

UC Berkeley

UC Berkeley Previously Published Works

Title

Interacting particle systems as stochastic social dynamics

Permalink

<https://escholarship.org/uc/item/8hq050fx>

Journal

Bernoulli, 19(4)

ISSN

1350-7265

Author

Aldous, David

Publication Date

2013

DOI

10.3150/12-bejsp04

Peer reviewed

Bernoulli **19**(4), 2013, 1122–1149
DOI: [10.3150/12-BEJSP04](https://doi.org/10.3150/12-BEJSP04)

Interacting particle systems as stochastic social dynamics

DAVID ALDOUS

*Statistics Dept., U.C. Berkeley, CA 94720, USA. E-mail: aldous@stat.berkeley.edu;
url: www.stat.berkeley.edu/~aldous/*

The style of mathematical models known to probabilists as Interacting Particle Systems and exemplified by the Voter, Exclusion and Contact processes have found use in many academic disciplines. In many such disciplines the underlying conceptual picture is of a social network, where individuals meet pairwise and update their “state” (opinion, activity etc) in a way depending on the two previous states. This picture motivates a precise general setup we call *Finite Markov Information Exchange (FMIE) processes*. We briefly describe a few less familiar models (Averaging, Compulsive Gambler, Deference, Fashionista) suggested by the social network picture, as well as a few familiar ones.

Keywords: epidemic; interacting particle system; Markov chain; social network; voter model

1. Introduction

What is the most broad-ranging currently active field of applied probability? Measuring breadth by the number of different academic disciplines where it appears, it seems hard to beat a field appearing in Physics, Computer Science and Electrical Engineering, Economics and Finance, Psychology and Sociology, Population Genetics and Epidemiology and Ecology. Curiously, the field has no good name, though readers of *Bernoulli* will most likely know it, under the continuing influence of Liggett’s 1985 book [37], as *Interacting Particle Systems*. It turns out that mathematically similar toy models of different real-world entities have been repeatedly re-invented in the different disciplines mentioned above, and literally thousands of papers have been written since 2000 on discipline-dependent variant models intended as more realistic.

Rather than attempting a brief overview of the whole field as it stands today (on which I am certainly not an expert), let me take an opposite approach, imagining starting with a blank slate. I will set out in the next section a particular conceptual setting – what sorts of things are we trying to model? – then a particular technical setup, and then indicate what sort of mathematical results one might seek. Given this manifesto, the rest of the article describes some basic models, gives pointers to the existing literature and presents

This is an electronic reprint of the original article published by the ISI/BS in <i>Bernoulli</i> , 2013, Vol. 19, No. 4, 1122–1149. This reprint differs from the original in pagination and typographic detail.

a few open problems suggested by our setting. The article is based on a short course given in summer 2012 at Warwick and Cornell, and more extensive informal material is available at [3]. I must emphasize that the article is intended as an introduction to the topics it mentions, not as an authoritative survey, which would require ten times the length! In particular none of the mathematics outlined here is both new and non-straightforward, except perhaps some open problems.

2. Conceptual and technical setup

2.1. Information flow through social networks via pairwise interaction

The particular setting we envisage is of the human social network. People meet or communicate, and explicitly or implicitly receive information from the other, and this interaction may change the “state” (opinion or possessions or ...) of the individuals. Now that sentence skips over many issues. “Information” is a complicated concept [33], and formalizing it as one of a set of prespecified states hardly captures the nuances of human experience. And obviously we receive information in many ways other than one-on-one interaction with other individuals. But the fact that humans do a lot of different things via interactions emphasizes the point that we are interested in a variety of different rules, a point that would be less visible if we took some real or hypothetical physical system as a prototype setting.

Social networks (whose mathematical modeling and analysis long predates the Facebook era) are usually modeled as unweighted graphs – vertices are individuals, edges indicate some specified relationship. In most contexts it would be more realistic to use *weighted* graphs, the weight indicating strength of relationship. This is seldom done because of increased model complexity and the fact that the meaning of “strength of relationship” is context-dependent. But we will give it the specific meaning of “frequency of meeting”, formalized via the rates ν_{ij} below.

The story above suggests the following technical framework, in which randomness arises primarily via assumed random times of meetings. Because the framework can be used outside the social network setting, I will use the impersonal word “agent” in what follows.

2.2. The technical framework

Consider n agents and a nonnegative array (ν_{ij}) , indexed by unordered pairs $\{i, j\}$, which is irreducible (i.e., the graph of edges corresponding to strictly positive entries is connected). Assume

- Each unordered pair i, j of agents with $\nu_{ij} > 0$ meets at the times of a rate- ν_{ij} Poisson process, independent for different pairs.

Call this collection of Poisson processes the *meeting process*. Where needed we write **Agents** for the set of agents.

Now introduce a set **States** of states, and consider a process $\mathbf{X}(t)$ in which each agent i is in some state $X_i(t)$ at time t , and the state can only change when agent i meets some other agent j , at which time the two agents' states are updated according to some rule (typically deterministic, but might be random) which depends only on the pre-meeting states $X_i(t-), X_j(t-)$. So in the deterministic case, there is a given arbitrary function $F : \mathbf{Agents} \times \mathbf{Agents} \rightarrow \mathbf{Agents}$ and when agents in states s_1 and s_2 meet their states are updated as

$$(s_1, s_2) \rightarrow (F(s_1, s_2), F(s_2, s_1)). \quad (2.1)$$

We call such a process an *FMIE (Finite Markov Information Exchange) process* and emphasize that a specific such process has two levels. The meeting process is determined by the meeting rate matrix (ν_{ij}) (which we loosely call the “geometry”, imagining that we tend to meet more frequently with individuals who are “close” in some sense). The information update process is determined by some rule. Then the different named FMIE models we shall describe (such as the Voter model) correspond to different rules for the informational superstructure; we regard such models as defined over arbitrary geometries.

2.3. Three disciplines using FMIE models

Of course much of the beauty and utility of mathematics is in providing abstractions that encompass apparently different things, and as stated at the start FMIE models have been used in many disciplines. Let me briefly mention three that are relevant to our models later.

Conceptually closest is the mathematical modeling of epidemics, whose long history is indicated by the 1957 monograph [11]. A typical classical model like the SIR (Susceptible–Infected–Recovered) model corresponds to a 3-state FMIE if the relevant distributions are Exponential.¹ Here of course the vertices are people or animals and the edges relate to physical proximity. Classical models considered homogeneously mixing populations (i.e., the complete graph geometry) but there is extensive post-2000 literature on more general geometries.

A quite different context, from theoretical computer science, concerns distributed randomized algorithms. Here agents are processors, communicating with nearby (direct connections in a network) processors, and performing computations based on data transferred. See [43] for the particular topic of *gossip algorithms*.

What are called *Interacting Particle Systems (IPS)* arose historically as statistical physics models of phenomena such as magnetism, and mathematicians became interested in particular in rigorous study of phase transitions, in the setting of infinite geometries, specifically the infinite d -dimensional lattice. This literature long ago branched out to study many “toy models” not so directly related to any specific physical phenomenon.

¹Capitalizing named distributions is a useful convention that should be more widely adopted.

Our development is in the same “toy models” spirit, and the models naturally suggested by the social networks context partly overlap and partly differ from those naturally suggested by statistical physics. In particular, the abstraction of what we are studying as “information flow through networks” seems more helpful than any physical visualization.

A broader but less focused view of the whole “statistical physics of social dynamics” topic can be found in [19].

2.4. Style of mathematical results

The concept of an *IPS* is generally discussed via example rather than explicit definition. For instance Liggett [37] takes as examples the Voter model, Contact process, Exclusion process, Glauber dynamics for the Ising model, and the Potlatch process. Our mathematical setup of *FMIE process* in Section 2.2 is intended as a more precisely circumscribed class of processes to study. For instance, only the first three processes on Liggett’s list are FMIE (for the contact process, add an artificial agent with rate-one meetings with each other agent, to effect recovery). What we call the meeting process is often called the *graphical representation* in IPS.

Our emphasis on a *finite* number of agents gives a somewhat different perspective than in classical IPS, in which the finite setting was considered only occasionally, e.g., [31]. Our viewpoint is by analogy with the modern literature on mixing and hitting times for finite Markov chains [4, 36]; study how the behavior of a model depends on the underlying geometry, where “behavior” means quantitative aspects of finite-time behavior. More precisely, we typically study $n \rightarrow \infty$ limits in a family of finite-agent processes, rather than $t \rightarrow \infty$ limits in a single infinite-agent process. The Averaging process studied in Section 4 provides a prototype for the style of results we would like to establish for a given FMIE model; see also the discussion of possible $t \rightarrow \infty$ limits in Section 3.3.

The FMIE setup allows an arbitrary update function F at (2.1) and an arbitrary “geometry” of meeting rates (ν_{ij}) . One can hardly hope to say anything of substance in this generality (we will mention various “general principles” but these are rather trite, once said). At a pragmatic level, if seeking analytic results one chooses an update function such that the FMIE process can be analyzed on the complete graph (“mean-field”, roughly) geometry, otherwise it will likely be difficult to handle other geometries. And obversely, one chooses a geometry in which one can analyze the two basic processes of Section 3, otherwise it will likely be difficult to handle more complicated update rules.

From our viewpoint, much of the literature since 2000 consists of inventing and investigating a large number of increasingly elaborate update rules, while considering only a small number of “standard geometries”, indicated below. Keep in mind that it is easy to invent and simulate models, but hard to give rigorous proofs or to relate convincingly to real-world data.²

²In talks I criticize a particularly egregious instance [44] of claims of real-world significance based on models with absolutely no data or evidence linking the model to the asserted real-world phenomena.

