

An Infinite Stochastic Model of Social Network Formation

Thomas M. Liggett^{1 2}

Silke W.W. Rolles³

January 21, 2004

Abstract

We consider an infinite interacting particle system in which individuals choose neighbors according to evolving sets of probabilities. If x chooses y at some time, the effect is to increase the probability that y chooses x at later times. We characterize the extremal invariant measures for this process. In an extremal equilibrium, the set of individuals is partitioned into finite sets called stars, each of which includes a “center” that is always chosen by the other individuals in that set.

1 Introduction

A number of recent papers have provided a rigorous analysis of stochastic models of social networks based on some form of behavior reinforcement. Examples are [BL03], [SP00], [PS03a], and [PS03b]. In the first of these, the model involves exchanges of gifts or rewards among finitely many individuals. When individual i rewards individual j , that action makes it more likely that j will reward i at some later time. The main result is that the system converges into an equilibrium that consists of randomly chosen collections of stars. A star has a center who rewards the other individuals in its star, while these other individuals reward only their center. The model in [SP00] is quite different, yet the conclusions are virtually the same.

To be more specific, the model analyzed in [BL03] is a discrete time Markov chain whose (uncountable) state space is the set of all $N \times N$ stochastic matrices $p = (p(i, j))$ with zero diagonal entries. The rows and columns correspond to N individuals and $p(i, j)$ represents the current probability that individual i decides to give a gift to individual j . If such a gift is made, then $p(j, i)$ is increased, and the other entries of the j th row of the matrix are decreased by a factor to retain the stochasticity of the matrix. At each time, i is chosen randomly and uniformly from $\{1, \dots, N\}$, and then the gift is made by i to a recipient chosen with probabilities given by the i th row of the matrix. Then the matrix is updated.

Initially, this model was studied by Sociologist P. Bonacich via computer simulations. He observed that there were typically many different limiting states for the chain, and asked what all possible limiting states are. The rigorous answer was provided in [BL03].

¹Department of Mathematics, University of California, Los Angeles, Box 951555, Los Angeles, CA 90095-1555, USA. Email: tml@math.ucla.edu

²Research supported in part by NSF Grants DMS-00-70465 and DMS-03-01795

³Department of Mathematics, University of California, Los Angeles, Box 951555, Los Angeles, CA 90095-1555, USA. Email: srolles@math.ucla.edu

We will give precise statements of the main results from that paper in Section 2 of the current paper in a form that facilitates their application here. As discussed in [BL03], this model is closely related to stochastic learning models, processes of randomly chosen maps, and random systems with complete connections.

In the present paper, we formulate a version of the [BL03] model with infinitely many individuals. Its form puts it in the general class of models studied in the area of interacting particle systems ([Lig85]). As is common in this subject, the uniform choice of individuals in discrete time is replaced by choices of individuals in continuous time at event times of independent Poisson processes. We determine the structure of the set of extremal invariant measures for the system. As in the finite case considered in [BL03], such a measure consists of a collection of finite stars. We borrow some tools from the finite case, but a number of aspects of the proof require modifications due to the presence of infinitely many individuals. We describe the model in the next section, and carry out its analysis in Sections 3 and 4.

2 Description of the model and result

2.1 Infinitely many individuals

Let $G = (V, E)$ be a simple graph with vertex set V and edge set E , i.e. we assume that there are no parallel edges with the same orientation and that every edge has two distinct endpoints. The graph G can be finite or infinite. All edges are directed. We assume that the vertices have uniformly bounded degree. If $(u, v) \in E$, we require that the reversed edge $(v, u) \in E$ as well. For a vertex $v \in V$, let E_v denote the set of edges with tail v , and let N_v be the set of neighbors of v :

$$\begin{aligned} E_v &:= \{(v, u) \in E\}, \\ N_v &:= \{u \in V : (v, u) \in E\}. \end{aligned}$$

For $v \in V$, we define the simplex

$$\Delta_v := \left\{ (x(e))_{e \in E_v} \in [0, 1]^{E_v} : \sum_{e \in E_v} x(e) = 1 \right\}.$$

The state space of our process $(\eta_t)_{t \geq 0}$ is $X = \prod_{v \in V} \Delta_v$, which is compact in the product topology. We have $\eta_t = (\eta_t(v))_{v \in V}$ and $\eta_t(v) = (\eta_t(e))_{e \in E_v}$ with $\sum_{e \in E_v} \eta_t(e) = 1$.

For $e = (u, v) \in E$ and $\eta \in X$, we define $\eta_e \in X$ as follows:

$$\eta_e(e') := \begin{cases} \frac{1 + \eta(e')}{2} & \text{if } e' = (v, u), \\ \frac{\eta(e')}{2} & \text{if } e' \in E_v \setminus \{(v, u)\}, \\ \eta(e') & \text{otherwise.} \end{cases} \quad (2.1)$$

Our process makes the transition $\eta \rightarrow \eta_e$ at rate $\eta(e)$. We can think of the process as follows: Attached to all sites are independent exponential clocks with rate 1. If the clock at site u rings, the individual at u randomly picks a neighbor $v \in N_u$ with probability $\eta(u, v)$ and the current configuration η is changed to η_e with $e = (u, v)$. We denote the distribution of the process started in the deterministic configuration η by P^η . The infinitesimal generator of the process is given by

$$\Omega f(\eta) := \sum_{e \in E} \eta(e) [f(\eta_e) - f(\eta)] \quad (2.2)$$

for continuous functions f depending on only finitely many coordinates $\eta(e)$. The fact that the Markov process with this generator is well defined is a consequence of Theorem 3.9 of Chapter I of [Lig85].

The extremal invariant measures for this process will be described in terms of “stars” and “constellations”. This leads us to the following definitions.

