

Leveraging collective intelligence in organizations

DANIEL M. ROMERO, School of Information, University of Michigan
ÁGNES HORVÁT AND BRIAN UZZI, NICO, Northwestern University

The concept of collective intelligence puts forth a key hypothesis: There is a cognitive surplus in crowds. Attempts to tap this cognitive surplus in science, business, and government are increasingly being employed [Hand 2010]. Yet, while there is relatively scarce evidence of its efficacy [Galton 1907], recent endeavors point to the potential of collective intelligence in improving people's performance in a wide range of tasks [Woolley et al. 2010]. However, these findings suggest at least one possible paradox: the individuals who contribute to the collective intelligence of the crowd may be unaware that it exists. We investigate the existence of collective intelligence signals in modern organizations that could result in better decisions than were actually made. Unlike crowds, organizations are the most pervasive human infrastructure for making consequential decisions. Moreover, the design of organizations creates new scientific challenges for understanding collective intelligence. The conditions associated with collective intelligence revolve around crowds that are made up of diverse decision-makers acting independently [Lorenz et al. 2011]. By contrast, the scale economics and stable membership of organizations dictate that decision-makers have correlated information and work towards consensual decisions. This points to the need to understand how collective intelligence can operate effectively in non-crowds.

In this study, we examine the resource allocation decisions of a large hedge fund. Each decision concerns the movement of funds in and out of stock holdings. Because hedge funds make a large number of these decisions and the accuracy of the decision can be clearly measured, even a single firm provides remarkably dense data to draw valid statistical tests. Our data covers two years during which the organization made approximately 800 decisions per week valued at an average of \$700,000 per decision. Financial organizations also have sophisticated communications that are now acquired in their totality and accurately through electronic records. To capture communications dynamics in the organization that may be indicative of collective intelligence mechanisms, we analyze the instant message communications of the organization, which include 22 million IMs exchanged among the 182 personnel of the hedge fund and their network of 8,646 outside contacts. From these IMs we identify two collective intelligence signals that we show to be predictive of the changes in the market's movement as recorded by public Dow Jones Industrial Average (DJIA) data.

The collective intelligence signals relay and concentration

One crucial aspect of communication dynamics is the speed at which information is relay from one person to the next. The urgency of relayed information is one mechanism we find to be related to collective intelligence. To measure the rapidity of market related information relay, we tag IMs that contain a valid stock symbol (e.g., AAPL). We compute the fraction of times a user u' who received an IM containing a given symbol from user u sends an IM with the same symbol to a contact u'' ($u'' \neq u$) within at most k hours after receiving the original message from u and denote it by $G_k(t)$. We use $k = 24$ to reflect the daily patterns of information. Then we further code IM communications for their sentiment, by using the standard Linguistic Inquiry Word Count dictionary [James et al. 1999]. The dictionary identifies words with positive and negative sentiment. IMs are short and are coded as having positive (negative) sentiment if they contain at least one positive (negative) word and no negative (positive) words. Requiring that the fraction of relayed messages $G_k(t)$ is computed only based on the IMs that contain sentiment, we obtain the *relay* measure $relay(t)$.

A second mechanism is the concentration of attention between employees who specialize in the same stocks. We use stock mentions in IMs as a proxy for specialization. For each user u , we let M_u be the number of messages sent by u and S_u be the set of stock symbols mentioned by u . For each pair of users (u_1, u_2) , we let $D(u_1 \rightarrow u_2)$ be the number of messages sent by u_1 to u_2 . We then ask whether users who mention many of the same stocks tend to concentrate their attention on each other. We let P_k be the set of pairs of users who share k stock symbols. That is, $P_k = \{(u_1, u_2) : |S_{u_1} \cap S_{u_2}| = k\}$. For each pair (u_1, u_2) in P_k we compute the ratio of messages exchanged between u_1 and

u_2 and the total number of messages sent by both users. We let F_k be the mean of these ratios. That is, $F_k = \frac{1}{|P_k|} \sum_{(u_1, u_2) \in P_k} \frac{D(u_1 \rightarrow u_2) + D(u_2 \rightarrow u_1)}{M_{u_1} + M_{u_2}}$. F_k increases as k increases, showing that pairs of users who mention the same stock symbols have high volumes of communication than other pairs. The rate of increase of the function F_k indicates how much members concentrate their attention on those similar to them. We compute the function $F_k(t)$ for each week t and perform a linear regression on the values $(k, F_k(t))$. We define *concentration*(t) to be the slope of the resulting approximation.

