A CLASS OF STOCHASTIC GAMES WITH INFINITELY MANY INTERACTING AGENTS RELATED TO GLAUBER DYNAMICS ON RANDOM GRAPHS

EMILIO DE SANTIS AND CARLO MARINELLI

ABSTRACT. We introduce and study a class of infinite-horizon non-zero-sum non-cooperative stochastic games with infinitely many interacting agents using ideas of statistical mechanics. First we show, in the general case of asymmetric interactions, the existence of a strategy that allows any player to eliminate losses after a finite random time. In the special case of symmetric interactions, we also prove that, as time goes to infinity, the game converges to a Nash equilibrium. Moreover, assuming that all agents adopt the same strategy, using arguments related to those leading to perfect simulation algorithms, spatial mixing and ergodicity are proved. In turn, ergodicity allows us to prove "fixation", i.e. that players will adopt a constant strategy after a finite time. The resulting dynamics is related to zero-temperature Glauber dynamics on random graphs of possibly infinite volume.

1. INTRODUCTION

The aim of this paper is to study a class of stochastic games with infinitely many interacting agents that is closely connected with a Glauber-type non-Markovian dynamics on random graphs. Let us briefly explain the setting and our contributions both from the point of view of game theory and of physics, referring to the next section for a precise construction of the model. Our central results are theorems 1, 2 and 3 below.

We consider an infinite number of agents located on the vertices of the two-dimensional lattice, where each agent is randomly linked with others, and has positive or negative feelings regarding them. Moreover, each agent is faced with the need of taking decisions that affect himself and all others to whom he is linked. The objective of each agent is to take (non-cooperative) decisions that ultimately do not affect him negatively. Under a specific

choice of the payoff function of each player, we shall prove that there exists a decision policy achieving this goal, and even more, that if each player adopts this strategy a non-cooperative Nash equilibrium is reached.

From the physical point of view, we study a Glauber-type dynamics on a random graph with the following features: the dynamics is non-Markovian and has long-range interactions, in the sense that the maximum distance between interacting particles is unbounded. For such dynamics we prove spatial mixing (hence ergodicity) and fixation. To the best of our knowledge, these problems are solved here for the first time, even in the simpler case of a standard Glauber dynamics on random graphs. Problems of dynamics on random graphs have attracted a lot of attention in recent years (see the monograph [4] for an extensive overview), as these structures are often more realistic models of several phenomena than classical deterministic structures (e.g. in network modeling, spread of epidemics, opinion formation, etc.). For instance, C. Cooper and A. M. Frieze [2] prove the existence of a critical coupling parameter at which the mixing time for the Swendsen-Wang process on random graphs of sufficiently high density is exponential in the number of vertices; M. Dyer and A. Frieze [5] study the rapid mixing (in time) of Glauber dynamics on random graphs with average degree satisfying a certain condition (see also A. Frieze and V. Juan [8] for a related result). In J. P. L. Hatchett et al. [10], the authors analyze the dynamics of finitely connected disordered Ising spin models on random connectivity graph, focusing on the thermodynamic limit. I. Pérez Castillo and N. S. Skantzos [17] study the Hopfield model on a random graph in scaling regimes with finite average number of connections per neuron and spin dynamics as in the Little-Hopfield model. On the other hand, as mentioned above, even though (spatial) mixing is one of the most natural questions to ask about stochastic models of interacting particle systems, it has not been discussed in the literature, to the best of our knowledge. It is probably important to recall that mixing is a key ingredient to obtain further results, such as ergodicity. Moreover, just to cite another important application, using Stein's methods (see e.g. [1]), mixing implies the central limit theorem, which gives qualitative estimates on the number of sites (or agents) with a positive spin (or opinion) in large regions of the graph.

We would like to stress that our results on mixing are quite general, and if one is only interested in the physical aspect of our work, they could essentially skip the part of the paper which deals with stochastic games, and concentrate only on the physical aspect.

Let us briefly discuss how the model and results of the present paper are related to the existing literature on using methods of the theory of interacting particle systems in economic modelling and game theory. One of

the first and still most cited works on the subject is a paper by H. Föllmer [6], who considered a pure exchange economy with (countably) infinitely many agents, each of which having random preferences and endowments. In particular, agents are located on the vertices of the d-dimensional lattice $\mathbb{Z}^d$, and their preferences can be influenced by all his neighbors (i.e. such that their euclidean distance is one). The author then considers the problem of existence of a price system stabilizing the economy. See also E. Nummelin [16] for further results in this connection, but with a finite number of agents. In U. Horst and J. Scheinkman [12] the authors study a system of social interactions where agents are located on the nodes of a subset of $\mathbb{Z}^d$, and each of them is provided with a utility function and a set of feasible actions. The behavior of an agent is assumed to depend on the choices of other agents in a reference group, which can be random and unbounded. The authors, in analogy to our case, work under the assumption that the probability of two agents being linked decays with distance, and are concerned with the existence of equilibrium (in the classical microeconomic sense). U. Horst [11] determines conditions such that non-zero-sum discounted stochastic games with agents interacting locally and weakly enough have a Nash equilibrium. While the set of feasible actions in this paper is much richer than in ours, we do not assume to have any knowledge on the reference group of each agent, apart of being finite almost surely. We only allow agents to be able to observe the dynamics of a (local) configuration around them. As a result of the structural differences in the settings, the optimal strategy in [11] is Markovian, while in our case it can never be Markovian.

In general, the following features of our setting and results could be particularly interesting from a game-theoretic perspective: we consider games where interactions among agents are not known a priori, and we explicitly construct a strategy that leads the game to equilibrium, while the typical result of game theory is the existence of equilibrium and a characterization of optimal strategies *at* equilibrium.

Let us also briefly recall that several other models of interacting particle systems admit a natural interpretation in terms of social interaction. Well-known examples are the voter process (see e.g. T. Liggett [13]), used in models for the formation and spread of opinions, or the Sherrington-Kirkpatrick model of spin glasses (see e.g. section 2.1 in M. Talagrand [18]). Infinite interacting particle systems have found applications in sociology as well: see, for instance, T. Liggett and S. Rolles [14] for a model of formation of social networks.

