# FORECASTING STOCK PRICES USING COLLECTIVE INTELLIGENCE: IMPLICATIONS FOR THE GLOBAL BRAIN

**CRAIG A. KAPLAN***

**iQ Company (www.iqco.com)**

## Abstract

The Global Brain is a metaphor for a global network of humans and technology that exhibits collective intelligence (CI). Heylighen (1999) defines CI as the ability of a group to solve more problems than its individual members. Expanding the scope of this definition allows inclusion of other types of cognition besides problem-solving. For example, in this paper, a group that makes better decisions than its individual members is considered to exhibit CI.

This paper describes the design and testing of a prototype system that uses CI to make stock forecasting and trading decisions. During an eleven trading-day test period, the system out-performed the NASDAQ, S&P 500, and DJIA stock indices by margins of 12.40%, 5.68%, and 2.25% respectively.

Statistical analysis showed that it was highly unlikely that a random sample of NASDAQ stock picks would have performed as well as our system ($p<.02$). We also found that the system performed better when more people participated, suggesting that the system's good performance was due to CI.

Further testing is needed to see if these results will hold up over a longer period of time and with more participants. Implications of this research for the Global Brain and for general decision-making and problem-solving systems based on collective intelligence are discussed.

## Keywords

Group decision-making, stock price forecasting, collective intelligence, problem solving, global brain.

## The Global Brain and Collective Intelligence

According to the Global Brain Workshop website: "*The 'Global Brain' is a metaphor for the emerging collectively intelligent network formed by the people of this planet together with the computers, knowledge bases, and communication links that connect them together.*"

Simply put, the Global Brain is a global network of humans and technology that exhibits collective intelligence.

Collective Intelligence (CI) has been defined as the ability of a group to solve more problems than its individual members (Heylighen 1999). This definition could certainly be expanded to include other types of cognition besides problem solving. For example, a group that makes better decisions than its individual members might be considered to exhibit CI. Similarly a group that makes more accurate stock forecasts than its individual members might also be said to exhibit CI.

Now an interesting question is: **Under what circumstances does a network of humans and technology exhibit CI?**

If we can answer this question, or at least make progress towards answering it, then we will have taken an important step towards developing architectures to support the Global Brain.

## Research Strategies

Global Brain research, like Cognitive Science, can benefit greatly from an inter-disciplinary approach (Simon & Kaplan 1989).

* The author would like to thank Erica Golden, Jason Laumeister, and Daniel Swid for their help in conducting this research. Portions of this paper are drawn from the article, Collective Intelligence: A new approach to stock price forecasting, which will appear in the *Proceedings of the 2001 IEEE Systems, Man, and Cybernetics Conference.*

For example, at this Global Brain workshop we have evolutionary approaches, ecological and biological approaches, AI architecture approaches, design analysis approaches, simulation approaches, Informological approaches, and web-based learning approaches, among others.

Some of these approaches arrive at theoretical claims relating to the Global Brain. Other approaches are more applied. The applied approaches tend to draw more heavily on computer science and related fields with the goal of building experimental systems that exhibit collective intelligence.

While I believe that both theoretical and experimental approaches are needed, this paper reflects an inductive, experimental approach.

As philosophers of science are fond of pointing out, it is necessary to establish the existence of a phenomenon before attempting to explain it. While it may seem obvious to many of the Global Brain researchers that CI exists, there is considerable skepticism and resistance to the concept of CI outside of the Global Brain research community. Therefore, my immediate objective has been to build a simple, small-scale, prototype system that exhibits CI.

To be relevant to Global Brain research, the CI system had to be composed of a network of humans and technology like the Global Brain, and the system needed to be capable of scaling (at least theoretically) to global size.

While the Global Brain is usually assumed to be a general collective intelligence, the prototype system necessarily has a much narrower scope of intelligence. For example, while the Global Brain conceivably would be able to tackle any sort of problem, the prototype is limited to tackling one kind of decision making in a limited domain.

The reason for the limited scope of the prototype is pragmatic. Researchers in the field of problem solving have discovered that problem representation (Heylighen 1988 & 1990, Kaplan & Simon 1990) and problem decomposition (Newell & Simon 1972) are tough nuts to crack. Although it may be possible to build an online distributed problem solving system capable of CI (Kaplan 2000), such an effort would require a significant investment of resources and time.

