



Nonlinear Dynamics for the Ising Model

Pietro Caputo¹ , Alistair Sinclair²

¹ Università Roma Tre, Rome, Italy. E-mail: pietro.caputo@uniroma3.it

² University of California, Berkeley, USA. E-mail: sinclair@cs.berkeley.edu

Received: 6 June 2023 / Accepted: 22 August 2024
Published online: 12 October 2024 – © The Author(s) 2024

Abstract: We introduce and analyze a natural class of nonlinear dynamics for spin systems such as the Ising model. This class of dynamics is based on the framework of mass action kinetics, which models the evolution of systems of entities under pairwise interactions, and captures a number of important nonlinear models from various fields, including chemical reaction networks, Boltzmann’s model of an ideal gas, recombination in population genetics and genetic algorithms. In the context of spin systems, it is a natural generalization of linear dynamics based on Markov chains, such as Glauber dynamics and block dynamics, which are by now well understood. However, the inherent nonlinearity makes the dynamics much harder to analyze, and rigorous quantitative results so far are limited to processes which converge to essentially trivial stationary distributions that are product measures. In this paper we provide the first quantitative convergence analysis for natural nonlinear dynamics in a combinatorial setting where the stationary distribution contains non-trivial correlations, namely spin systems at high temperatures. We prove that nonlinear versions of both the Glauber dynamics and the block dynamics converge to the Gibbs distribution of the Ising model (with given external fields) in times $O(n \log n)$ and $O(\log n)$ respectively, where n is the size of the underlying graph (number of spins). Given the lack of general analytical methods for such nonlinear systems, our analysis is unconventional, and combines tools such as information percolation (due in the linear setting to Lubetzky and Sly), a novel coupling of the Ising model with Erdős-Rényi random graphs, and non-traditional branching processes augmented by a “fragmentation” process. Our results extend immediately to any spin system with a finite number of spins and bounded interactions.

PC was supported in part by the Miller Institute at UC Berkeley. AS was supported in part by NSF grants CCF-1815328 and CCF-2231095. Part of this work was done while AS was visiting EPFL, Lausanne and while PC was visiting UC Berkeley.

1. Introduction

Mass action kinetics is a general framework for studying systems of interacting entities. The framework emerged in the study of chemical reaction networks, dating back at least to the seminal work of Horn and Jackson in the 1970s [29], and has seen a resurgence of activity in recent years; see the monograph [23]. However, it also captures a wide range of processes that are of interest in other fields, including Boltzmann's model of an ideal gas [7], classical models of population genetics [27,54], genetic algorithms in combinatorial optimization [25,38], and random sampling [44,47].

We describe mass action kinetics in the special case where all interactions are pairwise and homogeneous; this captures most of the complexity of general systems while keeping notation and technicalities to a minimum. Let Ω denote a finite set of *types*. A (quadratic) mass action system is described by a directed graph whose vertices are ordered pairs of types (σ, σ') , and a directed edge from (σ, σ') to (τ, τ') indicates the presence of a *reaction* in which types σ, σ' combine to produce types τ, τ' . Reactions involving a specific pair (σ, σ') are governed by a *collision kernel* $\mathcal{Q}(\sigma, \sigma'; \cdot, \cdot)$, where $\mathcal{Q}(\sigma, \sigma'; \tau, \tau')$ is the probability that the outcome of the reaction is the pair (τ, τ') . We assume throughout the symmetry property $\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \mathcal{Q}(\sigma', \sigma; \tau', \tau)$.

The state of the system at any time t is fully described by the vector p_t , where $p_t(\sigma)$ is the mass of type σ at time t , normalized so that $\sum_{\sigma \in \Omega} p_t(\sigma) = 1$ (i.e., the $p_t(\sigma)$ can be viewed as concentrations, or probabilities). The initial state is denoted p_0 . According to the so-called “mass action” principle, each reaction $(\sigma, \sigma') \rightarrow (\tau, \tau')$ takes place at a rate determined by the *product* of the current masses of types σ, σ' . The dynamics of the system is described by the following set of equations,¹ one for each type $\tau \in \Omega$:

$$p_{t+1}(\tau) = \sum_{\sigma, \sigma', \tau'} p_t(\sigma) p_t(\sigma') \mathcal{Q}(\sigma, \sigma'; \tau, \tau'). \quad (1.1)$$

At this level of generality such systems can be arbitrarily badly behaved (e.g., chaotic), so it is necessary to impose standard regularity conditions. A mass action system is said to be *reversible*² or *detailed balanced* if there exists a strictly positive mass vector $\mu = (\mu(\sigma)) > 0$ such that

$$\mu(\sigma) \mu(\sigma') \mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \mu(\tau) \mu(\tau') \mathcal{Q}(\tau, \tau'; \sigma, \sigma'), \quad \forall \sigma, \sigma', \tau, \tau'. \quad (1.2)$$

It is easy to check that any such μ is necessarily an *equilibrium* or *stationary* point for the dynamics (1.1). A mass action system may have many positive equilibrium points, but it is known (see, e.g., [23], and also Proposition 2.10 in this paper) that if any one of them satisfies the detailed balance condition then they all do. We stress that we do *not* require the kernel \mathcal{Q} to be irreducible (i.e., the directed graph describing it need not be strongly connected).

The mass action system defined in (1.1) above can be viewed as a natural nonlinear analog of a reversible Markov chain, whose dynamics takes the form $p_{t+1}(\tau) = \sum_{\sigma} p_t(\sigma) \mathcal{Q}(\sigma; \tau)$, where now $\mathcal{Q}(\sigma; \tau)$ is the transition matrix of the chain and the reversibility condition is $\mu(\sigma) \mathcal{Q}(\sigma; \tau) = \mu(\tau) \mathcal{Q}(\tau; \sigma)$ for all σ, τ . In the linear setting,

¹ We give the dynamics in discrete time here; a continuous time version, in which the $\mathcal{Q}(\sigma, \sigma'; \cdot, \cdot)$ are reaction *rates* can be defined analogously in the obvious way.

² In the mass action kinetics literature, the term “reversible” has unfortunately been used to denote the weaker property that $\mathcal{Q}(\sigma, \sigma'; \tau, \tau') \neq 0$ iff $\mathcal{Q}(\tau, \tau'; \sigma, \sigma') \neq 0$, whereas in physics reversibility is synonymous with detailed balance. In this paper, we shall use the terms “detailed balanced” and “reversible” interchangeably to denote the stronger condition (1.2).

there are well known criteria for convergence to stationarity and there is by now a vast literature on mixing times of reversible Markov chains and their algorithmic applications to sampling, approximate counting and integration, statistical physics, etc. By contrast, in the nonlinear setting even the most basic questions are still open: for example, the Global Attractor Conjecture [23] asserts that any detailed balanced³ mass action system converges to a stationary point when started from any initial point p_0 with full support. (In fact many stationary points may exist, but only one is consistent with any given initial condition p_0 .) And even in particular cases of interest where convergence has been proved, almost nothing is known about the *rate* of convergence (the analog of the *mixing time* for Markov chains).

In the discrete combinatorial setting, one of very few examples for which useful bounds on the convergence rate are known is the classical Hardy-Weinberg model of genetic recombination [27,54]. Here the types are bit strings $\Omega = \{0, 1\}^n$ (each bit representing an allele on a chromosome), and a reaction between two strings σ, σ' involves picking a “crossover” subset $\Lambda \subseteq \{1, \dots, n\}$ of positions according to some probabilistic rule and exchanging the bits in Λ between σ and σ' to obtain two new strings τ, τ' . (For example, one classical rule is to pick $i \in \{0, \dots, n\}$ u.a.r. and let Λ consist of the first i bits.) It is well known that this dynamics converges to the distribution μ in which all bits are *independent*, with the marginal probabilities of a 1 at each position given by those in the initial distribution p_0 . In [11,43], the rate of convergence was related precisely to the rate at which the strings are *fragmented* by the repeated random cuts Λ , thus enabling very precise estimates of the convergence time for any choice of crossover rule.