The organization of the paper is as follows: in section 2 we describe the model so we show how agents interact and what their aim is; in section 3 a general strategy achieving the goal of each agent is given, and section 4

proves spatial mixing, hence ergodicity, of the dynamics, when all agents adopt the same strategy. Finally, using the results on spatial mixing and ergodicity, we prove that the game "fixates", i.e. that agents will adopt a constant strategy after a finite time. This phenomenon is reminiscent of the fixation of zero-temperature dynamics (see e.g. E. De Santis and C. M. Newman [3], L. R. Fontes [7], O. Häggström [9], S. Nanda, C. M. Newman and D. L. Stein [15]).

2. Model and problem formulation

Let us first introduce some notation used throughout the paper. We consider the two-dimensional lattice $\mathbb{Z}^2$ with sites $x = (x_1, x_2)$ and distance d defined by

$$(1) \qquad d(x, y) = |x_1 - y_1| + |x_2 - y_2|.$$

The cardinality of a subset $\Gamma \subseteq \mathbb{Z}^2$ is denoted by $|\Gamma|$. We denote by Λ_M the set of all $x \in \mathbb{Z}^2$ such that $d(O, x) \leq M$, with $O = (0, 0)$. If $x \in \mathbb{Z}^2$, $\Lambda_M(x)$ stands for $\Lambda_M + x$. Our configuration space is $S = \{-1, +1\}^{\mathbb{Z}^2}$. The single spin space $\{-1, +1\}$ is endowed with the discrete topology, and S with the corresponding product topology. Given $\eta \in S$, or equivalently $\eta : \mathbb{Z}^2 \to \{-1, +1\}$, and $\Lambda \subseteq \mathbb{Z}^2$, we denote by η_Λ the restriction of η to Λ. Given a graph $G = (V, E)$, where V and E are the sets of its vertices and edges, respectively, we shall denote by $\{x, y\}$ an element of E connecting x, $y \in V$. For any $x \in V$, we shall denote by ρ_x the distance of the longest edge having x as endpoint, namely we define

$$\rho_x = \sup_{y:\,\{x,y\}\in E} d(x, y).$$

Recall that the distance in variation of two probability measures μ and ν on a discrete set Ω is defined as

$$\|\mu - \nu\| = \frac{1}{2} \sum_{\omega \in \Omega} |\mu(\omega) - \nu(\omega)|.$$

We shall now introduce an idealized model of a large ensemble of interacting individuals. The ingredients will be a random graph, a function on its edges (specifying an environment, roughly speaking), and a stochastic process with values in S describing the time evolution of the system.

Let $\mathcal{G} = (V, E)$ be a random graph, whose set of vertices V is given by all sites of the 2-dimensional lattice $\mathbb{Z}^2$, and whose set of edges E satisfy the following conditions: edges exist with probability one between each site x and all y such that $d(x, y) = 1$, and

$$(2) \qquad \mathbb{P}(|\{y : \{x, y\} \in E\}| < \infty) = 1 \quad \forall x \in V.$$

We suppose that each site is occupied by an individual (we shall often identify individuals with the sites they occupy, when no confusion will arise), and that relations among individuals are modeled by the edges of $\mathcal{G}$ and by a function $j : V \times V \to \{-1, 0, +1\}$; $j(x, y) = 0$ if $\{x, y\} \notin E$, otherwise $j(x, y) \in \{-1, +1\}$.

In particular, we shall say that individuals x and y are linked if $\{x, y\} \in E$, and the value $j(x, y)$ shall account for the "feelings" of x towards y: we set $j(x, y) = +1$ if x is a "friend" of y, and $j(x, y) = -1$ if x is an "enemy" of y. We do not assume symmetry of j, i.e. friendship of an individual towards another may not be reciprocal. Moreover, we assume that individuals do not know with whom they are connected, nor whether these individuals are friends or enemies. Note also that in this model x can be friend of y, y friend of z, but x and z can be either friends or enemies (a phenomenon also called frustration in physics).

Let us now introduce a stochastic process $\sigma : [0, \infty) \to S$ modelling the evolution of the "action" (or opinion) of the individuals. We shall use a graphical construction of the process, which provides a specific version of basic coupling, i.e. it provides versions of the whole family of stochastic processes on $\mathcal{G}$ (or on any finite subset of it), all on the same probability space. We assume that the initial configuration σ_0 is chosen from a symmetric Bernoulli product measure. Moreover, the continuous-time dynamics of σ_t is given by independent Poisson processes (with rate 1) at each site $x \in V$ corresponding to those times $(t_{x,n})_{n \in \mathbb{N}}$ when the individual x is asked to update his opinion. Before describing the set of feasible ways of opinion updating, let us introduce a reward for a generic individual x at time $t_{x,n}$, as a result of his action:

$$h_t(x) = \operatorname{sgn} \left(\sum_{y:\, \{x,y\} \in E} j(x, y) \sigma_t(x) \sigma_t(y) \right),$$

where we have set, for simplicity, $t \equiv t_{x,n}$.

We allow x to base his decision on the history of $\sigma_{\Xi_s}(s)$, $s \geq 0$, and $h(x)$, where Ξ_s are finite balls centered in x with random radius which is nondecreasing with respect to s, finite almost surely for all $s \geq 0$, and not 'exploding'. Formally, the decision of individual x at time $t_{x,n}$ is a $\{-1, +1\}$-valued random variable $u_{x,n}$ measurable with respect to the σ-algebra generated by $\{\sigma_{\Xi_s}(s),\ s \leq t_{x,n}\}$ and $\{h_s(x),\ s \leq t_{x,n}\}$, where Ξ_s are balls centered in x such that

$$\Xi_\infty = \lim_{s \to \infty} \Xi_s$$

exists and is finite with probability one. We shall denote by $\mathcal{E}_t^x$ the filtration just defined.

The dynamics of σ is then completely specified by the updating rule

$$\sigma_{t_{x,n}}(x) = u_{x,n}.$$

Several remarks are in order: the reward $h_t(x)$ obtained by individual x as a result of his decision at time $t = t_{x,n}$ is positive if the difference between pleased and damaged friends is bigger than the difference between pleased and damaged enemies, negative if the opposite happens, and zero if the value is the same. Since at a fixed arrival time $t = t_{x,n}$ of the Poisson clock of x no other clock is ringing, i.e. $\mathbb{P}(t_{y,m} = t) = 0$ for all $y \neq x$ and for all $m \in \mathbb{N}$, the dynamics of σ is well-defined (also using the graphical construction). Finally, at any positive time t, $\sigma_t(x)$ represents the last decision taken by individual x up to time t.

