

Inter-pattern speculation: beyond minority, majority and \$-games

Damien Challet

*Nomura Centre for Quantitative Finance, Mathematical Institute,
Oxford University, 24–29 St Giles’, Oxford OX1 3LB, United Kingdom**

(Dated: February 24, 2005)

A new model of financial market is proposed, based the sequential and inter-temporal nature of trader-trader interaction. In this pattern-based speculation model, the traders open and close their positions explicitly. Information ecology can be precisely characterised, and is strikingly similar to that of the Minority Game. Naive and sophisticated agents are shown to give rise to very different phenomenology.

PACS numbers:

Agent-based modelling is a way to mimic financial market phenomenology [1–9] that focuses on individual behaviour. Because this approach requires to find the right and delicate balance between simple and non-trivial rules, the complexity of the agents varies tremendously from model to model. At the same time, such models bring insights into financial market dynamics as long as they are amenable to analysis. As a consequence, the best strategy for building and studying this kind of models is to start from well-understood specimens, and either to push their boundaries, or to borrow their framework and create a new breed of models. For instance the Minority Game (MG) can be exactly solved despite being highly non-trivial [10, 11]. Many of its variants can still be solved exactly, including market-oriented games that are able to reproduce a subset of market phenomenology [5, 9, 12]. [29]

Although minority mechanisms are at work in financial markets [13, 14], as we shall see below, the latter are more complex, mostly because the gains of traders are intrinsically intertemporal. In real life, one first opens a position, waits, and then closes it. [30] Here I extend the discussion of Ref. [14] and show in what limits each mechanism is present in financial markets. Then, I propose a new simple model of pattern-based speculation that contains each of them. As one takes inspiration from MG framework, one can immediately implement some of its extensions known to be relevant to financial markets.

NATURE OF FINANCIAL MARKETS

The nature of financial markets is puzzling. On the one hand, they are very competitive, and not everybody can earn money in them, implying implicitly a minority game. On the other hand it is common knowledge that one wishes to anticipate the crowd; one might conclude from that statement that they are majority game, which is also incorrect: they are best described as an “escape game” [15], which contains two minority games, and one delayed majority game.

The main reason why this point is still somewhat un-

clear after so many years is due to the intertemporal nature of financial markets: all the traders are not active at the same time, hence perfect synchronisation cannot be achieved; even worse, all the trade orders arrive sequentially, hence simultaneity is a theoretician’s phantasm at best —which is rarely recognised in the literature. Let us review carefully the process of opening and closing a position.

Placing oneself in event time, which increases by one unit whenever an order is placed, makes the discussion easier. At time t_i , trader i decides to open a position $a_i = \pm 1$ by placing a market order. [31] His physical reaction time results in a delay in event time of δt_i which varies from transaction to transaction; the transaction takes place at time $t_i + \delta t_i$ and log-price $p(t_i + \delta t_i)$. [32] During that time period, the order book changes; as a consequence $p(t_i + \delta t_i)$ is likely to differ from $p(t_i)$:

$$p(t_i + \delta t_i) = p(t_i) + \sum_{t_j=t_i}^{t_i+\delta t_i} I[a(t_j), t_j] \quad (1)$$

where I is the price impact function, and $a(t_j)$ is the sign and size of the order at time t_j . In principle, $a(t_j)$ could be a limit order that would modify I . For the sake of simplicity, only market orders and linear impact $I(x) = x$ [16] will be considered here. In that case,

$$p(t_i + \delta t_i) = p(t_i) + A(t_i) \quad (2)$$

where

$$A(t_i) = \sum_{\tau=0}^{\delta t_i} \sum_{j=1}^N a_j \delta_{t_i+\tau, t_j+\tau_j}, \quad (3)$$

where the sum \sum_j is over all the traders possibly interested in trading this particular asset. Equation (3) means that the group of traders with which trader i interacts through market impact is different at each transaction, and that among this group, everyone has a different but partially overlapping interacting group.