The above analysis relies crucially on the fact that in the equilibrium distribution all bits are independent. When there is even a small amount of correlation, there appear to be no techniques available to obtain useful bounds on convergence rates. In this paper, we address this question for arguably the most natural example in which correlations arise, namely the Ising model of statistical physics. Here the types are spin configurations $\sigma \in \Omega = \{\pm 1\}^V$ which assign one of two possible spin values ± 1 to each vertex of a graph $G = (V, E)$. The *Gibbs distribution* is given by

$$\mu_{\mathbf{J}, \mathbf{h}}(\sigma) = \frac{1}{Z_{\mathbf{J}, \mathbf{h}}} \exp \left\{ \frac{1}{2} \sum_{x, y \in V} J_{xy} \sigma_x \sigma_y + \sum_{x \in V} h_x \sigma_x \right\}, \tag{1.3}$$

where $\mathbf{h} = \{h_x\}_{x \in V}$ is a vector of real numbers whose entry h_x represents the *external field* at vertex x , and $\mathbf{J} = \{J_{xy}\}_{x, y \in V}$ is a symmetric real matrix whose entry J_{xy} represents the *interaction* between spins at adjacent vertices x, y . (When there is no edge between x and y , $J_{xy} = 0$.) The normalizing factor $Z_{\mathbf{J}, \mathbf{h}}$ is the *partition function*. Note that we allow the interactions J_{xy} to be either positive (ferromagnetic) or negative (antiferromagnetic), and the fields h_x to be either positive (favoring +1 spins) or negative (favoring -1 spins). Setting $\mathbf{J} = \beta A$, where $\beta > 0$ and A is the adjacency matrix of G , corresponds to the standard ferromagnetic Ising model on G at inverse temperature β . We emphasize that, while for simplicity we develop our results for the specific case of the Ising model, they hold equally for any spin system with a constant number of different spins and bounded pairwise interactions (such as the q -state Potts model); see Sect. 5 for more detail.

³ Actually, this property is conjectured to hold under the weaker condition known as “complex balance” [23]; see also Sect. 2.6.

The classical *Glauber dynamics* for the Ising model picks a random vertex $x \in V$ at each step and resamples the spin at x according to the correct conditional distribution given its neighboring spins; this Markov chain converges to the Gibbs distribution (1.3) from any initial configuration. The analogous nonlinear mass action kinetics is defined by equation (1.1) with the following kernel: Given two configurations σ, σ' , pick a random vertex x and exchange the spins σ_x, σ'_x , obtaining two new configurations τ, τ' . The transition probabilities $\mathcal{Q}(\sigma, \sigma'; \tau, \tau')$ are chosen to satisfy the detailed balance condition (1.2), where $\mu = \mu_{\mathbf{J}, \mathbf{h}}$ is the Gibbs distribution. We emphasize that, in contrast to Glauber dynamics, here the system is evolving *endogenously* via pairwise interactions between configurations, rather than via exogenously applied spin updates. Our first result shows that this dynamics converges to the Gibbs distribution (1.3), where the fields \mathbf{h} are determined by the marginal probabilities of the spins at each vertex in the initial distribution. (The fact that the marginals determine a unique vector of fields \mathbf{h} follows from standard convexity arguments; see, e.g., [18].)

Theorem 1.1. *Let p_t denote the distribution at time t for the above mass action kinetics for the Ising model with interactions \mathbf{J} starting from any initial distribution p_0 , and let \mathbf{h} be the unique choice of external fields such that the marginal spin probabilities at each vertex $x \in V$ in $\mu_{\mathbf{J}, \mathbf{h}}$ are the same as those in p_0 . Then p_t converges to $\mu_{\mathbf{J}, \mathbf{h}}$ as $t \rightarrow \infty$.*

Note that, unlike the standard Glauber dynamics, the nonlinear dynamics has conserved quantities—namely, the marginal probabilities of the spins at each vertex—and the values of these conserved quantities determine which of the family of stationary points the dynamics converges to. This phenomenon is typical in mass action kinetics. The key to our proof of Theorem 1.1, which is based on an earlier argument in [44] (see Remark 2.12), is establishing an *irreducibility* property, namely that along any trajectory, the probability of any configuration eventually remains uniformly bounded away from zero.

We pause to briefly mention some features of mass action kinetics that make its analysis much more complex than that of Glauber dynamics, and which explain the lack of quantitative convergence results. First, as noted above, there are in general multiple equilibrium points, which are characterized by conserved quantities. Second, in contrast to the linear case, the total variation distance to stationarity is *not* monotonically decreasing (see [10, Remark 2.7]) and there are no simple coupling arguments to rely upon. Finally, the nonlinearity means that we do not have at our disposal a spectral theory and other functional analysis tools that have proved so powerful in the analysis of Markov chains. As usual in kinetic theory, a natural way to study convergence to stationarity here is to use relative entropy, which provides a monotonically decreasing functional; however, quantitative analysis of this quantity is a notoriously difficult problem in the nonlinear setting, which has so far been solved only in the non-interacting case (genetic recombination) [11, 12].

Our main result establishes tight bounds on the rate of convergence for this nonlinear dynamics in the so-called “high-temperature” regime, when the interactions are non-trivial but relatively weak. Specifically, the condition we require is that $\max_{x \in V} \sum_{y \in V} |J_{xy}| \leq \delta_0$ for some absolute constant $\delta_0 > 0$, i.e., the aggregated strength of all interactions at any given vertex is not too large. This condition mirrors the standard Dobrushin condition for Glauber dynamics, which gives a non-trivial sufficient condition for rapid mixing (see, e.g., [55]). We state this result in the following theorem.

Theorem 1.2. *In the scenario of Theorem 1.1, with the additional assumption that $\max_{x \in V} \sum_{y \in V} |J_{xy}| \leq \delta_0$ for an absolute constant $\delta_0 > 0$, the rate of convergence*

of p_t to $\mu_{\mathbf{J},\mathbf{h}}$ is given by

$$\|p_t - \mu_{\mathbf{J},\mathbf{h}}\|_{\text{TV}} \leq Cne^{-ct/n}$$

for absolute constants $C, c > 0$, where $\|\cdot\|_{\text{TV}}$ denotes total variation distance. Thus in particular the time required to achieve variation distance ε is $t = O(n \log(n/\varepsilon))$.

We note that this upper bound on convergence time is (up to constants) the same as for the analogous version of the genetic recombination model discussed above (where just a single allele is exchanged between the strings at each step) [43], and is therefore also tight by virtue of the lower bound in the same paper. That model is equivalent to the trivial case of the Ising model in which there are no interactions ($J_{xy} = 0$ for all x, y), with the h_x determined by the marginal probabilities at each allele x (and the spins ± 1 identified with the bits 1, 0). However, as we explain in Sect. 1.1 below, the correlations present in the Ising model make the analysis much more challenging.

We also consider a “block” version of the nonlinear dynamics, in which σ, σ' exchange their spins at a random subset $\Lambda \subseteq V$ of vertices (rather than just at a single randomly chosen vertex). The kernel $Q(\sigma, \sigma'; \tau, \tau')$ is again determined by the detailed balance condition (1.2), and the basic convergence result in Theorem 1.1 still holds. Under the same Dobrushin-type condition on the interactions as in Theorem 1.2, we again obtain a tight bound on the convergence rate:

Theorem 1.3. *With the same notation and assumptions as in Theorem 1.2, the variation distance of the block version of the mass action kinetics for the Ising model satisfies*

$$\|p_t - \mu_{\mathbf{J},\mathbf{h}}\|_{\text{TV}} \leq Cn^2e^{-ct}$$

for absolute constants $C, c > 0$. Thus in particular the time required to achieve variation distance ε is $t = O(\log(n/\varepsilon))$.

Note that convergence here is exponentially faster than in the single-vertex version of Theorem 1.2, reflecting the fact that this version is non-local and changes large portions of the configurations at each step. Again, the bound of Theorem 1.3 matches the lower bound for zero interaction [43].

We stress that the goal of this paper is *not* to design an efficient algorithm for sampling configurations of the Ising model. Such algorithms, based on standard linear Glauber dynamics, are already known throughout the high-temperature regime. Rather, our goal is to analyze the rate of convergence of a natural *nonlinear* dynamics, for the first time in a model with correlations. We view this as a first step towards a better understanding of such dynamics and the techniques needed to understand them; in addition to their inherent interest, these techniques may lead to algorithmic applications in future. Our work can be viewed as an extension of the successful application of a TCS lens in the analysis of mixing times of linear dynamics (Markov chains), which, as is well known, has seen both mathematical and algorithmic applications over many years. However, we should point out that our convergence analysis in Theorems 1.2 and 1.3 actually does yield polynomial time sampling algorithms based on simulation of the respective nonlinear dynamics. We outline these algorithms, together with some associated open questions, in Sect. 5 at the end of the paper.

