

COLLECTIVE INTELLIGENCE A NEW APPROACH TO STOCK PRICE FORECASTING

CRAIG A. KAPLAN

*Proceedings of the 2001 IEEE Systems, Man, and Cybernetics Conference
iQ Company (www.iqco.com)*

Abstract

A group that makes better decisions than its individual members is considered to exhibit collective intelligence (CI). This paper describes the design and testing of a prototype CI system that makes stock trading decisions based on group input. We hypothesized that the CI system would outperform the major stock indices, and that the performance of the system would improve as the size of the group increased.

During an eleven trading-day test period, the system outperformed 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 the system ($p < .02$).

We also found that the system performed better when trading decisions were based on input from larger groups, suggesting that CI rather than individual intelligence was responsible for the system's overall good performance. Further testing is needed to see if these results will hold up over a longer period of time and with more participants.

Keywords

Group decision-making, stock price forecasting, collective intelligence.

Introduction

Is it possible to forecast stock prices accurately?

Given that simply buying an index fund, such as the S&P 500, can outperform more than 80% of the professional analysts and stock traders on Wall

Street, some have argued that stock picking is a losing proposition [1]. Those who insist on trying anyway generally fall into one of two camps: value investors or technical investors.

Value investors, of whom Warren Buffet is probably the best-known example, buy companies that they feel are undervalued and then sell them when the stock prices adjust upwards [2]. Technical investors analyze charts and graphs in an attempt to find inflection points, market peaks and bottoms, or other patterns or trends that they think will give them an investment edge [3].

One common argument against the possibility of forecasting stock prices is the market efficiency argument. According to this point of view, if technical analysis, value investing, or any other method based on generally available information really worked then enough people would begin using it so that any advantage would disappear. However some researchers have pointed out that the very fact that traders are willing to pay for information indicates that markets cannot be informationally efficient [4, 5].

A potentially more serious problem – at least for value investors – is that the stock market does not seem to be rational. Rational expectation models of the stock market do not predict speculative bubbles and crashes, yet these are relatively commonplace phenomenon both in the real world and under experimental laboratory conditions [6, 7].

There is a growing body of evidence that the key to forecasting the stock market lies neither in value analysis nor in technical analysis. Rather, investor psychology seems to be the critical factor.

Instead of treating investor psychology as noise, we should recognize that it is actually the signal that drives much of the day-to-day price fluctuations. In fact, when researchers have conducted experimental studies of investors, they find that people do not seem to take a random walk down Wall Street. Instead they behave in predictable ways, which can be captured by models that incorporate psychological assumptions, such as momentum investing [7, 8].

People also tend to expect rising prices, miss price jumps, and learn from experience [6]. In fact one researcher has shown that the severity of price crashes has a very strong correlation ($R^2=0.98$) with the length of time since the last crash [9]. Apparently, over time, people forget that crashes can happen and eventually get carried away with “irrational exuberance.”

Of course the media and analysts have built a multi-billion dollar industry rationalizing and interpreting market psychology after the fact. People talk about whisper numbers and investor confidence, but the truth is most professionals haven't a clue what is about to happen. What almost all of them are quite good at doing is creating a story, after the fact, to explain what already happened.

But what if there were a way to instantly gather and synthesize all of the complex variables that comprise investor psychology, and then distill them into a single number – a forecast of a future stock price?

In fact, each individual investor is doing exactly such a synthesis in his or her own brain before making an investment decision. However the synthesis is imperfect. It is imperfect, in large measure, because:

- a. no one investor has access to all of the relevant information,
- b. no one investor has the brainpower to process all of the information rapidly enough to make optimal decisions, and

- c. individual investors make mistakes in judgment

Notice however that all three of these shortcomings might be overcome by groups. Groups, collectively, have access to more information than an individual does. Groups, collectively, have the brainpower to process much more information than an individual can. Groups, collectively, generally have better judgment than a single individual. (For example, the use of twelve jurors in a judicial system, or of peer-review in academia, relies on this principle.)

Collective Intelligence

Collective Intelligence (CI) has been defined as the ability of a group to solve more problems than its individual members [10]. 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.

Simply put, when CI is operating, many minds are better than one.

Applying CI to the domain of stock price forecasting, one could imagine a system that would combine the information, processing power, and judgment of many minds to create collective price forecasts that were more accurate than those produced by most individual investors. Since all investors buy or sell stock based on their future expectations (usually individual forecasts), a CI system that traded based on collective forecasts should outperform the major market indices. Specifically, we hypothesized:

1. A CI system would outperform the NASDAQ, S&P500, and DJIA.
2. A CI system would perform better as more people participated in the system.

3. 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.

