Consensus on networks

©A. J. Ganesh, University of Bristol

The spread of a rumour is one example of an absorbing Markov process on networks. It was a purely increasing process and so it reached the absorbing state rapidly, in time growing logarithmically in the number of nodes. We shall now consider two models related to agreement between interacting agents. Such models can be used to describe a number of situations - both those in which the agents actively seek to achieve consensus, and those in which consensus corresponds to the outcome of competition between different ideas.

Examples of the latter include competition between different video recording formats (VHS vs. Betamax), or operating systems (Windows vs. Apple Mac, or iPhone vs. Android), or social networks (Facebook vs. MySpace). In some of these cases, there is a substantial switching cost (buying new hardware), while in others the switching cost may be small (time to set up a new account and upload material). There are often "network externalities" - benefits to adopting one of the competing technologies that don't have to do with its intrinsic qualities, but only to do with who and how many have adopted it. In the examples above, developers are more likely to produce films or software for the more popular platform, where they can expect to sell more copies. In the social network example, you would prefer to be on the network that most of your friends are on, because that would make it easier to communicate with them. Thus, the question of which technology gets adopted may depend not just on its quality, but also on the random decisions of early adopters.

Examples of consensus seeking come from biological problems such as quorum sensing, which informs the decisions of bioluminescent bacteria, for example, or flocking, whereby a group of birds move collectively in a coordinated manner. These decisions involve some form of consensus regarding an action, or a direction of movement. It is more complex than simply following a leader, as there may be no leader involved, or the identity of

the leader may change over time without a formal process of leader election. Further examples and applications of consensus models come from computer science, where they form a building block of many distributed computing algorithms.

1 The de Groot model

The first model we look at involves agents who start out with an initial preference that is real-valued, and who want to reach a consensus on a real number that is a (possibly weighted) average of these initial values. An example from above is the case of a flock of birds agreeing on a direction. (In this case, the values are in $\mathbb{R}^3$, but the idea is the same.) The de Groot model is very simple, but was one of the first formal models of "social learning".

The model can be stated precisely as follows. There are n agents, who each have an initial preference or value in $\mathbb{R}$. We denote these initial values by the vector $\mathbf{x} = (x_1, x_2, \ldots x_n)$. The agents are located on the nodes of a directed graph $G = (V, E)$, and can only communicate directly with their neighbours in this graph. Time is discrete, and the agents update their values synchronously as follows. Agent i updates its value in time step $t+1$ to

$$x_i(t+1) = \sum_{j \in V : (i,j) \in E} p_{ij} x_j(t).$$

Here, the p_{ij} are fixed weights, describing how much weight i gives to the opinion of j in updating its own opinion. We assume that the p_{ij} are non-negative and sum to 1, so that $x_j(t+1)$ is a weighted average of the opinions of j's neighbours at time t. In other words, P is a stochastic matrix, and we can describe the evolution of opinions by the linear recursion

$$\mathbf{x}(t+1) = P\mathbf{x}(t). \tag{1}$$

The process is deterministic. We are interested in its long-term behaviour. Do the opinions of each of the agents converge in the long run, i.e., does $x_j(t)$ tend to some x_j as t tends to infinity? Morever, do the agents reach (or approach) agreement in that their opinions all converge to the same value x? If so, we can say that the agents asymptotically reach consensus. Another question we could ask is how quickly this consensus is reached.

The long-term behaviour of the process $\mathbf{x}(t)$ described by (1) is determined by the eigenvalues of the matrix P. Let us denote the eigenvalues by $\lambda_1, \lambda_2, \ldots, \lambda_n$, and order them so that $|\lambda_1| \leq |\lambda_2| \leq \ldots \leq |\lambda_n|$. Denote the corresponding eigenvectors by $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$. We are implicitly assuming that there are n linearly independent eigenvectors, i.e., that the matrix P is not defective. In that case, these eigenvalues form a basis of $\mathbb{R}^n$, and we can express the initial condition $\mathbf{x}(0)$ in this basis as

$$\mathbf{x}(0) = \sum_{i=1}^{n} \alpha_i \mathbf{v}_i,$$

where the α_i are constants. Suppose all other eigenvalues are strictly smaller in absolute value than the largest, i.e., $|\lambda_{n-1}| < |\lambda_n|$. Suppose, too, $\alpha_n \neq 0$, i.e., that the initial condition cannot be expressed just in terms of the remaining eigenvectors. Then,

$$\mathbf{x}(t) = P^t \sum_{i=1}^{n} \alpha_i \mathbf{v}_i = \sum_{i=1}^{n} \alpha_i P^t \mathbf{v}_i = \sum_{i=1}^{n} \lambda_i^t \alpha_i \mathbf{v}_i. \tag{2}$$

It follows that $\lambda_n^{-t}\mathbf{x}(t) \to c\mathbf{v}$, where c is some constant that depends only on the initial condition.