2.5. Some standard geometries

Here we briefly describe four of the “standard geometries” on which models are often studied. Recall n is the number of agents. Note that in a realistic social network there is no reason why $\sum_j \nu_{ij}$, which is the rate at which agent i meets other agents, should be constant in i . But many simple models do have this property, in which case we may rescale time to fit a *standardization convention*

$$\sum_j \nu_{ij} = 1 \quad \forall i \quad (2.2)$$

as in the first two examples below.

Complete graph. Here $\nu_{ij} = 1/(n-1), j \neq i$.

d -dimensional discrete torus \mathbb{Z}_m^d . Here $\nu_{ij} = 1/(2d)$ for adjacent agents i, j .

The next two examples involve networks modelled as *random* graphs. Formally this adds another level of randomness, though informally we can just think of being given a typical realization. In such (non-regular) graphs one typically cannot satisfy (2.2) and instead the default convention is to take rates $\nu_{ij} = 1$ for each edge (i, j) .

Small worlds. Here one starts with the d -dimensional discrete torus (or similar non-random graph) and then adds random other edges in some specified way, for instance independently with probabilities proportional to (Euclidean length) $^{-\gamma}$.

Random graphs with prescribed degree distribution. There are several ways (in particular the *configuration model* [16]) to define graphs corresponding to the intuitive notion of being “completely random” subject to the constraint that, as $n \rightarrow \infty$, the distribution of vertex-degrees converges to a prescribed distribution. Such models are tractable because their $n \rightarrow \infty$ local weak limit is just a modified Galton-Watson tree.

2.6. Some less standard geometries

There is a huge “complex networks” literature (see, e.g., the monographs [27, 40]) devoted to inventing and studying network models, and I will not write out a long list here. Let me instead mention three other mathematically natural geometries which have *not* been studied as much as one might have expected (and two more will be mentioned in Section 5.7).

Long range geometric interaction. Start with the torus \mathbb{Z}_m^d and add rates $\nu_{ij} = c_{d,m,\gamma} \|i - j\|_2^{-\gamma}$ for non-adjacent i, j . This is a way to interpolate between the torus and the complete graph. Note the distinction between this and the “small worlds” model, in which the rate is non-vanishing for a few random edges.

Proximity graphs. Given a model whose behavior is understood on the two-dimensional lattice, one could investigate the effect of “disorder” by taking instead a proximity graph over a Poisson process of points in the plane [9].

The Hamming cube $\{0, 1\}^d$. This is a standard example in Markov chain theory.

3. The two background models

Here we discuss two FMIE processes that are “basic” in two senses. They fit the definition of FMIE process but are much simpler than (and therefore unrepresentative of) a typical FMIE process; yet many other FMIE processes can be regarded as being built over one of these base processes, with extra structure added.

For each FMIE, we give only a verbal description, which the reader can readily formalize.

3.1. The Token process and the associated Markov chain

The Token process. There is one token. When the agent i holding the token meets another agent j , the token is passed to j .

This can evidently be viewed as a 2-state FMIE process. The natural aspect to study is $Z(t)$ = the agent holding the token at time t . This $Z(t)$ is the (continuous-time) Markov chain with transition rates (ν_{ij}) . We call this the *associated* Markov chain³ – associated with the underlying geometry.

A nuance. Let me digress to mention, by analogy, a nuance. A Google Scholar search on “exact phrase = Galton-Watson; year = 1965–1969” returns only papers studying Z_g = population size in generation g of a Galton-Watson branching process. Changing to “year = 2005–2009”, half the papers talk about the Galton-Watson **tree**, which of course has more structure than the population size process – one can study cousins, for instance. Analogously, the Token process has more structure than the associated Markov chain – one can study the time until agent i meets someone who has previously held the token, for instance.

For finite-state Markov chains, we have available both a classical general theory and a long catalog of explicit calculations in specific geometries. For our purposes, the main point is that several FMIE models (classically the Voter model and the exclusion process, but also the Averaging process) have close connections with the associated Markov chain, and parts of the “standard modern theory” [4, 36] of finite chains can be used in their analysis.

³Throughout, Markov chains are continuous-time.

For readers more familiar with discrete-time chains, let me emphasize that (because $\nu_{ij} = \nu_{ji}$) the stationary distribution of the associated chain is always uniform; this does not assume the standardization (2.2).

3.2. The Pandemic process and FPP

The Pandemic process. Initially one agent is infected. Whenever an infected agent meets another agent, the other agent becomes infected.

In classical language, this is the SI epidemic with Exponential(ν_{ij}) distributions for infection times. But again there is a nuance. Suppose that on each undirected edge (i, j) of the geometry we put a single Exponential(ν_{ij}) r.v. ξ_{ij} , independent over edges, and regard ξ_{ij} as the length of the edge. Then there is an induced distance (= minimum route length) between any pair of agents. This model of random distances is what I will call the *first passage percolation (FPP)* model. Note this differs from classical lattice FPP, which assumes i.i.d. variables with arbitrary distribution.

Here is the nuance. In the Pandemic process started with agent i , write T_{ij}^{pan} for the time when agent j is infected. In the FPP process write T_{ij}^{fpp} for the distance between i and j . From the memoryless property of the Exponential, we see that $T_{ij}^{\text{pan}} \stackrel{d}{=} T_{ij}^{\text{fpp}}$ and indeed that $(T_{ij}^{\text{pan}}, j \in \mathbf{Agents}) \stackrel{d}{=} (T_{ij}^{\text{fpp}}, j \in \mathbf{Agents})$ for each i . But Pandemic and FPP are different processes in that $T_{ij}^{\text{fpp}} = T_{ji}^{\text{fpp}}$ whereas we have only $T_{ij}^{\text{pan}} \stackrel{d}{=} T_{ji}^{\text{pan}}$. (One can fix this particular issue by re-interpreting, in FPP, the time ξ_{ij} as the elapsed time, after one end is infected, until the other end is infected, but this misses the overall point that the distributions of the whole arrays (T_{ij}) are different.)

The Pandemic model is “basic” in a very specific sense, in that it exhibits the fastest possible spread of information in any FMIE model. That is, in an arbitrary FMIE process $\mathbf{X}(t)$ over a given geometry, the value of $X_j(t)$ can only be influenced by the value of $X_i(0)$ if $T_{ij}^{\text{pan}} \leq t$.

Here of course we are exploiting the natural coupling of two arbitrary FMIE processes over the same geometry, which simply uses the same realization of the meeting process.⁴ Albeit obvious, one can view this as an instance of the following.

General Principle 1. *It is often useful to consider the natural coupling of two FMIE processes.*

Models of real-world epidemics are of course much more elaborate. Regarding our simple Pandemic model, it is analogous to the Markov chain in two ways:

- There is extensive literature doing calculations in specific geometries; I show a little in Section 7.
- Several other FMIEs can be “built over” the Pandemic model – see Section 8.

⁴If both update rules are random then we need also to specify some joint distribution.

But it differs it that there seems no existing “general theory”. See a conjecture in Section 7.2.

3.3. Remarks on time-asymptotics

In these remarks, we are envisaging the finite-state case.

General Principle 2. *Time-asymptotics for fixed n are not the issue; in most models it’s easy to see what’s going on.*

Informally there are four qualitatively different possible time-asymptotic behaviors, and each of the models we treat fits readily into one of these.

- Absorption in a random one (of a small number of) “ordered” configurations. [Pandemic, Averaging, Voter, Deference].
- Absorption in a random one (of a large number of) “disordered” configurations. [Compulsive Gambler].
-
- Convergence to the unique stationary distribution (perhaps on a subset of configurations). [Token, Interchange, Fashionista].

The elided third item is the most general possibility for a finite-state Markov chain (applied to the configurations of a FMIE process); there are one or more *strongly connected*⁵ components (SCCs) of the directed transition graph, and the limit distribution is a mixture over the stationary distributions on each SCC. While it’s easy to invent artificial FMIE models with this behavior, I don’t know any “natural” such model.

The following question is less easy that it first appears. The natural conjecture (involving subsets of **Agents** closed under F , and functions whose sum is conserved by the interaction) is false.

Open Problem 1. *Consider the FMIE process over the complete graph on n agents (for large n) defined by a deterministic update function F as at (2.1), What are the SCCs of the FMIE process?*

4. A prototype FMIE model: The Averaging process

As a prototype FMIE model – for which one can say something quantitative about its behavior over a general geometry – we take the following. Curiously, though it appears as an ingredient in more elaborate models, the only place I have seen the basic process itself is in the “gossip algorithms” literature [43] which derives a version of the “global bound”, Proposition 4.2. A careful treatment of the results below is written out in [5] – here I give a compressed version.

⁵Meaning there is a directed path between any ordered pair of states.

The Averaging process. Initially each agent has some (real-valued) amount of money. Whenever two agents meet, they divide their money equally.

It is rather obvious that the individual values $X_i(t)$ converge as $t \rightarrow \infty$ to the average of initial values, so what is there to say?

Write $\mathbf{1}_i$ for the initial configuration ($1_{(j=i)}, j \in \mathbf{Agents}$), that is agent i has unit money and other agents have none, and write $p_{ij}(t)$ for the transition probabilities of the associated Markov chain.

Proposition 4.1. *For the Averaging process with initial configuration $\mathbf{1}_i$ we have $\mathbb{E}X_j(t) = p_{ij}(t/2)$. More generally, from any deterministic initial configuration $\mathbf{x}(0)$, the expectations $\mathbf{x}(t) := \mathbb{E}\mathbf{X}(t)$ evolve exactly as the dynamical system*

$$\frac{d}{dt}\mathbf{x}(t) = \frac{1}{2}\mathbf{x}(t)\mathcal{N},$$

where \mathcal{N} is the generator of the associated Markov chain.