We formulate the following problem for the generic individual x: find a strategy $\pi_x = (u_{x,1}, u_{x,2}, \ldots)$ such that

$$h_t(x) \geq 0 \qquad \text{a.s.}$$

for all $t \geq T_x$, where T_x is a finite (random) time.

Remark 1. We built the random graph $\mathcal{G}$ on the two-dimensional lattice $\mathbb{Z}^2$ to give a "geographic" dimension to the problem and to have a simple notion of distance on the graph. However, all results in the next section still hold replacing $\mathbb{Z}^2$ with any higher dimensional lattice $\mathbb{Z}^d$, $d \geq 3$. We shall see below that choosing $d = 2$ also affects a constant appearing in an assumption used to prove spatial mixing.

3. Admissible strategies that eliminate losses

In this section we construct explicitly a strategy π_x for the generic individual x that asymptotically eliminate negative rewards, i.e. such that $\mathbb{P}(h_t(x) \geq 0) = 1$ for all t greater than a random time, which is finite with probability one. It will also be clear that this strategy is non-cooperative, that is π_x eventually eliminate negative rewards irrespectively of the strategies adopted by all other individuals.

For simplicity of notation let us describe the strategy $\pi \equiv \pi_0$ for the individual located at the origin O. The arrival times of his Poisson process and the corresponding decisions and rewards will be denoted by t_n, u_n, and h_n, $n \in \mathbb{N}$, respectively.

The strategy $\pi = (u_1, u_2, \ldots)$ is best defined algorithmically through a decision tree. We also need an additional "data structure", i.e. a collection $\mathcal{R}$ of ordered triples of the type (η, u, h), where $\eta \in S$ is supported on finite balls, $u \in \{-1, +1\}$, and $h \in \{-1, 0, +1\}$.

At the first arrival time t_1, u_1 is chosen accordingly to a Bernoulli law with parameter $1/2$ (a "fair coin toss"), and $(\sigma_{\Lambda_1}, u_1, h_1)$ is added to $\mathcal{R}$.

The description of the algorithm then follows inductively: at time t_{n+1}, let Λ_{M_n} be the support of the last configuration added to $\mathcal{R}$. Let $\sigma' := \sigma_{\Lambda_{M_n}}(t_{n+1}-)$ and check whether there exists $(\sigma', u', h') \in \mathcal{R}$.

- If yes, set
$$u_{n+1} = u' \frac{h'}{|h'|},$$
  with the convention $0/|0| := 1$. The reward h_{n+1} corresponding to u_{n+1} is now obtained.
  - If $h_{n+1} \geq 0$, no further action is needed.
  - If $h_{n+1} < 0$, then add to $\mathcal{R}$ the triplet $(\sigma_{\Lambda_{M_{n+1}}}, u_{n+1}, h_{n+1})$.
- Otherwise, set $u_{n+1} = u_n$, and add to $\mathcal{R}$ the triplet $(\sigma', u_{n+1}, h_{n+1})$.

The above algorithm formalizes the following heuristic procedure: the agent starts looking at the configuration on the smallest ball centered around him and plays tossing a coin. The next time his clock rings, he checks whether he has already seen such a configuration. If it is a new one, he will again memorize it and play by tossing a coin, while if it is a known one he will play as he did before if he got a positive reward, or the opposite way if he got a negative reward. Of course it could happen that this way of playing still does not guarantee a positive reward, in which case he will memorize the configuration on a larger ball around himself and its associated outcome.

Remark 2. One of the key steps of the algorithm requires one to look for a triplet (σ', u', h') in $\mathcal{R}$, given $\sigma' = \sigma_\Lambda$, for a certain $\Lambda \subset \mathbb{Z}^2$. This operation is uniquely determined, i.e. there can exist only one triplet $(\sigma', u', h') \in \mathcal{R}$ with a given σ'. This can be seen as a consequence of the structure of the algorithm itself. Namely, as soon as the player "observes" the same configuration $\sigma' = \sigma_\Lambda$ with a different associated outcome h, he will immediately enlarge the support of observed configurations Λ.

We shall now prove that the strategy just defined eliminates losses for large times.

Theorem 1. *For any individual x there exists a random time T_x, finite with probability one, such that*
$$\mathbb{P}\Big(h_{t \vee T_x}(x) \geq 0\Big) = 1.$$

Proof. Let us define a sequence of random times $(\tau_n)_{n \in \mathbb{N}}$ as follows:
$$\tau_n = \inf\{n \in \mathbb{N} \mid \exists (\overline{\sigma}, u, h) \in \mathcal{R}, \, \text{supp}\, \overline{\sigma} = \Lambda_{n+1}\},$$
with the convention that $\inf \emptyset = +\infty$. In other words, τ_n is the first time that individual x includes into his information set $\mathcal{R}$ the box Λ_{n+1} (and if this never happens, then $\tau_n = +\infty$). Let τ_k be the last finite element of the

sequence $(\tau_n)_{n \in \mathbb{N}}$. By assumption (2) we know that $|\{y : \{x, y\} \in E\}|$ is finite, hence $k \le \rho_x$ because $h_t(x)$ only depends on those y linked to x, for all times t. Therefore the biggest Λ_n observed by the agent in the origin is finite.

Define the family of sets

$$A_k(t) = \{\sigma_{\Lambda_k}(t_{x,\ell}) : t_{x,\ell} \in [\tau_k, t]\} .$$

It clearly holds $A_k(t_1) \subseteq A_k(t_2)$ for $t_1 < t_2$, hence we can define

$$A_k(\infty) = \lim_{t \uparrow \infty} A_k(t).$$

Since $A_k(t) \subset \{-1, +1\}^{\Lambda_k}$, and Λ_k is finite, then there exists $T_x > 0$ such that $A_k(T_x) = A_k(\infty)$, hence

$$A_k(t) = A_k(\infty) \qquad \forall t > T_x.$$

We claim that $h_{x,n} \ge 0$ for all $t_{x,n} > T_x$. In fact, for every $t_{x,n} > T_x$ there exists $(\overline{\sigma}, u, h) \in \mathcal{R}$ with $\overline{\sigma}(y) = \sigma_{t_{x,n}}(y)$ for all $y \in \Lambda_k$. But since $\tau_{k+1} = \infty$, the algorithm will give as output a u_n such that $h_n \ge 0$ (to convince oneself it is enough to "run" the algorithm). In a more suggestive way, one could say that after T_x individual x has already been faced at least once with all possible configurations that are relevant for him, and therefore knows how to take the right decision. $\qquad \square$

Remark 3. (i) Note that the strategies of other individuals never enter into the arguments used in the proof. Therefore individual x is sure to reach the goal of eliminating losses in finite time irrespectively of the strategies played by all other individuals.