Now, we know that P is a stochastic matrix. Hence, it has an eigenvalue 1 corresponding to the all-1 eigenvector, which we denote $\mathbf{1}$, i.e., $P\mathbf{1} = \mathbf{1}$. What can we say about the other eigenvalues and eigenvectors? The answer is given by the Perron-Frobenius theorem, which we state below. A matrix is said to be non-negative if all its elements are non-negative.

Theorem 1 (Perron-Frobenius theorem) *Suppose that $A \in \mathbb{R}^{n \times n}$ is a non-negative matrix and that A^k is strictly positive for some $k > 0$. Then, the following hold:*

(i) A has a positive eigenvalue λ such that $|\lambda_i| < \lambda$ for all other eigenvalues λ_i, which could be real or complex.

(ii) The eigenvector corresponding to λ is non-negative, and it is the only non-negative eigenvector.

If A is non-negative and non-zero, but there is no k such that A^k is strictly positive, then (i) holds with $|\lambda_i| \leq \lambda$ for all other eigenvalues λ_i. Moreover, the only non-negative eigenvectors are those corresponding to the positive eigenvalue λ, which could now have multiplicity bigger than one.

Remarks. The positive eigenvalue λ which has the largest absolute value among all eigenvalues of A is referred to as the Perron root or Perron eigenvalue of A. Part (ii) refers to right eigenvectors, but also applies to left eigenvectors, as can be seen by applying the theorem to A^T.

We will not prove this theorem. A proof can be found in any book on non-negative matrices, e.g., E. Seneta, *Non-negative matrices and Markov chains*, Springer, 2006, or R. B. Bapat and T. E. S. Raghavan, *Non-negative matrices and applications*, Cambridge University Press, 1997.

We now apply the theorem to our stochastic matrix P. Since the eigenvector corresponding to the eigenvalue 1 is the all-1 vector, which is non-negative, it follows that 1 must be the Perron eigenvalue of P. Consequently, all other eigenvalues are smaller than or equal to 1 in absolute value.

If the matrix P corresponds to an irreducible Markov chain, then the eigenvalue 1 is simple; otherwise, it is repeated with multiplicity equal to the number of closed communicating classes of the Markov chain. If the Markov chain is both irreducible and aperiodic, then it has a single eigenvalue of 1 and all other eigenvalues are strictly smaller than 1 in absolute value.

We now confine ourselves to the most interesting case, that of a stochastic matrix P corresponding to an irreducible, aperiodic Markov chain. In graphical terms, this corresponds to a strongly connected graph (one with a directed path between any two nodes), where the greatest common divisor of the lengths of all cycles is 1. In this case, the eigenvalue 1 is simple and all other eigenvalues are strictly smaller than 1 in absolute value. Hence, $\mathbf{x}(t) = P^t\mathbf{x}(0)$ converges to a constant multiple of the all-1 vector $\mathbf{1}$. Call this constant c. What this says is that the opinions of all agents converges to this same constant c as t tends to infinity, i.e., that they asymptotically reach consensus.

It remains to determine the value c to which the agents converge, and to say something about the speed of convergence. Let π denote the unique equilibrium distribution of the Markov chain with one-step transition probability matrix P. Observe that

$$\pi\mathbf{x}(t+1) = \pi(P\mathbf{x}(t)) = (\pi P)\mathbf{x}(t) = \pi\mathbf{x}(t),$$

i.e., that the inner product of the invariant distribution with the vector of agent preferences is conserved. Hence,

$$c = \pi(c\mathbf{1}) = \pi \lim_{t\to\infty} \mathbf{x}(t) = \lim_{t\to\infty} \pi\mathbf{x}(t) = \pi\mathbf{x}(0),$$

where the interchange of limits with inner products is justified because we are in finitely many dimensions. Hence,

$$\mathbf{x}(t) \to (\pi \mathbf{x}(0))\mathbf{1} \text{ as } t \to \infty.$$

The invariant distribution π is thus a measure of the *influence* of different individuals in determining the value on which consensus is reached; the bigger the value of π_j, the greater the weight assigned to agent j's initial opinion in the consensus outcome.