Definition 2.1 *We call a finite subset $A \subset V$ a star with center a if $A = \{a\} \cup \partial A$ for some non-empty set $\partial A \subseteq N_a$. Sometimes we denote the star A by $(A; a)$. A constellation $(A_k; a_k)_{k \in K}$ is a partition of the vertex set V into stars A_k with center a_k .*

Definition 2.2 *Let A be a star with center a , and let $\eta \in X$. We write $\eta = A$ if the following hold:*

$$\sum_{b \in \partial A} \eta(a, b) = 1, \quad (2.3)$$

$$\eta(b, a) = 1 \quad \text{for all } b \in \partial A, \quad (2.4)$$

$$\eta(v, b) = 0 \quad \text{for all } b \in \partial A, v \in N_b \setminus \{a\}, \text{ and} \quad (2.5)$$

$$\eta(v, a) = 0 \quad \text{for all } v \in N_a \setminus \partial A. \quad (2.6)$$

We write $\eta_t \rightarrow A$ if (2.3)–(2.6) hold as $t \rightarrow \infty$ with η replaced by η_t and “=” replaced by “ \rightarrow ”.

Let $(A_k; a_k)_{k \in K}$ be a constellation. We write $\eta = (A_k; a_k)_{k \in K}$ if $\eta = (A_k; a_k)$ for all $k \in K$. Similarly, we write $\eta_t \rightarrow (A_k; a_k)_{k \in K}$ if $\eta_t \rightarrow (A_k; a_k)$ for all $k \in K$.

2.2 Finitely many individuals

Before we state our main result, we describe results for a similar model which was studied in [BL03].

In this subsection, let $G = (V, E)$ be the complete graph without loops on N points, i.e. $V = \{1, 2, \dots, N\}$ and $E = \{(i, j) \in V^2 : i \neq j\}$. Let $\mathbb{N}_0 := \{0, 1, 2, \dots\}$ denote the set of non-negative integers. The model considered in [BL03] is a discrete time Markov chain $(p_n)_{n \in \mathbb{N}_0}$ whose state space is $\prod_{v \in V} \Delta_v$. We have $p_n = (p_n(i, j))_{1 \leq i, j \leq N}$ with $\sum_{j=1}^N p_n(i, j) = 1$ for all i . In other words, p_n is a stochastic $N \times N$ -matrix. Given the state $p_n = p$ of the chain at time n , p_{n+1} is obtained as follows: Choose $i \in \{1, 2, \dots, N\}$

with probability $1/N$ and then choose $j \in \{1, 2, \dots, N\}$ with probability $p_n(i, j)$. Then, using the notation (2.1), $p_{n+1} = p_{(i,j)}$.

Bonacich and Liggett studied a more general model. We describe only the case where their parameters $c_{i,j}$ are all equal to $1/2$. Below, we will need some of their results for this special case.

Let \mathcal{P}^p denote the distribution of the Markov chain $(p_n)_{n \in \mathbb{N}_0}$ with initial state p . The following convergence results were proved in [BL03]:

Theorem 2.3 (Theorem 1.1 [BL03]) (a) *We have*

$$\mathcal{P}^p(p_n \rightarrow (A_k; a_k)_{k=1, \dots, l} \text{ for some constellation } (A_k; a_k)_{k=1, \dots, l}) = 1.$$

(b) *A.s. on the event $\{p_n \rightarrow (A_k; a_k)_{k=1, \dots, l}\}$, p_n converges in distribution to a limit p_∞ where*

$$E p_\infty(a_k, i) = \frac{1}{|\partial A_k|}, \quad i \in \partial A_k.$$

(c) *For each constellation $(A_k; a_k)_{k=1, \dots, l}$,*

$$\mathcal{P}^p(p_n \rightarrow (A_k; a_k)_{k=1, \dots, l}) > 0$$

if and only if $p(i, a_k) + p(a_k, i) > 0$ for all $i \in \partial A_k$, $k = 1, \dots, l$.

For an edge (b, a) , we denote by $\delta_{(b,a)}$ the probability measure on Δ_b which gives mass 1 to the point $\eta(b, a) = 1$, $\eta(b, v) = 0$ for all $v \in N_b \setminus \{a\}$.

Let $(A_k; a_k)_{k=1, \dots, l}$ be a constellation, and consider the process $(p_n)_{n \in \mathbb{N}_0}$ restricted to the set of configurations η satisfying $\eta = (A_k; a_k)_{k=1, \dots, l}$. Bonacich and Liggett proved the following about the stationary distribution of this restricted process:

Proposition 2.4 (Proposition 2.1 [BL03]) *Let $(A_k; a_k)_{k=1, \dots, l}$ be a constellation. The Markov chain $(p_n)_{n \in \mathbb{N}_0}$ restricted to $\{p_n = (A_k; a_k)_{k=1, \dots, l} \text{ for all } n\}$ has a unique stationary distribution, and p_n converges weakly to it for any initial state p satisfying $p = (A_k; a_k)_{k=1, \dots, l}$. For the limiting distribution p_∞ ,*

$$E p_\infty(a_k, i) = \frac{1}{|\partial A_k|}, \quad i \in \partial A_k.$$

For a star $(A; a)$, let

$$\lambda_{A,a}(d\eta(a)) \otimes \bigotimes_{b \in \partial A} \delta_{(b,a)}(d\eta(b))$$

be the stationary distribution for the Markov chain $(p_n)_{n \in \mathbb{N}_0}$ in the case where the underlying graph has the vertex set A and the edge set $\{(a, v) : v \in \partial A\} \cup \{(v, a) : v \in \partial A\}$.

For $m_b \in \mathbb{N}_0$, let $M(m_b, b \in \partial A) := \int \prod_{b \in \partial A} \eta(a, b)^{m_b} \lambda_{A;a}(d\eta(a))$. Let $M_{b_0, j}(m_b, b \in \partial A) := M(m'_b, b \in \partial A)$ where $m'_{b_0} := j$ and $m'_b := m_b$ for all $b \neq b_0$. Then, the moments of $\lambda_{A;a}$ satisfy the following equations:

$$|\partial A| M(m_b, b \in \partial A) = \sum_{b \in \partial A} \frac{1}{2^m} \sum_{i=0}^{m_b} \binom{m_b}{i} M_{b,i}(m_{b'}, b' \in \partial A) \quad (2.7)$$

for all $m_b \in \mathbb{N}_0$; here $m := \sum_{b \in \partial A} m_b$.