We also point out a further interesting algorithmic aspect of our results. Recall that our nonlinear processes sample from an Ising Gibbs measure $\mu_{\mathbf{J},\mathbf{h}}$, where the fields \mathbf{h} are determined *implicitly* by the marginals at each site. It is these marginals (not the fields) that are specified by the initial distribution p_0 . (As far as we are aware, all existing

sampling algorithms for the Ising model require the specification of the fields \mathbf{h} rather than the marginals.) Thus our processes can be viewed as a novel, direct method for constructing a maximum entropy distribution subject to these marginal constraints, an important problem in its own right (see, e.g., [47]). Additionally, our processes can be used to learn the fields \mathbf{h} corresponding to given marginals, an inverse problem that is also of independent interest (see [40] for a survey of such inverse problems): given samples from $\mu_{\mathbf{J},\mathbf{h}}$ produced by the nonlinear dynamics, standard methods can be used to infer the field vector \mathbf{h} .

1.1. Techniques. We begin by describing the earlier approach of [43] to analyzing the rate of convergence of the simpler population genetics dynamics, which corresponds to the trivial case of the Ising model with no interactions ($\mathbf{J} = 0$). Since the equilibrium distribution here consists of independent bits, the analysis is relatively straightforward once one observes the following insight. The derivation of an individual σ at time t can be viewed as a binary tree going backwards in time, in which each individual inherits a random subset of its bits from each of its two parents according to the random crossover subset Λ . We may therefore follow the derivation of the n bits in σ back in time, until each of these bits is derived from a *distinct* individual at time 0. At that point we can deduce that the bits of σ are independently sampled from their respective marginal distributions, so σ is in equilibrium. The analysis therefore reduces to the question of how many steps are needed until all the n bits are separated, or “fragmented”, under the repeated action of partition by the random crossover subset Λ , which in turn is a straightforward combinatorial calculation. (See Sect. 2.4 for a more detailed description of this process.)

In the case where correlations are present in the equilibrium distribution, as in the Ising model, the above analysis breaks down because it is no longer sufficient to consider only fragmentation of the bits: indeed, the process must involve not only the breakdown of correlations in the initial distribution, but also, crucially, the *creation* of the correct equilibrium correlations as mandated by the Gibbs distribution (1.3). Moreover, the process by which an individual inherits bits from its parents is no longer independent of the parents, but dictated by a complex function of both parents.

To account for this, we appeal to the *information percolation* framework developed by Lubetzky and Sly [35] in the context of Glauber dynamics for the Ising model. This framework suggests that we keep track of a “dependence cluster” going back in time, which records the neighboring spins that have influenced each spin in our current configuration. In the linear setting of [35], it can be shown (under a similar high-temperature assumption to ours) that this cluster is dominated by a subcritical branching process and thus will die out with large probability: the equilibrium correlations are then implicitly encoded by the history of this process. The time until the process dies out gives a bound on the mixing time.

In our nonlinear setting, the dependence clusters are no longer describable in terms of a simple branching process, but rather by a new type of process that combines branching with fragmentation, a process we refer to as “fragmentation plus noise.” The first main ingredient of our analysis is the precise construction of such a process and the proof that it encodes the complex dependence structure of the nonlinear dynamics. The second main ingredient is the proof that, under the high-temperature assumption $\max_{x \in V} \sum_{y \in V} |J_{xy}| \leq \delta_0$, the fragmentation plus noise process has a subcritical behavior and therefore dies out with large probability on a suitable time scale. To establish this latter fact, we introduce a non-standard form of “high-temperature expansion” for

the dependence structure obtained by a coupling with non-uniform Erdős-Rényi random graphs. We refer the reader to Sect. 3.1 for a more technical high-level description of these ideas.

1.2. Related work. The framework of mass action kinetics was first introduced in the chemical reaction networks literature, most notably in the landmark paper of Horn and Jackson [29]. The dynamics defined in (1.1) above, generalized to allow interactions between arbitrary numbers of types, models chemical processes in an obvious way. The recent monograph of Feinberg [23] describes the state of the art in the area. While almost nothing is known about rates of convergence, much effort has been devoted to a proof of the Global Attractor Conjecture mentioned earlier; indeed, the original paper [29] contained this as a theorem, but the authors later retracted it and restated it as a conjecture [28]. Since then there have been numerous attempts at a proof, including recent papers that handle various special cases, typically based on rather severe structural conditions on the set of reactions (such as forming a single connected component)—see, e.g., [1, 2, 19, 21, 26, 41]. The Conjecture remains open.

Several other important classes of dynamics fit into the mass action framework, the most classical of which is Boltzmann's model of an ideal gas [7], where the types are momentum values of the molecules and interactions correspond to (randomized) collisions between pairs of them.⁴ In this case the dynamics (1.1) is the so-called *Boltzmann equation*, which remains a major object of study in mathematical physics today (see [51] for a survey). Another famous example is Hardy-Weinberg recombination in population genetics, as described earlier. The same dynamics can also be used to model the "recombination" step in genetic algorithms [25, 38], where two (or more) candidate solutions to a combinatorial optimization problem are combined to produce new solutions; here the mass action dynamics is typically combined with a "selection" operator that weeds out less desirable solutions. The viewpoint taken in this paper, where mass action kinetics are viewed as a nonlinear version of the Markov chain Monte Carlo method for sampling combinatorial structures from a given distribution, was first proposed in [44] and later explored in a different context in [47]. Negative results on the worst-case computational complexity of simulating mass action kinetics were derived in [4]. Mass action kinetics resembles certain more refined MCMC approaches that are used in practice for numerical simulations of spin systems such as the Ising model, including replica Monte Carlo [48], parallel tempering [31] and—most closely—cluster Monte Carlo [30]; these are in fact linear dynamics, but because they update a set of configurations in a dependent way, they can be viewed as approximate finite realizations of nonlinear dynamics.

As discussed earlier, there are almost no quantitative results on the rate of convergence of mass action kinetics in combinatorial settings, the main exception being genetic recombination [11, 12, 43]. There is a vast and still evolving literature on convergence rates of the Boltzmann equation (see, e.g., [20, 32, 37, 51]), which however is tailored to the specifics of that model and does not seem to generalize.

Finally we mention the more general class of so-called "nonlinear Markov chains", which (in discrete time) are stochastic processes (X_t) in which the distribution of X_t depends not only on the previous state X_{t-1} but also (in an arbitrarily complex way) on the *distribution* of X_{t-1} . (Mass action kinetics as discussed above is a particular example where the dependence is quadratic.) This class was formally introduced by McKean [36], who studied the continuous time version and its deep relationships to nonlinear

⁴ Here the set of types is continuous rather than discrete, as in our setting.

parabolic equations, including the Boltzmann equation, and other kinetic models. These nonlinear systems arise naturally as the limit of large finite mean-field particle systems through the so-called “propagation of chaos” [8, 13–17, 32, 33, 37, 49, 52]. Several works have been devoted to the development of nonlinear Markov chain Monte Carlo methods [3, 9, 22, 39, 45, 53], but in contrast with the classical (linear) Markov chain framework, quantitative results on convergence to stationarity are very limited and difficult to obtain.

1.3. Organization of the paper. In Sect. 2 we formally define both of our nonlinear dynamics and establish some of their basic properties, including a proof of a general convergence result, Theorem 2.8, which we then use to prove Theorem 1.1. In Sect. 3 we proceed with our quantitative analysis of the convergence rate for the nonlinear block dynamics, culminating in a proof of Theorem 1.3; prior to embarking on the details, we provide in Sect. 3.1 a more technical, high-level sketch of our approach. In Sect. 4 we apply a similar approach, though substantially different in detail, to analyze the nonlinear Glauber dynamics and prove Theorem 1.2. We conclude with some additional observations, extensions and open problems in Sect. 5.

2. Preliminaries

2.1. The Ising model. We recall from the introduction the definition (1.3) of the Ising model via its Gibbs distribution $\mu_{\mathbf{J}, \mathbf{h}}$. Note that we allow arbitrary edge-dependent interactions $\mathbf{J} = \{J_{xy}\}_{x, y \in V}$ and arbitrary external fields $\mathbf{h} = \{h_x\}_{x \in V}$. When all the external fields are zero we write simply $\mu_{\mathbf{J}}$. We identify the set of vertices (or sites) V with $[n]$, and denote the set of spin configurations by $\Omega = \{\pm 1\}^{[n]}$.