The key point behind this results is that we can rephrase the dynamics of the Averaging process as

when two agents meet, each gives half their money to the other.

In informal language, this implies that the motion of a random penny – which at a meeting of its owner agent is given to the other agent with probability $1/2$ – is as the associated Markov chain at half speed, that is with transition rates $\nu_{ij}/2$.

To continue, we need to recall some background facts (e.g., [4], Chap. 3) about reversible chains, here specialized to our continuous-time setting of the associated Markov chain, that is with uniform stationary distribution ($\nu_{ij} = \nu_{ji}$). A function $f : \mathbf{Agents} \rightarrow \mathbb{R}$ has (with respect to the uniform distribution) average \bar{f} , variance $\text{var } f$ and L^2 norm $\|f\|_2$ defined by

$$\bar{f} := n^{-1} \sum_i f_i, \quad \|f\|_2^2 := n^{-1} \sum_i f_i^2, \quad \text{var } f := \|f\|_2^2 - (\bar{f})^2.$$

The L^2 norm will be used in several different ways. For a possible time- t configuration $\mathbf{x}(t)$ of the averaging process, the quantity $\|\mathbf{x}(t)\|_2$ is a number, and so the quantity $\|\mathbf{X}(t)\|_2$ appearing in the proposition below is a random variable. The *Dirichlet form* of the associated Markov chain is

$$\mathcal{E}(f, f) := \frac{1}{2}n^{-1} \sum_i \sum_{j \neq i} (f_i - f_j)^2 \nu_{ij} = n^{-1} \sum_{\{i,j\}} (f_i - f_j)^2 \nu_{ij},$$

where $\sum_{\{i,j\}}$ indicates summation over *unordered* pairs. The spectral gap of the chain, defined as the gap between eigenvalue 0 and the second eigenvalue of \mathcal{N} , is characterized

as

$$\lambda = \inf_f \left\{ \frac{\mathcal{E}(f, f)}{\text{var}(f)} : \text{var}(f) \neq 0 \right\}. \quad (4.1)$$

The proof of the following result uses only the ingredients above and the Averaging process dynamics.

Proposition 4.2 (Global convergence theorem [5]). *From an initial configuration $\mathbf{x}(0) = (x_i)$ with average zero, the time- t configuration $\mathbf{X}(t)$ of the averaging process satisfies*

$$\mathbb{E} \|\mathbf{X}(t)\|_2 \leq \|\mathbf{x}(0)\|_2 \exp(-\lambda t/4), \quad 0 \leq t < \infty, \quad (4.2)$$

where λ is the spectral gap of the associated MC.

To introduce the next result, thinking heuristically of the agents who agent i most frequently meets as the “local” agents for i , it is natural to guess that the configuration of the averaging process might become “locally smooth” faster than the “global smoothness” rate implied by Proposition 4.2. In this context we may regard the Dirichlet form $\mathcal{E}(f, f)$ as measuring the “local smoothness”, more accurately the local roughness, of a function f , relative to the local structure of the particular meeting process. The next result implicitly bounds $\mathbb{E}\mathcal{E}(\mathbf{X}(t), \mathbf{X}(t))$ at finite times by giving an explicit bound for the integral over $0 \leq t < \infty$. Note that, from the fact that the spectral gap is strictly positive, we can see directly that $\mathbb{E}\mathcal{E}(\mathbf{X}(t), \mathbf{X}(t)) \rightarrow 0$ exponentially fast as $t \rightarrow \infty$; Proposition 4.3 is a complementary non-asymptotic result.

Proposition 4.3. *For the Averaging process with arbitrary initial configuration $\mathbf{x}(0)$,*

$$\mathbb{E} \int_0^\infty \mathcal{E}(\mathbf{X}(t), \mathbf{X}(t)) dt = 2 \text{var } \mathbf{x}(0).$$

This looks slightly magical because the identity does not depend on the particular rate matrix \mathcal{N} , but of course the definition of \mathcal{E} involves \mathcal{N} .

The next result illustrates the following.

General Principle 3. *Notions of duality are one of the interesting and useful tools in classical IPS, and equally so in the social dynamics models we are studying.*

There is a formal definition of duality in [37], Section 2.3, but we find the concept easier to understand via examples. The duality between the Voter model and coalescing chains (recalled in Section 5.1) is the simplest and most striking example. The relationship we give in Proposition 4.4 for the Averaging model is perhaps more representative of the general style of duality relationships. We repeat the “random penny” argument used in Proposition 4.1. Now there are two pennies, and at any meeting there are independent decisions to hold or pass each penny. Analyzing how the two pennies move leads to a

certain coupling ([5], Sec. 2.4) of two (half-speed) associated Markov chains $(Z_1(t), Z_2(t))$ – the motion of the two pennies is coupled because, at a meeting of their two owners, one or both may be passed. The coupled process has the following relationship with the Averaging process.

Proposition 4.4 (The duality relation for the Averaging process). *For the averaging model $\mathbf{X}(t)$ started from a configuration $\mathbf{x}(0)$ which is a probability distribution over agents, and for each t ,*

$$\mathbb{E}(X_i(t)X_j(t)) = \mathbb{P}(Z_1(t) = i, Z_2(t) = j),$$

where $(Z_1(t), Z_2(t))$ denotes the coupled process started from random agents $(Z_1(0), Z_2(0))$ chosen independently from $\mathbf{x}(0)$.

On a simple geometry where one can calculate $\mathbb{E}X_i(t)$ via Proposition 4.1, one could seek to use Proposition 4.4 to investigate second-order properties of the Averaging process, and some open problems of this type are mentioned in [5].

Finally, parallel to Proposition 4.2 which builds upon a standard L^2 bound for reversible Markov chains is the following result [5], inspired by the theory surrounding log-Sobolev inequalities [26]. For a configuration \mathbf{x} which is a probability distribution write

$$\text{Ent}(\mathbf{x}) := - \sum_i x_i \log x_i$$

for the entropy of the configuration. Consider the averaging process where the initial configuration is a probability distribution. By concavity of the function $-x \log x$ it is clear that in the averaging process $\text{Ent}(\mathbf{X}(t))$ can only increase, and hence $\text{Ent}(\mathbf{X}(t)) \uparrow \log n$ a.s. (recall $\log n$ is the entropy of the uniform distribution).

Proposition 4.5. *For the Averaging process whose initial configuration is a probability distribution $\mathbf{x}(0)$,*

$$\mathbb{E}(\log n - \text{Ent}(\mathbf{X}(t))) \leq (\log n - \text{Ent}(\mathbf{x}(0))) \exp(-\alpha t/2),$$

where α is the log-Sobolev constant of the associated Markov chain.

5. The Voter model

This model has a random update rule, most elegantly implemented via a “directed” convention in the meeting model: when agents i, j meet, choose a (uniform) random direction and indicate it using an arrow $i \rightarrow j$ or $j \rightarrow i$.

Voter model. Initially each agent has a different “opinion” – agent i has opinion i . When i and j meet at time t with direction $i \rightarrow j$, then agent j adopts the current opinion of agent i .

This model and many variants have acquired a substantial literature in the finite setting (as well as the classical infinite lattice setting) so what's written here is far from complete. In the classical setting, one traditionally assumed only two different initial opinions. In our finite-agent case, it seems more natural to take the initial opinions to be all different.

For notation, write

$$\mathcal{V}_i(t) := \text{the set of } j \text{ who have opinion } i \text{ at time } t.$$

So $\{\mathcal{V}_i(t), i \in \mathbf{Agents}\}$ is a random partition of **Agents**. Note that $\mathcal{V}_i(t)$ may be empty, or may be non-empty but not contain i . The number of different remaining opinions can only decrease with time. Ultimately the process must be absorbed in one of the n “everyone has same opinion” configurations. This fits one of our Section 3.3 general possibilities for a FMIE process: absorption in a random one (of a small number of) “ordered” configurations. A natural quantity of interest is this **consensus time**

$$T^{\text{voter}} := \min\{t : \mathcal{V}_i(t) = \mathbf{Agents} \text{ for some } i\}.$$

5.1. Duality with coalescing chains

Coalescing chains model. Initially each agent has a token – agent i has token i . At time t each agent i has a (maybe empty) collection (cluster) $\mathcal{C}_i(t)$ of tokens. When i and j meet at time t with direction $i \rightarrow j$, then agent i gives his tokens to agent j ; that is,

$$\mathcal{C}_j(t+) = \mathcal{C}_j(t-) \cup \mathcal{C}_i(t-), \quad \mathcal{C}_i(t+) = \emptyset.$$

So $\{\mathcal{C}_i(t), i \in \mathbf{Agents}\}$ is a random partition of **Agents**. A natural quantity of interest is the **coalescence time**

$$T^{\text{coal}} := \min\{t : \mathcal{C}_i(t) = \mathbf{Agents} \text{ for some } i\}.$$

Regarding each non-empty cluster as a particle, the particles move as the associated Markov chain at half-speed (rates $\nu_{ij}/2$), moving independently until two particles meet and thereby coalesce. Note this factor 1/2 in this section.

Proposition 5.1 (The duality relation for the Voter model). *For fixed t , $\{\mathcal{V}_i(t), i \in \mathbf{Agents}\} \stackrel{d}{=} \{\mathcal{C}_i(t), i \in \mathbf{Agents}\}$. In particular $T^{\text{voter}} \stackrel{d}{=} T^{\text{coal}}$.*

It is important to realize that they are different as processes. For fixed i , note that $|\mathcal{V}_i(t)|$ can only change by ± 1 , but $|\mathcal{C}_i(t)|$ jumps to and from 0.