In case $|\partial A| = 2$, $\lambda_{A;a}$ is the uniform distribution on Δ_a as observed in [BL03]. For larger ∂A there appears to be no similar description of $\lambda_{A;a}$.

2.3 Result

In this subsection, we state our main result for the process $(\eta_t)_{t \geq 0}$ with infinitely many individuals.

Definition 2.5 For a constellation $\mathcal{C} = (A_n; a_n)_{n \geq 1}$, we define

$$\nu_{\mathcal{C}}(d\eta) := \bigotimes_{n=1}^{\infty} \left(\lambda_{A_n; a_n}(d\eta(a_n)) \otimes \left(\bigotimes_{b \in \partial A_n} \delta_{(b, a_n)} \right) (d\eta(b)) \right).$$

Under $\nu_{\mathcal{C}}$, the coordinates $\eta(v)$, $v \in V$, are independent. Our main result is the following:

Theorem 2.6 The set \mathcal{I}_e of extremal invariant measures for the process $(\eta_t)_{t \geq 0}$ is given by

$$\mathcal{I}_e = \{\nu_{\mathcal{C}} : \mathcal{C} \text{ is a constellation}\}.$$

It would be interesting to study the convergence of the process $(\eta_t)_{t \geq 0}$.

3 Convergence to a star with positive probability

Throughout this section, we fix a star $A = \{a\} \cup \partial A$ with center a . In the following, usually b denotes a point of the boundary ∂A . The aim of this section is to prove the following proposition which states that $P^\eta(\eta_t \rightarrow A) > 0$ under appropriate conditions on η .

Proposition 3.1 If $\eta \in X$ satisfies

$$\eta(b, a) > 0 \quad \text{for all } b \in \partial A \quad \text{and} \quad (3.1)$$

$$\sum_{b \in A} \eta(v, b) < 1 \quad \text{for all } v \notin A, \quad (3.2)$$

then $P^\eta(\eta_t \rightarrow A) > 0$. Hence, for all $\eta \in X$, there exists $c(\eta) > 0$ such that the following inequality holds:

$$P^\eta(\eta_t \rightarrow A) \geq c(\eta) \left[\prod_{b \in \partial A} \eta(b, a) \right] \prod_{v \in V \setminus A} \left[1 - \sum_{b \in A} \eta(v, b) \right]. \quad (3.3)$$

In order to prove Proposition 3.1, we compare the Markov process $(\eta_t)_{t \geq 0}$ with a Markov process $(\zeta_t)_{t \geq 0}$ with slightly different rates.

Definition 3.2 *The Markov process $(\zeta_t)_{t \geq 0}$ makes the following transitions:*

$$\begin{aligned} \zeta &\rightarrow \zeta_{(v,b)} && \text{at rate 0 instead of rate } \zeta(v, b) \text{ for all } b \in \partial A, v \in N_b \setminus \{a\}, \\ \zeta &\rightarrow \zeta_{(v,a)} && \text{at rate 0 instead of rate } \zeta(v, a) \text{ for all } v \in N_a \setminus \partial A; \end{aligned}$$

all other transitions occur at the same rates as for the process $(\eta_t)_{t \geq 0}$. We denote the distribution of $(\zeta_t)_{t \geq 0}$ started in the deterministic configuration ζ by Q^ζ .

The effect of this choice is described by the following lemma:

Lemma 3.3 *The process $(\zeta_t(v); v \in A)$ is Markovian. Also, some of the coordinates are monotone in time:*

$$\zeta_t(b, v) \downarrow \quad \text{for } b \in \partial A, v \in N_b \setminus \{a\}, \quad (3.4)$$

$$\zeta_t(a, v) \downarrow \quad \text{for } v \in N_a \setminus \partial A, \quad (3.5)$$

$$\zeta_t(b, a) \uparrow \quad \text{for } b \in \partial A. \quad (3.6)$$

Proof. We note that by the definition of the process $(\zeta_t)_{t \geq 0}$, the following hold:

$$\begin{aligned} \zeta(a, b) &\rightarrow \frac{\zeta(a, b)}{2} && \text{at rate } \sum_{b' \in \partial A \setminus \{b\}} \zeta(b', a) \text{ for all } b \in \partial A, \\ \zeta(a, v) &\rightarrow \frac{\zeta(a, v)}{2} && \text{at rate } \sum_{b \in \partial A} \zeta(b, a) \text{ for all } v \in N_a \setminus \partial A, \text{ and} \\ \zeta(b, v) &\rightarrow \frac{\zeta(b, v)}{2} && \text{at rate } \zeta(a, b) \text{ for all } b \in \partial A \text{ and all } v \in N_b \setminus \{a\}. \end{aligned}$$

Hence, $(\zeta_t(v); v \in A)$ Markovian.

The statements (3.4) and (3.5) follow immediately from the definition of the process $(\zeta_t)_{t \geq 0}$. The monotonicity (3.6) follows from (3.4). ■

Lemmas 3.4–3.6 below will show that $Q^\eta(\zeta_t \rightarrow A) = 1$ if η satisfies (3.1) and (3.2). Later we will use an absolute continuity argument to deduce Proposition 3.1 from this fact.