Remark 2.1. For simplicity we have taken a model with no boundary conditions. However, there is no difficulty in extending our analysis and results to the case of a Gibbs measure with arbitrary boundary conditions, i.e., when the spins in some subset $V_0 \subseteq V$ are pinned to given values ± 1 . This generalization can be easily achieved by taking limits $h_x \rightarrow \pm\infty$ for all $x \in V_0$ that are pinned to the values ± 1 , respectively.

2.2. The nonlinear dynamics. Let $\mathcal{P}(\Omega)$ denote the set of probability measures on Ω , and $\mathcal{P}_+(\Omega) \subseteq \mathcal{P}(\Omega)$ the set of measures with full support, i.e. $p \in \mathcal{P}_+(\Omega)$ iff $p(\omega) > 0$ for all $\omega \in \Omega$. We consider nonlinear (mass action) dynamics $\{p_t\}_{t \geq 0}$ on the set of types Ω as defined in (1.1) of the introduction, with some kernel \mathcal{Q} satisfying the symmetry $\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \mathcal{Q}(\sigma', \sigma; \tau', \tau)$. We shall also assume that \mathcal{Q} satisfies the reversibility condition

$$\mu(\sigma)\mu(\sigma')\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \mu(\tau)\mu(\tau')\mathcal{Q}(\tau, \tau'; \sigma, \sigma'), \quad \forall \sigma, \sigma', \tau, \tau', \quad (2.1)$$

for some $\mu \in \mathcal{P}_+(\Omega)$.

We view this dynamics as a dynamical system $p \mapsto T_t(p)$, where $T_0(p) = p \in \mathcal{P}(\Omega)$ is the initial distribution, $T_t(p) \in \mathcal{P}(\Omega)$ is the distribution after t steps, and $T_t(p) = T_{t-1}(p) \circ T_{t-1}(p)$. Here one step of the dynamics is defined, as in (1.1), by

$$p \mapsto p \circ p := \sum_{\sigma, \sigma', \tau'} p(\sigma)p(\sigma')\mathcal{Q}(\sigma, \sigma'; \tau, \tau'). \quad (2.2)$$

It will be convenient later to write (2.2) in the equivalent form

$$p \mapsto p \circ p := \sum_{\sigma, \sigma'} p(\sigma) p(\sigma') Q(\cdot | \sigma, \sigma'), \tag{2.3}$$

where, for fixed $\sigma, \sigma' \in \Omega$, the distribution $Q(\cdot | \sigma, \sigma') \in \mathcal{P}(\Omega)$ is defined by

$$Q(\tau | \sigma, \sigma') := \sum_{\tau' \in \Omega} Q(\sigma, \sigma'; \tau, \tau'). \tag{2.4}$$

In this paper we take $\mu = \mu_{\mathbf{J}, \mathbf{h}}$ as the Ising measure (1.3) and consider two natural choices of the kernel Q that satisfy (2.1), which we now describe.

2.2.1. Nonlinear block dynamics The first model, which we refer to as the *nonlinear block dynamics*, corresponds to interactions in which a pair of configurations (σ, σ') exchange their spins at an arbitrary, randomly chosen subset $\Lambda \subseteq [n]$ of sites, i.e.,

$$(\sigma, \sigma') \mapsto (\sigma'_\Lambda \sigma_{\Lambda^c}, \sigma_\Lambda \sigma'_{\Lambda^c}), \tag{2.5}$$

where $\sigma_\Lambda \sigma'_{\Lambda^c}$ denotes the element of Ω with entries σ_x for $x \in \Lambda$ and σ'_x for $x \in \Lambda^c = [n] \setminus \Lambda$. Here, to ensure reversibility, the set Λ is chosen with probability proportional to $\mu(\sigma'_\Lambda \sigma_{\Lambda^c}) \mu(\sigma_\Lambda \sigma'_{\Lambda^c})$. Thus the associated kernel is defined as

$$Q_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') = \frac{\sum_{\Lambda \subseteq V} \mu(\sigma'_\Lambda \sigma_{\Lambda^c}) \mu(\sigma_\Lambda \sigma'_{\Lambda^c}) \mathbf{1}_{\tau = \sigma'_\Lambda \sigma_{\Lambda^c}} \mathbf{1}_{\tau' = \sigma_\Lambda \sigma'_{\Lambda^c}}}{\sum_{A \subseteq V} \mu(\sigma'_A \sigma_{A^c}) \mu(\sigma_A \sigma'_{A^c})}. \tag{2.6}$$

Note that transitions of the form (2.5) can only produce pairs (τ, τ') that belong to the equivalence class

$$\mathcal{C}(\sigma, \sigma') = \{(\sigma'_\Lambda \sigma_{\Lambda^c}, \sigma_\Lambda \sigma'_{\Lambda^c}), \Lambda \subseteq V\}.$$

Thus the kernel (2.6) defines a (linear) Markov chain on the pair space $\Omega \times \Omega$ that is in general not irreducible, and whose communicating classes are precisely $\mathcal{C}(\sigma, \sigma')$. Note also that the kernel $Q_{\mathbf{J}}$ depends on $\mu = \mu_{\mathbf{J}, \mathbf{h}}$ only through the interaction \mathbf{J} and is insensitive to the choice of fields \mathbf{h} . Indeed, once σ, σ' are given, then for any $(\tau, \tau') \in \mathcal{C}(\sigma, \sigma')$ and $(\eta, \eta') \in \mathcal{C}(\sigma, \sigma')$ one has

$$\frac{\mu_{\mathbf{J}, \mathbf{h}}(\tau) \mu_{\mathbf{J}, \mathbf{h}}(\tau')}{\mu_{\mathbf{J}, \mathbf{h}}(\eta) \mu_{\mathbf{J}, \mathbf{h}}(\eta')} = \frac{\mu_{\mathbf{J}, \mathbf{h}'}(\tau) \mu_{\mathbf{J}, \mathbf{h}'}(\tau')}{\mu_{\mathbf{J}, \mathbf{h}'}(\eta) \mu_{\mathbf{J}, \mathbf{h}'}(\eta')}, \quad \forall \mathbf{h}, \mathbf{h}' \in \mathbb{R}^n.$$

Thus, w.l.o.g., we may take $\mathbf{h} = 0$, and $\mu = \mu_{\mathbf{J}} = \mu_{\mathbf{J}, \mathbf{0}}$, in the definition of the kernel (2.6). The kernel $Q_{\mathbf{J}}$ in (2.6) is an example of a so-called “folding” transformation [50].

Observe that, for all $\mathbf{h} \in \mathbb{R}^n$, the reversibility condition (2.1) holds in the form

$$\mu_{\mathbf{J}, \mathbf{h}}(\sigma) \mu_{\mathbf{J}, \mathbf{h}}(\sigma') Q_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') = \mu_{\mathbf{J}, \mathbf{h}}(\tau) \mu_{\mathbf{J}, \mathbf{h}}(\tau') Q_{\mathbf{J}}(\tau, \tau'; \sigma, \sigma'), \tag{2.7}$$

for all $\sigma, \sigma', \tau, \tau' \in \Omega$. Thus, for a fixed interaction \mathbf{J} , the kernel $Q_{\mathbf{J}}$ is reversible w.r.t. all measures $\{\mu_{\mathbf{J}, \mathbf{h}}, \mathbf{h} \in \mathbb{R}^n\}$. In particular, all these measures are stationary for the dynamics (2.2), i.e.,

$$\mu_{\mathbf{J}, \mathbf{h}} \circ \mu_{\mathbf{J}, \mathbf{h}} = \mu_{\mathbf{J}, \mathbf{h}}, \quad \forall \mathbf{h} \in \mathbb{R}^n, \tag{2.8}$$

as can be easily checked from reversibility.

We note that in the case of zero interactions, i.e., $\mathbf{J} = 0$, the nonlinear block dynamics reduces to the uniform crossover model from population genetics [12, 43]. In this case, the stationary distributions $\mu_{\mathbf{J}, \mathbf{h}}$ are just product measures over spins with marginals determined by \mathbf{h} .

