

Structure-preserving desynchronization of minority games

G. Mosetti^{1,3}, D. Challet^{1,3,a}, and S. Solomon^{2,3}

¹ Département de physique, Université de Fribourg, Pérolles, 1700 Fribourg, Switzerland

² The Racah Institute of Physics, Hebrew University, Jerusalem, 91905, Israel

³ ISI Foundation, Vialle Settimio Severo 65, 14011 Turin, Italy

Received 14 January 2009 / Received in final form 23 June 2009

Published online 17 September 2009 – © EDP Sciences, Società Italiana di Fisica, Springer-Verlag 2009

Abstract. Perfect synchronicity in N -player games is a useful theoretical dream, but communication delays are inevitable and may result in asynchronous interactions. Some systems such as financial markets are asynchronous by design, and yet most theoretical models assume perfectly synchronized actions. We propose a general method to transform standard models of adaptive agents into asynchronous systems while preserving their global structure under some conditions. Using the minority game as an example, we find that the phase and fluctuations structure of the standard game subsists even in maximally asynchronous deterministic case, but that it disappears if too much stochasticity is added to the temporal structure of interaction. Allowing for heterogeneous communication speeds and activity patterns gives rise to a new information ecology that we study in details.

PACS. 89.65.Gh Economics; econophysics, financial markets, business and management – 89.75.Fb Structures and organization in complex systems – 64.60.De Statistical mechanics of model systems

1 Introduction

When a large number of agents taking part in a multi-player game submit their actions to a central authority (game server, financial market), the times at which their actions become effective are likely to differ because of reaction times, transmission delays or backlogs at the central server. If the resulting delay is sufficiently large, the synchronicity of actions and payoffs is not a reasonable assumption anymore. Yet the immense majority of the literature on games and agents assumes perfect synchronicity (see however e.g. [1–3] for some notable exceptions).

This is problematic in the modelling of many systems, among which the numerous financial markets where the actions of agents are discrete in time and asynchronous. Most financial market models aggregate traders' actions over a given period in one time step. Unless it corresponds to sensible time periods, such as one trading day, this approach is rather artificial. While time coarsening simplifies the description of market dynamics, speculation cannot be modelled by including all the agents' actions in one time step: one does not make money with a single transaction, i.e. in a single time step. In addition, the emergence of large price and volume fluctuations must also be explained in an asynchronous setting. Finally asynchronicity also originates from the heterogeneity of time scales of market participants, which is fat-tailed, possibly a power-law [4–6].

This raises fundamental issues regarding macroscopic synchronization, especially in the case of coordination and cooperation. Here we modify the well-understood minority game [7] by introducing a tunable time delay between the submission of a bid and its actual influence on the global outcome, and a tunable playing frequency.

Remarkably the structure of mean-field models such as the MG is preserved; as a consequence there is hope that the resulting asynchronous interaction somehow still belongs to the mean-field category, hence, that the powerful methods from statistical mechanics that solve the original model [8,9] can be generalized.

2 Desynchronizing global games

Let us consider a global synchronous game. Each agent $i = 1, \dots, N$ takes action $a_i(t)$ at time t . His payoff is a function of his own action and of the global action of all the agents $A(t) = \sum_{j=1}^N a_j(t)$: everything can be written as a function of time t only.

A way to desynchronize the agents while preserving the global structure of the game is to think in terms of delays: if agent i submits his action at time t , the latter becomes effective (i.e. is incorporated into A) only at time $t + \delta$. But at time $t + \delta$, $A(t + \delta)$ only contains the actions effective at that time, that is, does not contain the action of some other agents that are still being transmitted or thought about. Thus, the payoffs of the agents at this

^a e-mail: damien.challet@unifr.ch




Fig. 1. Schematic explanation of playing period T , time delay δ , bids $b(t)$ and actions $a(t + \delta)$.

time do not reflect the most recent actions, but only the last known actions. Differentiating between actions sent but not effective yet, thereafter called bids and denoted by $b_i(t)$, and actual actions $a_i(t)$ makes it easy to introduce asynchronous actions while keeping intact the payoff structure, as will be shown in the next section (see Fig. 1).