The proof in the classical lattice setting [34] (which refers to coalescing *random walks* extends unchanged to our setting of general symmetric rates (ν_{ij}) . One simply considers a realization of the (directed) meeting process on a fixed time interval $[0, t]$, uses this to construct the coalescing chains process at time t , then observes that reversing the direction of time and the arrows gives a construction of the Voter model at time t .

Almost all of the literature on the finite Voter model has focussed on estimating $T^{\text{voter}} \stackrel{d}{=} T^{\text{coal}}$ in terms of the geometry, and I will show some of this work in Sections 5.2–5.5. But there are other questions one can ask about the finite-time behavior, to be mentioned in Section 5.6.

5.2. The Voter model on the complete graph

There are two ways to analyze T_n^{voter} on the complete graph, both providing some bounds on other geometries. The first uses coalescing chains. One aspect of **Kingman's coalescent** [13] is the continuous-time MC on states $\{1, 2, 3, \dots\}$ with standardized rates $\lambda_{k,k-1} = \binom{k}{2}$, $k \geq 2$. For that chain, the mean hitting time to state 1 is

$$\mathbb{E}_m T_1^{\text{hit}} = \sum_{k=2}^m 1 / \binom{k}{2} = 2 \left(1 - \frac{1}{m}\right)$$

and in particular $\lim_{m \rightarrow \infty} \mathbb{E}_m T_1^{\text{hit}} = 2$. Now when we consider coalescing chains on the complete n -graph, the number of clusters evolves as the continuous-time MC on states $\{1, 2, 3, \dots, n\}$ with rates $\lambda_{k,k-1} = \frac{1}{n-1} \binom{k}{2}$. So $\mathbb{E} T_n^{\text{coal}} = (n-1) \times 2 \left(1 - \frac{1}{n}\right)$ and in particular

$$\mathbb{E} T_n^{\text{voter}} = \mathbb{E} T_n^{\text{coal}} \sim 2n. \quad (5.1)$$

The second way is to consider the Voter model directly, specifically the variant of the voter model with only 2 opinions. Consider the number $X(t)$ of agents with the first opinion. On the complete n -graph, $X(t)$ evolves as the continuous-time MC on states $\{0, 1, 2, \dots, n\}$ with rates

$$\lambda_{k,k+1} = \lambda_{k,k-1} = \frac{k(n-k)}{2(n-1)}.$$

This process arises in classical applied probability (e.g., as the Moran model in population genetics). We want to study the absorption time

$$T_{0,n}^{\text{hit}} := \min\{t : X(t) = 0 \text{ or } n\}.$$

By general birth-and-death formulas, or by comparison with simple RW,

$$\mathbb{E}_k T_{0,n}^{\text{hit}} = \frac{2(n-1)}{n} (k(h_{n-1} - h_{k+1}) + (n-k)(h_{n-1} - h_{n-k+1})) \quad \text{where } h_m := \sum_{i=1}^m 1/i.$$

This is maximized by $k = \lfloor n/2 \rfloor$, and

$$\max_k \mathbb{E}_k T_{0,n}^{\text{hit}} \sim (2 \log 2)n.$$

Now we can couple the original Voter model (n different initial opinions) with the variant with only 2 opinions, initially held by k and $n-k$ agents. (Just randomly assign

these two opinions, initially.) From this coupling, we see

$$\begin{aligned}\mathbb{P}_k(T_{0,n}^{\text{hit}} > t) &\leq \mathbb{P}(T_n^{\text{voter}} > t), \\ \mathbb{P}_k(T_{0,n}^{\text{hit}} > t) &\geq \frac{2k(n-k-1)}{n(n-1)} \mathbb{P}(T_n^{\text{voter}} > t),\end{aligned}$$

where the fraction is the chance that two specified agents are assigned different opinions, initially. In particular, the latter with $k = \lfloor n/2 \rfloor$ implies

$$\mathbb{E}T_n^{\text{voter}} \leq (4 \log 2 + o(1))n.$$

Comparing with the correct asymptotics (5.1), we see that although constant is not precise, the order of magnitude is correct.

5.3. A crude bound for general geometries

Consider a general geometry (ν_{ij}) , and write $\nu(A, A^c) := \sum_{i \in A, j \in A^c} \nu_{ij}$. Suppose the flow rates satisfy, for some constant $\kappa > 0$,

$$\nu(A, A^c) := \sum_{i \in A, j \in A^c} \nu_{ij} \geq \kappa \frac{|A|(n-|A|)}{n(n-1)}.$$

The right side is normalized so that on the complete graph this holds with $\kappa = 1$. We can repeat the analysis above – the process $X(t)$ now moves at least κ times as fast as on the complete graph, and so

$$\mathbb{E}T_n^{\text{voter}} \leq (4 \log 2 + o(1))n/\kappa.$$

This style of argument can be used for many FMIE processes.

General Principle 4. *Bottleneck statistics give crude general bounds.*

Often the speed at which the asymptotic behavior sets in depends on some quantification of the connectivity of the geometry. If for some m the quantity

$$\phi(m) = \min\{\nu(A, A^c) : |A| = m\}, \quad 1 \leq m \leq n-1$$

is small, it indicates a possible “bottleneck” subset of size m . It is convenient to combine the family $(\phi(m), 1 \leq m \leq n/2)$ into a single statistic, but the appropriate way to do this is (FMIE) model-dependent. In the Voter model case above, we used

$$\kappa := \min_A \frac{n(n-1)\nu(A, A^c)}{|A|(n-|A|)} = n(n-1) \min_m \frac{\phi(m)}{m(n-m)}.$$

Quantities like κ are descriptive statistics of a weighted graph, often given names like “isoperimetric constants”, but “bottleneck statistics” seems a more informative name.

The best known use of bottleneck statistics is Cheeger’s inequality (which bounds the spectral gap of a Markov chain in terms of κ) and its generalizations using other bottleneck statistics – see [38] for an detailed treatment. Such techniques are used extensively in the theory of algorithms, where κ is called *conductance* – see, e.g., [21] for analysis of a gossip algorithm.

5.4. Coalescing chains on general geometries

Estimating T^{coal} is clearly related to study of the *meeting time* T_{ij}^{meet} of two independent copies of the chain started at i and j , a topic that arises in other contexts. Recall that on the complete graph the mean coalescence time is asymptotically twice the mean meeting time. Also note that under enough symmetry (e.g., continuous-time random walk on the discrete torus) the relative displacement between the two copies evolves as the same random walk run at twice the speed, and study of T_{ij}^{meet} reduces to study of T_k^{hit} . All this suggests there is a close connection between the coalescence time and hitting times for the associated chain. The following result was conjectured long ago but only recently proved. For any chain, define a statistic

$$\tau^* := \max_{i,j} \mathbb{E}_i T_j^{\text{hit}}.$$

Theorem 5.1 ([41]). *There exist numerical constants $C_1, C_2 < \infty$ such that, for any finite irreducible reversible Markov chain,*

$$\max_{i,j} \mathbb{E} T_{ij}^{\text{meet}} \leq C_1 \tau^* \quad \text{and} \quad \mathbb{E} T^{\text{coal}} \leq C_2 \tau^*.$$

The proof involves intricate analysis that I will not try to outline here. If the stationary distribution π contains exponentially rare states then the theorem is uninformative. But in our setting of FMIE processes, the associated chain always has *uniform* stationary distribution. One can create examples of geometries in which the bound is not the correct order of magnitude, but I suspect that in any “natural” geometry Theorem 5.1 does give the correct order.

A different style of general result given recently in [22] shows that on a regular graph $\mathbb{E} T^{\text{coal}} = O(n/\lambda)$ where λ is the spectral gap of the associated MC. But this bound is of correct order only for expander-like graphs.

5.5. How general is the Kingman coalescent calculation?

Return to the setting of Section 5.1 – chains are associated chains run at half speed. Write τ_{meet} for the mean meeting time from independent uniform starts. In a sequence of chains with $n \rightarrow \infty$, impose a condition such as the following. For each $\varepsilon > 0$

$$n^{-2} |\{(i, j) : \mathbb{E} T_{ij}^{\text{meet}} \notin (1 \pm \varepsilon) \tau_{\text{meet}}\}| \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (5.2)$$

Open Problem 2. Assuming (5.2), under what further conditions can one prove

$$\mathbb{E}T^{\text{coal}} \sim 2\tau_{\text{meet}} \quad \text{as } n \rightarrow \infty? \quad (5.3)$$

That such a result is anticipated has been folklore since 1989, when Cox [23] proved (5.3) for the torus $[0, m-1]^d$ in dimension $d \geq 2$. (In that geometry, $\tau_{\text{meet}} \sim m^d R_d$ for $d \geq 3$, where R_d is the mean number of visits of standard random walk on \mathbb{Z}^d to the origin.) See [15] for a physicist’s view.

One could try to prove this in two stages.

Stage 1. For fixed m , show that the mean time for m initially independent uniform walkers to coalesce should be $\sim 2(1 - \frac{1}{m})\tau_{\text{meet}}$.

Stage 2. Show that for $m(n) \rightarrow \infty$ slowly, the time for the initial n walkers to coalesce into $m(n)$ clusters is $o(\tau_{\text{meet}})$.

It is likely that Stage 1 can be proved without further assumptions, as follows. From old results on mixing times [7], condition (5.2) is enough to show that the variation distance mixing time $\tau_{\text{mix}} = o(\tau_{\text{meet}})$. So – as a prototype use of τ_{mix} – by considering time intervals of length τ , for $\tau_{\text{mix}} \ll \tau \ll \tau_{\text{meet}}$, the events “a particular pair of walkers meets in the next τ -interval” are approximately independent. This makes the “number of clusters” process behave (to first order) as the Kingman coalescent as the number decreases from m to 1. This approach is implemented in [47] for \mathbb{Z}_m^2 , but it should hold for general geometries.

