From perceived risk to collective behavior in minority games

A. De Martino[1], I. Giardina[1], M. Marsili[2] and A. Tedeschi[1]

[1] INFM SMC and Dipartimento di Fisica, Università di Roma "La Sapienza", p.le A. Moro 2, 00185 Roma (Italy) Andrea.DeMartino@roma1.infn.it
[2] The Abdus Salam ICTP, Strada Costiera 11, 34014 Trieste (Italy) Matteo.Marsili@ictp.trieste.it

Summary. We analyze the effects induced by risk control in the Minority Game by allowing agents to modify their behavior according to the market conditions they perceive. First, we show that the macroscopic properties of the model are roughly unchanged if agents evaluate their trading strategies by balancing the expected profit and the associated risk averaged over sufficiently long times. Then we consider the case in which agents perceive the market as a Majority (rather than Minority) Game when the excess demand is below a given (fixed) threshold, mimicking a situation in which traders can switch from a fundamentalist to a trend-following conduct when the perceived risk is small. Now a crossover from a stable (fundamentalists dominated) to an unstable (trend-followers dominated) market regime appears as the agents' sensitivity to risk decreases. In the intermediate phase, Majority-like and Minority-like features coexist and the dynamics acquires several non-trivial traits.

1 Introduction

The Minority Game [1,2] was originally introduced along the lines of Arthur's El Farol problem [3] to model the inductive behavior of a population of heterogeneous agents competing for a limited resource. Despite its basic simplicity, even its most unpretentious versions required considerable analytic and numerical work to be elucidated in detail. As of today, one disposes of a coherent theoretical picture linking the key elements of the model's phenomenology – namely the production of exploitable information and the buildup of fluctuations – to static and dynamical phase transitions typical of disordered physical systems [4–7]. This mapping has allowed to clarify such issues as the role of 'market' impact [8,9] and to propose an explanation for the origin of several 'stylized facts' precisely in the vicinity of the critical point [10–12].

In the original setup, Minority Game players optimize the use of their trading strategies in time so as to maximize their expected profit subject to some fixed expectations about the game's outcome. In particular, it was shown in [13] that Minority Game's payoffs arise from a well defined market mechanism when agents behave as fundamentalists, i.e. assume that the price is close to a stationary state. It is reasonable to think that real traders may revise their expectations if they prove wrong or simply may want to weigh their decisions against other factors than the expected profit. For instance, in certain market regimes (e.g. bubbles) a trader could perceive the market as a Majority rather than Minority Game and consequently switch from a fundamentalist to a trend-following behavior. Similarly, in situations of high volatility traders would likely take into account the risk factor when choosing a trading strategy over another. An interesting question that naturally arises is how the macroscopic properties of the Minority Game would change if agents were allowed to modify their behavior and expectations according to the risk they perceive.

In this work, we introduce two extensions of the original setup that allow to tackle the problem above. We will first let agents evaluate the effectiveness of their trading strategies by confronting the expected profit with (a measure of) the associated risk. It turns out that if agents evaluate the expected risk over a sufficiently long time interval (i.e. if their memory is sufficiently long) the risk factor does not affect the global macroscopic picture. By contrast, a significant loss of efficiency may occur if the memory of agents is shorter. Then, we will study the case in which agents may switch from a fundamentalist to a herding conduct or vice versa, depending on the magnitude of price movements. The basic idea is that traders prefer to adopt a trend-following attitude, and thus perceive the market as a Majority Game, when fluctuations are small while they revert to fundamentals, and hence perceive the market as a Minority Game, when the price dynamics becomes more chaotic. In this case, one observes the formation of a crossover regime between a Minority-like market (where fundamentalist dominate) and a Majority-like market (where trend followers dominate) in which Minority-like and Majority-like features coexist. The former emerge when the number of players is large (more precisely: when the relative number of information patterns is small) while the latter are found when more agents interact. This picture differs significantly from that obtained for mixed Majority-Minority Games [14], whose macroscopic properties are to a large extent a linear combination of those of Minority and Majority Games, and suggests that a greater degree of behavioral adaptability may lead to a serious loss of collective efficiency.

