PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A
MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Self–organized segregation of traders within a market

N. F. Johnson, P. M. Hui and T. S. Lo

Phil. Trans. R. Soc. Lond. A 1999 **357**, 2013-2018
doi: 10.1098/rsta.1999.0413

THE ROYAL SOCIETY

Self-organized segregation of traders within a market

By N. F. Johnson[1], P. M. Hui[2] and T. S. Lo[2]

[1]*Department of Physics, University of Oxford, Parks Road, Oxford OX1 3PU, UK*
[2]*Department of Physics, The Chinese University of Hong Kong, Shatin, New Territories, Hong Kong*

We study a simple competitive market, in which individual traders adapt their trading strategies according to past experience. Because of the limited knowledge available to them, they are forced to make decisions based on inductive, rather than deductive, thinking. We show that a population of competing traders with similar capabilities and knowledge will tend to self-segregate into opposing groups characterized by extreme behaviour. To be successful, a trader should behave in an extreme way, by either always copying *or* rejecting past trends in the market's history. Cautious traders tend to perform poorly.

Keywords: evolutionary dynamics; competitive market; microstructure; strategies; bounded rationality; trading

A reasonable view of a financial market is that it comprises a population of individual members (i.e. traders, institutions) who adapt their interactions, and hence behaviour, according to their past experiences (Bak 1997; Kauffman 1993; Stewart 1998; Arthur 1994; Johnson *et al*. 1998). (For examples of 'microscopic' economics-based models in the physics literature, see Bak *et al*. (1997), Ilinski & Stepanenko (1998), Caldarelli *et al*. (1997), Cont & Bouchaud (1997), Savit *et al*. (1997) and Amaral *et al*. (1998)). In particular, traders often find themselves competing for a limited resource, or to be in a minority. A simple example is the situation where there are more buyers than sellers: this tends to increase the price and hence benefits the sellers. Hence it would be better for a trader to be in the minority group of sellers.

Here we introduce a simple model for such an evolving population containing traders who compete to be in the minority. Only *partial* information about the market is available to the traders and no *a priori* 'best' strategy exists: traders are hence forced to make decisions based on inductive, rather than deductive, thinking. Each trader tries to learn from his/her past mistakes and will adjust his/her strategy in order to survive. We find that a population of such traders with similar capabilities will tend to polarize itself into opposing groups. Although a large number of possible strategies exist, the most successful traders are those who behave in an extreme way, by either copying, or rejecting, past trends in the market.

Inspired by Challet & Zhang (1997, 1998) we consider the model of an odd number N of traders repeatedly choosing whether to be in group '0' or group '1': choosing group 0 denotes choosing to buy a given asset while 1 denotes choosing to sell the asset. After every trader has independently chosen a group, the winners are those in the minority group, i.e. the group with fewer traders. If the winning group is 0, there are more sellers than buyers and hence the price drops at that time-step. If the winning group is 1, there are more buyers than sellers and the price rises at that time-step. The 'output' for each time-step is a single binary digit denoting the winning group, or equivalently whether the market price went down or up. Each trader is given a bit-string of length m containing the previous m outcomes. Each trader also has access to a common register or 'memory' containing the outcomes from the most recent occurrences of all 2^m possible bit-strings of length m. Consider $m = 3$; denoting $(xyz)w$ as the $m = 3$ bit-string (xyz) and outcome w, an example memory would comprise (000)1, (001)0, (010)0, (011)1, (100)0, (101)1, (110)0, (111)1. Following a run of three wins for group 0 in the recent past, the winning group was subsequently 1. Faced with a given bit-string of length m, it might seem sensible for a trader simply to predict the same outcome as that registered in the memory. The trader will hence choose group 1 following the next 000 sequence. If 0 turns out to be the winning group, the entry (000)1 in the memory is replaced by (000)0. If all N traders act in this way, however, the system will be inefficient since all traders will choose the same group and will hence lose; all the traders are spotting the same trends and assuming that they will continue indefinitely. Because of this, the trend fails to continue. A critical quality of a successful financial trader, for example, is the ability to follow a trend as long as it is valid, but to predict correctly when it will end. Hence we assign each trader a single number or 'strategy' p: following a given m-bit sequence, p is the probability that the trader will choose the same outcome as that stored in the memory, i.e. he/she will follow the current prediction, while $1 - p$ is the probability he/she will choose the opposite, i.e. he/she will reject the current prediction. Using the example memory, the trader will choose 1 (i.e. sell) with probability p after spotting the sequence 000, and 0 (i.e. buy) with probability $1 - p$.