For the purposes of establishing the phenomenon of CI, it proved to be easier to implement a simple decision making system.

For example, in a decision-making task like deciding whether to buy, sell, or hold a stock for a 24-hour period, it is relatively easy to determine the effectiveness of the decision. Just wait 24 hours and see how much money the system made or lost due to its decision. In contrast, even the task of assigning credit or blame in a general problem solving system can be quite complex.

The goal of the prototyping effort was therefore to create a stock price forecasting and decision-making system that exhibits CI. The system would combine the information, processing power, and judgment of many minds (networked by technology) to create collective price forecasts. If trading decisions based on the collective forecasts were more profitable than the decisions made by individual investors (as reflected in the performance of major stock indices) then we would have evidence of CI. If the accuracy (and profitability) of the collective forecasts increased as more people participated in generating the collective forecasts, than this would be even stronger evidence of CI.

Specifically we hypothesized:

1. Our CI system would outperform the NASDAQ, S&P500, and DJIA indices.

2. Our CI system would perform better as more people participated in the system.

**Stock Forecasting Prototype**

iQ Company (iQ) built a prototype consisting of a website, a database, and proprietary information processing algorithms. Figure 1 illustrates the basic architecture of iQ's prototype system.

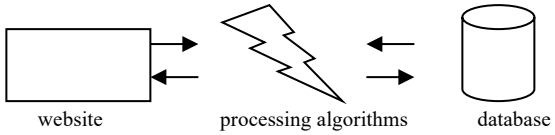


Figure 1. Basic Prototype Architecture

Investors went to the prototype's website where they forecast tomorrow's closing price for the stock(s) of their choice. To create collective forecasts, and to prevent unscrupulous individuals from gaming the system, we then processed the raw individual forecasts through proprietary and patent-pending technology.

The details of this processing system are described in our pending patent (Kaplan 2001). However, conceptually, our system weights each individual's forecast by the past track record of that individual on a particular stock. Information about the performance of each investor was stored in a database where it could be accessed by the processing algorithms as needed.

For example, suppose John and Sue are the only two individuals who have made forecasts about IBM's stock price. Our system tracks how accurate John and Sue are over time. If John is more accurate in his forecasts than Sue, then when the forecasts of John and Sue are combined into a single collective forecast, the system will give John's forecast more weight than Sue's forecast.

A powerful feature of the CI approach is that as more people use the system, the system becomes more accurate. This creates a positive feedback loop – sometimes called a "network effect" – in which more people using the system induces even more people to use the system.

It may be useful to describe how the system worked from the point of view of a participant in our test.

Participants began by going to a website and entering stock symbols to get 15-minute delayed quotes (Figure 2).


Figure 2. Initial Quote Request Screen

The prototype then returned a table of quotes, with blank fields where participants could enter their forecasts for tomorrow's closing price for these same stocks (Figure 3).


Figure 3. Enter Stock Forecasts Screen

Next, the system included the participant's individual forecasts in its collective intelligence calculations, and showed participants up-to-the-minute collective forecasts for each stock that the participant had forecast (Figure 4).


Figure 4. Display Collective Forecasts Screen

Note that this approach incorporated an informational trade: the participant entered an individual forecast and received in return a potentially more valuable collective forecast from the system.

At the top of every screen, participants saw iQ's patent-pending Tomorrow Ticker[tm], which displayed continuous updates of both stock quotes and collective forecasts (Figure 5).


Figure 5. The Tomorrow Ticker

The top line of the ticker showed a 15-minute-delayed quote for the symbol while the bottom line showed the latest collective forecast of tomorrow's closing price for the same stock symbol.

**Methodology**

iQ partnered with Save Our Shores (SOS) – a local environmental group with a mailing list of about 1,500 members. Interviews with members of SOS suggested that this group was representative of the general population with the exception that SOS members tended to be more environmentally conscious than the average citizen. Not all SOS members had Internet access, but enough did so that we could conduct a useful test of our prototype.

We wanted to target average people who did not possess any special knowledge and who were not professional traders. Our reasoning was that if we could get results by tapping the collective intelligence of members of the general population, then the results would only get better if we tapped stock trading professionals. We also wanted to generalize any results to the online investor population at large – and this population is made up mostly of people who do not invest professionally.