For the sake of simplicity, we shall consider here only games with random external information and price-taking agents. To begin with, we will shortly review some basic definitions and results on Majority/Minority Games and then discuss the model with risk-controlling agents, which requires passing from the usual 'single' learning dynamics to two coupled processes. Then we will move on to the model with adaptive trend-followers and fundamentalists, whose starting point is a simple modification of Minority Game's payoff function (a similar but more complex case was studied in [15]). To conclude, we shall address some of the implications of these results together with the possible generalizations to more complicated and realistic cases.

2 Some basic facts about Minority/Majority Games

Models of interacting fundamentalists and trend-followers have been extensively studied in the past [16–20]. It is well known that the latter cause market instabilities and that stylized facts often appear in connection with switches between fundamentalist-dominated and trend-followers-dominated market regimes. In the context of Minority Games, we will focus on the interplay between trader behavior, volatility and informational efficiency.

In a Minority Game, each of N agents must formulate at every time step $t = 0, 1, \ldots$ a bid $b_i(t)$ that for simplicity we assume to be binary, $b_i(t) \in \{-1, 1\}$ (e.g. 'buy/sell'), and as a consequence receives the payoff $u_i(t) = -b_i(t)A(t)$, where $A(t) = \frac{1}{\sqrt{N}}\sum_{i=1}^{N} b_i(t)$ represents the excess number of buyers (or sellers) at time t. (The normalizing $\sqrt{N}$ factor has been introduced for later convenience.) Clearly, every agent aims at taking the minority decision. In a Majority Game agents receive instead the payoff $u_i(t) = b_i(t)A(t)$ and aim at taking the majority decision. The naïve connection between the game's outcome $A(t)$ and a price $\pi(t)$ is established by

$$\log \pi(t+1) = \log \pi(t) + \frac{A(t)}{L} \tag{1}$$

where the 'excess demand' $A(t)$ is linked to the price evolution via the 'liquidity' L [21, 22]. Random traders, namely agents who at each time step choose independently whether to buy or sell at random with equal probability, lead to a state where[3] $\langle A \rangle = 0$ and $\sigma^2 \equiv \langle A^2 \rangle = 1$. For inductive agents, the simplest possibility is to employ a probabilistic rule of the form [23]

$$\mathrm{Prob}[b_i(t) = b] = C \exp\left[\Gamma b p_i(t)\right] \qquad b \in \{-1, 1\} \tag{2}$$

where C is a normalization constant, $\Gamma > 0$ represents the learning rate of agents and $p_i(t)$ is a function accounting for the agent's expectations about what will be the winning choice (if $p_i(t) > 0$ then he/she will choose $b_i(t) = +1$ with higher probability), which is updated according to

$$p_i(t+1) = p_i(t) + \epsilon_i A(t) \tag{3}$$

where $\epsilon_i = 1$ for a Majority Game and $\epsilon_i = -1$ for a Minority Game. This choice has been investigated in [13] upon varying the fraction f of trend-followers. The general conclusion (see however [13] for a more careful discussion) is that expectations of the majority group are self-fulfilled, as the system behaves like a Minority Game (with $\langle A \rangle = 0$ and $\langle A(t)A(t+1) \rangle$ negative) for $f < 1/2$ whereas it behaves like a Majority Game (with $\langle A \rangle \neq 0$ (trends!) and $\langle A(t)A(t+1) \rangle$ positive) for $f > 1/2$.

[3] Here and in what follows $\langle \cdots \rangle$ stands for a time average in the steady state of the dynamics.

The situation becomes more interesting when agents have access to public information. The basic setup is usually as follows. One considers a population of N agents labeled by i and a set of P possible information patterns or states of the world labeled by μ. At each time step t, an information pattern $\mu(t)$ is drawn at random with equal probability from $\{1,\dots,P\}$ and is presented to agents. Based on $\mu(t)$, agents have to formulate a binary bid $b_i(t)\in\{-1,1\}$. To this aim, each agent is endowed with S trading strategies $\boldsymbol{a}_{ig}=\{a_{ig}^{\mu}\}$ ($g=1,\dots,S$), namely P-component vectors that map informations $\mu\in\{1,\dots,P\}$ into actions $a_{ig}^{\mu}\in\{-1,1\}$. Each component a_{ig}^{μ} of every strategy is assumed to be selected randomly and independently from $\{-1,1\}$ with equal probability for every i, g and μ at the beginning of the game and kept fixed (in physical jargon, strategies are 'quenched disorder'). Furthermore, agents evaluate the performance of their strategies using a valuation function that is initialized to the value $p_{ig}(0)$ at time $t=0$ and is updated at the end of every round. At each time step, every agent picks the strategy to use at time t, denoted by $\widetilde{g}_i(t)$, according to the rule