The position is held for some time. At t' , the trader decides to close his position, obtaining for the same reason

as before $p(t' + \delta t'_i)$. His real payoff is

$$a_i[p(t' + \delta t'_i) - p(t + \delta t_i)] = -a_i A(t_i) + a_i \left[\sum_{t_i < t < t'_i}^t A(t) \right] - (-a_i) A(t'_i). \quad (4)$$

The first and the last terms come from market impact and reaction time. They are minority-game payoffs: one's market impact is reduced if one acts against the crowd (in this case, this means taking an action opposite to the majority of orders executed during the time delay). The central term represents the capital gain that could be achieved without reaction time nor market impact. It describes a *delayed* majority game, that is, a majority game to which trader i *does not take part*: whereas $A(t_i)$ and $A(t'_i)$ contain a contribution of trader i , the middle term does not.

The nature of financial markets depends therefore on the trading frequency and reaction time of each trader, and on the activity of the market: the relative importance of minority games decreases as the holding time $t'_i - t_i$ increases; reversely, somebody who trades at each time step plays a minority game. Interestingly, this is consistent with the behaviour of market makers who try and stabilise the price so as to minimise inventory risk, thereby inducing a mean-reverting behaviour. This is also precisely what minority players do.

A comparison between this market mechanism and the $\$$ -game [6, 7] is in order. In the $\$$ -game, the traders make a transaction at each time step and the payoff of trader i is $a_i(t)A(t+1)$. Replacing $t+1$ by $t+\delta t$ and assuming that trader i opens a position at time t and closes it at time $t+\delta t$, makes the $\$$ -game payoff look like Eq. (4) without market impact nor reaction time, i.e. without minority games. There are two possible ways for that kind of payoff to appear. First, if one knows in advance one's exact impact (that includes the reaction time), or if one's reaction time is negligible, which only happens for infrequently traded stocks and if the size of the market order is smaller than the volume available at the best price, provided that the best prices do not change during the submission of the order. The other possibility is that the holding time δt is very large, in which limit market impact becomes much smaller than the price return δt . In both cases, there is however an inconsistency: in the $\$$ -game, $A(t+1)$ also contains the contribution of trader i as it forces the traders to be active at each time step and not to hold and close their positions explicitly.[33] Finally, if the agents have expectations on the nature of the market, that is, on the middle term of Eq (4), the $\$$ -game payoff involves only one time step, and is a minority payoff for contrarians and a majority payoff for trend-followers [13].

All the above discussion clearly shows the need for a model whose strategy space would be as simple as that

of the Minority Game, while allowing the traders to hold positions.

THE BASIC MODEL

The agents base their trading decisions on patterns, such as mispricing (over/undervaluation), technical analysis, crossing of moving averages, etc. The ensemble the patterns comprises all kinds of information regarded as possibly relevant by all the traders. Each trader recognises only a few patterns; every time one of his patterns arises, he decides whether to open a position. How to close a position is a matter of more variations: one can assume a fixed-time horizon, stop gains, stop losses, etc. In this paper, a trader closes a position only when he recognises a pattern, which is how people using crossings of moving averages behave, for example. Accordingly, each trader measures the average return between all the pairs of patterns that he is able to recognise.

Defining what "average return" precisely means brings in the well-known problem of measuring returns of trading strategies in a backtesting, i.e., without actually using them. This is due to market impact and results usually in worse-than-expected gains when a strategy is used. Estimating correctly one's market impact is therefore a crucial but hard aspect of backtesting. There are two types of delusions in the present model. The first, temporal inaccuracy, is due to the over/underestimation of reaction time. Self-impact on the other hand is the impact that a real trade has on the price and market, which is not present in data used for backtesting. Both cause imprecision in estimating returns of strategies not being used and, accordingly, both are actively investigated by practitioners.