Each time a trader gets into the minority (majority) group, he/she gains (loses) one point. If the trader's score falls below a value $d < 0$, then his strategy is modified, i.e. the trader gets a new p value which is chosen with an equal probability from a range of values, centred on the old p, with a width equal to R. Hence d is the amount of money a trader is willing to lose before modifying his/her strategy. Although this is a fairly crude 'learning' rule as far as machines are concerned (Sutton & Barto 1998), in our experience it is not too dissimilar from the way that humans actually behave in practice. Since $0 \leqslant p \leqslant 1$, we can for simplicity enforce reflective boundary conditions. Our conclusions do not depend on the particular choice of boundary conditions (see figure 1). Upon strategy modification, the trader's score is reset to zero. Changing R allows the way in which the traders learn to be varied. For $R = 0$, the strategies will never change (though the memory will). If $R = 2$, the strategies before and after modification are uncorrelated. For small R, the new p value is close to the old value.

As traders are constantly attempting to do the opposite of the others, a reasonable expectation is that they should eventually organize themselves so that their strategies are evenly spread within $0 \leqslant p \leqslant 1$. Alternatively, given that no *a priori* best strategy exists, one might expect that traders would be ambivalent as to whether a present trend will continue, and hence cluster around $p = \frac{1}{2}$. Surprisingly, the opposite is true. Figure 1*a* shows the frequency distribution $P(p)$ at large time. The distribution $P(p)$ eventually becomes peaked around $p = 1$ and 0 regardless of the initial $P(p)$ distribution; these p values, respectively, correspond to always or never following

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

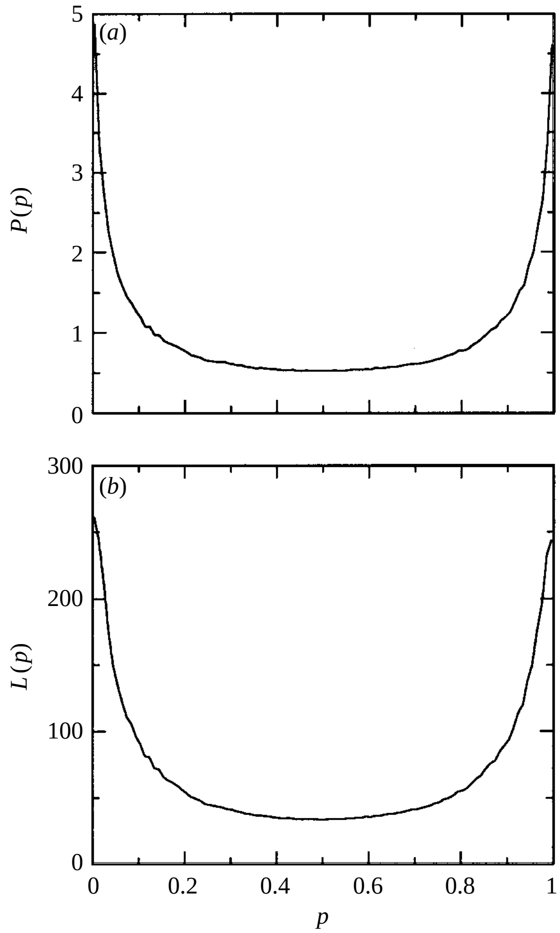

Figure 1. Distribution of (a) strategies $P(p)$ at large times. At $t = 0$, $P(p)$ is chosen to be flat. (b) Corresponding lifespans $L(p)$. $R = 0.2$, $N = 101$, $d = -4$ and $m = 3$.