Next, we look at the speed of convergence. Recall that we assumed P to be the transition matrix of an irreducible, aperiodic Markov chain, and hence that all its eigenvalues except one are strictly smaller than 1 in absolute value. Now, we can rewrite (2) as

$$\mathbf{x}(t) = c\mathbf{1} + \sum_{i=1}^{n-1} \lambda_i^t \alpha_i \mathbf{v}_i,$$

where $|\lambda_i| < 1$ for all $i = 1, 2, \ldots, n-1$. Hence, all terms in the sum on the RHS except the first are decaying geometrically in t. The term decaying most slowly corresponds to the largest in absolute value of these remaining eigenvalues, which is $|\lambda_{n-1}|$ by the way we ordered them. Hence, we conclude that

$$\|\mathbf{x}(t) - (\pi \mathbf{x}(0))\mathbf{1}\| \leq C|\lambda_{n-1}|^t,$$

for some constant C that depends only on the initial condition. Here $\|\cdot\|$ denotes the Euclidean norm of a vector. What this result says is that the "error", the distance between the opinions at any time t and their final consensus value, decays geometrically over time, and the rate of decay is determined by the "spectral gap", namely how far any of the other eigenvalues is, in absolute value, from the largest one, 1.

2 The classical voter model on complete graphs

Consider a population of n individuals, each of whom has a preference for one of two political parties, which we shall denote by 0 and 1. They interact in some way, and change their opinion as a result of the interaction. In practice, their opinion would also be influenced by external factors, such as the policies or performance of the parties, but we don't consider that in our

models. We want to know how the interaction alone influences the evolution of preferences over time.

There are many possible ways to model interactions. It may be that individuals change their opinion if some fraction of their friends or family have a different opinion and influence them. Individuals may differ in how big this fraction needs to be before they change their mind, and also in the relative weights they ascribe to the opinions of different people in their social circle. While it is possible to build models incorporating many of these features, we shall consider a much simpler model, described below.

We can think of the n individuals as the nodes of a complete graph, K_n. The nodes are started in an arbitrary initial state $\mathbf{X}(0) = \{X_v(0), v \in V\}$, where $X_v(0) \in \{0, 1\}$ specifies the initial preference of node v. There are n independent unit rate Poisson processes, one associated with each node. At every time that there is an increment of the Poisson process associated with node v, node v becomes active, chooses a node w uniformly at random from the set of all nodes (including itself), and copies the preference of node w at that time. Thus, individuals in this model are easily persuaded and show no resistance to changing their mind, which is obviously unrealistic. However, it leads to tractable models.

It should be clear from the description above that $\mathbf{X}(t), t \geq 0$ is a continuous-time Markov process as, given the state $\mathbf{X}(t)$ at time t, both the time to the next jump, and the state reached after that jump are independent of the past of the process before time t. Letting e_v denote the unit vector with a 1 corresponding to node v and zeros for all other elements, we can write the transition rates for this Markov process as follows:

$$\begin{aligned} q(\mathbf{x}, \mathbf{x} + e_v) &= \frac{1}{n}(1 - x_v) \sum_{w \in V} x_w, \\ q(\mathbf{x}, \mathbf{x} - e_v) &= \frac{1}{n} x_v \sum_{w \in V} (1 - x_w). \end{aligned}$$

The $1 - x_v$ term in the first equation says that an increment in the v^{th} element of $\mathbf{x}$ is possible only if this element is 0. In that case, it changes to 1 at rate $1/n$, the contact rate between any two nodes, times the number of nodes which are in state 1, which is given by $\sum_{w \in V} x_w$. Likewise, the second equation gives the rate for node v to move from state 1 to state 0.

The state space of the Markov process $\mathbf{X}(t), t \geq 0$ is $\{0, 1\}^V$, the set of $0-1$ valued vectors indexed by the vertex set. Clearly, the all-0 and all-1 vectors,

which we'll denote **0** and **1** are absorbing states; no vertex can change state if the system has reached either of these states. Moreover, all other states form a single communicating class as it is possible to move from any of them to any other. Hence, the Markov process eventually hits one of these two absorbing states, and becomes absorbed. We say that consensus is reached, either on the value 0 or the value 1, which is adopted by all nodes. The questions we want to address are:

- How likely are we to reach consensus on, say, the value 1, given the initial state?
- How long does it take to reach consensus?