Mathematically, N traders can submit buy orders (+1) or sell orders (-1). They base their decisions on patterns $\mu = 1, \dots, P$, which classify the state of the market. Each trader $i = 1, \dots, N$ is able to recognise S patterns $\mu_{i,1}, \dots, \mu_{i,S}$, and is active, i.e., may wish to buy or sell one share of stock, only when $\mu(t) \in \{\mu_{i,1}, \dots, \mu_{i,S}\}$. The kind of position taken ($a_i(t) \in \{0, \pm 1\}$) is determined by the cumulative price return between two consecutive occurrences of patterns. The time unit is that of pattern change, i.e., at each time step, $\mu(t)$ is drawn at random and uniformly from $1, \dots, P$. The duration of a time step is assumed to be larger than the time needed to placed an order. The order in which agents' actions arrive is disregarded here, although it would be straightforward to take it into account. Therefore, at time t , the excess return $A(t) = \sum_{i=1}^N a_i(t)$ results in a (log-) price change of

$$p(t+1) = p(t) + A(t). \quad (5)$$

$p(t+1)$, not $p(t)$ is the price actually obtained by the people who placed their order at time t . There are therefore

various ways to compute returns between two patterns. Assume that $\mu(t) = \mu$ and that $t' > t$ is the first subsequent occurrence of pattern ν : $p(t') - p(t)$ is the price difference that does not take into account price impact, whereas $p(t' + 1) - p(t + 1)$ is the realistic price difference. The cumulative price return $U_{i,\mu \rightarrow \nu}$ between pattern μ and ν evolves according to

$$U_{i,\mu \rightarrow \nu}(t' + 1) = U_{i,\mu \rightarrow \nu}(t) + p(t' + 1) - p(t + 1) - (1 - \alpha_i(t))\zeta_i[A(t') - A(t)] \quad (6)$$

ζ is the naivety factor: *simplex*, or naive, agents have $\zeta_i = 1$ and fail to take reaction time into account properly, while *sapiens*, or sophisticated, agents have $\zeta_i = 0$ and compute perfectly the average price return. Finally, the variable $\alpha_i(t)$ states whether the agent has opened a position at time t ($\alpha_i(t) = 1$), or stayed outside of the market ($\alpha_i(t) = 0$). When an agent trades, he perceives perfectly the effect of his reaction time whatever ζ_i . In practice, $\zeta_i \neq 0$, and can be of any sign and value. This is because estimating reaction time exactly is impossible: even if real traders are often acutely aware of its importance, they always over or underestimate it.

An agent only trades between his E best pairs of patterns, where $E \leq S - 1$ is his maximal exposure. In the following, the simplest case $S = 2$ is analysed, and $E = 1$. When $\mu(t) = \mu_{i,1}$, trader i closes his position if he has an opened one. Then depending on the sign of $U_{i,\mu_1 \rightarrow \mu_2}(t)$, he buys one share ($U_{i,\mu_1 \rightarrow \mu_2}(t) > 0$) or sells one share ($U_{i,\mu_1 \rightarrow \mu_2}(t) < 0$) [34], and holds his position until $\mu(t) = \mu_{i,2}$. [?]

Thus the basic model has four parameters so far: N , S , P , and E , and retains from the MG the heterogeneity of the agents, limited cognition abilities, and the distinction between naive and sophisticated agents.

Relevant quantities include the variance of A

$$\sigma^2 = \frac{\langle A^2 \rangle - \langle A \rangle^2}{P}, \quad (7)$$

the naive predictability

$$J = \frac{1}{P^2} \sum_{\mu, \nu, \mu \neq \nu} \langle A(t) | \mu \rightarrow \nu \rangle^2, \quad (8)$$

where $\langle A(t) | \mu \rightarrow \nu \rangle$ stands for the average price difference between the occurrence of μ at time t and the next occurrence of ν . J measures predictability that the naive agents hope to exploit. Another closely related quantity is

$$K = \frac{1}{P^2} \sum_{\mu, \nu} \langle A(t + 1) | \mu \rightarrow \nu \rangle^2 \quad (9)$$

is the exploitable predictability that takes into account price impact. Finally, a measure of price impact predictability is given by the average return conditional to