Lemma 3.4 *If (3.1) holds, then Q^η -a.s.*

$$\lim_{t \rightarrow \infty} \sum_{b \in \partial A} \zeta_t(a, b) = 1 \quad \text{exponentially rapidly.}$$

Proof. The sum $S(t) = \sum_{v \in N_a \setminus \partial A} \zeta_t(a, v)$ can only decrease. It decreases by an amount $\frac{1}{2}S(t)$ at rate $\sum_{b \in \partial A} \zeta_t(b, a)$. Therefore

$$\frac{d}{dt} E^\eta S(t) = -\frac{1}{2} E^\eta \left[S(t) \sum_{b \in \partial A} \zeta_t(b, a) \right];$$

here E^η denotes the expectation with respect to Q^η . On the other hand, $\zeta_t(b, a)$ can only increase for $b \in \partial A$, so that

$$\frac{d}{dt} E^\eta S(t) \leq -\frac{1}{2} \left[\sum_{b \in \partial A} \eta(b, a) \right] E^\eta S(t).$$

Solving this differential inequality yields

$$E^\eta S(t) \leq \left[\sum_{v \in N_a \setminus \partial A} \eta(a, v) \right] \exp \left(-\frac{t}{2} \left[\sum_{b \in \partial A} \eta(b, a) \right] \right).$$

Therefore, since $S(t)$ is monotone, $S(t) \rightarrow 0$ exponentially rapidly Q^η -a.s., so

$$\sum_{b \in \partial A} \zeta_t(a, b) = 1 - S(t) \rightarrow 1$$

exponentially rapidly Q^η -a.s. ■

Lemma 3.5 *If (3.1) holds, then Q^η -a.s.*

$$\lim_{t \rightarrow \infty} \zeta_t(b, a) = 1 \quad \text{exponentially rapidly for all } b \in \partial A.$$

Proof. Fix a $b \in \partial A$. Then

$$\zeta_t(b, a) \rightarrow \frac{1 + \zeta_t(b, a)}{2} \quad \text{at rate } \zeta_t(a, b),$$

and

$$\zeta_t(a, b) \rightarrow \begin{cases} \frac{1}{2} \zeta_t(a, b) & \text{at rate } \sum_{b' \in \partial A \setminus \{b\}} \zeta_t(b', a) \\ \frac{1 + \zeta_t(a, b)}{2} & \text{at rate } \zeta_t(b, a). \end{cases}$$

Therefore, we may couple $(\zeta_t(b, a), \zeta_t(a, b))$ with a process (X_t, Y_t) so that $X_0 = \eta(b, a)$, $Y_0 = \eta(a, b)$,

$$\zeta_t(b, a) \geq X_t, \quad \zeta_t(a, b) \geq Y_t,$$

and

$$X_t \rightarrow \frac{1 + X_t}{2} \quad \text{at rate } Y_t, \quad Y_t \rightarrow \begin{cases} 0 & \text{at rate } |\partial A| - 1 \\ \frac{1}{2} & \text{at rate } \eta(b, a). \end{cases}$$

The fact that this is possible follows from

$$\sum_{b' \in \partial A \setminus \{b\}} \zeta_t(b', a) \leq |\partial A| - 1 \quad \text{and} \quad \zeta_t(b, a) \geq \eta(b, a).$$

Note that Y_t is a two state Markov chain. Also, since $(1 - X_t) \rightarrow (1 - X_t)/2$ at rate Y_t , the conditional distribution of $1 - X_t$ given the process $(Y_s)_{s \geq 0}$ is that of $(1 - X_0)/2^N$, where N is Poisson with parameter $\lambda = \int_0^t Y_s ds$. Therefore,

$$E(1 - X_t) = (1 - X_0)E \exp \left[-\frac{1}{2} \int_0^t Y_s ds \right].$$

The right side above tends to 0 exponentially rapidly, so since X_t is monotone, $X_t \rightarrow 1$ exponentially rapidly a.s. The same is true for $\zeta_t(b, a)$, since it is bounded below by X_t .

■

Lemma 3.6 *If (3.1) and (3.2) hold, then Q^η -a.s.*

$$\begin{aligned} \int_0^\infty \zeta_t(v, b) dt &< \infty \quad \text{for all } b \in \partial A, v \in N_b \setminus \{a\} \text{ and} \\ \int_0^\infty \zeta_t(v, a) dt &< \infty \quad \text{for all } v \in N_a \setminus \partial A. \end{aligned} \quad (3.7)$$

Proof. If $v \in \partial A$, (3.7) is immediate from Lemma 3.5 and the fact that $\sum_w \zeta_t(v, w) = 1$. So, we may assume that $v \notin A$. Fix such a v . Let b_1, \dots, b_k be the members of A such that $(v, b_i) \in E$. We need to show that

$$\int_0^\infty \zeta_t(v, b_i) dt < \infty \quad Q^\eta - a.s.$$

for all $i \in \{1, 2, \dots, k\}$. Recall that $(\zeta_t(u); u \in A)$ is Markov by Lemma 3.3, so we can and will consider the process obtained by conditioning on $(\zeta_t(u); u \in A)_{t \geq 0}$. For $1 \leq i \leq k$, let M_t^i be (conditionally) independent temporally inhomogeneous Poisson processes with rates $\zeta_t(b_i, v)$ at time t . By Lemmas 3.4 and 3.5, $M = \sum_{i=1}^k M_\infty^i < \infty$ a.s. In what follows, we will also condition on the processes M_t^i for $1 \leq i \leq k$.

Now we proceed somewhat like we did in the proof of Lemma 3.5. There are a few extra complications resulting from the fact that all processes corresponding to X_t and Y_t can increase and decrease. Let

$$K = \sum_{v' \in N_v \setminus A} |N_{v'}|.$$

Now let $(X_t^1, \dots, X_t^k, Y_t)$ be coupled to the process $(\zeta_t(u), u \notin A)$ so that $X_0^i = \eta(v, b_i)$, $Y_0 = 0$,

$$\zeta_t(v, b_i) \leq X_t^i, \quad \sum_{v' \in N_v \setminus A} \zeta_t(v', v) \geq Y_t,$$

$$X_t^i \rightarrow \frac{1}{2}X_t^i \quad \text{together at rate } Y_t,$$

and

$$Y_t \rightarrow \begin{cases} 0 & \text{at rate } K \\ \frac{1}{2} & \text{at rate } \left(1 - \sum_{i=1}^k \eta(v, b_i)\right) 2^{-M}. \end{cases}$$

In addition, $X_t^i \rightarrow (1 + X_t^i)/2$ at the event times of the process M_t^i . The fact that this coupling is possible comes from the following considerations:

(a) the rate at which the $\zeta_t(v, b_i)$'s decrease (together) is

$$\sum_{v' \in N_v \setminus A} \zeta_t(v', v) \geq Y_t,$$

(b) the rate at which $\sum_{v' \in N_v \setminus A} \zeta_t(v', v)$ decreases is at most K , and

(c) the rate at which $\sum_{v' \in N_v \setminus A} \zeta_t(v', v)$ increases is

$$\sum_{v' \in N_v \setminus A} \zeta_t(v, v') = 1 - \sum_{i=1}^k \zeta_t(v, b_i) \geq \left(1 - \sum_{i=1}^k \eta(v, b_i)\right) 2^{-M}, \quad (3.8)$$

and if it does increase, its new value is at least $\frac{1}{2}$.