Predicting the movement of the market

To discover whether future change in the market is associated with relay and concentration, we test for the utility of these measures to predict the movements of the DJIA. We start by looking at the values of *relay*(t) and *concentration*(t) when the DJIA went up and down during week $t + 1$. Table 1 shows basic statistics of the relay and concentration values, given the direction of the market. The values indicate that the behavior of the organization with regard to relay and concentration at time t differ significantly ($P < 0.01$ and $P < 0.003$ respectively) depending on whether the market rises or falls at time $t + 1$, according to a Wilcoxon–Mann–Whitney test for difference in medians. A relatively high relay and a relatively low concentration at time t is associated with the market falling at time $t + 1$. One interpretation for the rise in relay and fall in concentration before the market goes down is that these dynamics positively correlate with the level of uncertainty about the future state of the market. Information is circulated faster and reaches also employees who are not experts for a given stock as members of the firm attempt to disambiguate the uncertainty.

	# weeks	<i>relay</i> (t)		<i>concentration</i> (t)	
		Mean	STD	Mean	STD
DJIA goes up ($t + 1$)	54	0.036	0.014	0.010	0.008
DJIA goes down ($t + 1$)	44	0.046	0.018	0.007	0.005

Table 1: Mean and standard deviation of relay and concentration values when DJIA goes up and down during the following week

To understand whether relay and concentration are associated with, not just the direction of the market, but also the magnitude of the change, we perform a linear regression using Granger causality [Granger 1969]. As shown in Table 2, increases in relay or decreases in concentration at time t are associated with greater magnitudes of change in market at $t + 1$.

Variable	Coefficient	SE
Relay	-64.98**	25.31
Concentration	64.14*	27.79

Table 2: Regression model that predicts the change in the DJIA from relay and concentration. *: $P < 0.05$ and **: $P < 0.01$

These results suggest that relay and concentration, two variables derived from the collective behavior of the firm, are correlated with future movements of the market. Many observers of markets, as well as well-developed economic theory, suggest that consistent prediction of the market is impossible. Yet, we find that these signals consistently predict the direction and size of the change of the market. This suggests that using communications in an organization can provide a residual of cognitive surplus for responding to problems that otherwise overwhelm individual decision-makers.

Using collective intelligence signals to improve performance

Using the stock trades performed by the employees of the hedge fund, we explore whether the predictions derived from the collective intelligence signals could have resulted in better decisions than the ones actually made by the firm. Our tests hold the human capital of the firm, the timing, stocks, and size of trades constant to create a control group of altered trades¹. The only difference

¹ There are four kinds of transactions in the hedge fund's trading data: buy, sell, short, and cover. In buy/sell transactions the trader is hoping to buy a stock for a low price and sell it later for a higher price, hence making a profit. When a trader 'shorts' a stock, she obtains the stock and pays for it in the future, at a future price. When a trader 'covers' a stock, she is paying for a

between the actual trades and the altered trades is that the altered trades are based on the predictions made from the collective intelligence signals.

We measure the weekly profit earned by the hedge fund by trading stocks from the 30 companies included in the DJIA. Since the same stocks are bought and shorted at different prices, the profit of the firm can be computed variously based on which of the previously bought (shorted) shares are assumed to be sold (covered). To cover a broad range of possibilities, we compute the profit based on three methods: a.) FIFO: we sell (cover) the shares that were *first* bought (shorted); b.) FILO: we sell (cover) the shares that were *last* bought (shorted); c.) Random: we choose a random share to sell or cover. For each week t , we then train the logistic regression model on the data from weeks $1, \dots, t$ and use $relay(t)$ and $concentration(t)$ to predict the direction of the DJIA during week $t + 1$. Based on this prediction, we modify the trades performed by the hedge during week $t + 1$. The alterations only affect 'buy' transactions that later 'sell' within the week and 'short' transactions that are later 'covered' within the week, i.e. they only affect the short-term decisions of the company. We alter the trades in two ways: 1. If we predict the DJIA is going down we switch every 'buy'-'sell' pair of transactions to a 'short'-'cover' pair; 2. If we predict the DJIA is going up we switch every 'short'-'cover' pair to a 'buy'-'sell' pair.