$$\mathrm{Prob}[\widetilde{g}_i(t)=g]=C\exp\left[\mathit{\Gamma} p_{ig}(t)\right]\qquad g=1,\dots,S \tag{4}$$

and formulates the corresponding bid $b_i(t)=a_{i\widetilde{g}_i(t)}^{\mu(t)}$. Finally, every agent updates his/her strategy valuations according to

$$p_{ig}(t+1)=p_{ig}(t)+\epsilon_i a_{ig}^{\mu(t)}A(t) \tag{5}$$

where as before $A(t)=\frac{1}{\sqrt{N}}\sum_{i=1}^{N}b_i(t)$ and $\epsilon_i=1$ for trend-followers and $\epsilon_i=-1$ for fundamentalists. In this case a further macroscopic quantity of interest arises besides $\langle A\rangle$ and σ^2, namely the amount of exploitable information[4] $H=\frac{1}{P}\sum_{\mu=1}^{P}\langle A|\mu\rangle^2$. It is easily understood that $H=0$ implies the absence of exploitable information, as the winning action is not predictable on the basis of the state of the world alone. By contrast, if $H>0$ the game's outcome time series contains information an external agent could exploit. Just like the original Minority Game [6], this model can be studied analytically in the limit $N\to\infty$ using methods of statistical mechanics (this is done in [14]). The most interesting case is that where the number of possible information patterns scales linearly with N. The macroscopic properties of the stationary state turn out to depend only on the relative number of information patterns $\alpha=\lim_{N\to\infty}\frac{P}{N}$ and on the fraction f of trend-followers. The typical picture obtained for $\mathit{\Gamma}=\infty$ (i.e. for $\widetilde{g}_i(t)=\arg\max_g p_{ig}(t)$) and $S=2$ is shown in Fig. 1. For $f<1/2$, a Minority-Game type of behavior is recovered, with an informationally inefficient phase ($H>0$) at high α separated by an efficient one ($H=0$) at low α. Fluctuations are small in the supercritical phase while for $\alpha<\alpha_c$ the stationary state depends strongly on the initial conditions of the

[4]The symbol $\langle\cdots|\mu\rangle$ denotes a time average in the steady state conditioned on the occurrence of the pattern μ.

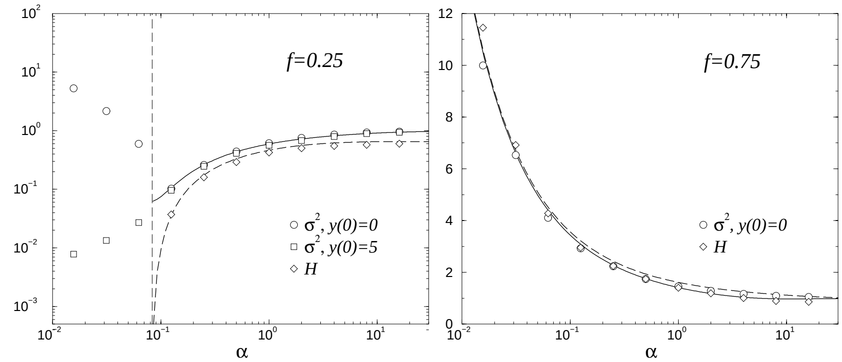

Fig. 1. Behavior of σ^2 and H vs α for $f = 0.25$ (left) and $f = 0.75$ (right). Markers represent results from numerical simulations with homogeneous initial conditions ($y(0) \equiv y_i(0) = \frac{1}{2}[p_{i1}(0) - p_{i2}(0)]$), averaged over 200 realizations of the trading strategies. The dashed vertical lines give the location of α_c (from theory). Continuous (resp. dashed) lines represent analytical approximations for σ^2 (resp. H), valid only for $\alpha > \alpha_c$. See [14] for details.

learning dynamics and both high-volatility and low-volatility states can occur. The transition point α_c decreases as f increases, as trend followers inject more and more exploitable information in the game. For $f > 1/2$, the system produces exploitable information ($H > 0$) for all α and fluctuations increase with f. The difference between σ^2 and H diminishes as f increases.and for $f = 1$ one has $\sigma^2 = H$. No dependence on initial conditions is observed. Again, the expectations of the majority group turn out to be fulfilled.