On the other hand, Stage 2 requires some different assumptions to control short-time behavior. See [42] for a recent extensive study of this problem, and the current best results.

5.6. Pre-asymptotic behavior

As mentioned earlier, almost all of the literature on the finite Voter model has focussed on estimating T^{voter} in terms of the geometry. But there are other aspects of finite-time behavior one might study, and we can uncover some such aspects by comparison with what we have studied in the Averaging process. Recall we start with all opinions different.

1. If the proportions of agents with the various opinions are written as $\mathbf{x} = (x_i)$, the statistic $q := \sum_i x_i^2$ is one measure of concentration versus diversity of opinion. This suggests studying $Q(t) := \sum_i (n^{-1}|\mathcal{V}_i(t)|)^2$. Duality implies

$$\mathbb{E}Q(t) = \mathbb{P}(T^{\text{meet}} \leq t),$$

where T^{meet} is the meeting time for independent chains with uniform starts. One can study the distribution of T^{meet} in special geometries, and more generally can often apply the explicit Exponential distribution approximation from [4], Prop. 3.23.

2. A corresponding “local” measure of concentration versus diversity is the probability that agents (I, J) chosen with probability $\propto \nu_{ij}$ (“random neighbors”) have same opinion at time t . Refinements of this “clustering” behavior have been studied in the classical two-dimensional lattice setting [24].

3. The diversity statistic $q := \sum_i x_i^2$ emphasizes large clusters (large time); the statistic $\text{Ent}(\mathbf{x}) = -\sum_i x_i \log x_i$ emphasizes small clusters (small time). So one could consider

$$\mathbb{E}(t) := -\sum_i (n^{-1}|\mathcal{V}_i(t)|) \log(n^{-1}|\mathcal{V}_i(t)|).$$

This has apparently not been studied.

5.7. Approximating finite graphs by infinite graphs

General Principle 5. *Certain special families of geometries have $n \rightarrow \infty$ local weak limits (which are infinite rooted random networks), and aspects of FMIE models on the finite geometries can often be related to the behavior of the model on the infinite geometry.*

Two familiar examples of such geometries are

- (i) For the torus \mathbb{Z}_m^d , the $m \rightarrow \infty$ limit is the infinite lattice \mathbb{Z}^d ;
- (ii) For the “random graphs with prescribed degree distribution” geometry, the limit is a Galton-Watson tree with a size-biased offspring distribution after the first generation;

and two less familiar examples are

- (iii) For the complete graph with random edge-weights, the limit is the PWIT [6], an infinite-degree tree where the weights on the edges at a given vertex form a rate-1 Poisson process on $[0, \infty)$;
- (iv) A more elaborate graph [8] built over (iii), designed to have a $n \rightarrow \infty$ local weak limit which is more complex (not just a tree) yet still fairly tractable.

The relation between finite and infinite geometries is simplest for the “epidemic” (FPP) type models later, but also can be used for Markov chain related models, starting from the following.

Local transience principle. For a large finite-state Markov chain whose behavior near a state i can be approximated by a transient infinite-state chain, we have

$$\mathbb{E}_\pi T_i^{\text{hit}} \approx R_i / \pi_i,$$

where R_i is defined in terms of the approximating infinite-state chain as $\int_0^\infty p_{ii}(t) dt$. The approximation (an instance of the Poisson clumping heuristic [1]) comes from the finite-state mean hitting time formula ([4], Lemma 2.11) via a “interchange of limits” procedure which requires ad hoc justification.

In the context of the Voter model, if there is a local weak limit of the finite geometries, the local transience principle implies that the property $\mathbb{E}T^{\text{voter}} = \Theta(n)$ is essentially equivalent to transience of the chain on the limit infinite geometry.

6. Two further models

6.1. The Compulsive Gambler process

The Compulsive Gambler process. Initially each agent has some (nonnegative real-valued) amount of money. Whenever two agents meet, they instantly play a fair game in which one agent acquires the combined money. In other words, if one has a and the other has b then the first acquires all $a + b$ with chance $a/(a + b)$.

Note that on the complete graph geometry, this process is just an augmentation of the Kingman coalescent process. On a general geometry, a configuration in which the set of agents with non-zero money forms an “independent set” (no two are adjacent in the weighted graph) is obviously an absorbing configuration, and conversely.

This process has apparently not been studied. I mention it because it provides a simple example of the “disordered” limit behavior mentioned in Section 3.3.

6.2. The Interchange (Exclusion) process

The Interchange process. Initially each agent has a different token. When two agents meet, they exchange tokens.

I will call this the Interchange process; the “two type” version is the classical Exclusion process. Though this model is not particularly natural in the social networks context, the result below (which established a longstanding conjecture) illustrates a result for which our general “weighted graphs” setting is the (mathematically) correct setting.

The motion of an individual token is as the associated Markov chain, so has some spectral gap $\lambda_{MC}(\mathcal{N}) > 0$. The whole Interchange process is itself a reversible Markov chain, so has a spectral gap $\lambda_{IP}(\mathcal{N})$ which by a general “contraction” fact ([4], Sec. 4.6.1) satisfies $\lambda_{IP}(\mathcal{N}) \leq \lambda_{MC}(\mathcal{N})$.

Proposition 6.1 ([18]). *For every geometry $\mathcal{N} = (\nu_{ij})$ we have $\lambda_{IP}(\mathcal{N}) = \lambda_{MC}(\mathcal{N})$.*

The proof involves intricate analysis that I will not try to outline here.

7. The Pandemic process

Recall the discussion in Section 3.2. For the purposes of this section, the Pandemic process is the same as the FPP process; we will use the language of “infectives” and “time” (rather than “percolation distance”) and write T_{ij}^{fpp} for the time taken for an epidemic started at i to reach j .

In this section, we will recall the well-understood behaviors on the lattice (Section 7.1) and the complete graph (Section 7.3). Each type of behavior raises questions about the extent to which similar behavior remains true in more general geometries (Sections 7.2 and 7.4) which have not been systematically studied.

7.1. FPP on the torus

We quote the classical result for FPP on the infinite lattice, with brief remarks.

Theorem 7.1 (Shape theorem for lattice FPP [35]). *Consider the FPP process on the edges of \mathbb{Z}^d with Exponential(1) edge-lengths, starting from the origin. Write $\mathcal{S}(t) \subset \mathbb{Z}^d$ for the infected set at time t , and $\bar{\mathcal{S}}(t) \subset \mathbb{R}^d$ for its fattening (replace each point by a cube of side 1). There is a non-random closed convex set $B = B_d$ such that, for each $0 < \varepsilon < 1$,*

$$\mathbb{P}((1 - \varepsilon)tB \subseteq \bar{\mathcal{S}}(t) \subseteq (1 + \varepsilon)tB) \rightarrow 1.$$

- The key proof ingredient is the *subadditive ergodic theorem*.
- The result hold for more general IID edge-times.
- B_d is *not* the unit ball in Euclidean norm.
- For later use write b_d for the (d -dimensional) “volume” of B_d .
- Understanding the asymptotic variance of T_{ij}^{fpp} is a famous hard problem; see, e.g., [12].

Consider now the FPP process on the discrete torus \mathbb{Z}_m^d . We use our standardization convention (2.2), so that rates $\nu_{ij} = 1/(2d)$ for each edge (i, j) . The following three “first-order” results are immediate from the shape theorem. Asymptotics are as $m \rightarrow \infty$ for fixed d .

$$\mathbb{E}|\mathcal{S}(t)| \sim (2d)^{-1} b_d t^d \text{ over } 1 \ll t \ll m. \quad (7.1)$$

The same holds (in L^1) for $|\mathcal{S}(t)|$ itself. Second, the time $T_*^{\text{pan}}(m)$ until all agents are infected satisfies

$$m^{-1} T_*^{\text{pan}}(m) \rightarrow c_d \quad \text{a.s.} \quad (7.2)$$

for a constant c_d defined in terms of the shape of B_d . Third, the proportion $X_m(t)$ of infected agents satisfies

$$\sup_{0 \leq s < \infty} |X_m(ms) - F_d(s)| \rightarrow_p 0 \quad (7.3)$$

for a certain function F_d defined in terms of the shape of B_d .

7.2. Toward a general WLLN

Having a limit “shape” represented by B_d is very special to the lattice setting. What about the general setting of rates (ν_{ij}) ? Suppose we have a result of the kind

(*) provided i and j are not close, then T_{ij}^{fpp} is close to its mean.

Then the random set $\mathcal{S}_i(t)$ of infectives at t from initial i is approximately the deterministic set $\{j : \mathbb{E}T_{ij}^{\text{fpp}} \leq t\}$. So, if the essence of a “shape theorem” is a deterministic

approximation to the infected set, it is essentially just a WLLN (weak law of large numbers).

A precise conjecture is formulated below. Everything depends on n (not written) and limits are as $n \rightarrow \infty$. The conclusion we want is

$$\frac{T_{ij}^{\text{fpp}}}{\mathbb{E}T_{ij}^{\text{fpp}}} \rightarrow_p 1. \quad (7.4)$$

We work in the FPP model, writing ξ_{ab} for the time to traverse edge (a, b) .

The obvious obstacle to the desired conclusion is that there may be a set A (with $i \in A, j \notin A$) such that, in the random percolation path π_{ij} , the time taken to traverse the edge $A \rightarrow A^c$ is not $o(\mathbb{E}T_{ij}^{\text{fpp}})$. One could conjecture this is the only obstacle; here is a slightly weaker conjecture.