Mathematically, agent i submits his bid $b_i(t)$; between times t and $t + \delta - 1$, his last known action is unchanged; at time $t + \delta$, the last bid becomes effective, i.e. $a_i(t + \delta) = b_i(t)$ and of course $A(t + \delta) = \sum_i a_i(t + \delta)$; agent i then receives his payoff that depends on $a_i(t + \delta)$ and $A(t + \delta)$. In this way, a non-trivial structure of overlapping bids and actions can be built.

The temporal structure of the game is completely specified by assuming that agent i is active at times $t_i^{(n)}$, $n > 0$. For instance, $t_i^{(n)} = nT_i + \phi_i$ where n is an integer. The maximally asynchronous case corresponds to $T_i = N$ and $\phi_i = i$: only one agent is active at each time step. Alternatively, an agent may be active with some probability at each time step, thereby removing the rigid structure imposed by periodic $t_i^{(n)}$.

3 Example: minority games

The minority game (MG thereafter) is a prototype model of global competition between adaptive agents [7,8]. Well-understood [8,9], it provides an ideal test-bed for new ideas and extensions. At the same time, it has highly non-trivial and characteristic fluctuations structure and phase transition. The existence of a phase transition with symmetry breaking [10,11] is robust with respect to a surprising number and types of modifications [8,12]; as such, if present in modified games, it is a signature that the original dynamics has not been overly altered.

The aim of the agents is, as the name of the game implies, to be in the minority: in the original game, agent i takes action $a_i(t) \in \{-1, 1\}$; $A(t)$ is defined as above and is positive if the majority chose +1 and vice-versa. The payoff of agent i is $-a_i(t)A(t)$: those who happen to be in the minority are rewarded.

The various types of minority games found in the literature differ mostly in their learning and decision mechanisms. The original one is defined as follows: agents are fed with the last m winning decisions, a bitstring called history and denoted by $\mu(t)$. Each agent has a set of S strategies, i.e., of predefined ways to react to all possible public pieces of information. Denoting the strategies of agent i by $a_{i,s}$, $s = 1, \dots, S$, one can rewrite

$A(t) = \sum_i a_{i,s_i(t)}^{\mu(t)}$ where $s_i(t)$ is the strategy trusted by agent i at time t .

Which strategy to choose is determined by reinforcement learning. To this effect, since the agents gather information about the world through the use of their strategies, they store experience about the past in virtual performance scores of their strategies that evolve according to

$$U_{i,s}(t+1) = U_{i,s}(t) - a_{i,s}^{\mu(t)} A(t). \quad (1)$$

The agents choose their best strategy at time t . In other words, $s_i(t) = \arg \max_s U_{i,s}(t)$.

When desynchronizing the game, a slight complication with respect to game histories arises: whereas in perfectly synchronized games $\mu(t)$ is the same for all the agents, this cannot hold anymore in asynchronous settings since all the agents do not see the same A when they receive a payoff. Thus each agent has his own history of the game, which encodes the past m right choices for him.

By construction $\langle A \rangle = 0$. We shall be interested in the fluctuations of the global outcome $\sigma^2 = \langle A^2 \rangle$ which quantify the degree of coordination of the players, the benchmark being the random outcome $\sigma^2 = N$. Predictability must be measured at the individual level: one defines the conditional average of the attendance from the point of view of agent i , that is, conditional to his histories, which we denote as $\langle A^{(i)} | \mu_i \rangle$, and the predictability as seen by this agent $H_i = \sum_{\mu_i} \langle A^{(i)} | \mu_i \rangle^2 / P$, where $P = 2^M$, while the average predictability is $H = \sum_i H_i / N$.

Predictability $H > 0$ corresponds to a (conditional) symmetry breaking between the two choices. The standard game is characterized by a predictable phase ($H > 0$) for $\alpha = P/N > \alpha_c$, an unpredictable phase ($H = 0$, $\alpha < \alpha_c$) and critical point $\alpha_c = 0.3374 \dots$ where $H \rightarrow 0$ and σ^2/N reaches a minimum [8–11]. The presence of this phase transition is acknowledged to be robust with respect to many modifications of the game, except when all the agents take into account their impact on the game by removing their contribution to A , i.e., by replacing A with $A_{-i} = A - a_i(t)$ in equation (1) [12,13]. However, little is known about the importance of the synchronization of histories with respect to the existence of the phase transition. Local games are synchronized MGs where an agent plays against his neighbors, giving rise to partially spatially overlapping games, hence, histories [14,15], whereas the overlaps are in time in our case.