3 Agents controlling the risk in the Minority Game

Let us now consider a standard Minority Game with random external information and substitute the learning dynamics (5) with the following coupled processes:

$$p_{ig}(t+1) = p_{ig}(t) - a_{ig}^{\mu(t)} A(t) \tag{6}$$

$$r_{ig}(t+1) = r_{ig}(t) + A(t)^2 \tag{7}$$

Let the strategy choice rule be now given by

$$\mathrm{Prob}[\widetilde{g}_i(t) = g] = \widetilde{C} \exp\left\{\Gamma\left[p_{ig}(t) - \lambda\left(r_{ig}(t) - p_{ig}(t)^2\right)\right]\right\} \tag{8}$$

where $g = 1, \ldots, S$ and λ is a constant. This learning/choice dynamics is easily understood. The function p_{ig} accounts for the profit of agents while r_{ig}

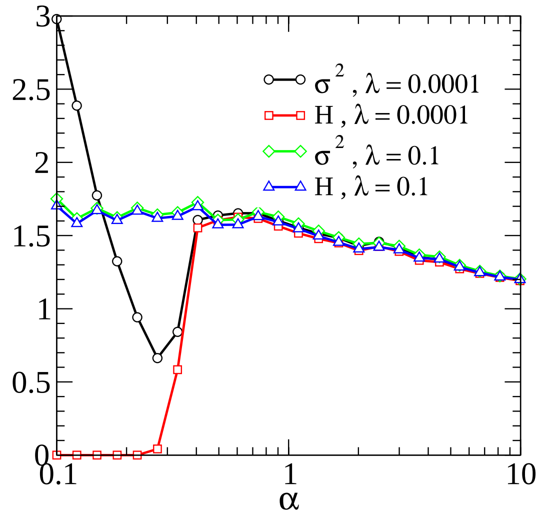

Fig. 2. Behavior of σ^2 and H vs α for $\lambda = 0.0001$ and $\lambda = 0.1$. Simulations performed with $\alpha N^2 = 16000$, with averages over 100 realizations of the trading strategies.

accounts for the square of the profit. In the long run, we expect the former to be related to the expected profit and the latter to be related to the second moment of the profit (i.e., in this case, to the volatility). Hence the term proportional to λ in (8) represents the variance of the profit. In their decision-making process, agents try to select the strategy that ensures the highest expected profit with the minimum variance. If the latter is taken as a measure of the risk associated to a certain strategy, we see that agents weigh the expected profit against the risk so that, in principle a profitable strategy may be discarded in favor of a less profitable strategy if the former is accompanied by a high risk. The parameter λ modulates the agents' sensitivity to risk: for $\lambda > 0$ agents are risk-averse, for $\lambda < 0$ they are instead risk-prone. A simple interpretation of λ is easily obtained if one assumes that agents learn the expected profit and the risk at different rates, say Γ and Γ'. Then λ represents the ratio Γ'/Γ of the learning rates if Γ is finite or the limit to which the ratio tends when $\Gamma \to \infty$. In principle, λ should be agent-dependent. We however analyze the case of fixed λ in order to keep the focus on the effects induced by the risk term.

The stationary state turns out to depend strongly on λ (see Fig. 2). One observes that for $\Gamma = \infty$ and any $\lambda > 0$ the magnitude of fluctuations increases drastically with respect to the standard Minority Game (which corresponds to $\lambda = 0$). However if λ is small enough the system is informationally efficient and Minority-Game like for α smaller than $\alpha_c \simeq 0.2$. Above α_c one finds that $H \simeq \sigma^2$, as in a Majority Game. As λ increases fluctuations increase but