(ii) However, we would like to stress that the random time T_x is not a stopping time (i.e. it is not adapted to the filtration $\mathcal{E}_t^x$). In fact, T_x depends in general on the decisions of other individuals, whose policies are not necessarily adapted to $\mathcal{E}_t^x$. In general, even if all policies were adapted, the random times $\{T_x\}_{x \in \mathbb{Z}^2}$ would not be stopping times.

(iii) Let us also observe that although we formally allowed the strategy π_x to be adapted to $\mathcal{E}_t^x$, the information used by the strategy constructed in the proof of Theorem 1 is much smaller. Similarly, one could refine the way the memory structure $\mathcal{R}$ is constructed, for instance by eliminating configurations on smaller balls, when one starts to add new configurations on balls of higher radius. However, we preferred to keep the construction of $\mathcal{R}$ as it is to avoid non-essential complications.

As a consequence of theorem 1 and of observation (i) in the above remark, one has the following result, which essentially states that the games admits an "asymptotic" Nash equilibrium.

Proposition 1. *Let $M \in \mathbb{N}$ and assume that each player $x \in \Lambda_M :=$ $[-M, M] \times [-M, M]$ adopts the strategy π_x defined above. Then there exists a finite random time T_M after which no agent can gain by any change in their strategy given the strategies currently pursued by other players.*

It is important to observe that in the above proposition we implicitly assume that each player only cares about "not loosing", or equivalently he distinguishes only between "loosing" ($h_t(x) < 0$) and "not loosing" ($h_t(x) \geq 0$). In this sense, after T_M, there is no point for any player $x \in \Lambda_M$ to change his strategy, as proved in theorem 1. The statement of the proposition is in general false if the player distinguishes between $h_t(x) > 0$, $h_t(x) = 0$, and $h_t(x) < 0$.

We think that one can prove (and we leave it as a conjecture), that this asymptotic equilibrium is not Pareto. This could be done adapting ideas of O. Häggström [9], who proved that zero-temperature dynamics on a random graph does not reach the minimum energy configuration.

4. Spatial mixing and ergodicity

The main result of this section, which plays an essential role in the results about fixation of the next section, is that a spatial mixing property holds. We shall work under the following hypothesis, which states that the probability of two agents being linked decays algebraically with their distance.

Standing assumption. It holds that

$$(3) \qquad \mathbb{P}(\{x, y\} \in E) \leq \frac{C}{d(x, y)^9},$$

for all y such that $d(x, y) > 1$, where C is a positive constant.

Note that assumption (3) implies (2). Moreover, the exponent appearing on the right-hand side of (3) depends on the dimension of the lattice and it is needed in order to use well-known combinatorial estimates on path counting in $\mathbb{Z}^2$ in the proofs to follow. However, it would not be difficult to generalize our arguments to any higher dimensional lattice $\mathbb{Z}^d$, $d \geq 3$, at the expense of replacing the exponent 9 with a (higher) constant depending on the dimension d, and of using more complicated estimates in the proofs. Since this point is not essential and would only add technical complications, we preferred to fix $d = 2$.

Before stating the main theorem of this section, we need to introduce the following set of conditions.

Hypothesis H. *The process random graph $\mathcal{G} = (V, E)$ and the process $\sigma : [0, \infty) \to \{-1, 1\}^{\mathbb{Z}^2}$, satisfy the following conditions:*

(i) *For each vertex $x \in \mathbb{Z}^2$ there exists a Poisson process P_x, and the Poisson processes $\{P_x\}_{x \in \mathbb{Z}^2}$ are mutually independent. Denoting by $\Upsilon_x = \{t_{x,n}\}$ the set of arrival times of P_x, the value of $\sigma_t(x)$ is allowed to change only at times $t \in \Upsilon_x$.*

(ii) *Given any couple $(x, y) \in \mathbb{Z}^2 \times \mathbb{Z}^2$, the probability $\mathbb{P}(\{x, y\} \in E)$ is defined and it can depend on $d(x, y)$. Moreover, for any choice of (x, y), $(v, w) \in \mathbb{Z}^2 \times \mathbb{Z}^2$ with $(x, y) \neq (v, w)$, the events $\{x, y\} \in E$ and $\{v, w\} \in E$ are mutually independent.*

(iii) *The evolution of the process is local, i.e. $\sigma_{t_{x,n}}(x)$ is measurable with respect to $\mathfrak{F}^x_{t_{x,n}}$, where $\mathfrak{F}^x_t$ denotes the σ-algebra generated by $\{\sigma_s(y) : \{x, y\} \in E \text{ or } y = x, \ s < t\}$. We denote by $\mathcal{F}^V_t$ the σ-algebra generated by $\cup_{x \in V} \mathfrak{F}^x_t$.*

(iv) *Both the probability of two agents being linked and the evolution of σ are translation invariant, i.e. $\mathbb{P}(\{x, y\} \in E) = \mathbb{P}(\{x+v, y+v\} \in E)$ and $\mathbb{P}(\sigma_t \in A | \sigma_0 = \eta) = \mathbb{P}(\sigma_t \in A + v | \sigma_0(\cdot) = \eta(\cdot + v))$.*

We can now state the main theorem of this section.

Theorem 2. *If Hypothesis H holds true, then σ satisfies the spatial mixing property*

$$(4) \qquad \lim_{\Lambda \to \mathbb{Z}^2} \mathbb{P}\Big(\sigma_{\Lambda_0}(t) = \eta | \mathcal{F}^{\Lambda^c}_t\Big) = \mathbb{P}(\sigma_{\Lambda_0}(t) = \eta),$$

where Λ_0 is any finite region in $\mathbb{Z}^2$.