2.2.2. *Nonlinear Glauber dynamics* In our second model, the configurations σ, σ' exchange their spins at a *single* randomly chosen site $x \in [n]$, i.e.,

$$(\sigma, \sigma') \mapsto (\sigma'_x \sigma_{[n] \setminus \{x\}}, \sigma_x \sigma'_{[n] \setminus \{x\}}).$$

By analogy with the familiar Glauber dynamics (a Markov chain that updates the spin at one site in each step), we refer to this as the *nonlinear Glauber dynamics*. As usual, to ensure reversibility w.r.t. $\mu_{\mathbf{J}, \mathbf{h}}$, we need to perform such an exchange with an appropriate probability $\alpha_x(\sigma, \sigma')$. Specifically, we use the generic dynamics (2.2) with the kernel

$$\mathcal{Q}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') = \frac{1}{n} \sum_{x \in V} \mathcal{Q}_{\mathbf{J}, x}(\sigma, \sigma'; \tau, \tau'), \tag{2.9}$$

where

$$\mathcal{Q}_{\mathbf{J}, x}(\sigma, \sigma'; \tau, \tau') := \alpha_x(\sigma, \sigma') \mathbf{1}_{\tau = \sigma'_x \sigma_{[n] \setminus \{x\}}} \mathbf{1}_{\tau' = \sigma_x \sigma'_{[n] \setminus \{x\}}} + (1 - \alpha_x(\sigma, \sigma')) \mathbf{1}_{\tau = \sigma} \mathbf{1}_{\tau' = \sigma'},$$

and

$$\alpha_x(\sigma, \sigma') := \frac{\mu(\sigma'_x \sigma_{[n] \setminus \{x\}}) \mu(\sigma_x \sigma'_{[n] \setminus \{x\}})}{\mu(\sigma'_x \sigma_{[n] \setminus \{x\}}) \mu(\sigma_x \sigma'_{[n] \setminus \{x\}}) + \mu(\sigma) \mu(\sigma')}.$$

Once again, the Markov chain on pairs $\Omega \times \Omega$ defined by the kernel (2.9) is not irreducible, and the kernel $\mathcal{Q}_{\mathbf{J}}$ depends on $\mu = \mu_{\mathbf{J}, \mathbf{h}}$ only through the interaction \mathbf{J} . As in (2.7)-(2.8), reversibility and stationarity of all measures $\mu_{\mathbf{J}, \mathbf{h}}$ can be easily checked.

2.3. *Conservation laws.* In both dynamics defined above, the map $p \mapsto p \circ p$ conserves the marginal probabilities of spins at every vertex, i.e., for every $x \in [n]$, and for any $p \in \mathcal{P}(\Omega)$, one has

$$(p \circ p)_x = p_x, \tag{2.10}$$

where $p_x(a) := p(\sigma_x = a), a \in \{-1, 1\}$, denotes the marginal of p at x . It is convenient to state the following stronger property. Let us define the commutative convolution product of two distributions $p, q \in \mathcal{P}(\Omega)$ by

$$p \circ q := \frac{1}{2} \sum_{\sigma, \sigma'} (p(\sigma)q(\sigma') + p(\sigma')q(\sigma)) \mathcal{Q}(\cdot | \sigma, \sigma'), \tag{2.11}$$

where \mathcal{Q} is defined by (2.4) and (2.6) for the nonlinear block dynamics and by (2.4) and (2.9) for the nonlinear Glauber dynamics, respectively. Note that the notation (2.11) is consistent with (2.3).

Lemma 2.2. *Both of the above dynamics satisfy*

$$(p \circ q)_x = \frac{1}{2} (p_x + q_x).$$

In particular, the conservation law (2.10) holds.

Proof. The proof is a consequence of the fact that, in both dynamics, spins are simply exchanged between σ, σ' , as well as the symmetry of $\mathcal{Q}_{\mathbf{J}}$. For the kernel (2.6) we compute

$$\begin{aligned} \sum_{\tau, \tau'} \mathcal{Q}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') \mathbf{1}_{\tau_x=a} &= \frac{\sum_{\Lambda \subseteq V} \mu(\sigma'_\Lambda \sigma_{\Lambda^c}) \mu(\sigma_\Lambda \sigma'_{\Lambda^c}) (\mathbf{1}_{\sigma'_x=a, \Lambda \ni x} + \mathbf{1}_{\sigma_x=a, \Lambda \not\ni x})}{\sum_{A \subseteq V} \mu(\sigma'_A \sigma_{A^c}) \mu(\sigma_A \sigma'_{A^c})} \\ &= \mathbf{1}_{\sigma_x=a} + \frac{\sum_{\Lambda \subseteq V} \mu(\sigma'_\Lambda \sigma_{\Lambda^c}) \mu(\sigma_\Lambda \sigma'_{\Lambda^c}) (\mathbf{1}_{\sigma'_x=a} - \mathbf{1}_{\sigma_x=a}) \mathbf{1}_{\Lambda \ni x}}{\sum_{A \subseteq V} \mu(\sigma'_A \sigma_{A^c}) \mu(\sigma_A \sigma'_{A^c})}. \end{aligned}$$

By symmetry,

$$\frac{1}{2} \sum_{\tau, \tau'} (\mathcal{Q}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') + \mathcal{Q}_{\mathbf{J}}(\sigma', \sigma; \tau, \tau')) \mathbf{1}_{\tau_x=a} = \frac{1}{2} (\mathbf{1}_{\sigma_x=a} + \mathbf{1}_{\sigma'_x=a}).$$

In conclusion,

$$\begin{aligned} (p \circ q)_x(a) &= \frac{1}{2} \sum_{\sigma, \sigma'} (p(\sigma)q(\sigma') + p(\sigma')q(\sigma)) \mathcal{Q}(\cdot | \sigma, \sigma') \\ &= \frac{1}{2} \sum_{\sigma, \sigma'} p(\sigma)q(\sigma') \sum_{\tau, \tau'} (\mathcal{Q}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') + \mathcal{Q}_{\mathbf{J}}(\sigma', \sigma; \tau, \tau')) \mathbf{1}_{\tau_x=a} \\ &= \frac{1}{2} \sum_{\sigma, \sigma'} p(\sigma)q(\sigma') (\mathbf{1}_{\sigma_x=a} + \mathbf{1}_{\sigma'_x=a}) = \frac{1}{2} (p_x(a) + q_x(a)). \end{aligned} \tag{2.12}$$

This proves the lemma for the nonlinear block dynamics. To prove it for the nonlinear Glauber dynamics, observe that in this case by (2.9) one has

$$\sum_{\tau, \tau'} \mathcal{Q}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') \mathbf{1}_{\tau_x=a} = \mathbf{1}_{\sigma_x=a} + \frac{1}{n} \alpha_x(\sigma, \sigma') (\mathbf{1}_{\sigma'_x=a} - \mathbf{1}_{\sigma_x=a}).$$

By the symmetry $\alpha_x(\sigma, \sigma') = \alpha_x(\sigma', \sigma)$,

$$\frac{1}{2} \sum_{\tau, \tau'} (\mathcal{Q}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') + \mathcal{Q}_{\mathbf{J}}(\sigma', \sigma; \tau, \tau')) \mathbf{1}_{\tau_x=a} = \frac{1}{2} (\mathbf{1}_{\sigma_x=a} + \mathbf{1}_{\sigma'_x=a}),$$

and the conclusion follows as in (2.12). \square

2.4. The derivation tree and fragmentation. Throughout the paper, the following view of the nonlinear dynamics will be central. By definition, $T_t(p)$ is the result of repeated pairwise interactions and can be represented as the distribution at the root of a binary “derivation” tree, where each leaf is equipped with the distribution $p \in \mathcal{P}(\Omega)$, and recursively, starting from the leaves, each internal node is assigned the distribution $p_1 \circ p_2$ where p_1, p_2 represent the distributions assigned to the left and right descendants of that node; see Fig. 1 for a schematic picture of the case $t = 2$.