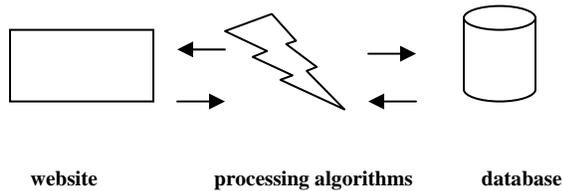


Figure 1. Basic Prototype Architecture

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 [11]. 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™, which displayed continuous updates of both stock quotes and collective forecasts (Figure 5).

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.

Current Stock Quotes: IBM 94.96 MSFT 54 ORCL 16.06 ARBA 1
 Collective Forecasts: IBM 95 MSFT 53 ORCL 16 ARBA 1

Figure 5. The Tomorrow Ticker

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.

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.

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 NASDAQ, 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.

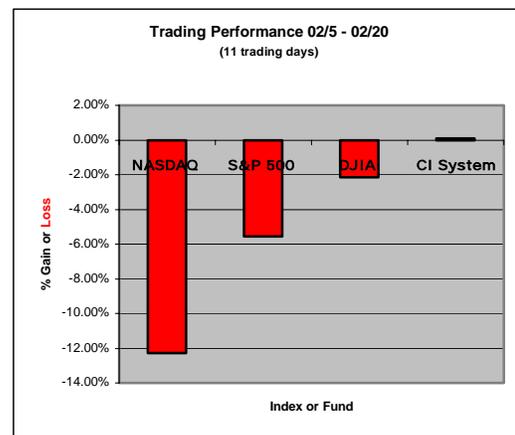


Figure 6. System vs. Major Indices

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.

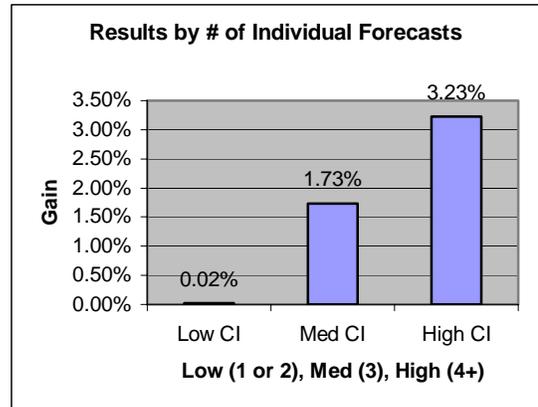


Figure 7. Results by Number of Individual Forecasts

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 the 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 scale, and experimenting with variations of the CI processing algorithms to identify those that are most effective. If the preliminary results reported here continue to hold over longer periods of time and with larger numbers of participants, then systems incorporating CI may become essential tools for investors of the future.

References

- [1] B. G. Malkiel, *A Random Walk Down Wall Street*, New York, NY: Norton, 1996.
- [2] B. Graham and W. Buffet (Preface), *The Intelligent Investor*, New York, NY: Harper Collins, 1985.
- [3] R. D. Edwards and J. F. Magee, *Technical Analysis of Stock Trends*, New York, NY: AMACOM, 1997.
- [4] G. Caginalp and D. Balenovich, Trend-Based Asset Flow in Technical Analysis and Securities Marketing, *Psychology & Marketing*, Vol. 13, 4, pp. 407-444, 1996.
- [5] S. Grossman and J. Stiglitz, On the Impossibility of Informationally Efficient Markets, *American Economic Review*, Vol. 70, pp. 393-407, 1980.
- [6] D. P. Porter and V. L. Smith, Stock Market Bubbles in the Laboratory, *Applied Mathematical Finance*, Vol. 1, pp. 111-127, 1994.
- [7] G. Caginalp, D. Porter, and V. L. Smith, Overreactions, Momentum, Liquidity, and Price Bubbles in Laboratory and Field Asset Markets, *The Journal of Psychology and Financial Markets*, Vol. 1, No. 1, pp 24-28, 2000.
- [8] G. Caginalp, D. Porter, and V. Smith, Momentum and Overreaction in Experimental Financial Markets, *International Journal of Industrial Organization*, Vol. 18, No. 1, pp. 187-204, 2000.
- [9] E. Renshaw, The Crash of October 19 in Retrospect, *The Market Chronicle*, Vol. 22, 1, 1988.
- [10] F. Heylighen, Collective Intelligence and Its Implications on the Web: Algorithms to Develop a Collective Mental Map, *Computational and Mathematical Theory of Organizations*, Vol. 5, 3, pp. 253-280, 1999.
- [11] C. A. Kaplan, Tomorrow Ticker™: *A System for Using Collective Intelligence to Forecast Future Values of Financial or Other Quantifiable Information*, Patent-Pending, USPTO, 2001.

The author would like to thank Erica Golden and Jason Laumeister for their help in conducting this research.