Let us explain why the inequality (3.8) holds: Clearly, Q^η -a.s., $1 - \sum_{i=1}^k \zeta_0(v, b_i) = 1 - \sum_{i=1}^k \eta(v, b_i)$. The quantity

$$1 - \sum_{i=1}^k \zeta_t(v, b_i)$$

decreases by a factor of $\frac{1}{2}$ each time there is an event time for one of the processes M^i . All other transitions only decrease the values of the $\zeta_t(v, b_i)$'s.

The proof that $X_t^i \rightarrow 0$ exponentially rapidly is essentially the same as in the proof of Lemma 3.5. The main point is that the transition $X_t^i \rightarrow (1 + X_t^i)/2$ occurs only finitely many times a.s. since $M_\infty^i < \infty$ a.s., while the transition $X_t^i \rightarrow X_t^i/2$ occurs at rate $1/2$ a positive fraction of the time. The finitely many increases in X_t^i does not change its exponential decay to 0. Since $\zeta_t(v, b_i) \leq X_t^i$, it follows that $\zeta_t(v, b_i) \rightarrow 0$ exponentially rapidly as well. ■

The main property of Q^ζ that we will need is its absolute continuity with respect to the distribution P^ζ of the process $(\eta_t)_{t \geq 0}$ starting at ζ . This issue has been studied in various contexts – primarily diffusion processes in Euclidean spaces. (See Chapter 4 of [Sko65], for example.) Since none of them quite fit our situation, we provide a proof of the result we need. Let Q_t^ζ and P_t^ζ be the distributions of the two processes up to time t .

Lemma 3.7 *For any $t > 0$, $Q_t^\zeta \ll P_t^\zeta$ and*

$$\log \left(\frac{dQ_t^\zeta}{dP_t^\zeta}(\zeta \bullet) \right) = \int_0^t \left[\sum_{b \in \partial A} \sum_{v \in N_b \setminus \{a\}} \zeta_s(v, b) + \sum_{v \in N_a \setminus \partial A} \zeta_s(v, a) \right] ds \quad Q_t^\zeta - a.s. \quad (3.9)$$

Proof. Note that the right side of (3.9) is a continuous function of the path ζ_\bullet , and depends on the values of $\zeta_s(e)$ for only finitely many edges e . Therefore, as we will see, it is enough to prove the lemma for the processes obtained by suppressing all transitions outside of a large finite subgraph of G .

To see this, we will abstract the situation a bit. Suppose that μ_n and ν_n are probability measures such that $\nu_n \ll \mu_n$, C is a closed set that supports ν_n for each n , h is a nonnegative continuous function so that $d\nu_n/d\mu_n = h$ on C for each n , $\mu_n \rightarrow \mu$ and $\nu_n \rightarrow \nu$ weakly, and $\mu_n(C) \rightarrow \mu(C)$. Then for any continuous function f taking values in $[0, 1]$,

$$\int f d\nu_n = \int_C f h d\mu_n.$$

We may pass to the limit as $n \rightarrow \infty$ on the left side with no difficulty. For the right side, note that $f h 1_C$ is an upper semicontinuous function, so

$$\int_C f h d\mu \geq \limsup_{n \rightarrow \infty} \int_C f h d\mu_n. \quad (3.10)$$

Inequality (3.10) also holds with h replaced by 1 and/or with f replaced by $1 - f$. Therefore

$$\int_C f d\mu \geq \limsup_{n \rightarrow \infty} \int_C f d\mu_n \quad \text{and} \quad \int_C (1 - f) d\mu \geq \limsup_{n \rightarrow \infty} \int_C (1 - f) d\mu_n. \quad (3.11)$$

The sum of the left sides of (3.11) is $\mu(C)$, while the sum of the integrals that appear on the right of (3.11) is $\mu_n(C)$. Since $\mu_n(C) \rightarrow \mu(C)$, (3.11) holds with equality in both cases, and with each \limsup replaced by \lim . Therefore, $\int_C f d\mu = \lim_{n \rightarrow \infty} \int_C f d\mu_n$ for any bounded continuous function f . Applying this to $f = \min(h, k)$ and letting $k \rightarrow \infty$ gives

$$\int_C h d\mu \leq \liminf_{n \rightarrow \infty} \int_C h d\mu_n.$$

Combining this with (3.10) with $f \equiv 1$, we see that

$$\int_C h d\mu = \lim_{n \rightarrow \infty} \int_C h d\mu_n = \lim_{n \rightarrow \infty} \nu_n(C) = 1.$$

A similar argument using this fact and (3.10) for both f and $1 - f$ gives

$$\int_C f h d\mu = \lim_{n \rightarrow \infty} \int_C f h d\mu_n,$$

so that

$$\int f d\nu = \int_C f h d\mu,$$

and hence $\nu \ll \mu$ and $d\nu/d\mu = h$ on C . In our application, C is the set of paths $(\zeta_s)_{0 \leq s \leq t}$ so that $\zeta_s(b, v)$ is nonincreasing for $b \in \partial A, v \in N_b \setminus \{a\}$, and $\zeta_s(a, v)$ is nonincreasing for $v \in N_a \setminus \partial A$. The fact that distributions of the processes with transitions suppressed

outside subgraphs G_n with $G_n \uparrow G$ converge to the original processes is a consequence of Theorem 2.12 of Chapter I of [Lig85].

When transitions are restricted to a large finite subgraph of G , both processes are countable state continuous time Markov chains. (While Δ_v is not countable, the set of states that can be reached from a given initial configuration is.) For a path $(\zeta_s)_{0 \leq s \leq t}$ in the support of Q_t^ζ , the rates of the transitions that occur in that path are the same for both processes. One process is obtained from the other by setting some of the rates to zero.