To encourage participation, iQ offered to donate a minimum of one cent for each individual forecast anybody made during the test period. iQ also offered to increase this donation if the system outperformed the S&P 500 and/or if participants entered forecasts on more than one day during the test period.

The test period was scheduled solely for the convenience of SOS. The test period began on February 5, 2001 and ended on February 20, 2001. Due to weekends and holidays, this period contained eleven trading days when the major US stock markets were open. iQ contacted SOS members via mail and Email to promote participation in the experimental test. All together, 63 individuals participated – about a 4% response rate from the 1,500 members contacted. These 63 participants entered a total of 785 individual forecasts. From these 785 individual forecasts, the collective intelligence algorithms generated 401 unique collective forecasts. The results reported below are based on various analyses of these 401 collective forecasts.

Here is how the CI system operated during the test:

As individuals entered their stock forecasts, the system processed these forecasts using its computer algorithms and generated collective forecasts. Each collective forecast represented the processed intelligence of one or more individuals regarding the closing price of a particular stock the next day. For example, the collective forecast, IBM 114.25, means that based on the input of all individual forecasts for IBM, the system has produced a single collective forecast that IBM will close at 114.25 tomorrow.

Each day, just before the market closed, we simulated trading on all available collective forecasts for the next day. Trading on the collective forecasts means we followed these rules:

1. If today's closing price is lower than the collective forecast, then buy the stock because the system is saying it will go up tomorrow.

2. If today's closing price is higher than the collective forecast, then sell the stock because the system is saying it will go down tomorrow.

3. If today's closing price is identical to the collective forecast, do nothing because the system is saying there will be no change tomorrow.

4. The amount of stock bought or sold is proportional to the % difference between the collective forecast and today's closing price. This means you buy more of a stock that is forecast to go up a lot than you do of a stock that is forecast to go up only a little. Similarly, you sell more if the stock is forecast to go down more.

5. If the collective forecast was correct in predicting the direction that the stock moved, then you made the difference between the closing price on the day the collective forecast was generated and the next day – when the collective forecast could be tested. If the collective forecast was wrong about the direction, then you lost the difference.

The overall return from trading on the collective forecasts is simply the average of the percentage

gains and losses from all trades made during the test period.

**Results**

Our first hypothesis was that a CI system would outperform (i.e. yield better returns than) the NADAQ, S&P 500, or DJIA.

We compared the actual return that would have been generated by trading on the system's collective forecasts with the actual return that would have been generated by investing in the NASDAQ, S&P 500, and the DJIA stock indices. Figure 6 shows the results of this comparison graphically. Performance of the indices assumed we bought each index on Feb 5th and sold it on Feb 20th – the same period covered by our test of the CI system.

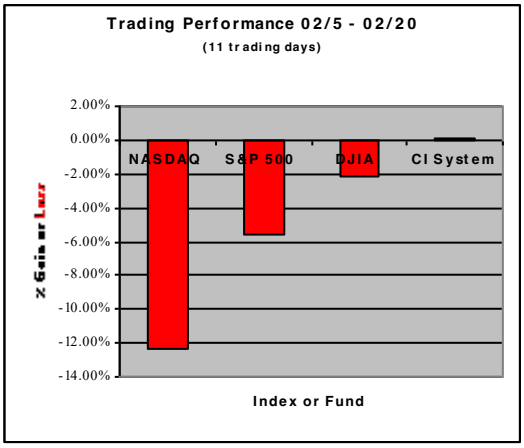


Figure 6. System vs. Major Indices

The CI system gained .11% during the eleven trading-day period while the other indices suffered losses. The system outperformed the NASDAQ, S&P 500, and DJIA stock indices by margins of 12.40%, 5.68%, and 2.25% respectively.

A t-test showed that the probability that the CI system would outperform the NASDAQ by as much as it did simply by chance is less than two in one hundred ($p<.02$, two-tailed t-test). This result is statistically significant.

If collective intelligence is responsible for the better performance of the system, then as more people use the system it should perform better. The best way to test this hypothesis would be to run a follow-up test with a very large number of participants to see if the CI system performs even better. However it is also possible to divide the data we have already collected into groups of collective forecasts that were based on different numbers of individual forecasts.