FIG. 1: Price time series with $\epsilon = 0$ ($N = 1000$, $P = 32$, $S = 2$, $\zeta = 0$)

a given pattern

$$H = \frac{\sum_{\mu} \langle A | \mu \rangle^2}{P}. \quad (10)$$

Before carrying out numerical simulations, one should keep in mind that the price is bounded between $-N$ and $+N$, since the traders are not allowed to have an exposure larger than 1. Assume that $\epsilon = 0$, and that all the scores are have random and small initial valuation (otherwise nobody trades in the first place). One observes in such case beginnings of bubbles or anti-bubbles, the price rising or decreasing to $\pm N$, and then staying constant. Indeed, the price increase/decrease is echoed in the scores of all the agents, which have all the same sign, therefore all stipulate the same action. The price is stuck (Fig 1), as nobody wants to close its position, because everybody is convinced that the price should carry on on its way up/down. This phenomenon is found for all values of ζ .

The way to break the bubble is to make the agents compare the average returns to a risk-free interest rate $\epsilon > 0$ [6, 9, 12]. When an agent detects a pattern at time t , he opens a position only if $|U| > \epsilon t$. If the price has reached $\pm N$, the time needed to break the bubble is at least of order $N/(\epsilon)$.

Figure 2 illustrates the typical price time series for $\epsilon > 0$: first a bubble, then a waiting time until some traders begin to withdraw from the market. The price goes back to 0 and then fluctuates for a while. How these fluctuations are interpreted by the agents depends on ζ and ϵ : increasing ζ makes it more difficult for the agents to understand that they should refrain from trading, because they are not capable of disentangling their own contribution from these fluctuations. Accordingly, the larger ζ , the later the agents withdraw from the market, and the smaller ϵ , the longer it takes (Fig. 3). In this figure, the maximum number of iteration was capped at 10^6 ; naive agents need therefore a very long time before withdrawing if ϵ is small. The scaling $T_w \propto N/\epsilon$ holds only for small ζ and ϵ . For large ϵ , ζ does not matter much.




FIG. 2: Price time series with $\epsilon = 0.01$ ($N = 1000$, $P = 32$, $S = 2$). Black lines are for *simplex* traders and green lines are for *sapiens* traders.




FIG. 3: Scaled time to withdraw as a function of ϵ . $P = 10$, $N = 50$ (full symbols) and 100 (empty symbols), $\zeta = 0$ (circles), 0.75 (squares), 0.85 (diamonds) and 1 (triangles); average over 500 samples

All the agents eventually withdraw from the market. This makes complete sense, as there is no reason to trade. Naivety results in a diminished ability to withdraw rapidly enough from a non-predictable market, and, as a by-product, larger price fluctuations. This is consistent with the fact that volatility in real markets is much larger than if the traders were as rational as mainstream Economics theory assumes [20]. Naivety, an unavoidable deviation from rationality, is suggested here as one possible cause of excess volatility.

Noise traders

First, as the traders try and measure average returns, adding noise to the price evolution ($A(t) \rightarrow A(t) + N_n \eta(t)$ where $\eta(t)$ is uniformly distributed between -1 and 1) does not provide any trading opportunity, but makes it more difficult to estimate precisely average returns. Accordingly, the agents should withdraw later, and the larger ζ , the later. This is precisely what happens: Fig 4 reports the average behaviour of T_w for sophisticated and




FIG. 4: Scaled time to withdraw as a function of N_n . $P = 10$, $N = 100$, $\epsilon = 0.1$, $\zeta = 0$ (circles) and 1 (squares); average over 500 samples

naive agents. One can therefore reinterpret this figure as a clue that naive agents are blinded by the fluctuations that they produce themselves.

Market impact heterogeneity

Real traders are heterogeneous in more than a way. In a population where each trader has his own ζ_i , people with a small ζ_i evaluate gain opportunities better. This also means that the impact of people with large ζ_i provides predictability to the agents with a lower ζ_i , and therefore the former are exploited by the latter, giving a good reason to trade to sophisticated agents as long as naive agents actively trade in the market.

MARKET INFORMATION STRUCTURE

Up to now, the model showed how hard it is not to trade. But how hard is it to make the agents want to trade?