These are the questions we shall address in the remainder of this section.

2.1 Hitting probabilities

We begin by observing that the process representation above yields a Markov chain on 2^n states, where $n = |V|$ is the number of nodes. This is unnecessarily large. Given the symmetry in the model, it suffices to keep track of the number of nodes in each state. If we let $Y(t) = \sum_{v \in V} X_v(t)$ denote the number of nodes in state 1 at time t, then $Y(t), t \geq 0$ is a Markov process on the much smaller state space $\{0, 1, 2, \ldots, n\}$. The states 0 and n are absorbing, and our questions above reduce to determining the htting probabilities of each of these states, as well as the absorption time. First, note that the transition rates for the new Markov process $Y(t)$ are given by

$$q(k, k+1) = \frac{(n-k)k}{n}, \quad q(k, k-1) = \frac{k(n-k)}{n}. \tag{3}$$

To see this, note that a transition from k to $k+1$ happens if any of the $n-k$ nodes in state 0 becomes active (which happens at a total rate of $n-k$) and chooses one of the nodes in state 1 to contact (which has probability k/n). Likewise, the second term corresponds to one of the k nodes in state 1 becoming active and choosing a node in state 0 as its random contact. In order to compute the hitting probabilities of the two absorbing states 0 and n, starting from an arbitrary initial state, we shall make use of results from the theory of martingales, which we now introduce, first in discrete time and then in continuous time.

Definition A stochastic process $X_t, t \in \mathbb{N} \cup \{0\}$ is called a martingale if $\mathbb{E}[X_{t+1}|X_t, X_{t-1}, \ldots, X_0] = X_t$ for all t.

Note that this is different from the Markov property. It doesn't say that the probability distribution of X_{t+1} given the past depends only on X_t; it only says that the mean only depends on X_t. However, it is very restrictive about the form of this dependence, as it says that the mean has to be equal to X_t. Intuitively, we think of a martingale as representing one's fortune in a fair game of chance. The fortune after the next play (next roll of dice or spin of the roulette wheel or whatever) is random, but is equal in expectation to the current fortune. An example would be a game of dice in which you win five times your bet (plus your stake) if the die comes up 6, and lose your stake otherwise. Note that this game would be a martingale whatever fraction of your wealth you decided to stake each time.

It is clear, by induction, that $\mathbb{E}[X_{t+s}|X_t, X_{t-1}, \ldots, X_0] = X_t$ for all $s \geq 1$. Taking $t = 0$, $\mathbb{E}[X_s] = \mathbb{E}[X_0]$ for all $s \geq 1$. In fact, it turns out this relationship not only holds for all fixed (deterministic) times s, but also for random times satisfying certain conditions.

Definition A random time T is called a stopping time for a stochastic process $X_t, t \in \mathbb{N} \cup \{0\}$ if the event $T = t$ is measurable with respect to $\{X_s, s \leq t\}$.

In less measure-theoretic language, the random T is a stopping time if you can decide whether $T = t$ just by observing $X_s, s \leq t$. In other words, T is a function of the past and present, not of the future.

Example. Let X_t be the fortune after t time steps in a game where you bet 1 pound repeatedly on rolls of a die where you get 6 pounds (including your stake) if the die comes up 6. Suppose your initial fortune X_0 is 10 pounds. Let T_1 be the first time that your fortune is either 20 pounds or zero. Let T_2 be the time one time step before your fortune hits zero. Then T_1 is a stopping time, whereas T_2 is not. (Both are perfectly well-defined random variables on the sample space of infinite sequences of outcomes of rolls of the die.)

Theorem 2 *(Optional Stopping Theorem) Let $X_t, t \in \mathbb{N} \cup \{0\}$ be a bounded martingale (i.e., there is a finite constant M such that $|X_t| \leq M$ for all $t \geq 0$), and let T be a stopping time for it. Suppose that T is finite almost surely, i.e., $\mathbb{P}(T < \infty) = 1$. Then $\mathbb{E}[X_T] = \mathbb{E}[X_0]$.*

The requirement that the martingale be bounded can be relaxed, but needs to be replaced with conditions that are more complicated to state, and to check in applications. It will suffice for us to confine ourselves to the bounded case.