Note that the process σ is translation invariant if each agent adopts the same strategy at each decision time (the strategy does *not* need to be the one defined in section 3). Before giving the proof of the theorem, we establish some auxiliary results.

We shall use the following terminology: by "box of side length L" we mean the set $[-L/2, L/2]^2 \subset \mathbb{Z}^2$. For $\rho < 1$, we call "subbox of side length L^ρ" any one of the $L^{1-\rho}$ square sets into which a box of side length L can be subdivided. We always assume L^ρ, $L^{1-\rho} \in \mathbb{N}$ (without loss of generality, as it will be clear). Furthermore, we shall say that two subboxes R and S are "neighbors" if $d(R, S) \leq \sqrt{2}$, so every subbox has 8 neighbor subboxes. We shall call "path of subboxes" a sequence of subboxes $(R_k)_{k=1,\dots,K}$ such that R_k and R_{k+1} are neighbors for each $k = 1, \dots, K - 1$. Two subboxes R, S are "linked" if there exist $x \in R$, $y \in S$ such that $\{x, y\} \in E$.

In the following lemma we introduce a sequence of boxes increasing to $\mathbb{Z}^2$, each of one further subdivided into a variable number of boxes also increasing to $\mathbb{Z}^2$, but at a lower rate.

Lemma 1. *There exist a sequence of integer numbers $L_n \uparrow +\infty$, a sequence of square boxes Q_{L_n} of side length L_n, each of them partitioned into subboxes*

of side L_n^ρ, $\rho = 13/42$, such that only a finite number of the boxes Q_{L_n} will contain linked non-neighbor subboxes.

Proof. We use a Borel-Cantelli argument on a suitable sequence of box side lengths L_n. In particular, let L be a positive integer, Q_L a square of side L, subdivided into subboxes of side L^ρ. The probability of an agent x to be linked with some other agent of a non-neighbor subbox is bounded by

$$\sum_{y:d(x,y)\geq L^\rho} \frac{C}{d(x,y)^9} \leq C_1 \int_{L^\rho}^\infty \frac{1}{v^9}\, 2\pi v\, dv = C_2 \frac{1}{L^{7\rho}},$$

where C, C_1, C_2 are positive constants. Therefore two agents in non-neighbor subboxes exist with probability not larger than

$$L^2 \sum_{y:d(x,y)\geq L^\rho} \frac{C}{d(x,y)^9} \leq C_2 \frac{1}{L^{7\rho-2}} \to 0, \quad \text{as } L \to \infty.$$

Taking now a subsequence L_n growing to infinity rapidly enough,

$$\sum \mathbb{P}(A_{L_n}) < \infty,$$

where A_{L_n} denotes the event that Q_{L_n} contains linked non-neighbor subboxes. By Borel-Cantelli lemma, only a finite number of occurrences of A_{L_n} can happen, which finishes the proof. $\qquad\square$

Recall that for a sequence of i.i.d. standard exponential random variables $\{X_i\}$ one has

$$(5) \qquad \mathbb{P}\Big(\sum_{i=1}^n X_i < n\alpha\Big) \leq e^{-\Phi(\alpha)n} \quad \forall \alpha < \mathbb{E}X_1,$$

where the so-called rate function Φ is given by

$$\Phi(\alpha) = \alpha - 1 - \log\alpha.$$

Proof of Theorem 2. We use a coupling argument to show that

$$(6) \quad \sup_{\zeta',\zeta''} |\mathbb{P}(\sigma_{\Lambda_0}(t) = \eta|\sigma_{\Lambda^c}(0) = \zeta') - \mathbb{P}(\sigma_{\Lambda_0}(t) = \eta|\sigma_{\Lambda^c}(0) = \zeta'')| \to 0$$

and hence, by the inequality

$$\sup_{\zeta',\zeta''} |\mathbb{P}(\sigma_{\Lambda_0}(t) = \eta|\sigma_{\Lambda^c}(0) = \zeta') - \mathbb{P}(\sigma_{\Lambda_0}(t) = \eta|\sigma_{\Lambda^c}(0) = \zeta'')|$$

$$\geq |\mathbb{P}(\sigma_{\Lambda_0}(t) = \eta|\sigma_{\Lambda^c}(0) = \zeta) - \mathbb{P}(\sigma_{\Lambda_0}(t) = \eta)| \quad \forall \zeta,$$

that (4) holds.

We construct two coupled systems σ', σ'' on the same probability space supporting σ in the following way: $\sigma'_x(0) = \sigma''_x(0) = \sigma_x(0)$ for all $x \in \Lambda$; σ' and σ'' update their state according to the same translation-invariant rule of σ; all other randomness in the system (the random graph, the Poisson

processes, the "coin tosses" needed for the decision rules) coincide. Define, for any $x \in V$, the random time

$$\tau_x = \inf\{t \geq 0 : \sigma'_x(t) \neq \sigma''_x(t)\},$$

and introduce the process

$$[0, \infty) \times V \ni (t, x) \mapsto \nu_x(t) = 1(t \geq \tau_x) \in \{0, 1\}.$$

Using a pictorial language, we shall say that we color x with black as soon as the two processes σ' and σ'' differ at x. Let us also introduce another process $\tilde{\nu} : [0, \infty) \times V \to \{0, 1\}$ with the property $\tilde{\nu}_x(t) \geq \nu_x(t)$ a.s. for all x and all t. The dynamics of $\tilde{\nu}$ is specified as follows: $\tilde{\nu}_x(0) = 0$ for all $x \in V$, and $\tilde{\nu}_x$ can turn to one as a consequence of two classes of events. In particular, (i) $\tilde{\nu}_x(t) = 1$ if there exists x' belonging to the same subbox of x such that $\nu_{x'}(t) = 1$, and (ii) $\tilde{\nu}_x(\tau) = 1$ if there exists y belonging to a neighbor subbox such that $\tilde{\nu}_y(\tau) = 1$, where τ is any arrival time of the Poisson process relative to x. Moreover, we assume that 1 is an absorbing state for $\tilde{\nu}_x$, for all x. Using again a pictorial analogy, we could say that the black area generated by $\tilde{\nu}$ is bigger than the black area generated by ν. In particular, as soon as a site x turns black, (i) implies that the whole subbox to which it belongs becomes black as well.