As in the MG [23], and following the picture proposed in Ref. [21], one can distinguish between speculators and producers. The latter do not care much about market timing and adjust more slowly their strategies because they need the market for other purposes than mere speculation. A simple way to include N_p producers in the model is to assume that they induce a bias in the excess demand that depends on $\mu(t)$, i.e., $A(t) = A_{\text{prod}}^{\mu(t)} + A_{\text{spec}}(t)$. Each producer has a fixed contribution ± 1 to A_{prod}^{μ} , hence $A_{\text{prod}}^{\mu} \propto \sqrt{N_p}$. In that way, the amount of market impact predictability is well controlled.

If there is no information structure in the market, i.e., if $\mu(t)$ does not depend at all on past patterns, the effect of producers is akin to that of noise traders, hence, the speculators cannot exploit the predictability left by the producers. This is because the speculators need some




FIG. 5: Volatility σ^2 (circles), price impact H/P (squares), naive (down triangles) and sophisticated (up triangles) predictability, scaled fraction of active speculators (diamonds), and average gain per agent (right triangles); $P = 10$, $N_p = 100$, $\epsilon = 0.05$; empty symbols: average over 10000 samples of speculators; filled symbols: average over 1000 samples

temporal correlation between occurring biases in order to exploit them, which happens when the transitions between market states are not equiprobable, i.e., when the transition matrix between patterns W is such that $W_{\mu \rightarrow \nu} \neq 1/P$. This assumption is supported by an empirical study [22] which determined states of the market with a clustering method, and found that the transitions between the states is highly non-random and has long-term memory which we neglect here. Numerically, we chose to fix $W_{\mu \rightarrow \nu}$ to a random number between 0 and 1 and then normalised the transition probabilities; the variance of $W_{\mu \rightarrow \nu}$ is a parameter of the game, which controls the amount of correlation between biases induced by the producers.

Adding more producers increases linearly the values of predictability J and K , as well as H and σ^2 . Then, keeping fixed the number of producers and their actions, adding more and more sophisticated speculators first decreases σ^2 , H , J and K , which reach a minimum and then increase again (Fig 5); this seems to occur at $N \simeq P^2$, where on average every pair of patterns is attributed to some speculator. However, the average total number of speculators in the market reaches a plateau. The fact that the fluctuations σ^2 increase for $N > P^2$ means that the agents with the same pair of patterns enter and withdraw from the market in a synchronous way. Avoiding such synchronicity can be realized by letting the agents have heterogeneous ϵ_i or ζ_i . The average gain of the speculators becomes negative for $N < P^2$; if evolving capitals were implemented, less speculators would survive, which would lower N , thus keeping the average gain at or above zero.

The relationship between the producers and the spec-




FIG. 6: Number of active speculators as a function of time ($P = 10$, $N = 1000$, $N_p = 100000$, $\epsilon = 0.001$)

ulators can be described as a symbiosis: without producers, the speculators do not trade; without speculators, the producers lose more on average, as the speculators (with evolving capital) reduce H . This picture, remarkably similar to that of the Minority Game with producers [23–25], justifies fully in retrospect the study of information ecology in Minority Games. But in contrast to the MG speculators, the agents here are not able to remove much predictability. A more refined mapping between the state of the world and the perceived patterns such as the one proposed in Ref [26] might improve information efficiency.

Guided by the conceptual similarity between the grand canonical MG with producers and the present model, we expect to see volatility clustering when the number of agents is large, and if $\epsilon > 0$; more generally, as argued in Ref. [27], any mechanism responsible for switching between being in and out of the market should produce volatility clustering. This is exactly what happens, as shown by Fig. 6, where the volume N_{act} has clearly a long term memory. Whether this model is able to reproduce faithfully real-market phenomenology is still unclear.

CONCLUSION

This model provides a new simple yet remarkably rich market modelling framework that builds on the Minority Game. It is readily extendable, and many relevant modifications are to be studied. The similarity of information ecology between the present model and the MG is striking. Whether this model is exactly solvable is under scrutiny.

I thank T. Bochud and W. Shaw for their useful comments, and Wadham College for support.