The definitions and theorem above extend to continuous time martingales. We won't restate the theorem, or the definition of a stopping time, which are identical, but just the definition of a martingale, which extends in the obvious way.

Definition A continuous-time stochastic process $X_t, t \geq 0$ is called a martingale if $\mathbb{E}[X_{t+s}|(X_u, u \leq t)] = X_t$ for all $s, t \geq 0$.

Let us now go back to the continuous-time Markov process $Y_t, t \geq 0$ representing the number of nodes in state 1. We saw in equation (3) that the rates for going from k to $k+1$ and $k-1$ are equal. Hence, it is equally likely that the first jump from state k is to either of these states, which implies the following lemma.

Lemma 1 *The stochastic process $Y_t, t \geq 0$ is a martingale.*

Proof. Fix $t \geq 0$. If $Y_t = 0$ or $Y_t = n$, then Y_t remains constants at all subsequent times, and hence $\mathbb{E}[Y_s|(Y_u, u \leq t)] = Y_t$ for all $s \geq t$. Hence, it only remains to verify this equality if $Y_t = k$ for some $k \in \{1, 2, \ldots, n-1\}$. We suppose from now that this is the case.

Clearly, Y_s is constant and equal to Y_t for all $s \in (t, \tau)$ where τ is the random time of the first jump after time t. Hence $\mathbb{E}[Y_s|(Y_u, u \leq t)] = Y_t$ for all $s \in (t, \tau)$. Moreover,

$$\mathbb{E}[Y_\tau|(Y_u, u \leq t)] = \frac{1}{2}(Y_t + 1) + \frac{1}{2}(Y_t - 1) = Y_t.$$

Thus, the equality $\mathbb{E}[Y_s|(Y_u, u \leq t)] = Y_t$ holds for all s up to and including the first jump time after t. By induction on the sequence of jump times, it holds for all $s \geq t$.

This way of proving that Y_t is a martingale basically reduces the continuous-time process to a discrete-time process watched at the jump times. Another approach is to look at the change over infinitesimal time intervals. Observe from equation (3) that

$$\mathbb{E}[Y_{t+dt} - Y_t|(Y_u, u \leq t)] = (+1) \cdot \frac{(n-k)k}{n} dt + (-1) \cdot \frac{(n-k)k}{n} dt = 0,$$

since the possible values of $Y_{t+dt} - Y_t$ are $+1$ or -1, and we multiply them by the corresponding rates and take the average. (Jumps of 2 or more have probability $o(dt)$, which we ignore.) The fact that the change in conditional expectation is zero over infinitesimal time intervals implies that the conditional expectation is constant, and hence that Y_t is a martingale. $\square$

It is now straighforward to compute the hitting probability of each of the absorbing states starting from any given initial state k. Let $T = \inf\{t \geq 0 : Y_t = 0 \text{ or } n\}$ denote the random time to absorption. Then, T is clearly a stopping time, and we have by the Optional Stopping Theorem that

$$\mathbb{E}[Y_T] = \mathbb{E}[Y_0] = k.$$

But

$$\mathbb{E}[Y_T] = n\mathbb{P}(Y_T = n) + 0\mathbb{P}(Y_T = 0),$$

and so

$$\mathbb{P}(Y_T = n) = \frac{\mathbb{E}[Y_T]}{n} = \frac{k}{n}.$$

This gives us the probability of hitting n before 0 as a function of the initial state k.

2.2 Hitting times

We shall derive bounds on the time to consensus by establishing a duality with coalescing random walks. Suppose we want to know whether, on a specific sample path of the random process, consensus has been reached by a given time, τ. We can determine this by following the evolution of the process, given its initial condition, from time 0 to time τ. Alternatively, we can do so by following the evolution backwards from τ. Imagine that, at time τ, each node is occupied by a single particle.

Let $\tau - T_1$ denote the last time before τ (first time looking backwards) that there is a contact between any two nodes. Suppose that, at time $\tau - T_1$, a node denoted v_1 copies a node denoted u_1. As v_1 is involved in no further communications after time $\tau - T_1$, and neither is any other node, it is clear that the state of v_1 at time τ, and consequently the state of all nodes at time τ, is fully determined by the state of *all nodes other than* v_1 at time $\tau - T_1$. We represent this by saying that the particle at v_1 has moved to u_1 and coalesced with the particle there at time $\tau - T_1$.