Conjecture 1. *With arbitrary rates (ν_{ij}) , if*

$$\frac{\max\{\xi_{ab} : (a, b) \text{ edge in } \pi_{ij}\}}{\mathbb{E}T_{ij}^{\text{fpp}}} \rightarrow_p 0 \quad (7.5)$$

then (7.4) holds.

It is intuitively clear (and not hard to prove) that condition (7.5) is necessary. Proving the conjecture would hardly be a definitive result (because condition (7.5) is not a condition directly on (ν_{ij})) but would represent a start on understanding when the WLLN holds.

7.3. FPP on the complete n -vertex graph

Here we have rates

$$\nu_{ij} = 1/(n-1), \quad j \neq i$$

and we start with one infective. This is a fundamental stochastic process whose basic properties have been rediscovered many times from different viewpoints, but there seems no authoritative survey. Our focus is on the randomly-shifted logistic limit result, Proposition 7.1. We give an outline taken from [2] (a similar treatment is in [29]), though essentially the same result can be found in older genetics literature. There are two reasons for this particular focus. From this one result, we can read off many results like (7.14, 7.15), a fact which has not been appreciated in the combinatorics-inspired literature. And secondly, it plays a central role in investigating FMIE processes build over Pandemic, as we shall see in Section 8.

Write $D(k) = D_n(k) =$ time at which k people are infected. The observation

$$\text{the r.v.'s } D(k+1) - D(k) \text{ are independent with Exponential}\left(\frac{k(n-k)}{n-1}\right) \text{ distribution} \quad (7.6)$$

is a basic starting place for analysis. For instance, the time T_*^{pan} until the whole population is infected is just $D(n)$, and so

$$\mathbb{E}T_*^{\text{pan}} = \sum_{k=1}^{n-1} \mathbb{E}(D(k+1) - D(k)) = (n-1) \sum_{k=1}^{n-1} \frac{1}{k(n-k)} \sim 2 \log n.$$

And the time $T_{\text{rand}}^{\text{fpp}}$ until a random person (amongst the $n-1$ initially uninfected) is infected is

$$\begin{aligned} \mathbb{E}T_{\text{rand}}^{\text{fpp}} &= \mathbb{E}D(U), \quad U \text{ uniform on } \{2, 3, \dots, n\} \\ &= \frac{1}{n-1} \sum_{k=2}^n \mathbb{E}D(k) = \frac{1}{n-1} \sum_{k=1}^{n-1} (n-k) \mathbb{E}(D(k+1) - D(k)) \\ &= \sum_{k=1}^{n-1} \frac{1}{k} \sim \log n. \end{aligned} \tag{7.7}$$

This analysis can be pursued to get distributional results, but let us first look at different methods that indicate the behavior without extensive calculation. The idea is to analyze separately the *initial phase* when $o(n)$ people are infected, and the *pandemic phase* when $\Theta(n)$ people are infected.

The “deterministic, continuous” analog of our “stochastic, discrete” model of an epidemic is the *logistic equation*

$$F'(t) = F(t)(1 - F(t))$$

for the proportion $F(t)$ of a population infected at time t . A solution is a shift of the basic solution

$$F(t) = \frac{e^t}{1 + e^t}, \quad -\infty < t < \infty \quad (\text{the logistic function}).$$

Setting $M_n(t)$ for the number of individuals infected in our stochastic process, then scaling to the proportion $X_n(t) = M_n(t)/n$, we find

$$\mathbb{E}(dX_n(t) | \mathcal{F}_n(t)) = \frac{n}{n-1} X_n(t)(1 - X_n(t)) dt.$$

By including variance estimates, one could prove the following formalization of the idea that, during the pandemic phase, $X_n(t)$ behaves as $F(t)$ to first order. Precisely, for fixed ε and t_0 we have (in probability)

$$(X_n(D_n(n\varepsilon) + t), 0 \leq t \leq t_0) \rightarrow (F(F^{-1}(\varepsilon) + t), 0 \leq t \leq t_0). \tag{7.8}$$

That is, we “start the clock” when a proportion ε of the population is infected. This can now be reformulated more cleanly in terms of the time $G_n = D_n(n/2)$ at which half the

population is infected:

$$\sup_{-G_n \leq t < \infty} |n^{-1} M_n(G_n + t) - F(t)| \rightarrow 0 \quad \text{in probability.}$$

For later use, consider the times at which each process equals k_n/n , where $k_n/n \rightarrow 0$ slowly. We find

$$G_n + \log(k_n/n) = D_n(k_n) + o(1)$$

which rearranges to

$$D_n(k_n) - \log k_n = G_n - \log n + o(1). \quad (7.9)$$

Now turn to the initial phase of the epidemic. On a fixed initial time interval $[0, t_0]$, the process $M_n(t)$ of number of infectives converges in distribution to the process $M_\infty(t)$ for which the times $D(k)$ satisfy

the r.v.'s $D(k+1) - D(k)$ are independent with Exponential(k) distribution

and this is the classic *Yule process*, for which it is well known that

$$\begin{aligned} M_\infty(t) &\text{ has Geometric } (e^{-t}) \text{ distribution} \\ e^{-t} M_\infty(t) &\rightarrow \mathcal{E}_\infty \text{ a.s. as } t \rightarrow \infty, \text{ where } \mathcal{E}_\infty \text{ has Exponential}(1) \text{ distribution.} \end{aligned} \quad (7.10)$$

Now calculate informally, for $1 \ll k \ll n$,

$$\begin{aligned} \mathbb{P}(D_n(k) \leq t) &= \mathbb{P}(M_n(t) \geq k) \approx \mathbb{P}(M_\infty(t) \geq k) \\ &= \mathbb{P}(e^{-t} M_\infty(t) \geq k e^{-t}) \approx \mathbb{P}(\mathcal{E}_\infty \geq k e^{-t}) = \exp(-k e^{-t}). \end{aligned}$$

In other words

$$D_n(k) - \log k \approx_d G, \quad 1 \ll k \ll n, \quad (7.11)$$

where G has the Gumbel distribution $\mathbb{P}(G \leq x) = \exp(-e^{-x})$. Comparing with (7.9) gives (7.12) below.

Proposition 7.1 (The randomly-shifted logistic limit). *For the simple epidemic process on the complete n -vertex graph, there exist random G_n such that*

$$\begin{aligned} \sup_{-G_n \leq t < \infty} |n^{-1} M_n(G_n + t) - F(t)| &\rightarrow 0 \quad \text{in probability,} \\ G_n - \log n &\xrightarrow{d} G, \end{aligned} \quad (7.12)$$

where F is the logistic function and G has Gumbel distribution.

One could prove this by formalizing the arguments above, but there is a more efficient though less illuminating way. The result is essentially equivalent to the assertion

$$D_n(\lfloor un \rfloor) - \log n \xrightarrow{d} F^{-1}(u) + G, \quad 0 < u < 1. \quad (7.13)$$

The basic observation (7.6) allows one to write down an explicit expression for the Fourier transform of the left side, and one just needs to work through the calculus to check it converges to the Fourier transform of the right side. (Transform methods are of course very classical in applied probability, and their use in epidemic models goes back at least to [10].)

The same transform argument (then taking a different limit regime) can be used to prove the following result (e.g., [45]) for the time $D_n(n)$ until the entire population is infected:

$$D_n(n) - 2 \log n \xrightarrow{d} G_1 + G_2, \quad (7.14)$$

where G_1, G_2 are independent with Gumbel distribution. While not obvious from the statement of the proposition above, it can in fact be deduced as follows. From (7.13), the time $D_n(n/2)$ until half the population is infected satisfies

$$D_n(n/2) - \log n \xrightarrow{d} G.$$

Then symmetry (under $k \rightarrow n - k$) of the transition rates (7.6) implies that $D_n(n)$ is essentially distributed as the sum of two independent copies of $D_n(n/2)$.

Reconsider the time $D_n(U)$ until a random person is infected. Proposition 7.1 implies it has asymptotic distribution

$$D_n(U) - \log n \xrightarrow{d} G + L, \quad (7.15)$$

where L has logistic distribution function, independent of the Gumbel time-shift G . So we expect

$$\mathbb{E}D_n(U) - \log n \rightarrow \mathbb{E}G + \mathbb{E}L.$$

Now $\mathbb{E}L = 0$ by symmetry ($L \stackrel{d}{=} -L$) and one can calculate (or look up) that $\mathbb{E}G = \text{Euler's constant}$, so the limit theorem is consistent with the exact formula (7.7).

7.4. In what other geometries does similar behavior hold?

The extent to which different aspects of the “mean-field” behavior in Section 7.3 extend to more general geometries is not well understood. For the “random graphs with prescribed degree distributions” (with appropriate moments finite) geometry, one expects to see qualitatively similar behavior (the pandemic phase to have duration $O(1)$, centered at time $O(\log n)$), and almost precise analogs have indeed been proved – e.g., [14], Theorem 3.1(b) gives the analog of (7.15). There has been extensive subsequent work on related problems and related “locally tree-like” geometries, but these proofs are sub-

stantially more technical than for the complete graph. And even on the Hamming cube $\{0, 1\}^d$ the behavior of FPP is not well understood [32].

To think about possible “general theory”, let me suggest a loose analogy with the established theory of random walk on infinite graphs [46]. In that theory, most geometries fall into one of two categories, exemplified by \mathbb{Z}^d and d -trees, according to whether

$$\mathbb{P}_i(X(t) = i) \rightarrow 0 \quad \text{as } t^{-\alpha} \text{ or } e^{-\beta t}.$$

Properties involving spectral radius, isoperimetric inequalities, amenability, (sub)linear rate of escape can then be related to this categorization. That theory formalizes a distinction (informally, between “finite dimensional” and “infinite dimensional”) from the random walk viewpoint.