By Lemma 1, there exists a positive integer N and a sequence L_n such that for all $n > N$ the boxes Q_{L_n} contain no linked non-neighbor subboxes. The shortest path of subboxes from the boundary of the box Q_{L_n} to its center has length $L_n^{1-\rho}/2$ (therefore, for n large enough, the shortest path of subboxes from the boundary of the box Q_{L_n} to Λ_0 has length greater or equal than $L_n^{1-\rho}/2 - 1$). Setting $T_{Q_{L^\rho}} = \inf_{x \in Q_{L^\rho}} t_{x,1}$ (recall that $t_{x,1}$ is the first arrival time of the Poisson processes relative to x), one has that the distribution of $T_{Q_{L^\rho}}$ is $\mathrm{Exp}(L^{2\rho})$, where $\mathrm{Exp}(\lambda)$ stands for the law of an exponential random variable with parameter λ. The minimum time for the formation of a path of k "black" subboxes along a fixed path (of sites) from the boundary of Q_{L_n} to the origin is given by

$$T = \sum_{i=1}^{k} T_i$$

for all $n > N$ (from now we shall tacitly assume $n > N$), where $T_1, \ldots, T_k$ are i.i.d. exponential random variables with parameter $L_n^{2\rho}$ (independence and the value $L_n^{2\rho}$ follow by the memoryless property of the exponential distribution).

Note that the sequence of subboxes in a path turning black does not influence the minimum time needed for the formation of such path, which is a sum of independent exponential random variables of parameter $L_n^{2\rho}$, using

again the memoryless property of exponential distributions. It follows by (5) that, for $0 < \alpha < 1$, one has

$$\mathbb{P}(T \leq k\alpha L_n^{-2\rho}) \leq e^{-(\alpha-1-\log\alpha)k}.$$

Denoting by $T_{\partial Q_{L_n} \to O}$ the (random) time needed to form a path of black subboxes from the boundary of Q_{L_n} to the origin O, we obtain the estimate

$$\mathbb{P}\Big(T_{\partial Q_{L_n} \to O} \leq \frac{\alpha}{2} L_n^{1-3\rho}\Big) \quad \leq \quad 4L_n^{1-\rho} \sum_{k \geq \frac{L_n^{1-\rho}}{2}} 8 \cdot 7^{k-1} \exp\Big(-(\alpha-1-\log\alpha)k\Big),$$

hence, for $0 < \alpha < \frac{1}{7e}$,

$$\lim_{n \to \infty} \mathbb{P}\Big(T_{\partial Q_{L_n} \to O} \leq \frac{\alpha}{2} L_n^{1-3\rho}\Big) = 0.$$

Here the term $4L_n^{1-\rho}$ accounts for the possible initial subbox on the boundary of Q_L, and $8 \cdot 7^{k-1}$ is an upper bound for the number of paths (of subboxes) of length k starting in a given subbox. We obtain that, as $n \to \infty$, the term on the right hand side goes to zero like $e^{-\beta L_n^{1-\rho}}$ (modulo polynomial terms), with β a positive constant. Again by a Borel-Cantelli argument we obtain

$$\mathbb{P}\Big(\lim_{L \to \infty} T_{\partial Q_L \to O} = \infty\Big) = 1.$$

Moreover, the evolution of the central subbox is completely independent on the configuration outside Λ until it turns black, and so the theorem is proved. □

Remark 4. Although (6) has been proved only for a particular choice of a sequence of increasing boxes Λ_n, one can easily show that any increasing sequence of boxes will do. In fact, the supremum appearing in (6) is decreasing with respect to Λ, hence it is enough to prove the theorem for any (fixed) subsequence.

5. Fixation

In this section we shall work under the general assumptions introduced in section 2 and 4, and furthermore we assume that each player adopts the same strategy (hence the dynamics is translation invariant), and that interactions are symmetric, i.e. that $j(x,y) = j(y,x)$ for any $x, y \in V$. The latter hypothesis is essential, as it would be possible to find counterexamples to our results in the case of asymmetric interactions. As before, we shall denote by x an arbitrary agent, fixed throughout this section. Let us define the random time T_x as

(7) $T_x = \sup\{t : \text{ at time } t \text{ agent } x \text{ sees a new configuration or loses}\}.$

As it follows from Theorem 1, T_x is finite with probability one. Moreover, by definition, agent x will not loose at any time after T_x. Let us also define the random variable M_x as the number of times agent x changes his state (i.e. updates his opinion) during the time interval $(0, +\infty)$.

The main result of this section is the following:

Theorem 3. *Assume that each agent adopts the strategy constructed in section 3. Then each agent $x \in V$ updates his opinion only a finite number of times, i.e.*

$$\mathbb{P}(M_x < \infty) = 1.$$

Before proving theorem 3, we shall need some more definitions and preparatory results.

Let us recall the definition of ρ_x:

$$(8) \qquad \rho_x = \sup_{y:\, \{x,y\}\in E} d(x,y),$$

the distance from x of his farthest connected agent. Note that one has, as follows by the standing assumption (3),

$$(9) \qquad \mathbb{P}(\rho_x \geq r) \leq \sum_{s \geq r} 4s \frac{C}{s^{8+\varepsilon}} \leq \frac{K}{r^{6+\varepsilon}}$$

where C and K are constants depending on x. Therefore $\mathbb{E}\rho_x^k$, $1 \leq k \leq 5$ are finite:

$$(10) \quad \mathbb{E}\rho_x^k \leq 1 + \sum_{r=2}^{\infty} r^k \mathbb{P}(\rho = r) \leq 1 + \sum_{r=2}^{\infty} r^k \mathbb{P}(\rho \geq r) \leq 1 + \sum_{r=2}^{\infty} r^k \frac{K}{r^{6+\varepsilon}} < \infty.$$

Let us also define the energy (or Lyapunov) function on a finite set $\Lambda \subset \mathbb{Z}^2$ as

$$(11) \qquad H_\Lambda(\sigma) = -\sum_{u \in \Lambda} \tilde{h}_u(\sigma),$$

where

$$(12) \qquad \tilde{h}_u(\sigma) = \sum_{v:\, \{u,v\}\in E} j(u,v)\sigma_u\sigma_v.$$

In the following we shall denote by Λ_n the square box $[-n,n] \times [-n,n]$.