Source code for this model is available at

-
- * Electronic address: challet@maths.ox.ac.uk
- [1] H. Levy, M. Levy, and S. Solomon, *Microscopic Simulation of Financial Markets: From Investor Behavior to Market Phenomena* (Academic Press, London, 2000).
- [2] T. Lux and M. Marchesi, *Nature* **397**, 498 (1999).
- [3] R. Cont and J.-P. Bouchaud, *Macroecon. Dyn.* **4**, 170 (2000).
- [4] D. Challet and Y.-C. Zhang, *Physica A* **246**, 407 (1997).
- [5] N. F. Johnson, M. Hart, P. M. Hui, and D. Zheng, *ITJFA* **3** (2000).
- [6] I. Giardina and J.-P. Bouchaud (2002).
- [7] J. V. Andersen and D. Sornette, *Eur. Phys. J. B* **31**, 141 (2003), [cond-mat/0205423](#).
- [8] R. Cont, F. Ghoulmie, and J.-P. Nadal, *Journal of Physics: Condensed Matter* (2005).
- [9] D. Challet, M. Marsili, and Y.-C. Zhang, *Physica A* **294**, 514 (2001).
- [10] D. Challet, M. Marsili, and R. Zecchina, *Phys. Rev. Lett.* **84**, 1824 (2000).
- [11] J. A. F. Heimerl and A. C. C. Coolen, *Phys. Rev. E* **63**, 056121 (2001).
- [12] D. Challet and M. Marsili, *Phys. Rev. E* **68**, 036132 (2003).
- [13] M. Marsili, *Physica A* **299**, 93 (2001).
- [14] D. Challet and T. Galla, *Quant. Fin* (2005), in print, [cond-mat/0404264](#).
- [15] F. Slanina and Y.-C. Zhang, *Physica A* **289**, 290 (2001).
- [16] J. D. Farmer, *Tech. Rep. 98-12-117*, Santa Fe Institute (1999).
- [17] J.-P. Bouchaud and M. Potters, *Theory of Financial Risks* (Cambridge University Press, Cambridge, 2000).
- [18] J. Y. Campbell, A. W. Lo, and A. C. MacKinlay, *The Econometrics of Financial Markets* (Princeton University Press, 1996).
- [19] M. M. Dacorogna, R. Geny, U. A. Miller, R. B. Olsen, and O. V. Pictet, *An Introduction to High-Frequency Finance* (Academic Press, London, 2001).
- [20] R. J. Shiller, *Am. Econ. Rev.* **71**, 421 (1981).
- [21] Y.-C. Zhang, *Physica A* **269**, 30 (1999).
- [22] M. Marsili, *Quant. Fin.* **2** (2002).
- [23] D. Challet, M. Marsili, and Y.-C. Zhang, *Physica A* **276**, 284 (2000), [cond-mat/9909265](#).
- [24] D. Challet, A. Chessa, M. Marsili, and Y.-C. Zhang, *Quant. Fin.* **1**, 168 (2000), [cond-mat/0011042](#).
- [25] Y.-C. Zhang (2001), [cond-mat/0105373](#).
- [26] J. Berg et al., *Quant. Fin* **1** (2001), [cond-mat/0101351](#).
- [27] J.-P. Bouchaud, I. Giardina, and M. Mézard, *Quant. Fin.* **1**, 212 (2001).
- [28] J.-P. Bouchaud, *private communication*.
- [29] However, market phenomenology in MGs still escapes analytical understanding.
- [30] Some authors have proposed other kinds of payoffs which attempt to remedy the shortcomings of pure minority games [6, 7, 14, 15]
- [31] It is an order to buy/sell immediately at the best price. More patient traders place limit orders at or beyond best prices. thus obtaining a better deal.
- [32] The smallest reaction time is around 1s for the DAX [28].
- [33] Since the payoff that involves two time steps, this model is not exactly solvable.
- [34] Selling shares that one does not own is called short-selling. Closing such a position consists in buying back the stock later at a hopefully lower price.
- [] The agents are allowed to close and open exactly the same kind of position at the same time step; for the sake of simplicity, we let them act that way, but discount the resulting volume.