As we continue following the process backwards from time $\tau - T_1$, there may be a time $\tau - T_2$ at which some node $v_2 \neq v_1$ copies some other node, denoted u_2. We are not interested in times before $\tau - T_1$ at which v_1 copies some other node, because that won't affect the final state of v_1. Again, we represent this by the particle at v_2 moving to u_2. If $u_2 \neq v_1$, then the moving particle coalesces with the one occupying u_2. If $u_2 = v_1$, then the particle at u_2 moves to v_1 but there is no particle there to coalesce with.

The above is a verbal description of the process, looking backwards in time from τ, but we would like a probabilistic model. What can we say about the random time T_1, which is the first time looking back from τ that a contact occured between two nodes? Nodes becomes active at the points of independent unit rate Poisson processes. It is worth noting that the time reversal of a Poisson process is also a Poisson process of the same rate. Hence, looking back from τ, node activation times are again independent unit rate Poisson processes. When a node v becomes active, it contacts a node u chosen uniformly at random (again, time reversal doesn't change this) and copies its state. In our description, this corresponds to the particle at v moving to u, and coalescing with the particle already there, if any. Thus, the process backwards from τ corresponds to particles performing independent continuous time random walks on the complete graph until they meet another particle and coalesce. Coalesced particles behave just like any other particle. We thus arrive at the following probabilistic model for the process looking back from τ.

Initially, there are n particles, one at each node of the complete graph. Particles move independently of each other. Each particle waits for a random time exponentially distributed with unit mean, then moves to a node chosen uniformly at random (including itself). If there is a particle at that node, the two particles coalesce, and henceforth behave as a single particle obeying the same rules. It is clear from this description that the process can be modelled as a Markov chain $Z_t, t \geq 0$ on the state space $\mathcal{P}(V)$, the set of all subsets of the vertex set. The state at time t denotes the set of nodes occupied by a particle. The event that consensus has occured by time τ is then equivalent to the event that the set Z_τ of occupied nodes at time τ of the backwards-time process is a set of nodes which all have the same initial condition in the forward-time process. Thus, this event depends on the initial condition in general. We would like to obtain a bound on the consensus time that doesn't depend on the initial condition. The only way to ensure that all nodes in Z_τ have the same initial state, irrespective of the initial condition,

is if Z_τ is a singleton set, i.e., consists of a single node. This is the case if all n particles have coalesced into a single particle. We will now analyse the distribution of the random time for this event to occur.

In a general graph, we would have to keep track of the locations of all extant particles, but as we are working on the complete graph, all nodes are identical. Hence, we only need to keep track of the number of particles at any time t, which we shall denote by W_t. Let T_k denote the first time that $W_t = k$. We have $T_n = 0$ as we start with n particles. We want to estimate T_1, the random time that all particles have coalesced to a single one. Now, what can we say about $T_{k-1} - T_k$? Each of the k particles alive at time T_k becomes active according to a unit-rate Poisson process, and moves to a node chosen uniformly at random from all n nodes (including its current location). Using the fact that Bernoulli splittings of Poisson processes are Poisson, we can model this by associating independent Poisson processes of rate $1/n$ with each directed edge of a complete directed graph on n nodes. If the Poisson clock on the directed edge (u, v) goes off, then the particle at u moves to v and coalesces with the particle there, if any. As we are only interested in coalescences, it is enough to look at those directed edges which link occupied nodes. As there are k occupied nodes, there are $k(k-1)$ such edges. Each of these has an independent rate $1/n$ Poisson clock associated with it. Using the fact that superpositions of independent Poisson processes are Poisson, the time until the clock on some edge between occupied nodes rings is given by an $Exp(k(k-1)/n)$ random variable. When this happens, two particles coalesce. Hence,

$$T_{k-1} - T_k \sim \mathrm{Exp}\Big(\frac{k(k-1)}{n}\Big), \quad \mathbb{E}[T_{k-1} - T_k] = \frac{n}{k(k-1)} = n\Big(\frac{1}{k-1} - \frac{1}{k}\Big).$$

Recalling that $T_n = 0$, we obtain that

$$\mathbb{E}[T_1] = \sum_{k=2}^{n} \mathbb{E}[T_{k-1} - T_k] = n - 1.$$

As we argued above, the random time T_1 is an upper bound on the time to consensus. Thus, the above result tells us that the mean time to consensus is bounded above by n, the number of nodes. While we did not derive a lower bound, this is in fact the correct scaling relationship. A more detailed but tedious calculation shows that the mean time to consensus, starting from an initial condition in which a fraction α of nodes are in state 0, is given by $nh(\alpha)$, where $h(\alpha) = -\alpha \log \alpha - (1-\alpha)\log(1-\alpha)$ denotes the binary

entropy function evaluated at α. Thus, the time to reach consensus on the complete graph scales linearly in the number of nodes, in contrast to the rumour spreading time, which scales logarithmically.