So by analogy, we have seen that the finite Pandemic process behaves differently on \mathbb{Z}_m^d and on the complete graph. Can we select one specific aspect of the different behavior as a useful basis for classifying geometries? Write $t_{ij} = \mathbb{E}T_{ij}^{\text{fpp}}$. One obvious distinction between these two geometries is that on \mathbb{Z}_m^d the means t_{ij} vary (depending on distance between i and j) but on the complete graph, the t_{ij} are the same (to first order) for most pairs (i, j) .

So we can define a property (for a sequence of weighted graphs (ν_{ij})) which combines the idea of the WLLN (7.4) with the idea of the t_{ij} not varying much:

$$\begin{aligned} &\text{there exist } t^{(n)} \text{ such that as } n \rightarrow \infty \text{ for fixed } \varepsilon > 0 \\ &n^{-2} |\{(i, j) : \mathbb{P}(T_{ij}^{\text{fpp}} \notin (1 \pm \varepsilon)t^{(n)}) > \varepsilon\}| \rightarrow 0. \end{aligned} \quad (7.16)$$

This property (which holds for the “random graphs with prescribed degree distribution” geometry [14]) is an attempt to formalize the notion of a geometry being “locally infinite-dimensional” from the FPP viewpoint.

Open Problem 3. *Study equivalences between this property and the “geometry” of (ν_{ij}) .*

Two brief comments.

(i) We can’t hope to establish (7.16) using only isoperimetric inequalities – they can only give order of magnitude.

(ii) There is a loose analogy with the cutoff phenomenon in finite Markov chains [25], in that we expect the property to hold in “highly symmetric” geometries. But that analogy is discouraging in that there is no useful general condition for cut-off.

7.5. An epidemic process on the discrete torus \mathbb{Z}_m^2 with short- and long-range interactions

Consider the geometry \mathcal{N} with agents \mathbb{Z}_m^2 and meeting rates

- (i) at rate 1 an agent meets a uniform random neighbor;
- (ii) at rate $m^{-\alpha}$ an agent meets a uniform non-neighbor.

Here $0 < \alpha < 3$. Over this range the geometry qualitatively interpolates between the complete-graph case and the nearest-neighbor lattice case. So we expect the center of the pandemic window to interpolate between orders $\log m$ and m , and the width of the pandemic window to interpolate between orders 1 and m .

For the Pandemic model over this geometry, it turns out that an analog of the “randomly-shifted logistic limit” holds, albeit with a different deterministic function F in place of the logistic, and different scaling. This is rather counter-intuitive – we envisage infectives as growing patches of mold in a Petri dish, one might expect the detailed geometry of overlaps to play a role. But the key qualitative property is that, during the pandemic phase, the new infectives arise from “colonies” started earlier in the pandemic phase, not in the initial phase. This enables us to derive the following equation for the standardized function F describing proportion of infectives r over the pandemic window:

$$1 - F(t) = \exp\left(-\int_{-\infty}^t F(s)(t-s)^2 ds\right).$$

See [2] for heuristics and [20] for careful treatment of a closely related model.

8. Some social analogs of epidemics

There is a large literature modeling spread of opinions, adoption of innovations etc via social networks, and some of the models fit our FMIE setting. Instead of describing some standard model let me describe two new FMIE models suggested by the “social networks” context. Mathematically they are variants of Pandemic, and we will see in Sections 8.1–8.2 that on the complete graph their behavior can be derived from the precise understanding of Pandemic given by Proposition 7.1.

8.1. The Deference model

The Deference process. The agents are labelled 1 through n , and agent i initially has opinion i . When two agents meet, they adopt the same opinion, the smaller of the two opinion-labels.

Clearly opinion 1 spreads as Pandemic, so the “ultimate” behavior of Deference is not a new question.

We will consider the complete graph geometry, and study $(X_1^n(t), \dots, X_k^n(t))$, where $X_k^n(t)$ is the proportion of the population with opinion k at time t . We can rewrite Proposition 7.1 as

$$(X_1^n(\log n + s), -\infty < s < \infty) \xrightarrow{d} (F(C_1 + s), -\infty < s < \infty), \quad (8.1)$$

where F is the logistic function and $C_1 = \log(\xi_1)$, where $\xi_1 \stackrel{d}{=} \text{Exponential}(1)$ arises as the limit r.v. (7.10) associated with the Yule process.

The key insight is that opinions 1 and 2 and \dots and k combined behave as one infection in Pandemic, except for a different initial phase, and so the sum $\sum_{i=1}^k X_i^{(n)}(\cdot)$ behaves as in (8.1) but with a different random shift C_k . Because the initial phase is just a collection of k independent Yule processes, we see that the sum has a representation of format (8.1) with

$$C_k = \log(\xi_1 + \dots + \xi_k), \quad k \geq 1. \quad (8.2)$$

Here the (ξ_i) are i.i.d. Exponential(1).

So the $n \rightarrow \infty$ limit behavior of the Deference model in the complete graph geometry is, over the pandemic phase,

$$\begin{aligned} & ((X_1^n(\log n + s), X_2^n(\log n + s), \dots, X_k^n(\log n + s)), -\infty < s < \infty) \\ & \rightarrow ((F(C_1 + s), F(C_2 + s) - F(C_1 + s), \dots, F(C_k + s) - F(C_{k-1} + s)), -\infty < s < \infty). \end{aligned} \quad (8.3)$$

8.2. The Fashionista model

The Deference model envisages agents as “slaves to authority”. Here is a conceptually opposite “slaves to fashion” model, whose analysis is mathematically surprisingly similar.

The Fashionista model. At the times of a rate- λ Poisson process, a new fashion originates with a uniform random agent, and is time-stamped. When two agents meet, they each adopt the latest (most recent time-stamp) fashion.

It is easy to see that, for the random partition of agents into “same fashion”, there is a stationary distribution.

For the complete graph geometry, we can copy the analysis of the Deference model. Combining all the fashions appearing after a given time, these behave as one infection in Pandemic (over the pandemic window; newly-originating fashions have negligible effect), hence their proportion behaves as a random time-shift of the logistic curve F . So when we study the vector $(X_k^n(t), -\infty < k < \infty)$ of proportions of agents adopting different fashions k , we expect $n \rightarrow \infty$ limit behavior of the form

$$(X_k^n(\log n + s), -\infty < k < \infty) \xrightarrow{d} (F(C_k + s) - F(C_{k-1} + s), -\infty < k < \infty), \quad (8.4)$$

where $(C_k, -\infty < k < \infty)$ are the points of some stationary process on $(-\infty, \infty)$. Knowing this form for the $n \rightarrow \infty$ asymptotics, we can again determine the distribution of (C_i) by considering the initial stage of spread of a new fashion. It turns out that

$$C_i = \log \left(\sum_{j \leq i} \exp(\gamma_j) \right) = \gamma_i + \log \left(\sum_{k \geq 1} \exp(\gamma_{i-k} - \gamma_i) \right), \quad (8.5)$$

where $(\gamma_j, -\infty < j < \infty)$ are the times of a rate- λ Poisson process on $(-\infty, \infty)$. The second expression makes it clear that (C_i) is a stationary process.

Here is the outline argument for (8.5). Consider the recent fashions at time $t = 0$ adopted by small but non-negligible proportions of the population. More precisely, con-

sider fashions originating during the time interval $[-\log n + t_n, -\log n + 2t_n]$, where $t_n \rightarrow \infty$ slowly. For a fashion originating at time $-\log n + \eta$, the time-0 set of adopting agents will be a subset of the corresponding epidemic process, which we know has proportional size $\xi \exp(-\eta) = \exp(-\eta + \log \xi)$ where ξ has Exponential(1) distribution.

The times $-\log n + \eta_j$ of origination of different fashions form by assumption a rate- λ Poisson process, and after we impose i.i.d. shifts $\log \xi_j$ we note (as an elementary property of Poisson processes) that the shifted points $-\log n + \eta_j + \log \xi_j$ still form a rate- λ Poisson process, say γ_j , on $(-\infty, \infty)$. So the sizes of small recent fashion groups (that is letting $j \rightarrow -\infty$), for which overlap between fashions becomes negligible, are approximately $\exp(\gamma_j)$. Summing over $j \leq i$ gives

$$\sum_{j \leq i} \exp(\gamma_j) \approx F(C_i) \approx \exp(C_i)$$

and we end up with the representation (8.5).

8.3. Other geometries

On the torus \mathbb{Z}_m^2 the stationary distribution of Fashionista is qualitatively similar to the *dead leaves model* [17] and, as outlined in [3], one can apply simple scaling arguments to analyze the diversity statistic

$$s(m, \lambda) := \mathbb{P}(\text{two random agents share same fashion})$$

and show

$$s(m, \lambda) \sim c \lambda^{-2/3} m^{-2/3} \quad \text{for } m^{-1} \ll \lambda \ll m^2.$$

The corresponding result on the complete graph is

$$s(n, \lambda) \rightarrow s(\infty, \lambda) := \mathbb{E} \sum_k (F(\lambda^{-1} C_k) - F(\lambda^{-1} C_{k-1}))^2$$

and (8.5) provides a (complicated) explicit expression for the right side. In particular, we can see the $\lambda \rightarrow \infty$ behavior is

$$s(\infty, \lambda) \sim c \lambda^{-1}$$

for constant c .

Studying (n, λ) in general geometries is an **Open Topic**, which opens up other questions to study about short-term behavior of Pandemic.