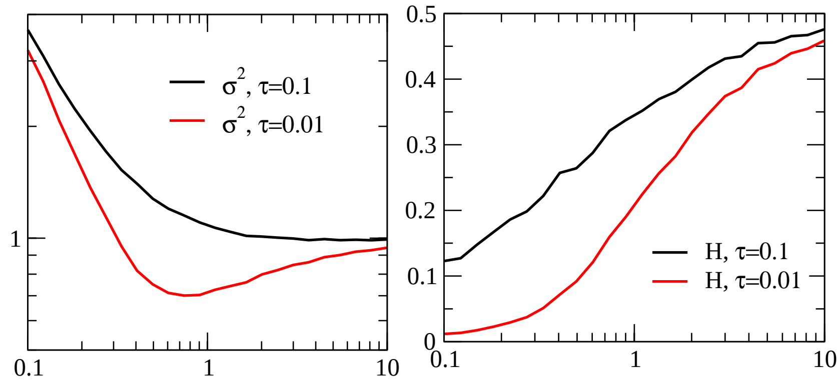

Fig. 3. Behavior of σ^2 and H vs α for $\lambda = 0.1$. Simulations parameters are as in Fig. 2.

the informationally efficient phase persists (with the same α_c) until $\lambda \simeq 0.1$, at which point it disappears. (It should however be said that the difficulty of performing reliable numerics for very small values of α prevents us from drawing sharp conclusions concerning the existence of an informationally efficient phase at low α.) We conclude that as agents learn the risk and the profit at increasingly comparable rates it becomes harder and harder for them to process the information efficiently and the steady state becomes more and more inefficient.

One could argue that real traders do not assess the risk of their trading strategies over the whole time span of their experience (it would result in a prohibitive inter-temporal optimization problem) but rather tend to perform time averages over restricted, though perhaps large, time intervals. A naïve way to implement these considerations consists in modifying (6,7) to

$$p_{ig}(t+1) = (1-\tau)p_{ig}(t) - \tau a_{ig}^{\mu(t)} A(t) \tag{9}$$

$$r_{ig}(t+1) = (1-\tau)r_{ig}(t) + \tau A(t)^2 \tag{10}$$

where the parameter $\tau > 0$ can be seen as the inverse time span over which agents assess the validity of their strategies. Notice that $\tau > 0$ corresponds also to a finite memory of initial conditions [24]. One sees in Fig. 3 that for $\lambda = 0.1$, i.e. when the learning rates for profit and risk are close (resulting in a highly inefficient state in the previous case), the stationary state develops a remarkable dependence on τ. If τ is small, the volatility σ^2 and the predictability H behave as in a standard Minority Game, although the strictly speaking efficient phase with $H = 0$ disappears. The system becomes more and more inefficient as τ increases, i.e. as agents evaluate averages over shorter and shorter time intervals and ultimately fluctuations become Majority-like. These simple models highlight on the one hand the robustness of Minority

Game's macroscopic picture with respect to the introduction of risk managing (hence more 'volatile') agents and on the other hand the relevance of the time scales over which risk is assessed. While several other generalizations may be introduced, models with risk control as the ones we introduce here are definitely worth a more detailed investigation.

4 Adaptive trend-followers and contrarians

Now let us introduce the case in which agents may switch from one group to the other. There are of course many ways to allow for this. One is to couple the learning process (5) with a dynamics involving ϵ_i (and thus f). However, if the time scales over which ϵ_i varies are sufficiently long compared to those required for (5) to reach a steady state, then the only states with self-fulfilling expectations will be those with $f = 0$ and $f = 1$. Another possibility is to generalize (5) as

$$p_{ig}(t+1) = p_{ig}(t) + a_{ig}^{\mu(t)} F_i[A(t)] \tag{11}$$

(clearly, $F(A) = -A$ for a Minority Game whereas $F(A) = A$ for a Majority Game) and consider functions F_i embodying the way in which agent i perceives the performance of their g-th trading strategy in the market. For simplicity we shall henceforth assume that $F_i = F$ for all i. Let us choose

$$F(A) = [\Theta(\eta - |A|) - \Theta(|A| - \eta)]\, A \tag{12}$$

with $\eta \geq 0$ a constant. ($\Theta(x)$ denotes Heaviside's step function.) Agents will now perceive the market as a Majority Game, hence behave as trend-followers, when $|A| < \eta$ while they will perceive it as a Minority Game, hence behave as fundamentalists, when $|A| > \eta$. The idea is that when the excess demand is large the risk perceived by agents is large and a fundamentalist attitude is preferred. Thus η mimics a risk threshold: the smaller it is, the more agents are sensitive to risk or the larger the risk they perceive. Clearly, a pure Minority (resp. Majority) Game is recovered for $\eta = 0$ (resp. $\eta \to \infty$).