So, consider this situation. Suppose X_t is a Markov chain with transition rates $q(x, y)$, and let

$$c(x) = \sum_{y: y \neq x} q(x, y)$$

be the total rate for transitions out of x . Consider another Markov chain X_t^* whose transition rates are $q^*(x, y)$, where for each x, y , $q^*(x, y)$ is either $q(x, y)$ or zero, and let $c^*(x)$ be the corresponding total rate out of x for this process. Consider a path $\gamma(s), 0 \leq s \leq t$ that X_s^* can follow: $\gamma(s) = x_i$ for $t_i \leq s < t_{i+1}$, where $0 = t_0 < t_1 < \dots < t_n < t_{n+1} = t$. (the fact that X_s^* can follow this path means that $q^*(x_i, x_{i+1}) > 0$ for each i .) The density of this path for the process X_s is

$$\left[\prod_{i=0}^{n-1} e^{-c(x_i)(t_{i+1}-t_i)} q(x_i, x_{i+1}) \right] e^{-c(x_n)(t_{n+1}-t_n)},$$

while its density for X_s^* is

$$\left[\prod_{i=0}^{n-1} e^{-c^*(x_i)(t_{i+1}-t_i)} q(x_i, x_{i+1}) \right] e^{-c^*(x_n)(t_{n+1}-t_n)}.$$

The ratio of the second of these to the first (which is the appropriate Radon-Nikodym derivative of the distributions of the two processes on path space) is

$$\exp \left[\int_0^t [c(\gamma(s)) - c^*(\gamma(s))] ds \right].$$

Returning to our models, note that in configuration ζ , the difference between the total transition rates for the two processes is

$$\sum_{b \in \partial A} \sum_{v \in N_b \setminus \{a\}} \zeta(v, b) + \sum_{v \in N_a \setminus \partial A} \zeta(v, a).$$

This gives (3.9) in this case, and thus completes the proof of the lemma. ■

Proof of Proposition 3.1. Suppose $\eta \in X$ satisfies (3.1) and (3.2). Lemmas 3.4–3.6 show that $Q^\eta(\zeta_t \rightarrow A) = 1$. By Lemma 3.7, as $t \rightarrow \infty$,

$$\log \left[\frac{dQ_t^\eta}{dP_t^\eta}(\zeta_\bullet) \right] \uparrow \int_0^\infty \left[\sum_{b \in \partial A} \sum_{v \in N_b \setminus \{a\}} \zeta_s(v, b) + \sum_{v \in N_a \setminus \partial A} \zeta_s(v, a) \right] ds$$

which is finite Q^η -a.s. by Lemma 3.6. Hence, by Theorem (3.3) on page 242 of [Dur96] it follows that $Q^\eta \ll P^\eta$. Consequently, $P^\eta(\eta_t \rightarrow A) > 0$. If $\eta \in X$ does not satisfy (3.1)–(3.2), then the right-hand side of (3.3) equals 0. This completes the proof of the proposition. ■

4 The extremal invariant measures

In this section, we prove Theorem 2.6.

We denote by \mathcal{I} the set of invariant probability measures for the process $(\eta_t)_{t \geq 0}$.

Lemma 4.1 *If $\mu \in \mathcal{I}_e$ and $\mu(\eta(u, v) = 0) > 0$, then $\mu(\eta(u, v) = 0, \eta(v, u) = 0) = 1$.*

Proof. Let $\mu \in \mathcal{I}_e$ with $\mu(\eta(u, v) = 0) > 0$. Note that $\eta_t(u, v) = 0$ only if $\eta_s(u, v) = 0$ for all $s \in [0, t]$. Since $\eta(u, v) \rightarrow \frac{1+\eta(u, v)}{2}$ at rate $\eta(v, u)$, we must have $\mu(\eta(u, v) = 0, \eta(v, u) > 0) = 0$.

Let $B := \{\eta(u, v) = 0, \eta(v, u) = 0\}$. We have just shown that $\mu(B) > 0$. Since $\eta_t \in B$ iff $\eta_0 \in B$ we have for all measurable sets C

$$\begin{aligned} P^{\mu(\cdot|B)}(\eta_t \in C) &= P^\mu(\eta_t \in C | \eta_0 \in B) = P^\mu(\eta_t \in C | \eta_t \in B) \\ &= P^\mu(\eta_0 \in C | \eta_0 \in B) = P^{\mu(\cdot|B)}(\eta_0 \in C). \end{aligned}$$

Consequently, $\mu(\cdot|B) \in \mathcal{I}$. Analogously, if $\mu(B) < 1$, $\mu(\cdot|B^c) \in \mathcal{I}$. Hence, we can write μ as a linear combination of two invariant measures:

$$\mu(\cdot) = \mu(B)\mu(\cdot|B) + \mu(B^c)\mu(\cdot|B^c).$$

Since we assumed μ to be an extremal invariant measure, it follows that $\mu(B) \in \{0, 1\}$ and we conclude $\mu(B) = 1$. ■

Lemma 4.2 *If $\mu \in \mathcal{I}_e$, then $\mu(\eta(u, v) = 0) \in \{0, 1\}$ and $\mu(\eta(u, v) = 1) \in \{0, 1\}$.*

Proof. The statement $\mu(\eta(u, v) = 0) \in \{0, 1\}$ is an immediate consequence of Lemma 4.1.

