Sociobiol. 49, 416–427

[6] Marshall, James. Bogacz, Rafal. Dornhaus, Anna. Planqué, Robert. Kovacs, Tim. Franks, Nigel. "On optimal decision-making in brains and social insect colonies." Soc. Interface 2009.

[7] Seeley, Thomas D., et al. "Stop signals provide cross inhibition in collective decision-making by honeybee swarms." Science 335.6064 (2012): 108-111.

[8] Seeley, Thomas D. Honeybee Democracy. Princeton Univ. Press, 2010.

[9] Seeley, Thomas D., Visscher, P. Kirk. "Choosing a home: How the scouts in a honey bee swarm perceive the completion of their group decision making." Behavioural Ecology and Sociobiology 54 (5) 511-520.

[10] Usher, M. McClelland J.L 2001 "The time course of perceptual choice: the leaky, competing accumulator model." Psychol. Rev. 108, 550–592