A few remarks are in order.

1. This mechanism is expected to induce a feed-back in the dynamics of the excess demand: when it is small, trend-followers dominate and drive it to larger values until fundamentalists eventually take over and drive it back to smaller values.

2. It is reasonable to think that η should fluctuate in time and possibly be coupled to the system's performance. A possible microscopic mechanism is the following. When η is large a high volatility is to be expected as agents are more likely to behave as trend-followers. As a consequence, they should likely reduce their threshold since the market is risky; however, for small η fundamentalists are expected to dominate and the game should acquire a Minority character. Hence the predictability will be smaller and there will be less profit opportunities. Agents may then decide to adopt a larger threshold to

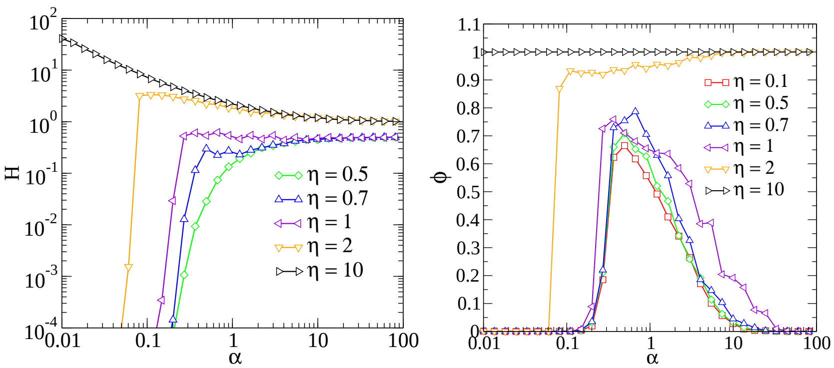

Fig. 4. Behavior of H (left) and ϕ (right) vs α for different η (top figures) and the reverse (bottom figures). Simulations performed with $\alpha N^2 = 16000$, with averages over 100 realizations of the trading strategies.

seek for convenient speculations on a wider scale. If these two competing effects are appropriately described by an evolution equation for η, the system should self-organize around an 'optimal' value of the parameter. However such a time evolution should take place on time-scales much longer than those which the model addresses (intra-day/daily trading) and hence it is reasonable to study the case of fixed η.

3. Similarly, η should be agent-dependent. While we do not consider this case here, we notice that the existence of such risk thresholds can in principle be verified from the analysis of high-frequency order book data, for instance from the probability that an agent places a certain order conditioned on the occurrence of a given price increment. Some work along these lines is at present in progress.

Numerical results for the relevant macroscopic quantities have been obtained for $\Gamma = \infty$ and $S = 2$, with initial conditions of 11 set at $p_{ig}(0) = 0$ for all i and g. In Fig. 4 we show the behavior of the predictability H and of the fraction ϕ of 'frozen' agents[5] as a function of α for different values of η. For small η the system reproduces qualitatively a pure Minority Game, although the informationally efficient phase shrinks as η increases and more information is injected into the market. Increasing η further (e.g. $\eta = 2$) one sees a clear crossover from a Minority Game like system at low α to a Majority Game like system at high α, which seems even more striking if one observes that ϕ jumps sharply from 0 (as in a Minority Game) to a large value (in a

[5] An agent is 'frozen' when he/she ends up using only one of his strategies. For $S = 2$ this happens when the difference between the strategy valuations, $y_i(t) = \frac{1}{2}[p_{i1}(t) - p_{i2}(t)]$ grows as $y_i(t) \sim v_i t$, with v_i a constant. Clearly, in such cases $\widetilde{g}_i(t) \to (3 - \mathrm{sign}[v_i])/2$ as $t \to \infty$.