Suppose $\mu(\eta(u, v) = 1) > 0$. Since $\sum_{v' \in N_u} \eta(u, v') = 1$ we have $\mu(\eta(u, v') = 0) > 0$ and thus $\mu(\eta(u, v') = 0) = 1$ for all $v' \in N_u \setminus \{v\}$. It follows from $\sum_{v' \in N_u} \eta(u, v') = 1$ that $\mu(\eta(u, v) = 1) = 1$. ■

Lemma 4.3 *If $\mu \in \mathcal{I}_e$ and A is a star, then $\mu(\eta = A) \in \{0, 1\}$.*

Proof. Recall the definition of $\eta = A$ from Definition 2.2. Suppose that $\mu(\eta = A) > 0$. Then, by Lemma 4.2, (2.4)–(2.6) hold μ -a.s. Lemma 4.1 implies that $\mu(\eta(v, a) = 0, \eta(a, v) = 0) = 1$ for all $v \in N_a \setminus \partial A$. Hence, (2.3) holds μ -a.s. as well, and we have shown that $\mu(\eta = A) = 1$. ■

Proof of Theorem 2.6. Let A be a star with center a . By Proposition 2.4,

$$\lambda_{A;a}(d\eta(a)) \otimes \left(\bigotimes_{b \in \partial A} \delta_{(b,a)} \right) (d\eta(b))$$

is an invariant measure for the process $(\eta_t)_{t \geq 0}$ restricted to the graph with vertex set A and edge set $\{(a, v) : v \in \partial A\} \cup \{(v, a) : v \in \partial A\}$. Hence, $\nu_{\mathcal{C}} \in \mathcal{I}$ whenever $\mathcal{C} = (A_n; a_n)_{n \geq 1}$ is a constellation.

Suppose that $\nu_{\mathcal{C}} = \alpha\nu_1 + (1 - \alpha)\nu_2$ for some $\alpha \in [0, 1]$ and $\nu_1, \nu_2 \in \mathcal{I}$. By the definition of $\nu_{\mathcal{C}}$, we have

$$\nu_{\mathcal{C}}(\eta = A_n) = 1 \quad \text{for all } n \geq 1.$$

Consequently, $\nu_i(\eta = A_n) = 1$ must hold for $i = 1, 2$ and all $n \geq 1$. Hence, for all i and n , ν_i is an invariant measure for the process $(\eta_t)_{t \geq 0}$ restricted to the set A_n . By Proposition 2.4, the restriction of ν_i to the set $\prod_{b \in A_n} \Delta_b$ equals

$$\lambda_{A_n; a_n}(d\eta(a_n)) \otimes \left(\bigotimes_{b \in \partial A_n} \delta_{(b, a_n)} \right) (d\eta(b)).$$

Consequently, $\nu_i = \nu_{\mathcal{C}}$ for $i = 1, 2$, and we have shown that $\nu_{\mathcal{C}}$ is an extremal invariant measure.

It remains to show that all extremal invariant measures are of the form $\nu_{\mathcal{C}}$ with a constellation $\mathcal{C} = (A_n; a_n)_{n \geq 1}$. Let $\mu \in \mathcal{I}_e$, and let $a \in V$.

Claim: There exists a star A with $\mu(\eta = A) = 1$ having either a as its center or $a \in \partial A$.

Let us explain why the claim implies the theorem: Suppose $\mu(\eta = A) = 1$ for a star A with center a_0 . Then μ -a.s. $\eta(b, a_0) = 1$ for all $b \in \partial A$. Since $\sum_{v \in N_b} \eta(b, v) = 1$ it follows that $\eta(b, v) = 0$ μ -a.s. for all $v \in N_b \setminus \{a_0\}$. Consequently, the distribution of $\eta(b)$ under μ equals $\delta_{(b, a_0)}$ for all $b \in \partial A$. Furthermore, $\mu(\eta = A) = 1$ implies that $\eta(v, b) = 0$ μ -a.s. for all $v \in N_b \setminus \{a_0\}$, $\eta(v, a_0) = 0 = \eta(a_0, v)$ μ -a.s. for all $v \in N_{a_0} \setminus \partial A$. This shows that μ -a.s. the sites in A interact only with sites in A . The dynamics for $\eta_t(a_0, b)$ with $b \in \partial A$ is as follows:

$$\eta_t(a_0, b) \rightarrow \begin{cases} \frac{1 + \eta_t(a_0, b)}{2} & \text{at rate } \eta_t(b, a_0) = 1, \\ \frac{\eta_t(a_0, b)}{2} & \text{at rate } \sum_{v \in N_{a_0} \setminus \{b\}} \eta_t(v, a_0) = |\partial A| - 1. \end{cases}$$

Hence, by Proposition 2.4, the moments of $(\eta_t(a_0, b))_{b \in \partial A}$ under μ satisfy the equation (2.7). Hence $\eta(a_0)\mu = \lambda_{A; a_0}$. Since every site is contained in a star A such that $\mu(\eta = A) = 1$, the statement of the theorem follows.

We turn to the proof of the claim.

If there exists a star A with center a such that $\mu(\eta = A) = 1$, we are done. Otherwise, by Lemma 4.3, we know that $\mu(\eta = A) = 0$ for all stars A with center a . Since μ is invariant, $P^\mu(\eta_t \rightarrow A) = \mu(\eta = A) = 0$. Integrating (3.3) with respect to μ , yields

$$0 = P^\mu(\eta_t \rightarrow A) \geq \int c(\eta) \left[\prod_{b \in \partial A} \eta(b, a) \right] \prod_{v \in V \setminus A} \left[1 - \sum_{b \in A} \eta(v, b) \right] \mu(d\eta).$$

Consequently, for all stars A with center a and μ -almost all $\eta \in X$, the following holds:

$$\eta(b, a) = 0 \quad \text{for some } b \in \partial A \quad \text{or} \quad (4.1)$$

$$\sum_{b \in A} \eta(v, b) = 1 \quad \text{for some } v \notin A. \quad (4.2)$$

Using Lemma 4.2, we see that $\mu(\eta(b, a) = 0) \in \{0, 1\}$. Furthermore,

$$\mu \left(\sum_{b \in A} \eta(v, b) = 1 \right) = \mu(\eta(v, b) = 0 \text{ for all } b \in V \setminus A) \in \{0, 1\}$$

by Lemma 4.2. We enumerate the elements of

$$\{b \in N_a : \eta(b, a) > 0\} :$$

b_1, b_2, \dots, b_k . (The set is μ -a.s. defined.) Let $A_i := \{a, b_i\}$ denote the star with center a and boundary $\{b_i\}$. Since $\eta(b_i, a) > 0$ μ -a.s., an application of (4.1) and (4.2) for the star A_i yields

$$\eta(v_i, a) + \eta(v_i, b_i) = 1 \quad \text{for some } v_i \notin \{a, b_i\}. \quad (4.3)$$

Case 1: $\eta(v_i, b_i) = 1$ for some $i \in \{1, 2, \dots, k\}$.