3 The voter model on general graphs

Let $G = (V, E)$ be a directed graph. Each node $v \in V$ can be in one of two states, 0 or 1. We denote by $X_v(t)$ the state of node v at time t, and by $\mathbf{X}(t)$ the vector, $(X_v(t), v \in V)$. The process by which nodes change their state is as follows. Associated with each directed edge $(v, w) \in E$ is a Poisson process of rate $q_{vw} > 0$, at the points of which node v copies the state of node w. The Poisson processes associated with distinct edges are mutually independent. It follows from this description that $\mathbf{X}(t)$ is a Markov process.

We assume that the graph G is *strongly connected*, i.e., that there is a *directed path* from v to w for every pair of nodes v, w. If there is such a directed path, then node w can influence node v. Clearly, the all-0 and all-1 states, which we denote by $\mathbf{0}$ and $\mathbf{1}$ respectively, are absorbing. The assumption says that every node can influence every other node, and hence that all states other than $\mathbf{0}$ and $\mathbf{1}$ form a single communicating class, from which both these states are accessible. Hence, the Markov chain eventually reaches one of these absorbing states, and we want to know the probability of reaching $\mathbf{0}$ and $\mathbf{1}$ starting from an arbitrary initially state.

Note that the assumption of strong connectivity is essential for absorption in one of these two states to be guaranteed. To see this, we consider a counterexample. First, consider a graph consisting of 3 nodes: $V = \{u, v, w\}$ and $E = \{(v, u), (v, w)\}$. Consider the initial state, $X_u(0) = 0$, $X_w(0) = 1$, and arbitrary $X_v(0)$. In this graph, both u and w can influence v but v can't influence them, and they can't influence each other. Hence, $X_u(t) = 0$ for all $t \geq 0$, $X_w(t) = 1$ for all $t \geq 0$, while $X_v(t)$ keeps oscillating between 0 and 1. Hence, there is no absorption in this example. Equally, if the graph is disconnected, there may be multiple absorbing states, one for each connected component of the graph.

Let Q be a rate matrix (infinitesimal generator) with off-diagonal elements q_{vw} as specified above; the diagonal elements are then given by the requirement that the row sums should all be zero. Since Q is the rate matrix of a finite state Markov chain on the state space V, it has an invariant distri-

bution π, namely a probability vector on the vertex set V which solves the global balance equations $\pi Q = \mathbf{0}$. Under the assumption that the graph is strongly connected, the invariant distribution is unique, and strictly positive (i.e., π_v is not zero for any $v \in V$). Define $M(t) = \pi \mathbf{X}(t)$, the product of the row vector π and the column vector $\mathbf{X}(t)$. We now claim the following.

Lemma 2 *The stochastic process $M(t) = \pi\mathbf{X}(t) = \sum_{v\in V} \pi_v X_v(t)$ is a martingale.*

Proof. Observe that

$$\mathbb{E}[M(t+dt) - M(t)|(\mathbf{X}(u), u \le t)] = \sum_{v\in V: X_v(t)=0} \pi_v \mathbb{P}(X_v(t+dt) = 1|\mathbf{X}(t))$$
$$- \sum_{v\in V: X_v(t)=1} \pi_v \mathbb{P}(X_v(t+dt) = 0|\mathbf{X}(t)) + o(dt), \tag{4}$$

since the possible changes in the state $\mathbf{X}(t)$ over an infinitesimal time interval dt are the change in state of some single node from 0 to 1 or 1 to 0; by the definition of $M(t)$, a change in state of X_v from 0 to 1 increases $M(t)$ by π_v, while a change from 1 to 0 decreases it by π_v. Now, by the description of the model,

$$\begin{aligned}
\mathbb{P}(X_v(t+dt) = 1|\mathbf{X}(t), X_v(t) = 0) &= \sum_{w:(v,w)\in E, X_w=1} q_{vw}dt, \\
\mathbb{P}(X_v(t+dt) = 0|\mathbf{X}(t), X_v(t) = 1) &= \sum_{w:(v,w)\in E, X_w=0} q_{vw}dt.
\end{aligned}$$