Lemma 2. *There exists a continuous function $e : \mathbb{R}_+ \to [-\mathbb{E}\rho^2, \mathbb{E}\rho^2]$ such that*

$$(13) \qquad \lim_{n \to \infty} \frac{H_{\Lambda_n}(\sigma(t))}{|\Lambda_n|} = e(t) \quad \text{a.s.}$$

Proof. By the definitions of $\tilde{h}_x(\sigma(t))$, ρ_x, it follows that for each time t

$$-\rho_x^2 \leq \tilde{h}_x(\sigma(t)) \leq \rho_x^2,$$

hence, taking expectations, recalling (10), and using translation invariance

$$-\infty < -\mathbb{E}\rho_O^2 \leq \mathbb{E}\tilde{h}_x(\sigma(t)) \leq E\rho_O^2 < \infty.$$

At any time t, using the space ergodicity of the system (implied by the spatial mixing property proved in Theorem 2), we obtain

$$(14) \qquad \lim_{n\to\infty} \frac{H_{\Lambda_n}(\sigma(t))}{|\Lambda_n|} = \mathbb{E}\tilde{h}_O(\sigma(t)) \quad \text{a.s.}$$

Setting $e(t) = \mathbb{E}\tilde{h}_O(\sigma(t))$, we just have to prove that e is continuous. Using again the spatial ergodicity of σ, the proportion of agents in Λ_n taking at least a decision in the time interval $]t_1, t_2[$ tends to $1 - e^{-(t_2-t_1)} \leq t_2 - t_1$ as $n \to \infty$. Since each agent is endowed with a Poisson process that is independent from all other processes and random variables describing the dynamics of the system, the mean energy variation of each agent is bounded by $\mathbb{E}\rho^2$. Therefore we also have

$$(15) \qquad \begin{aligned} |e(t_2) - e(t_1)| &= \lim_{n\to\infty} \frac{|H_{\Lambda_n}(\sigma(t_2)) - H_{\Lambda_n}(\sigma(t_1))|}{|\Lambda_n|} \\ &\leq (1 - e^{-(t_2-t_1)})\mathbb{E}\rho^2 \leq (t_2 - t_1)\mathbb{E}\rho^2, \end{aligned}$$

i.e. the function e is Lipschitz continuous. $\qquad\square$

Let us now define the following discrete random sets for agent x, which are subsets of the set of arrival times of his Poisson process:
$N_1(x) =$
$= \{t : t \leq T_x, \text{the agent in } x \text{ sees a known configuration at time } t \text{ and loses}\}$
$N_2(x) =$
$\quad = \{t : t \leq T_x, \text{there is an arrival of the Poisson process in } x\} \setminus N_1(x)$
$N_3(x) = \{t : t > T_x, \text{the agent in } x \text{ changes opinion}\}$

Note that by definition of T_x, at any time $t > T_x$ agent x can only see known configurations, and can only win.
We also define, for every $t > 0$ and $x \in \mathbb{Z}^2$, the random sets

$$N_i(t, x) = N_i(x) \cap [0, t],$$

for $i = 1,\ 2,\ 3$.
Moreover, for $\Lambda \subset \mathbb{Z}^2$, we set

$$N_i(t, \Lambda) = \bigcup_{x\in\Lambda} N_i(t, x).$$

The dynamics of the system and the definition of $e(t)$ imply that $e(t)$ is determined only by the changes of $\sigma_\tau(\cdot)$, $\tau \in \{N_i(t,x)\}_{x\in\mathbb{Z}^2}$, $i = 1,2,3$. We can therefore write

$$e(t) = e_1(t) + e_2(t) + e_3(t),$$

where $e_i(t)$ denotes the component of $e(t)$ determined by changes of $\sigma_\tau(\cdot)$ for $\tau \in \{N_i(t,x)\}_{x\in\mathbb{Z}^2}$. Moreover one has $e_2(t) \leq 0$ because we are eliminating the arrivals where the agent lost, and in this case the energy can only decrease.

We are now in the position to prove the theorem on the fixation of the stochastic dynamics.

Proof of Theorem 3. In virtue of the translation invariance of the system, it is enough to prove the result for the agent in the origin. First observe that $M_O \leq |N_1(O)| + |N_2(O)| + |N_3(O)|$, because $N_1(O)$ and $N_2(O)$ may contain Poisson arrival times in which the agent O does not change his opinion. Denoting $N_i(O)$ by N_i for simplicity of notation, we shall prove that $|N_i| < \infty$ almost surely for $i = 1,2,3$. Let us first observe that the following inclusion relations hold:

$$\{|N_1(O)| + |N_2(O)| = \infty\} \subset \{T_O = \infty\} \cup \left(\bigcup_{n\geq 1}\{N_1(O) = \infty,\; T_O < n\}\right),$$

and

$$\{|N_1(O)| + |N_2(O)| = \infty,\; T_O < n\} \subset$$
$$\{|\text{Arrivals in } (0,n) \text{ of the Poisson process in the origin}| = \infty,\; T_O < n\}.$$

Recalling that $\mathbb{P}(\{T_O = \infty\}) = 0$ we obtain

$$P(|N_1(O)| + |N_2(O)| = \infty) \leq$$
$$\sum_{n=1}^{\infty}\mathbb{P}(|\text{Arrivals in } (0,n) \text{ of the Poisson process in the origin}| = \infty) = 0.$$

Thus we only need to show that $|N_3(O)|$ is almost surely finite. First we observe that one has

$$(16)\quad e_1(t) =$$
$$= \lim_{n\to\infty}\frac{1}{|\Lambda_n|}\sum_{\tau\in N_1(t,\Lambda_n)} H_{\Lambda_n}(\sigma(\tau)) - H_{\Lambda_n}(\sigma(\tau^-)) \leq \mathbb{E}\rho_O^3 \quad\text{a.s.},$$

because the number of changes in the origin $N_1(t,O)$ is at most ρ_O (the maximum number of enlargements of the box observed by the agent x), and in any change the energy can increase at most by ρ_O^2. Finally, the spatial ergodicity yields the almost sure upper bound in (16).