9. Other epidemic models

There is a huge literature attempting more realistic epidemic models, which I will not try to outline here. The FMIE-type model that has been most extensively studied in

the mathematical literature is the Contact process (the SIS model with Exponential infection times), studied classically on \mathbb{Z}^d [28]. A nice survey, in the style of this article, of recent rigorous work on the Contact process on preferential attachment, small worlds and power-law degree geometries can be found in [30], and the latest technical work in [39].

Acknowledgements

The current form of Conjecture 1 emerged from conversations with Jian Ding. And I thank an anonymous referee for many helpful expository suggestions. Research supported by NSF Grant DMS-07-04159.

References

- [1] ALDOUS, D. (1989). *Probability Approximations via the Poisson Clumping Heuristic. Applied Mathematical Sciences* **77**. New York: Springer. [MR0969362](#)
- [2] ALDOUS, D. (2010). When knowing early matters: Gossip, percolation and Nash equilibria. Available at [arXiv:1005.4846](#).
- [3] ALDOUS, D. (2012). Finite Markov information-exchange processes. Beamer slides of five lectures. Available at http://www.stat.berkeley.edu/~aldous/FMIE/2012_FMIE_lectures.html.
- [4] ALDOUS, D. and FILL, J. (2000). Reversible Markov chains and random walks on graphs. Available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- [5] ALDOUS, D. and LANOUE, D. (2012). A lecture on the averaging process. *Probab. Surv.* **9** 90–102. [MR2908618](#)
- [6] ALDOUS, D. and STEELE, J.M. (2004). The objective method: Probabilistic combinatorial optimization and local weak convergence. In *Probability on Discrete Structures. Encyclopaedia Math. Sci.* **110** 1–72. Berlin: Springer. [MR2023650](#)
- [7] ALDOUS, D.J. (1982). Some inequalities for reversible Markov chains. *J. London Math. Soc.* (2) **25** 564–576. [MR0657512](#)
- [8] ALDOUS, D.J. (2003). A stochastic complex network model. *Electron. Res. Announc. Amer. Math. Soc.* **9** 152–161 (electronic). [MR2029476](#)
- [9] ALDOUS, D.J. and SHUN, J. (2010). Connected spatial networks over random points and a route-length statistic. *Statist. Sci.* **25** 275–288. [MR2791668](#)
- [10] BAILEY, N.T.J. (1950). A simple stochastic epidemic. *Biometrika* **37** 193–202. [MR0038636](#)
- [11] BAILEY, N.T.J. (1957). *The Mathematical Theory of Epidemics*. New York: Hafner Publishing Co. [MR0095085](#)
- [12] BENAÏM, M. and ROSSIGNOL, R. (2008). Exponential concentration for first passage percolation through modified Poincaré inequalities. *Ann. Inst. Henri Poincaré Probab. Stat.* **44** 544–573. [MR2451057](#)
- [13] BERESTYCKI, N. (2009). *Recent Progress in Coalescent Theory. Ensaios Matemáticos [Mathematical Surveys]* **16**. Rio de Janeiro: Sociedade Brasileira de Matemática. [MR2574323](#)
- [14] BHAMIDI, S., VAN DER HOFSTAD, R. and HOOGHIEMSTRA, G. (2010). First passage percolation on random graphs with finite mean degrees. *Ann. Appl. Probab.* **20** 1907–1965. [MR2724425](#)

- [15] BLYTHE, R.A. (2010). Ordering in voter models on networks: Exact reduction to a single-coordinate diffusion. *J. Phys. A* **43** 385003, 33. [MR2718319](#)
- [16] BOLLOBÁS, B. (2001). *Random Graphs*, 2nd ed. *Cambridge Studies in Advanced Mathematics* **73**. Cambridge: Cambridge Univ. Press. [MR1864966](#)
- [17] BORDENAVE, C., GOUSSEAU, Y. and ROUEFF, F. (2006). The dead leaves model: A general tessellation modeling occlusion. *Adv. in Appl. Probab.* **38** 31–46. [MR2213962](#)
- [18] CAPUTO, P., LIGGETT, T.M. and RICHTHAMMER, T. (2010). Proof of Aldous’ spectral gap conjecture. *J. Amer. Math. Soc.* **23** 831–851. [MR2629990](#)
- [19] CASTELLANO, C., FORTUNATO, S. and LORETO, V. (2009). Statistical physics of social dynamics. *Rev. Mod. Phys.* **81** 591–646.
- [20] CHATTERJEE, S. and DURRETT, R. (2011). Asymptotic behavior of Aldous’ gossip process. *Ann. Appl. Probab.* **21** 2447–2482. [MR2895421](#)
- [21] CHERICHETTI, F., LATTANZI, S. and PANCONESI, A. (2010). Rumour spreading and graph conductance. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms* 1657–1663. Philadelphia, PA: SIAM. [MR2809775](#)
- [22] COOPER, C., ELSÄSSER, R., ONO, H. and RADZIK, T. (2012). Coalescing random walks and voting on graphs. In *Proceedings of the 2012 ACM symposium on Principles of Distributed Computing* 37–56. ACM: New York.
- [23] COX, J.T. (1989). Coalescing random walks and voter model consensus times on the torus in \mathbf{Z}^d . *Ann. Probab.* **17** 1333–1366. [MR1048930](#)
- [24] COX, J.T. and GRIFFEATH, D. (1986). Diffusive clustering in the two-dimensional voter model. *Ann. Probab.* **14** 347–370. [MR0832014](#)
- [25] DIACONIS, P. (1996). The cutoff phenomenon in finite Markov chains. *Proc. Natl. Acad. Sci. USA* **93** 1659–1664. [MR1374011](#)
- [26] DIACONIS, P. and SALOFF-COSTE, L. (1996). Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Probab.* **6** 695–750. [MR1410112](#)
- [27] DOROGOVTSEV, S.N. and MENDES, J.F.F. (2003). *Evolution of Networks*. Oxford: Oxford Univ. Press. [MR1993912](#)
- [28] DURRETT, R. (1991). The contact process, 1974–1989. In *Mathematics of Random Media (Blacksburg, VA, 1989). Lectures in Applied Mathematics* **27** 1–18. Providence, RI: Amer. Math. Soc. [MR1117232](#)
- [29] DURRETT, R. (2008). *Probability Models for DNA Sequence Evolution*, 2nd ed. *Probability and Its Applications (New York)*. New York: Springer. [MR2439767](#)
- [30] DURRETT, R. (2010). Some features of the spread of epidemics and opinions on a random graph. *Proc. Natl. Acad. Sci. USA* **107** 4491–4498.
- [31] DURRETT, R. and SCHONMANN, R.H. (1988). The contact process on a finite set. II. *Ann. Probab.* **16** 1570–1583. [MR0958203](#)
- [32] FILL, J.A. and PEMANTLE, R. (1993). Percolation, first-passage percolation and covering times for Richardson’s model on the n -cube. *Ann. Appl. Probab.* **3** 593–629. [MR1221168](#)
- [33] FLORIDI, L. (2010). *Information: A Very Short Introduction*. Oxford, UK: Oxford Univ. Press.
- [34] GRIFFEATH, D. (1979). *Additive and Cancellative Interacting Particle Systems. Lecture Notes in Math.* **724**. Berlin: Springer. [MR0538077](#)
- [35] KESTEN, H. (2003). First-passage percolation. In *From Classical to Modern Probability. Progress in Probability* **54** 93–143. Basel: Birkhäuser. [MR2045986](#)
- [36] LEVIN, D.A., PERES, Y. and WILMER, E.L. (2009). *Markov Chains and Mixing Times*. Providence, RI: Amer. Math. Soc. With a chapter by James G. Propp and David B. Wilson. [MR2466937](#)

- [37] LIGGETT, T.M. (1985). *Interacting Particle Systems. Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]* **276**. New York: Springer. [MR0776231](#)
- [38] MONTENEGRO, R. and TETALI, P. (2006). Mathematical aspects of mixing times in Markov chains. *Found. Trends Theor. Comput. Sci.* **1** 1–121. [MR2341319](#)
- [39] MOUNTFORD, T., MOURRAT, J.-C., VALESIN, D. and YAO, Q. (2012). Exponential extinction time of the contact process on finite graphs. Available at [arXiv:1203.2972](#).
- [40] NEWMAN, M.E.J. (2003). The structure and function of complex networks. *SIAM Rev.* **45** 167–256 (electronic). [MR2010377](#)
- [41] OLIVEIRA, R.I. (2012). On the coalescence time of reversible random walks. *Trans. Amer. Math. Soc.* **364** 2109–2128. [MR2869200](#)
- [42] OLIVEIRA, R.I. (2012). Mean field conditions for coalescing random walks. Available at [arXiv:1109.5684](#).
- [43] SHAH, D. (2008). Gossip algorithms. *Foundations and Trends in Networking* **3** 1–125.
- [44] Social Cognitive Networks Academic Research Center (2011). Minority rules: Scientists discover tipping point for the spread of ideas. Available at <http://scnarc.rpi.edu/content/minority-rules-scientists-discover-tipping-point-spread-ideas>.
- [45] VAN DER HOFSTAD, R., HOOGHIEMSTRA, G. and VAN MIEGHEM, P. (2002). The flooding time in random graphs. *Extremes* **5** 111–129 (2003). [MR1965974](#)
- [46] WOESS, W. (2000). *Random Walks on Infinite Graphs and Groups. Cambridge Tracts in Mathematics* **138**. Cambridge: Cambridge Univ. Press. [MR1743100](#)
- [47] ZÄHLE, I., COX, J.T. and DURRETT, R. (2005). The stepping stone model. II. Genealogies and the infinite sites model. *Ann. Appl. Probab.* **15** 671–699. [MR2114986](#)