Another issue is the information ecology: given the fact that predictability is easily measured, MGs give insights on who exploits whom [16], that is, in this case, on the risk associated with delays, for instance. Since the MG is a negative sum game, making an average positive gain is hard; it is only achieved in the original game by some agents that can exploit very efficiently the predictability left by other agents, such as e.g. some of those able to settle on one strategy in the standard game (frozen agents) [10,11,16], speculators feeding on producers, insider trading or a longer history length deep in the unpredictable phase [16]. Recently two of us designed and studied synchronous MGs where the agents had heterogeneous time scales [17]. Here




Fig. 2. (Color online) Fluctuations σ^2/N (red squares; right scale) and predictability H/N (black circles; left scale) of the maximally asynchronous MG as a function of $\alpha = P/N$. $M = 4$ ($P = 16$), $S = 2$, $400NP$ iterations, $200NP$ time steps until equilibration, averages over 200 samples.

we shall characterize the importance of communication delays and frequency of play.

3.1 Results

Numerical simulations are about N times slower than those of the usual MGs, as one effective time-step from a measurement point of view ends when all the agents have updated their actions. Even worse, the interesting regions are found for quite low $\alpha = P/N$, which makes computations even slower. This unfortunately limits the system sizes one can study with current computers to $M = 4$, i.e. $P = 16$. Indeed one run for the maximally asynchronous case at $P = 32$ and $P/N = 0.01$ needs 150 min a modern computer (Core 2 duo, 2 GHz), hence averaging over 200 samples requires 20 days just for this point. Fortunately, $M = 4$ yields good enough results, as discussed in the concluding section.

3.1.1 Maximally asynchronous game

Assuming that $t_i^{(n)} = nT_i + \phi_i$ as above and setting $T_i = T = N$ and $\phi_i = i$, there is only one active agent at each time step, hence A changes at most by 2 in a time step. In this case, the histories μ_i of agents 1 and $N/2$ will likely differ, unless the dynamics of A has a memory longer than $N/2$.

Plotting as usual σ^2/N and H/N reveals that this system undergoes the same phase transition as the original minority game (see Fig. 2). This means that the desynchronization we propose has the remarkable property to keep the global structure of the game unchanged, while allowing for extreme asynchronicity. It is notable that the predictability is orders of magnitude smaller than σ^2/N , which is to be expected since the very long time delays T_i of the agents dilutes information; the order of magnitude of the fluctuations on the other hand is unchanged, since $A \propto \sqrt{N}$.




Fig. 3. (Color online) Autocorrelation of $A(t)$ in a population with heterogeneous time delays T_i ; $M = 4$ ($P = 16$), $N = 51$, $S = 2$, $400NP$ iterations, $200NP$ time steps until equilibration.

The auto-correlation of A reveals a complex pattern, similar to that seen in Figure 3: since A is only possibly changed by only one agent at a time, it displays persistence for $O(N/2)$ time steps. After this decay, it shows on average negative auto-correlation, as can be expected in a minority game where every deviation from $A = 0$ tends to be cancelled by adaptive agents. The additional oscillations are of period N .

3.1.2 Heterogeneous time delays

The pattern of auto-correlation of A suggest that playing with a high frequency allows one to take profit from the persistence of A .

Drawing at random T_i from $\{1, \dots, \rho N\}$ and setting $\delta_i = T_i$, allows one to study the respective gains associated with a time scale denoted by $\langle g|T \rangle = \langle g|\delta \rangle$, in order words, the information ecology that arises from being active more often and having a shorted delay in an asynchronous setting. Figure 4 shows that in asynchronous situations being faster is an clear advantage, a few of the agents reaping even positive gains.

Positive gains come quite peculiarly for minority games from the unconditional persistence of A . Computing the autocorrelation. $\langle A(t)A(t+\tau) \rangle$ (Fig. 3) shows a memory of A lasting $N/2$ time steps, which corresponds exactly to the point where the losses of the agents saturates in Figure 4. Earning a positive gain, though, requires enough persistence to overcome the cost of play the MG, a negative sum game.