what happened last time. The lifespan $L(p)$, defined as the average length of time a strategy p survives between modifications, shows similar behaviour. Henceforth we denote $P(p)$ and $L(p)$ as representing the long-time limits. If we consider the game simply as a random walk, with individual traders deciding randomly which group to choose, we would expect the mean number in group 0 or 1 to be $\frac{1}{2}N$ with a standard deviation of $\sqrt{N/4}$. At each time-step, the net number of points awarded will therefore be $-\sqrt{N}$. The average lifespan would be $d\sqrt{N}$. The observed average lifespan is indeed proportional to $d\sqrt{N}$. However, the average value of the $L(p)$ in figure 1b is larger than $d\sqrt{N}$ by a factor of approximately 2 for $d = -4$, confirming that the traders are organizing themselves better than randomly. Furthermore, the root-mean-square (RMS) separation of the strategies is higher than the value for uniform $P(p)$, indicating the desire of traders to do the opposite of the majority. It increases with N due to increased possibilities for self-organization. Even when R is large, and the strategy values are hence picked randomly upon modification, the RMS strategy separation remains high. The RMS strategy separation and the average value of $L(p)$ are typically maximal at $R \sim 0.5$; this is a compromise between a lack of learning when $R \sim 0$ and excessive strategy modification for large R. We also note that the standard deviation of the actual attendance time-series for group 0 (or group 1) is less than that obtained for traders choosing via independent coin-tosses: this again confirms that the system is organizing itself better than random.

Varying the length of the bit-string m has little effect on $P(p)$ and $L(p)$: since all traders have similar capabilities and available information, these benefits tend to cancel out. It is what each trader decides to do with the common knowledge that matters ($p = 0, 1$ traders outperform $p = \frac{1}{2}$ traders). Similarly if the memory is not updated dynamically according to the recent outcomes as discussed earlier, but is instead kept constant (i.e. time independent) or is randomly chosen at each time-step, then $P(p)$ and $L(p)$ are also essentially unchanged. Once again, the memory is common to all traders and hence all traders agree on the current prediction: no trader hence has any relative advantage in terms of available information (Cavagna (1998) proposes a similar result for the model of Challet & Zhang (1997, 1998)). It has been shown for the basic minority game (Manuca *et al*. 1998), in contrast to the claim in Cavagna (1998), that the memory is relevant since it can introduce hidden correlations into the winning-group time-series. This point will be discussed in detail for the present model elsewhere.

We now provide some analytic analysis. The simplest example of our system contains $N = 3$ traders i, j, k with memory m and three discrete p values $p = 0, \frac{1}{2}, 1$. (The fact that $N < 3$ is impossible suggests that our system contains the level of complexity typically associated with three-body, versus two-body, problems). All traders agree on the current prediction (say 0). Trader i will choose 0 or 1 with probability p_i and $1 - p_i$, respectively. Likewise for j (p_j) and k (p_k). The 2^3 possible decisions for ijk are 000, 001, 010, 100, 110, 101, 011, 111. There are $3^3 = 27$ possible configurations (p_i, p_j, p_k). For a given (p_i, p_j, p_k), the eight possible decisions yield the expected gain for the traders. For example, for $(p_i, p_j, p_k) = (0, 0, \frac{1}{2})$, i and j both choose 1 while k chooses 0 with probability $\frac{1}{2}$. Hence k wins with probability $\frac{1}{2}$ whereas i and j both lose. The net number of points gained per trader per turn, given by the points awarded minus the points deducted, is -1 for i, -1 for j and 0 for k. The total is hence -2. Given that the maximum is -1 (there is a maximum of one winner) we see that $(0, 0, \frac{1}{2})$ is not optimal.

Table 1 shows the various configuration types, or classes. The last column shows the average points per trader: $[-\frac{1}{2}]$ for class (i) implies the average trader loses $-\frac{1}{2}$ points per turn, and would hence modify his/her strategy after time $2d$. Such strategy modification allows the system to sample the 27 configurations. Classes (vi)–(viii) are optimal, having maximum points. To obtain the average distribution $P(p)$ and $L(p)$, we must average over all 27 configurations. Since some classes are more favourable (i.e. more points), we should weight the distributions in an appropriate way. In the extreme case of large weighting, we include only the optimal classes (vi)–(viii), yielding

$$P(0) : P(\tfrac{1}{2}) : P(1) = 2.5 : 1 : 2.5 \quad \text{and} \quad L(0) : L(\tfrac{1}{2}) : L(1) = 5 : 1 : 5.$$

For zero weighting, we consider instead the system as visiting all configurations with equal probability regardless of points gained per trader; such a zero-weight averaging is similar to that for the microstates in a gas within the microcanonical ensemble

PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY A MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Table 1. *Configuration classes showing the distribution of the three traders (each denoted by x) and the average points awarded per time-step for each strategy-value p*

(Also given are the number of distinct configurations per class, and the average number of points per trader per time-step.)