In each case, we add up the rates q_{vw} of contacting nodes w in the opposite state to node v. Substituting these transition probabilities in (4), we get

$$\begin{aligned}
&\mathbb{E}[M(t+dt) - M(t)|(\mathbf{X}(u), u \le t)] \\
&= \sum_{\substack{(v,w)\in E: \\ X_v(t)=0, X_w(t)=1}} \pi_v q_{vw}dt - \sum_{\substack{(v,w)\in E: \\ X_v(t)=1, X_w(t)=0}} \pi_v q_{vw}dt + o(dt) \\
&= \sum_{(v,w)\in E: X_v(t)=0, X_w(t)=1} (\pi_v q_{vw} - \pi_w q_{wv})dt + o(dt).
\end{aligned}$$

It is convenient to extend the sum to all vertex pairs (v, w) rather than just $(v, w) \in E$; we may do this as $q_{vw} = 0$ by definition whenever $(v, w) \notin E$.

Thus, we can rewrite the above equation as

$$\begin{aligned}&\mathbb{E}[M(t+dt) - M(t)|(\mathbf{X}(u), u \le t)] \\ &= \sum_{v,w \in V: X_v(t)=0, X_w(t)=1} (\pi_v q_{vw} - \pi_w q_{wv})dt + o(dt).\end{aligned}$$

Now, let us add and subtract

$$\sum_{v,w \in V: X_v(t)=0, X_w(t)=0} \pi_v q_{vw} \equiv \sum_{v,w \in V: X_v(t)=0, X_w(t)=0} \pi_w q_{wv}$$

on the RHS of above expression, to obtain that

$$\begin{aligned}&\mathbb{E}[M(t+dt) - M(t)|(\mathbf{X}(u), u \le t)] = \\ &\sum_{v \in V: X_v=0} \sum_{w \in V} \pi_v q_{vw} dt - \sum_{v \in V: X_v=0} \sum_{w \in V} \pi_w q_{wv} dt + o(dt). \qquad (5)\end{aligned}$$

Now, for each $v \in V$, observe that $\sum_{w \in V} \pi_v q_{vw} = \pi_v \sum_{w \in V} q_{vw} = 0$, since the row sums of the rate matrix Q are zero. Moreover, $\sum_{w \in V} \pi_w q_{wv} = 0$ because π was assumed to be an invariant distribution, i.e., it satisfies the global balance equations $\pi Q = \mathbf{0}$. Substituting in (5), we conclude that

$$\mathbb{E}[M(t+dt) - M(t)|(\mathbf{X}(u), u \le t)] = o(dt),$$

which implies that $M(t)$ is a martingale, as claimed. □

It is now straightforward to compute the hitting probability of the all-zero and all-one states, starting from an arbitrary initial condition $\mathbf{X}(0)$. Let $M(0) = \pi\mathbf{X}(0)$ denote the value of the martingale corresponding to this initial condition. Let T denote the absorption time in one of the two absorbing states, and note that it is a stopping time. Hence, by the Optional Stopping Theorem, $\mathbb{E}[M(T)] = M(0) = \pi\mathbf{X}(0)$. But

$$\begin{aligned}\mathbb{E}[M_T] &= \mathbb{P}(\mathbf{X}(T) = \mathbf{1}) \cdot \pi\mathbf{1} + \mathbb{P}(\mathbf{X}(T) = \mathbf{0}) \cdot \pi\mathbf{0} \\ &= \mathbb{P}(\mathbf{X}(T) = \mathbf{1}) \cdot 1 + \mathbb{P}(\mathbf{X}(T) = \mathbf{0}) \cdot 0.\end{aligned}$$

It follows that

$$\mathbb{P}(\mathbf{X}(T) = \mathbf{1}) = \mathbb{E}[M(T)] = M(0) = \pi\mathbf{X}(0).$$

Notice the similarity of the final result with that of the de Groot model. There, the value on which consensus was reached was $\pi\mathbf{X}(0)$. Here, consensus is reached on either $\mathbf{0}$ or $\mathbf{1}$, but the expected value is still the same, $\pi\mathbf{X}(0)$. In particular, the invariant distribution π still determines the influence of different nodes or agents in determining the final outcome.