3.1.3 Activity frequency and time delays

In the two previous subsections, one assumed that the frequency of activity of a given agent is exactly equal to his time delay. This restriction is unrealistic and can be lifted easily: assume that an agent plays every T_i time steps and has a time delay of δ_i . The delay δ_i may happen




Fig. 4. Average gain of agents with a given time delay T_i versus T_i . $M = 4$ ($P = 16$), $N = 101$, $S = 2$, $400NP$ iterations, $200NP$ time steps until equilibration, averages over 10^5 samples.




Fig. 5. (Color online) Contour plot $M = 4$ ($P = 16$), $N = 101$, $S = 2$, $400NP$ iterations, $200NP$ time steps until equilibration, averages over 10^5 samples.

to be larger than T_i : if an agent is unfortunate enough to have a larger delay than activity frequency, he waits for δ_i time steps to receive his payoff and sends a new bid at the next possible activity time $t_i^{(n)}$.

The average gains conditional on δ and T , denoted by $\langle g|\delta, T \rangle$, reveal in more details the information ecology of asynchronicity: Figure 5 reports a contour plot of $\langle g|\delta, T \rangle$; T and δ play a similar role: the more frequently one is active, the more one can profit from the persistence of A , i.e., from the slowness of other players, as before. But δ , which can be seen as the quickness of reaction to new information, is expectedly the most relevant parameter: it mainly controls whether one obtains a positive or negative gain, except around $\delta = 4$, where T plays this role; the gain decreases as T increases, which is consistent with the results of the previous subsection. The last feature of the plot is the slight bump at $\delta = T$, which is due to the




Fig. 6. (Color online) Predictability H/N and fluctuations σ^2/N as a function of $\alpha = P/N$ for Poissonian activity ($\theta = 1$). $M = 4$ ($P = 16$), $S = 2$, $400NP$ iterations, $200NP$ time steps until equilibration, averages over 200 samples.

fact that the players who have $\delta = T + 1$ have effectively $T' = 2T$, hence the slight decrease of average gain.

3.1.4 Poissonian activity

Finally, let us relax completely the periodicity of agent activity: each agent i plays now at time t with probability θ/N ; the standard MG is recovered when $\theta = N$. We first take $\theta = 1$, which is stochastically equivalent to the setup of Section 3.1.1, except that the temporal interaction has no more structure. Figure 6 suggests that the phase transition disappears completely, at least for the range of parameters we swept over (we tried also individual runs at $\alpha = 0.001$ that confirm this result).

This means that at least some structure of overlap between the histories of the agents is a crucial ingredient of the phase transition. If the phase transition is resilient to some stochastic perturbation of the personal history update temporal structure, one should find a critical θ_c for which a system placed in the symmetric phase undergoes a phase transition when sweeping over θ . The best way to test for a phase transition is to plot the fraction of frozen agents, ϕ , defined as the fraction of agents which played the same strategy during the last half of the time steps [11]. The only situation for which $\phi = 0$ can occur is when $H = 0$. Figure 7 shows that the system stays in the symmetric phase as long as $\theta/N \geq \theta_c/N \simeq 0.5$ and then ϕ increases rapidly. The fluctuations drop discontinuously at $\theta = 1$: anergordicity is broken by stochasticity, but the system stays in a symmetric phase.

4 Discussion

We have proposed a powerful and generic way to desynchronize global games while preserving their structure. The resulting asynchronicity allows for the study of the




Fig. 7. (Color online) Fraction of frozen agents ϕ and fluctuations σ^2/N as a function of θ/N for Poissonian activity for a single realisation of the disorder. $M = 4$ ($P = 16$), $S = 2$, $80\,000NP/\theta$ iterations, $4000NP/\theta$ time steps until equilibration.

effect of time delays and playing frequency in principle in any N -player game.

Desynchronized MGs provide yet another example of the robustness of the phase transition in MGs (see the many other phase transition-preserving modifications listed in [8]). We did not study in detail the location of this phase transition; additional lengthy numerical simulations are needed to study this point in detail, in particular as a function of T .

The fact that the phase transition disappears in the presence of strong enough stochastic desynchronization is a clue that this robustness relies on the temporal structure of individual history updates, as also confirmed by the fact that history-less games ($P = 1$) [18,19] do not reproduce the delay information ecology found in the present study. Interestingly, recent work on the Prisoner's Dilemma showed that similar Poissonian desynchronization leads to a first order phase transition at finite activity frequency [20].