Fix such a b_i . We define

$$B := \{b_i\} \cup \partial B \quad \text{with} \quad \partial B := \{b \in N_{b_i} : \eta(b, b_i) = 1\}$$

to be the star with center b_i . By assumption, $\partial B \neq \emptyset$. Suppose $\mu(\eta = B) = 0$. Then, by (4.1) and (4.2) for the star B , there exists $v \notin B$ such that $\sum_{b \in B} \eta(v, b) = 1$. Since $\eta(b, b_i) = 1$ μ -a.s. for $b \in \partial B$, we have $\eta(v, b) = 0$ μ -a.s. and Lemma 4.1 implies $\eta(v, b) = 0$ μ -a.s. Hence, $\eta(v, b_i) = 1$ which means $v \in \partial B$. But this contradicts $v \notin B$. Hence $\mu(\eta = B) > 0$ and by Lemma 4.3, $\mu(\eta = B) = 1$. Since $\eta(b_i, a) > 0$, the site a must be in ∂B and we found a star B containing a in its boundary with $\mu(\eta = B) = 1$.

Case 2: $\eta(v_i, b_i) < 1$ for all $i \in \{1, 2, \dots, k\}$.

Then by (4.3), $\eta(v_i, a) > 0$ for all i . Since $\{b_i : 1 \leq i \leq k\} = \{v : \eta(v, a) > 0\}$, we have $v_i = b_{j_i}$ for some $j_i \in \{1, 2, \dots, k\} \setminus \{i\}$ and (4.3) becomes

$$\eta(b_{j_i}, a) + \eta(b_{j_i}, b_i) = 1. \quad (4.4)$$

Suppose that

$$\text{there exists } i \neq i' \text{ with } j_i = j_{i'}. \quad (4.5)$$

Then, (4.4) applied for i' gives

$$\eta(b_{j_i}, a) + \eta(b_{j_i}, b_{i'}) = \eta(b_{j_{i'}}, a) + \eta(b_{j_{i'}}, b_{i'}) = 1. \quad (4.6)$$

Using the identity $\sum_{v \in V} \eta(b_{j_i}, v) = 1$ and adding up (4.4) and (4.6) yields

$$1 \geq \eta(b_{j_i}, a) + \eta(b_{j_i}, b_i) + \eta(b_{j_i}, b_{i'}) = 2 - \eta(b_{j_i}, a).$$

Thus, $\eta(b_{j_i}, a) \geq 1$, and we conclude that

$$\eta(b_{j_i}, a) = 1. \quad (4.7)$$

Consider the star $C := \{a\} \cup \{b_i \in \{b_1, b_2, \dots, b_k\} : \eta(b_i, a) = 1\}$ with center a . By (4.7), $\partial C \neq \emptyset$. An application of (4.1/4.2) shows that there exists $v \notin C$ such that

$$\eta(v, a) + \sum_{\{i: \eta(b_i, a)=1\}} \eta(v, b_i) = 1. \quad (4.8)$$

Consider b_i with $\eta(b_i, a) = 1$. Then $\eta(b_i, v) = 0$ and thus, by Lemma 4.1, $\eta(v, b_i) = 0$ follows. Consequently, (4.8) implies that

$$\eta(v, a) = 1.$$

Since $\{b_i : 1 \leq i \leq k\} = \{v : \eta(v, a) > 0\}$, we have $v = b_i$ for some i . Hence, $v \in C$, which contradicts $v \notin C$. This shows that our assumption (4.5) was wrong. Consequently, $j_i \neq j_{i'}$ whenever $i \neq i'$, i.e. the map $\{1, \dots, k\} \rightarrow \{1, \dots, k\}, i \mapsto j_i$ is one-to-one and thus bijective.

We define $D := \{a, b_1, b_2, \dots, b_k\}$. Equation (4.4) implies that $\eta(b_i, v) = 0$ for all $v \notin D$. Thus, by Lemma 4.1,

$$\eta(v, b_i) = 0 \quad \text{for all } v \notin D.$$

By the definition of the b_i 's,

$$\eta(v, a) = 0 \quad \text{for all } v \notin D,$$

and by Lemma 4.1, $\eta(a, v) = 0$ for all $v \notin D$. This means that the sites in D do not interact with sites outside of D and vice versa. By Theorem 2.3a), we know that there exists a star A containing a such that $\eta_t \rightarrow A$ with positive P^μ -probability. Hence for this star, $\mu(\eta = A) > 0$, and consequently, by Lemma 4.3, $\mu(\eta = A) = 1$. This completes the proof of the claim. ■

Acknowledgement: We would like to thank the referee for carefully reading the article.

References

- [BL03] Phillip Bonacich and Thomas M. Liggett. Asymptotics of a matrix valued Markov chain arising in sociology. *Stochastic Process. Appl.*, 104(1):155–171, 2003.
- [Dur96] R. Durrett. *Probability: Theory and Examples*. Duxbury Press, Second edition, 1996.
- [Lig85] Thomas M. Liggett. *Interacting Particle Systems*, volume 276 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, New York, 1985.
- [PS03a] R. Pemantle and B. Skyrms. Network formation by reinforcement learning: the long and medium run. Preprint, 2003.
- [PS03b] R. Pemantle and B. Skyrms. Time of absorption in discounted reinforcement models. *Stochastic Process. Appl.*, to appear.
- [Sko65] A. V. Skorokhod. *Studies in the theory of random processes*. Translated from the Russian by Scripta Technica, Inc. Addison-Wesley Publishing Co., Inc., Reading, Mass., 1965.
- [SP00] B. Skyrms and R. Pemantle. A dynamic model of social network formation. *Proc. Nat. Acad. Sci.*, 97:9340–9346, 2000.