We focus now on the simple case $\mathbf{J} = 0$, i.e., no correlations between spins. Under block dynamics, the configuration τ at the root of the tree (at time t) is constructed according to the random partition (Λ, Λ^c) of V , which is equivalent to drawing each spin τ_x from the configuration at the left or right child node with probability $\frac{1}{2}$, independently for each site $x \in V$. Continuing down the tree in the same fashion, we see that each spin




Fig. 1. Graphical representation of the distribution $(p_1 \circ p_2) \circ (p_3 \circ p_4)$ at time $t = 2$ when each leaf $i = 1, \dots, 4$ is equipped with distribution p_i . When $p_i = p \forall i$, the distribution at the root is $T_2(p)$

at the root is drawn from one of the 2^t leaves (at time 0), independently and uniformly at random. Thus this process induces a *partition* of the sites V into 2^t disjoint subsets (some of which may be empty), where the ℓ th subset consists of those sites that draw their spin from the configuration at leaf ℓ .

Now let \mathcal{A} denote the event that none of the 2^t subsets in this partition contains more than one site; equivalently, each spin in τ is drawn from a *distinct* leaf. We call \mathcal{A} the “complete fragmentation” event. Note that, conditional on \mathcal{A} , the distribution of the configuration τ at the root is just the product $\pi := \otimes_{x \in V} p_x$, since there are no remaining correlations between spins. Hence we may write

$$T_t(p) = \nu(\mathcal{A}) \pi + \nu(\mathcal{A}^c) q(t), \tag{2.13}$$

for some other distribution $q(t)$, where ν denotes the uniform distribution over all $2^t - 1$ independent random subsets Λ occurring in the tree.

We can use (2.13) to obtain an upper bound on the convergence time for the nonlinear block dynamics when $\mathbf{J} = 0$, as was done in [43]. First, we claim that $\nu(\mathcal{A}^c) \leq \binom{n}{2} 2^{-t}$. To see this, note that for any given pair of distinct sites $x, y \in V$, the probability that x, y are not separated after t levels of the successive partitioning process is 2^{-t} , and then take a union bound over pairs. Hence by (2.13), taking $t = O(\log(n/\varepsilon))$ ensures that $\|T_t(p) - \pi\|_{TV} \leq \varepsilon$, so the convergence time is $O(\log(n/\varepsilon))$.

For the nonlinear Glauber dynamics with $\mathbf{J} = 0$ a similar analysis applies, except that now each node in the tree chooses one random spin from the left child and the remainder from the right child. The complete fragmentation event \mathcal{A} now corresponds to isolating each of the n bits in this process, which is just a coupon-collecting event for $n = |V|$ coupons. Thus we have $\nu(\mathcal{A}^c) \leq n(1 - \frac{1}{n})^t \leq ne^{-t/n}$ (where now ν denotes the uniform distribution over all $2^t - 1$ independent choices of random spins occurring in the tree), which by (2.13) implies a convergence time of $O(n \log(n/\varepsilon))$.

When $\mathbf{J} \neq 0$, so that correlations are present, it is no longer possible to reduce the analysis of convergence to the above simple fragmentation process, because the mechanism by which a configuration inherits spins from its parents depends on the actual configurations at the parent nodes. Thus to prove Theorems 1.2 and 1.3 we will need to augment the simple derivation process above to obtain a more complex process that we call “fragmentation with noise” (see Sects. 3 and 4).

2.5. Irreducibility. Next we establish a rough lower bound on the probability of any configuration after a sufficiently long time. This observation will be key to our proof

of convergence in the next subsection. We note that the lack of such a lower bound is the principle obstacle to proving the Global Attractor Conjecture for general reversible mass action systems.

The following definition captures the desired property.

Definition 2.3. We say that an initial distribution $p \in \mathcal{P}(\Omega)$ is *irreducible* for a given mass action system with kernel \mathcal{Q} if there exist $\varepsilon > 0$ and t_0 such that $T_t(p)(\tau) \geq \varepsilon$ for all $t \geq t_0$ and all $\tau \in \Omega$.

Thus irreducibility says that the trajectory of the dynamics starting from p eventually remains bounded away from the boundary of the simplex. Note that, under this definition, irreducibility is a property of initial distributions: i.e., for a given kernel \mathcal{Q} , some initial distributions may be irreducible while others are not.

In what follows we shall assume that the initial distribution p of our dynamical system has *nondegenerate marginals*, by which we mean that there exists $\delta > 0$ such that

$$\min_{x \in V} \min_{a \in \{-1,+1\}} p_x(a) \geq \delta > 0. \tag{2.14}$$

This is actually no loss of generality since one can otherwise restrict to the nondegenerate spins and consider the degenerate spins as a fixed boundary condition, or pinning; see Remarks 2.1 and 2.11.

Lemma 2.4. *For any interaction matrix \mathbf{J} , any initial distribution $p \in \mathcal{P}(\Omega)$ with nondegenerate marginals is irreducible for both the nonlinear block dynamics and the nonlinear Glauber dynamics.*

Proof. Let us first consider the nonlinear block dynamics (2.6). We note that for a fixed interaction matrix $\mathbf{J} \in \mathbb{R}^{n \times n}$, there exists a constant $\delta_{\mathbf{J}} > 0$ such that

$$\mu_{\mathbf{J}}(\sigma)\mu_{\mathbf{J}}(\sigma') \geq \delta_{\mathbf{J}}, \quad \forall \sigma, \sigma' \in \Omega.$$

(Recall that $\mathcal{Q}_{\mathbf{J},\mathbf{h}} = \mathcal{Q}_{\mathbf{J}}$ does not depend on \mathbf{h} , so since we are discussing a property of $\mathcal{Q}_{\mathbf{J}}$ we may take $\mathbf{h} = 0$.) It follows from (2.6) that

$$\mathcal{Q}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau') = \delta_{\mathbf{J}} 2^{-n} \sum_{\Lambda \subseteq V} \mathbf{1}_{\tau = \sigma'_{\Lambda} \sigma_{\Lambda^c}} \mathbf{1}_{\tau' = \sigma_{\Lambda} \sigma'_{\Lambda^c}} + (1 - \delta_{\mathbf{J}}) \widehat{\mathcal{Q}}_{\mathbf{J}}(\sigma, \sigma'; \tau, \tau'),$$

where, for each $\sigma, \sigma', \widehat{\mathcal{Q}}_{\mathbf{J}}(\sigma, \sigma'; \cdot, \cdot) \in \mathcal{P}(\Omega \times \Omega)$ is some probability distribution that we do not need to describe explicitly. Therefore, from (2.3),

$$(p \circ p) = \delta_{\mathbf{J}} 2^{-n} \sum_{\Lambda \subseteq V} (p_{\Lambda} \otimes p_{\Lambda^c}) + (1 - \delta_{\mathbf{J}}) \widehat{\Phi}(p), \tag{2.15}$$

where $(p_{\Lambda} \otimes p_{\Lambda^c})(\tau) = p_{\Lambda}(\tau_{\Lambda})p_{\Lambda^c}(\tau_{\Lambda^c})$ denotes the product of marginals of p on Λ , Λ^c , and $\widehat{\Phi}(p) \in \mathcal{P}(\Omega)$ is some new distribution. We may interpret (2.15) as saying that the outcome of each interaction is, with probability (at least) $\delta_{\mathbf{J}}$, equal to the factorized distribution $p_{\Lambda} \otimes p_{\Lambda^c}$, where Λ is chosen u.a.r. among all subsets of V . Now note that the map $p \mapsto 2^{-n} \sum_{\Lambda \subseteq V} p_{\Lambda} \otimes p_{\Lambda^c}$ corresponds to one step of the block dynamics when $\mathbf{J} = 0$. Writing $\widetilde{T}_t(p)$ for the t -step evolution of this $\mathbf{J} = 0$ dynamics, and noting from the tree representation that the construction of $T_t(p)$ involves $N := 2^t - 1$ interactions, we may rewrite (2.15) as

$$T_t(p) = \delta_{\mathbf{J}}^N \widetilde{T}_t(p) + (1 - \delta_{\mathbf{J}}^N) \widehat{T}_t(p), \tag{2.16}$$

where $\widehat{T}_t(p) \in \mathcal{P}(\Omega)$ is again some other distribution that we will not describe.