Although we had to consider quite small systems ($M = 4$), the results presented here will not change qualitatively when simulating much larger systems is doable. Finite size effects of minority games are well-studied [8]; two points are crucial: (i) all the macroscopic variables depend only on $\alpha = 2^M/N$, up to finite size effects; (ii) the existence of the phase transition is found for all $M > 1$. It should be noted that macroscopic variables (H/N , σ^2/N) of the MG with real histories may have peculiar finite size effects for very small M because of the De Bruijn graph on which bitstrings of length M live; when $M \geq 4$, its complexity is large enough to prevent the histories of the game to be stuck in a small trivial subgraph.

The familiar phase structure of asynchronous games suggests to solve this kind of games with mathematical methods that have been successfully applied to the original MGs and some of its extensions [8,9]. Two new problems arise: first actions are delayed, which makes the whole calculus more complex; the second problem comes from the fact that each agent has his own real history: solving the standard MG with real (global) histories was a mathematical *tour de force* [9,21]; solving systems with individual histories will need an even more impressive feat. Nevertheless, there is nothing *in principle* that makes the computation infeasible, although it will be much more complex. We hope that such future work will reveal the inner dynamical effects of asynchronicity.

References

1. B.A. Huberman, N.S. Glance, Proc. Nat. Acad. Sci. **90**, 7716 (1993)
2. R. Lagunoff, A. Matsui, Econometrica **65**, 1467 (1997)
3. V. Bhaskar, F. Vega-Redondo, J. Econ. Theory **103**, 334 (2002)
4. P. Lynch, G. Zumbach, Quant. Financ. **3**, 320 (2003)
5. F. Lillo, Eur. Phys. J. B **55**, 453 (2007), e-print [arXiv:physics/0612016](https://arxiv.org/abs/physics/0612016)
6. D. Challet, D. Morton, *Trader heterogeneity and market efficiency* (2009) in preparation
7. D. Challet, Y.-C. Zhang, Physica A **246**, 407 (1997), e-print [arXiv:adap-org/9708006](https://arxiv.org/abs/adap-org/9708006)
8. D. Challet, M. Marsili, Y.-C. Zhang, *Minority Games* (Oxford University Press, Oxford, 2005)
9. A.A.C. Coolen, *The Mathematical Theory of Minority Games* (Oxford University Press, Oxford, 2005)
10. R. Savit, R. Manuca, R. Riolo, Phys. Rev. Lett. **82**, 2203 (1999)
11. D. Challet, M. Marsili, Phys. Rev. E **60**, R6271 (1999), e-print [arXiv:cond-mat/9904392](https://arxiv.org/abs/cond-mat/9904392)
12. A. de Martino, T. Galla, J. Phys. A: Math. Gen. **41**, 324003 (2008)
13. D. Challet, M. Marsili, R. Zecchina, Phys. Rev. Lett. **84**, 1824 (2000), e-print [arXiv:cond-mat/9904392](https://arxiv.org/abs/cond-mat/9904392)
14. T. Kalinowski, H.-J. Schulz, M. Brieese, Physica A **277**, 502 (2000)
15. S. Moelbert, P. De Los Rios, Physica A: Statistical Mechanics and its Applications **303**, 217 (2002)
16. D. Challet, M. Marsili, Y.-C. Zhang, Physica A **276**, 284 (2000), e-print [arXiv:cond-mat/9909265](https://arxiv.org/abs/cond-mat/9909265).
17. G. Mosetti, D. Challet, Y.-C. Zhang, J. Stat. Mech.: Exp. Theory (2005), e-print [arXiv:physics/0509078](https://arxiv.org/abs/physics/0509078), submitted
18. M. Marsili, D. Challet, Adv. Complex Systems **3**, 3 (2001), e-print [arXiv:cond-mat/0011042](https://arxiv.org/abs/cond-mat/0011042)
19. D. Challet, A. De Martino, M. Marsili, J. Stat. Mech.: Exp. Theory (2008)
20. A. Saif, P.M. Gade, J. Stat. Mech.: Exp. Theory P07023 (2009)
21. A.C.C. Coolen, J. Phys. A: Math. Gen. **38**, 2311 (2005)