For example, on one day we might find that the collective forecast for IBM was based on only two individual forecasts, while on another day, the collective forecast for IBM might have been based on eight individual forecasts. If the collective forecast based on eight individual forecasts proved to be more accurate than the forecast based on only two individual forecasts, then this would be evidence in support of CI.

We divided the 401 collective forecasts into three groups – those based on only one or two individual forecasts (Low CI, n=368), those based on three individual forecasts (Med CI, n=17) and those based on four or more individual forecasts (High CI, n=16). We then compared the gains made by trading on collective forecasts in the three groups. Figure 7 shows the results of this comparison.

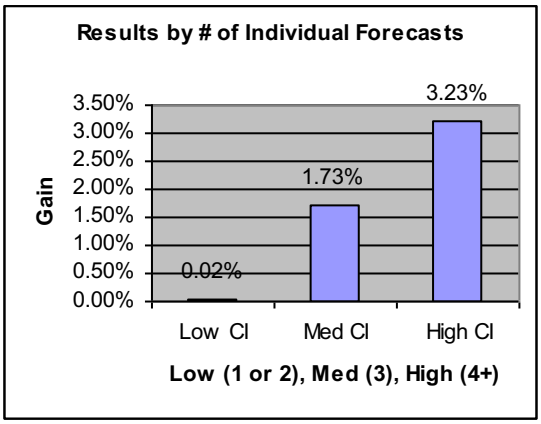


Figure 7. Results by # of Individual Forecasts

Figure 7 shows clearly that when the system's collective predictions are based on more forecasts, they are more accurate, and translate into higher % gains. It is unlikely that the better gains achieved by the collective forecasts based on three or more individual forecasts are simply due to chance ($p<.07$, one tailed t-test).

Perhaps more importantly, from an investors point of view, Figure 7 shows we could have made 3.23% profit (instead of the .11% overall profit shown in Figure 6) during the eleven-day test period, if we used only collective forecasts based on four or more

individual forecasts. Assuming 250 trading days, this equates to a 73.4% annual return on investment.

**Conclusions**

First, the system for forecasting stock prices outperformed the NASDAQ, S&P 500, and DJIA during our test period. Second, statistical analysis showed that it was highly unlikely that this exceptional performance was due to chance. Third, detailed analysis showed that the system's performance improved as more people participated (i.e. when collective forecasts were based on more individual forecasts), suggesting that CI was responsible for the overall good performance of the system.

Together, these preliminary results make a strong case that the CI system works and that it offers significant gains over simply investing in index funds. However since the test was conducted over only eleven trading days, and because a total of only 63 people were involved in the test, more testing is needed to prove the advantage of the CI approach conclusively.

Next steps include improving the reliability and usability of the CI prototype, conducting a test on a much larger (global) scale, and experimenting with variations of the CI processing algorithms to identify those that are most effective.

As the factors that influence CI are understood more clearly, it should be possible to develop a range of architectures that could help support the Global Brain.

**References**

Heylighen F. (1988): "Formulating the Problem of Problem-Formation", in: Cybernetics and Systems '88, Trappl R. (ed.), (Kluwer Academic Publishers, Dordrecht), p. 949-957.

Heylighen F. (1990): Representation and Change. A Metarepresentational Framework for the Foundations of Physical and Cognitive Science, (Communication & Cognition, Gent), 200p.

Heylighen F. (1999): "Collective Intelligence and its Implications on the Web: algorithms to develop a collective mental map", Computational and Mathematical Theory of Organizations 5(3), 253-280.

Kaplan, C.A. (2001) Tomorrow Ticker: A system for using collective intelligence to forecast future values of financial or other quantifiable information. Patent Pending, USPTO.

Kaplan, C.A. (2000) Online Distributed Problem Solving (ODPS) System. Patent Pending USPTO.

Kaplan C.A. & Simon H.A. (1990). In Search of Insight. *Cognitive Psychology*, 22, 374-419.

Newell A. & Simon H.A. (1972) Human Problem Solving. Englewood Cliffs, NJ: Prentice Hall

Simon H. A., & Kaplan C. A. (1989). Foundations of cognitive science. In M. Posner (Ed.), *Foundations of Cognitive Science*. Cambridge, MA: The MIT Press