At any time $\tau \in N_3$ the energy $H_\Lambda(\sigma(t))$ decreases at least of one unit, i.e. $H_\Lambda(\sigma(t)) \leq H_\Lambda(\sigma(t^-)) - 1$, otherwise the agent does not change opinion. Thus

$$
\begin{aligned}
(17) \quad e_3(t) &= \lim_{n\to\infty} \frac{1}{|\Lambda_n|} \sum_{\tau \in N_3(t,\Lambda_n)} H_{\Lambda_n}(\sigma(\tau)) - H_{\Lambda_n}(\sigma(\tau^-)) \\
&\leq \lim_{n\to\infty} \frac{-|N_3(t,\Lambda_n)|}{|\Lambda_n|} = -\mathbb{E}|N_3(t,O)|,
\end{aligned}
$$

where we have used once more the spatial ergodicity.

By Lemma 2 and noting that the energy is initially zero (because agents choose $+1$ or -1 with probability $1/2$), one has the following inequality

$$
-\mathbb{E}\rho^2 \leq e(t) = e_1(t) + e_2(t) + e_3(t) \leq e_1(t) + e_3(t),
$$

which holds uniformly in time t. Using inequalities (16) and (17) we obtain $\mathbb{E}|N_3(t,O)| \leq \mathbb{E}\rho_O^3 + \mathbb{E}\rho_O^2 \leq \infty$ uniformly in t, hence also in the limit as $t \to \infty$. But $\mathbb{E}|N_3(O)| < \infty$ obviously implies $|N_3(O)| < \infty$ a.s., so we have shown that $M_O \leq |N_1(O)| + |N_2(O)| + |N_3(O)| < \infty$ and the proof is complete. $\qquad\square$

Remark 5. We can also deduce, following the proof of Theorem 3, that

$$
P(N_3 > C) \leq \frac{E(\rho_O^2) + E(\rho_O^3)}{C},
$$

as an immediate consequence of Markov's inequality.

Remark 6. Let us briefly comment on the connection between the fixation result just proved and the results of De Santis and Newman [3]. The improvement is twofold: namely, the dynamics considered here does not coincide (locally) with zero-temperature dynamics. It is immediate to prove that at any given time there is at least an agent which does not follow the zero-temperature dynamics. This implies that on any time interval the zero-temperature dynamics and our dynamics are almost surely different. Our could say, perhaps somewhat informally, that our dynamics is a perturbation of zero-temperature dynamics with the property of preserving fixation. Moreover, as already mentioned several times, our dynamics is non-Markovian, while the arguments used in [3] hold only for Markovian dynamics.

<div align="center">REFERENCES</div>

[1] P. Billingsley. *Probability and measure*. John Wiley & Sons Inc., New York, third edition, 1995.

[2] C. Cooper and A. M. Frieze. Mixing properties of the Swendsen-Wang process on classes of graphs. *Random Structures Algorithms*, 15(3-4):242–261, 1999.

[3] E. De Santis and C. M. Newman. Convergence in energy-lowering (disordered) sto-
    chastic spin systems. *J. Statist. Phys.*, 110(1-2):431–442, 2003.
[4] R. Durrett. *Random graph dynamics.* Cambridge University Press, Cambridge, 2007.
[5] M. Dyer and A. Frieze. Randomly coloring graphs with lower bounds on girth and
    maximum degree. *Random Structures Algorithms*, 23(2):167–179, 2003.
[6] H. Föllmer. Random economies with many interacting agents. *J. Math. Econom.*,
    1(1):51–62, 1974.
[7] L. R. Fontes. *Spin dynamics at zero temperature.* Publicações Matemáticas do IMPA.
    IMPA, Rio de Janeiro, 2003.
[8] A. Frieze and J. Vera. On randomly colouring locally sparse graphs. *Discrete Math.
    Theor. Comput. Sci.*, 8(1):121–127 (electronic), 2006.
[9] O. Häggström. Zero-temperature dynamics for the ferromagnetic Ising model on
    random graphs. *Phys. A*, 310(3-4):275–284, 2002.
[10] J. P. L. Hatchett, B. Wemmenhove, I. Pérez Castillo, T. Nikoletopoulos, N. S. Skant-
    zos, and A. C. C. Coolen. Parallel dynamics of disordered Ising spin systems on
    finitely connected random graphs. *J. Phys. A*, 37(24):6201–6220, 2004.
[11] U. Horst. Stationary equilibria in discounted stochastic games with weakly interact-
    ing players. *Games Econom. Behav.*, 51(1):83–108, 2005.
[12] U. Horst and J. Scheinkman. Equilibria in systems of social interactions. *J. Econom.
    Theory*, 130(1):44–77, 2006.
[13] T. M. Liggett. *Interacting particle systems.* Springer-Verlag, New York, 1985.
[14] T. M. Liggett and S. W. W. Rolles. An infinite stochastic model of social network
    formation. *Stochastic Process. Appl.*, 113(1):65–80, 2004.
[15] S. Nanda, C. M. Newman, and D. L. Stein. Dynamics of Ising spin systems at zero
    temperature. In *On Dobrushin's way. From probability theory to statistical physics*,
    volume 198 of *Amer. Math. Soc. Transl. Ser. 2*, pages 183–194. Amer. Math. Soc.,
    Providence, RI, 2000.
[16] Esa Nummelin. On the existence and convergence of price equilibria for random
    economies. *Ann. Appl. Probab.*, 10(1):268–282, 2000.
[17] I. Pérez Castillo and N. S. Skantzos. The Little-Hopfield model on a sparse random
    graph. *J. Phys. A*, 37(39):9087–9099, 2004.
[18] M. Talagrand. *Spin glasses: a challenge for mathematicians.* Springer-Verlag,
    Berlin, 2003.

EMILIO DE SANTIS, DIPARTIMENTO DI MATEMATICA "G. CASTELNUOVO", UNIVERSITÀ
DI ROMA "LA SAPIENZA", P.LE ALDO MORO 2, 00185 ROMA, ITALY,
E-MAIL: DESANTIS@MAT.UNIROMA1.IT

CARLO MARINELLI, INSTITUT FÜR ANGEWANDTE MATHEMATIK, RHEINISCHE FRIEDRICH-
WILHELMS-UNIVERSITÄT BONN, WEGELERSTR. 6, D-53115 BONN, GERMANY,
E-MAIL: CM788@UNI-BONN.DE