To complete the argument, we appeal to the analysis of the $\mathbf{J} = 0$ case from the previous subsection. Specifically, we use equation (2.13) together with the analysis of the complete fragmentation event \mathcal{A} , which implies that $\nu(\mathcal{A}) \geq \frac{1}{2}$ for $t \geq t_0 = \lceil 2 \log_2 n \rceil$, to deduce that $\widetilde{T}_{t_0}(p)(\tau) \geq \frac{1}{2}\pi(\tau)$ for all $\tau \in \Omega$. Here π is the stationary distribution for the $\mathbf{J} = 0$ case, which is just the product $\pi := \otimes_{x \in V} p_x$, and the assumption (2.14) implies $\pi(\tau) \geq \delta^n$ for all τ . Plugging these observations into (2.16) gives

$$T_{t_0}(p)(\tau) \geq \frac{1}{2} \delta_{\mathbf{J}}^{N_0} \delta^n, \quad \forall \tau \in \Omega, \tag{2.17}$$

where $N_0 = 2^{t_0} - 1$. This proves the claim for $t = t_0$ with $\varepsilon = \frac{1}{2} \delta_{\mathbf{J}}^{N_0} \delta^n$. To prove it for all $t \geq t_0$, observe that if $t \geq t_0$ then $T_t(p) = T_{t_0}(T_{t-t_0}(p))$. Since by Lemma 2.2, $T_{t-t_0}(p)$ has the same marginals as p , we can apply the bound (2.14) to $T_{t-t_0}(p)$ with the same constant δ . Using (2.17) with p replaced by $T_{t-t_0}(p)$ yields $T_t(p)(\tau) \geq \varepsilon$ with the same ε for all $t \geq t_0$. This completes the proof of the lemma for the nonlinear block dynamics.

To prove it for the nonlinear Glauber dynamics we follow the same reasoning. We first observe that for some $\delta_{\mathbf{J}} > 0$ one has

$$\alpha_x(\sigma, \sigma') \geq \delta_{\mathbf{J}}, \quad \forall x \in V, \sigma, \sigma' \in \Omega$$

and therefore (2.15) now takes the form

$$p \circ p = \delta_{\mathbf{J}} \frac{1}{n} \sum_{x \in V} p_x \otimes p_{[n] \setminus \{x\}} + (1 - \delta_{\mathbf{J}}) \widehat{\Phi}(p),$$

for some new measure $\widehat{\Phi}(p) \in \mathcal{P}(\Omega)$. We may again write the expression (2.16), where now $\widetilde{T}_t(p)$ is the t -step evolution of the $\mathbf{J} = 0$ nonlinear Glauber dynamics, $p \mapsto p_x \otimes p_{[n] \setminus \{x\}}$ with $x \in V$ chosen uniformly at random. (Incidentally, it is interesting to note that, in contrast with the block dynamics, the process $\widetilde{T}_t(p)$ here is actually linear in p , since the marginals p_x are constants of the motion.)

Following the same reasoning as above, again using equation (2.13) from the previous subsection together with the analysis of the complete fragmentation event \mathcal{A} (which in this case corresponds to coupon-collecting) leads to $\nu(\mathcal{A}) \geq \frac{1}{2}$ for $t \geq t_0 = n \lceil 1 + \log n \rceil$, which in turn with the assumption (2.14) implies $\widetilde{T}_{t_0}(p)(\tau) \geq \delta^n/2$ for all $\tau \in \Omega$. It follows that

$$T_{t_0}(p)(\tau) \geq \frac{1}{2} \delta_{\mathbf{J}}^{t_0} \delta^n, \quad \forall \tau \in \Omega.$$

The desired conclusion for all $t \geq t_0$ follows as in the block dynamics case. \square

Remark 2.5. For (linear) Markov chains, an irreducibility statement as in Lemma 2.4 immediately implies exponential convergence to stationarity (at some possibly very slow rate). For the nonlinear dynamics considered here, however, it is an open question whether this holds. We will see in the next subsection that Lemma 2.4 is enough to ensure convergence, though without any information about the rate; one reason why this may be delicate is the fact that, for our nonlinear dynamics, the total variation distance to stationarity is in general not monotonically decreasing [10, Remark 2.7]. However, our main results (Theorems 1.2 and 1.3 in the introduction) do imply exponential convergence in the high-temperature regime (i.e., when the interactions in \mathbf{J} are sufficiently weak).

2.6. *Convergence to stationarity.* As we have seen, reversibility implies that, for a fixed interaction matrix \mathbf{J} , the Ising measures $\mu_{\mathbf{J},\mathbf{h}}$ defined in (1.3) are all stationary, regardless of the choice of \mathbf{h} . In fact, these are the *only* stationary distributions, as proved in [11]:

Lemma 2.6. [11, Lemma 3.2] *For both the above dynamics, for any fixed interaction matrix \mathbf{J} , a distribution $\mu \in \mathcal{P}(\Omega)$ is stationary for (2.3) if and only if μ has the form (1.3) for some choice of the fields \mathbf{h} .*

We now address the convergence result claimed in Theorem 1.1 in the introduction, which we restate here more formally.

Theorem 2.7. *Fix an interaction matrix \mathbf{J} . For any initial distribution $p \in \mathcal{P}(\Omega)$ with nondegenerate marginals, both the nonlinear block dynamics and the nonlinear Glauber dynamics satisfy the convergence*

$$\|T_t(p) - \mu_{\mathbf{J},\mathbf{h}}\|_{\text{TV}} \rightarrow 0 \quad \text{as } t \rightarrow \infty,$$

where \mathbf{h} is the unique choice of external fields such that $\mu_{\mathbf{J},\mathbf{h}}$ and p have the same marginals, i.e., $(\mu_{\mathbf{J},\mathbf{h}})_x = p_x$ for all $x \in V$.

We will prove Theorem 2.7 as a consequence of a more general convergence theorem (Theorem 2.8 below), together with the irreducibility property established in Lemma 2.4.

Our general convergence theorem is based on the following more abstract framework for mass action kinetics. Given a finite space of types Ω , consider the nonlinear mass action dynamics $T_t(p) \in \mathcal{P}(\Omega)$ defined by $T_t(p) = T_{t-1}(p) \circ T_{t-1}(p)$, $T_0(p) = p$, where the collision operator \circ is defined as in (2.2), and \mathcal{Q} is now a generic probability kernel on $\Omega \times \Omega$, i.e., $\mathcal{Q}(\sigma, \sigma'; \tau, \tau') \geq 0$ for all $\sigma, \sigma', \tau, \tau' \in \Omega$ and

$$\sum_{\tau, \tau' \in \Omega} \mathcal{Q}(\sigma, \sigma'; \tau, \tau') = 1, \quad \forall \sigma, \sigma' \in \Omega.$$

As always we assume the mild exchange symmetry property

$$\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \mathcal{Q}(\sigma', \sigma; \tau', \tau). \tag{2.18}$$

Indeed, as a consequence of (2.18), we may actually assume w.l.o.g. that \mathcal{Q} satisfies the additional symmetries

$$\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \mathcal{Q}(\sigma, \sigma'; \tau', \tau) = \mathcal{Q}(\sigma', \sigma; \tau, \tau'), \quad \forall \sigma, \sigma', \tau, \tau' \in \Omega. \tag{2.19}$$

To see this, note that because of the symmetry of the factor $p(\sigma)p(\sigma')$ in the dynamics 2.2, one can always replace \mathcal{Q} in (2.2) by the symmetrized kernel

$$\bar{\mathcal{Q}}(\sigma, \sigma'; \tau, \tau') = \frac{1}{2}(\mathcal{Q}(\sigma, \sigma'; \tau, \tau') + \mathcal{Q}(\sigma', \sigma; \tau, \tau'))$$

without altering the dynamics. Using (2.18), we immediately see that this symmetrized kernel satisfies (2.19). Accordingly, we will assume (2.19) from now on.

We also assume that the kernel has *positive diagonal elements*, i.e.,

$$\mathcal{Q}(\sigma, \sigma'; \sigma, \sigma') > 0, \quad \forall \sigma, \sigma' \in \Omega; \tag{2.20}$$

in particular, this rules out periodic behavior. However, note that, as usual, we do *not* assume that \mathcal{Q} is irreducible. We will also not assume that \mathcal{Q} is reversible (detailed

balanced), as it is in our Ising model systems; rather, it will be enough to assume the weaker property that there exists a strictly positive distribution $\mu \in \mathcal{P}_+(\Omega)$ such that

$$\sum_{\sigma, \sigma' \in \Omega} \mu(\sigma)\mu(\sigma')\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \mu(\tau)\mu(\tau'), \quad \forall \tau, \tau' \in \Omega. \quad (2.21)$$

This is equivalent to saying that the product distribution $\mu \otimes \mu \in \mathcal{P}(\Omega \times \Omega)$ is \mathcal{Q} -invariant, i.e., $(\mu \otimes \mu)\mathcal{Q} = \mu \otimes \mu$. If (2.21) holds for some positive μ , we say that \mathcal{Q} is *balanced*.⁵ Note that if \mathcal{Q} is detailed balanced (reversible), as defined in (2.1), then it is also balanced. Moreover, it is easy to check that any μ satisfying (2.21) must necessarily also be stationary for the mass-action dynamics defined by \mathcal{Q} . Indeed, for any balanced system, it will follow from the proof of Theorem 2.8 below that the converse is also true: any stationary μ must satisfy (2.21) (see Proposition 2.10).