class	$p=0$	$p=\frac{1}{2}$	$p=1$	no. configs.	average pts/trader
(i)	—	$xxx[-\frac{1}{2}][-\frac{1}{2}][-\frac{1}{2}]$	—	1	$[-\frac{1}{2}]$
(ii)	$x[-\frac{1}{2}]$	$xx[-\frac{1}{2}][-\frac{1}{2}]$	—	3	$[-\frac{1}{2}]$
(iii)	$xx[-1][-1]$	$x[0]$	—	3	$[-\frac{2}{3}]$
(iv)	$xxx[-1][-1][-1]$	—	—	1	$[-1]$
(v)	—	—	$xxx[-1][-1][-1]$	1	$[-1]$
(vi)	$x[1]$	—	$xx[-1][-1]$	3	$[-\frac{1}{3}]$
(vii)	$xx[-1][-1]$	—	$x[1]$	3	$[-\frac{1}{3}]$
(viii)	$x[0]$	$x[-1]$	$x[0]$	6	$[-\frac{1}{3}]$
(ix)	—	$xx[-\frac{1}{2}][-\frac{1}{2}]$	$x[-\frac{1}{2}]$	3	$[-\frac{1}{2}]$
(x)	—	$x[0]$	$xx[-1][-1]$	3	$[-\frac{2}{3}]$

and yields

$$P(0):P(\tfrac{1}{2}):P(1)=1:1:1 \quad \text{and} \quad L(0):L(\tfrac{1}{2}):L(1)=1:1:1.$$

For an intermediate case, whereby all classes are weighted by the average points per trader, we obtain

$$P(0):P(\tfrac{1}{2}):P(1)=1.1:1:1.1 \quad \text{and} \quad L(0):L(\tfrac{1}{2}):L(1)=1.5:1:1.5.$$

In fact, any sensible weighting that favours the more profitable configurations yields a non-uniform $P(p)$ and $L(p)$ as observed numerically. This implies that the population, by self-segregating, has also managed to self-organize itself around the most profitable configurations. We emphasize that the system is dynamic since the membership of the various configurations is constantly changing (i, j and k interdiffuse) but $P(p)$ remains essentially constant. For general N we can loosely think of i, j, k as three equal-size groups of like-minded traders.

In summary, we have shown that an evolving population of traders with similar capabilities and information will self-segregate. To flourish in such a population, a trader should behave in an extreme way ($p=0$ or $p=1$).

We thank R. Jonson, D. Challet and P. Binder for discussions.

References

Amaral, L. A. N., Buldyrev, S. V., Halvin, S., Salinger, M. A. & Stanley, H. E. 1998 *Phys. Rev. Lett.* **80**, 1385.

Arthur, W. B. 1994 *Am. Econ. Assoc. Papers Proc.* **84**, 406.

Bak, P. 1997 *How nature works.* Oxford University Press.

Bak, P., Paczuski, M. & Shubik, M. 1997 *Physica* A **246**, 430.

Caldarelli, G., Marsili, M. & Zhang, Y. C. 1997 *Europhys. Lett.* **40**, 479.

Cavagna, A. 1998 Preprint, cond-mat/9812215. Available on e-print server at http://xxx.lanl.gov.

Challet, D. & Zhang, Y. C. 1997 *Physica* A **246**, 407.

Challet, D. & Zhang, Y. C. 1998 *Physica* A **256**, 514.

Cont, R. & Bouchaud, J. P. 1997 Preprint, cond-mat/9712318. Available on e-print server at http://xxx.lanl.gov.

Ilinski, K. N. & Stepanenko, A. S. 1998 Preprint, cond-mat/9806138. Available on e-print server at http://xxx.lanl.gov.

Johnson, N. F., Jarvis, S., Jonson, R., Cheung, P., Kwong, Y. R. & Hui, P. M. 1998 *Physica* A **258**, 230.

Kauffman, S. 1993 *The origins of order: self-organization and selection in evolution*. Oxford University press.

Manuca, R., Li, Y., Riolo, R. & Savit, R. 1998 Preprint, adap-org/9811005. Available on e-print server at http://xxx.lanl.gov.

Savit, R., Manuca, R. & Riolo, R. 1997 Preprint, adap-org/9712006. Available on e-print server at http://xxx.lanl.gov.

Stewart, I. 1998 *Life's other secret*. Wiley.

Sutton, R. S. & Barto, A. G. 1998 *Reinforcement learning: an introduction (adaptive computation and machine learning)*. MIT Press.