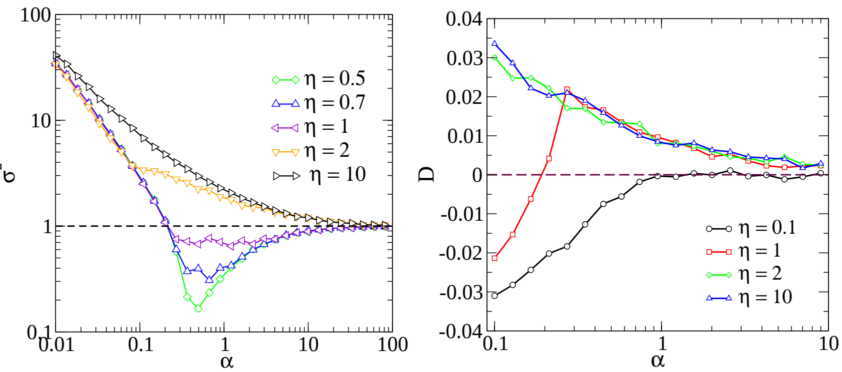

Fig. 5. Behavior of σ^2 (left) and $D = \langle A(t)A(t+1)\rangle/\sigma^2$ (right) vs α for different η. Simulations parameters are as in Fig. 4.

Majority Game all agents ultimately freeze) as α increases. Finally, for large η the system reproduces qualitatively a pure Majority Game. Notice that ϕ decreases as η decreases, indicating that as agents become more sensitive to risk it becomes harder for them to identify an optimal strategy. The crossover from the fundamentalist-dominated regime to the trend-followers-dominated regime gets sharper and sharper as α increases and a detailed analysis shows that H and ϕ develop jump discontinuities at $\eta \simeq 1.4$ when $\alpha \to \infty$. While the magnitude of fluctuations σ^2 increases smoothly with η (as expected, Fig. 5), the existence of a non-trivial crossover for intermediate values of η is evident from the behavior of the (normalized) autocorrelation function of the excess demand $D = \langle A(t)A(t+1)\rangle/\sigma^2$. For small η, $D < 0$ as the dynamics is completely dominated by anticorrelations. As η increases and agents become more risk prone the fundamentalist regime shrinks until the market is completely dominated by trend followers for large η.

For values of η that allow the coexistence of Minority-like and Majority-like features the dynamics acquires surprisingly rich traits. One of these is shown in Fig. 6, where we report the time series of the excess demand at fixed state of the world for three different patterns together with the time series of the 'price' $R(t) = \sum_{t' \le t} A(t')$. One sees that the system tends to generate large excess demands (or price fluctuations are large) in response to a particular information pattern for long time stretches and then quiesce. Several though not all patterns however have this property[6]. At the same time, the profile of $R(t)$ clearly shows the formation of well-defined trends. (Notice however that

[6]A similar phenomenon has been observed in at least another version of the Minority Game, namely that with finite score memory [24].

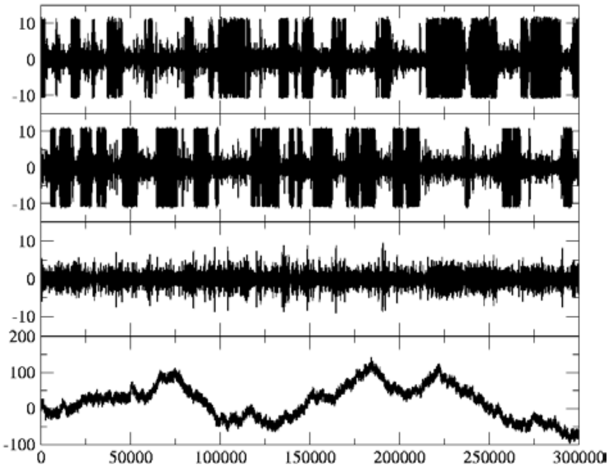

Fig. 6. Typical realization of market dynamics at $\alpha = 0.05$ and $\eta = 1$. Top three rows: $A(t)$ vs t at fixed $\mu(t) = \mu_\ell$ ($\ell = 1, 2, 3$). Bottom: $R(t)$ vs t.

the dynamics becomes strongly sample-dependent and trivial samples may also occur.)