Figure 1 shows the cumulative profit for the actual trades and the altered trades. We find that regardless of the way we compute profit, the profit generated by the actual trades is much lower than the one generated by the altered trades, validating the potential usefulness of the relay and concentration measures. These results imply that by aggregating the communication patterns of the employees of the hedge fund we are able to not only extract predictive information of the market, but also improve the performance of the firm.

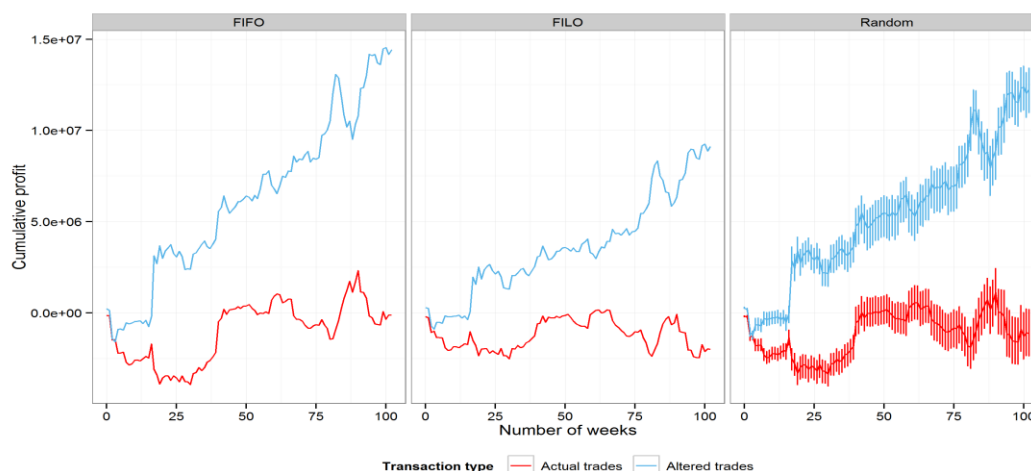


Figure 1: Cumulative profit earned by the actual trades performed by the hedge fund (red) and the altered trades that reflect the prediction of the direction of the DJIA (blue). The last curves are based on 1,000 random draws

Conclusion

The powerful idea of the wisdom of the crowds applies in the setting of an organization. In this case the number of people in “the crowd” tends to be small and there is constant interaction among people. We find that certain aggregate signals that come from the collective interactions among the personnel of an organization are predictive of the movement of the market. When the crowd urgently relays information, the market is likely to go down in the future. On the other hand, when they concentrate their attention on other users who discuss the same stock symbol, the market is more likely to up. Furthermore, the signals allowed us to alter the firm’s trades, improving its performance.

previously 'shorted' stock. Hence, in short/cover transactions the trader is hoping to make a profit by obtaining a stock of high price and covering it for a low price.

References

- Eric Hand. 2010. Citizen science: people power. *Nature* 466: 685–687
- Francis Galton. 1907. Vox populi. *Nature* 75 (1949): 450–451
- Anita Williams Woolley, Christopher F. Chabris, Alex Pentland, Nada Hashmi and Thomas W. Malone. 2010. Evidence for a collective intelligence factor in the performance of human groups. *Science* 330 (6004): 686–688
- Jan Lorenz, Heiko Rauhut, Frank Schweitzer and Dirk Helbing. 2011. How social influence can undermine the wisdom of crowd effect. *PNAS* 108(22): 9020–9025
- James W Pennebaker, Roger J. Booth, Martha E. Francis. 1999. Linguistic inquiry and word count. URL: <http://www.liwc.net/index.php>
- Clive Granger. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37(3): 424–438