Before moving on, we remark that (12) is perhaps the simplest of several possible choices describing the desired effect. An alternative having similar though not identical properties, namely $F(A) = A - \epsilon A^3$ with ϵ a real positive parameter, was studied in [15]. We just mention that the main difference lies in the absence of an informationally efficient phase for any value of the risk threshold ϵ. Remarkably, in this case market-like phenomenology arises exactly in the crossover regime, as the probability distribution of excess demands spontaneously develops 'heavy tails' for small values of α. We refer the reader to [15] for details.

5 Epilogue

The phenomenology of Minority Games is remarkably robust to changes of the agents' microscopic descriptions such as a broadening of the set of strategies at their disposal or of their possible actions. However, significant new phenomena appear when agents are allowed to modify their behavior (expectations, choice rules, etc.) according to the market conditions they perceive. In this paper we have analyzed perhaps the simplest realistic cases, namely Minority Games in which agents can switch from a fundamentalist to a trend-following attitude in market conditions of low risk and Minority Games where agents keep tracks of the risk associated to a given strategy and use it to counter the expected profit in their decision making process. The general conclusions are that (a) risk does not affect the macroscopic properties of the Minority Game provided the memory of agents is sufficiently long and (b) if agents modify their expectations switching from a fundamentalist to a trend following conduct in conditions of low risk both the stationary and the dynamical properties undergo serious changes, and the global efficiency decreases as agents become less and less risk sensitive. The analysis presented here is but a start, and a more detailed investigation would definitely be welcome.

Acknowledgments – We would like to thank D. Challet, E. Marinari and M.A. Virasoro for stimulating discussions. This work was supported partially by the EU Human Potential Program under contract HPRN-CT-2002-00319 (STIPCO), by the EU EVERGROW Project and by the INFM-MIUR Strategic Project on 'High frequency dynamics of financial markets'.

References

1. D. Challet and Y.-C. Zhang (1997) Physica A 246:407
2. Y.-C. Zhang (1998) Europhys. News 29:51
3. W.B. Arthur (1994) Am. Econ. Assoc. Papers and Proc 84:406
4. D. Challet and M. Marsili (1999) Phys. Rev. E 60:R6271
5. A. Cavagna, J.P. Garrahan, I. Giardina and D. Sherrington (1999) Phys. Rev. Lett. 83:4429
6. D. Challet, M. Marsili and R. Zecchina (2000) Phys. Rev. Lett. 84:1824
7. J.A.F. Heimel and A.C.C. Coolen (2001) Phys. Rev. E 63:056121
8. M. Marsili, D. Challet and R. Zecchina (2000) Physica A 280:522
9. J.A.F. Heimel and A. De Martino (2001) J. Phys.A: Math. Gen. 34:L539
10. D. Challet, M. Marsili and Y.-C. Zhang (2001) Physica A 276:284
11. P. Jefferies, M.L. Hart, P.M. Hui and N.F. Johnson (2001) Eur. Phys. J. B 20:493
12. D. Challet and M. Marsili (2003) Phys. Rev. E 68:036132
13. M. Marsili (2001) Physica A 299:93
14. A. De Martino, I. Giardina and G. Mosetti (2003) J. Phys. A: Math. Gen. 36:8935
15. A. De Martino, I. Giardina, A. Tedeschi and M. Marsili (2004) Phys. Rev. E 70 (in press)
16. C. Chiarella (1992) Ann. Oper. Res. 37:101
17. W.B. Arthur, J.H. Holland, B. LeBaron, R. Palmer and P. Tyler, in W.B. Arthur, S.N. Durlauf and D.A. Lane (Eds.), *The economy as an evolving complex system II* (Addison Wesley, Reading, MA, 1997)
18. J.D. Farmer and A.W. Lo (1999) Proc. Nat. Acad. Sci. 96:9991
19. T. Lux and M. Marchesi (1999) Nature 397:498
20. I. Giardina and J.-Ph. Bouchaud (2003) Eur. Phys. J. B 31:421
21. J.D. Farmer (2002) Ind. Corp. Change 11:895
22. P. Jefferies, M.L. Hart, P.M. Hui and N.F. Johnson (2003) Int. J. Th. Appl. Fin. 3:3
23. C. Camerer and T.-H. Ho (1999) Econometrica 67:827
24. D. Challet, A. De Martino and M. Marsili (2004) Physica A 338:143