The following general theorem says that, for any mass action kernel satisfying the above properties, irreducibility of the initial distribution (as specified in Definition 2.3) is sufficient to guarantee convergence.

Theorem 2.8. *Suppose the kernel \mathcal{Q} satisfies (2.19) and (2.20), and is also balanced. Then, for any irreducible initial distribution p , we have that $T_t(p) \rightarrow \nu$ as $t \rightarrow \infty$ for some stationary $\nu \in \mathcal{P}_+(\Omega)$. Moreover, $\nu \otimes \nu$ is \mathcal{Q} -invariant.*

Note that, in general, the limit point ν will depend on the initial distribution p .

The first ingredient in the proof of Theorem 2.8 is decay of relative entropy. Let $\mathcal{D}(\cdot\|\cdot)$ denote the relative entropy, or KL-divergence, for probability measures on Ω . I.e., $\mathcal{D}(p\|q) := \sum_{\omega \in \Omega} p(\omega) \log \frac{p(\omega)}{q(\omega)}$.

Lemma 2.9. *Suppose the kernel \mathcal{Q} satisfies (2.19) and (2.20). If \mathcal{Q} is balanced w.r.t. μ , then*

$$\mathcal{D}(p \circ p\|\mu) < \mathcal{D}(p\|\mu), \quad (2.22)$$

for any non-stationary $p \in \mathcal{P}(\Omega)$.

Proof. Define $\rho := p \otimes p$ and $\pi := \mu \otimes \mu$. The assumption that \mathcal{Q} is balanced w.r.t. μ implies that $\pi\mathcal{Q} = \pi$, while the fact that p is not stationary implies that $\rho\mathcal{Q} \neq \rho$. Now we may equivalently write the operation (2.2) in the form

$$(p \circ p)(\tau) = \sum_{\sigma, \sigma', \tau'} \rho(\sigma, \sigma')\mathcal{Q}(\sigma, \sigma'; \tau, \tau') = \sum_{\tau'} (\rho\mathcal{Q})(\tau, \tau'). \quad (2.23)$$

Equation (2.23) suggests a two-step decomposition of $p \mapsto p \circ p$. The first step is a mapping on the pair space $\Omega \times \Omega$, which takes the product distribution $\rho = p \otimes p$ to the distribution $\rho\mathcal{Q}$ (which is not typically a product). The second step is the mapping back down to Ω obtained by marginalizing out the second element τ' of the pair. We argue that each of these steps separately decreases the relative entropy w.r.t. μ , the first step yielding a strict decrease whenever p is not stationary.

For the first step, note that \mathcal{Q} defines one step of a Markov chain on $\Omega \times \Omega$. Simple convexity considerations imply the inequality $\mathcal{D}(\rho\mathcal{Q}\|\pi) \leq \mathcal{D}(\rho\|\pi)$. We claim the stronger property that

$$\mathcal{D}(\rho\mathcal{Q}\|\pi) < \mathcal{D}(\rho\|\pi) \quad (2.24)$$

⁵ In the chemical reaction networks literature, this property is known as ‘‘complex balance’’.

whenever $\rho_Q \neq \rho$. To prove this, we use the following basic fact that applies to any (not necessarily irreducible or reversible) Markov chain with a finite state space, with positive diagonal entries, and with an everywhere positive stationary distribution π . Let $\mathcal{C}_1, \dots, \mathcal{C}_k$ denote the communicating classes (irreducible components) associated with Q . Note that the assumption $\pi > 0$ implies that there is no transient state, so that the \mathcal{C}_i partition the state space $\Omega \times \Omega$. Let $F = \rho/\pi$ denote the density of ρ w.r.t. π and note that $\rho_Q = \rho$ iff F is constant within each \mathcal{C}_i . Thus, to prove (2.24) it is sufficient to show that the identity

$$\mathcal{D}(\rho_Q \parallel \pi) = \mathcal{D}(\rho \parallel \pi) \tag{2.25}$$

holds iff F is constant within each \mathcal{C}_i . Direct computation shows that

$$\mathcal{D}(\rho \parallel \pi) - \mathcal{D}(\rho_Q \parallel \pi) = \pi(F \log F) - \pi(Q^* F \log Q^* F), \tag{2.26}$$

where $Q^*(u, v) = \pi(v)Q(v, u)/\pi(u)$ is the time-reversal of Q . Clearly, if F is constant within each \mathcal{C}_i then $Q^* F = F$ and (2.25) holds. To see the converse, observe that the strict convexity of $x \mapsto x \log x$ implies that

$$H := Q^*(F \log F) - Q^* F \log Q^* F \geq 0$$

and that $H(u) = 0$ for some state u if and only if $F(v') = F(v)$ for all v, v' in the neighborhood of u in the Markov chain graph (i.e., the set of v such that $Q(u, v) > 0$). By the assumption (2.20) of positive diagonal elements for Q , this is equivalent to $F(v) = F(u)$ for all v such that $Q(u, v) > 0$. Therefore, $H = 0$ everywhere implies that F is constant within each communicating class \mathcal{C}_i . On the other hand, it follows from (2.26) and the invariance $\pi Q^* = \pi$ that (2.25) is equivalent to $\pi(H) = 0$, which by the positivity assumption on π , and the fact that $H \geq 0$, is equivalent to $H = 0$ everywhere. This ends the proof of (2.24).

For the second (marginalization) step, we appeal to the well-known fact that, among all distributions on $\Omega \times \Omega$ with fixed marginals, the relative entropy w.r.t. a product measure $\pi = \mu \otimes \mu$ is minimized when that distribution is a product distribution. Namely, if a probability $\nu \in \mathcal{P}(\Omega \times \Omega)$ has marginals ν_1, ν_2 on the first and second element respectively, one has

$$\mathcal{D}(\nu \parallel \mu_1 \otimes \mu_2) \geq \mathcal{D}(\nu_1 \parallel \mu_1) + \mathcal{D}(\nu_2 \parallel \mu_2), \tag{2.27}$$

for all probability measures $\mu_1, \mu_2 \in \mathcal{P}_+(\Omega)$. To see this, recall the variational principle for relative entropy, asserting that

$$\mathcal{D}(\nu \parallel \zeta) = \sup\{\nu(F) - \log \zeta(e^F)\},$$

where ν, ζ are arbitrary probability measures and F ranges over all real functions. Now, take $\zeta = \mu_1 \otimes \mu_2$ and $F(\sigma, \sigma') = \log f_1(\sigma) + \log f_2(\sigma')$ where $f_1 = \nu_1/\mu_1$ and $f_2 = \nu_2/\mu_2$, and observe that $\nu(F) = \mathcal{D}(\nu_1 \parallel \mu_1) + \mathcal{D}(\nu_2 \parallel \mu_2)$ while $\zeta(e^F) = 1$ because of the product structure. This proves (2.27).

Now, in view of the symmetry (2.19), the marginal of ρ_Q on the first element of the pair equals the marginal of ρ_Q on the second element of the pair, and they are both equal to $p \circ \rho$. It follows that

$$\mathcal{D}(\rho_Q \parallel \mu \otimes \mu) \geq 2\mathcal{D}(p \circ \rho \parallel \mu).$$

In conclusion, if p is not stationary we have shown that

$$\mathcal{D}(p\|\mu) = \frac{1}{2} \mathcal{D}(\rho\|\pi) > \frac{1}{2} \mathcal{D}(\rho\mathcal{Q}\|\pi) \geq \mathcal{D}(p \circ p\|\mu).$$

This completes the proof of (2.22) and hence the lemma. \square

Proof of Theorem 2.8. Let \mathcal{Q} be balanced w.r.t. $\mu \in \mathcal{P}_+(\Omega)$. Lemma 2.9 shows that $\mathcal{D}(T_t(p)\|\mu)$ is monotonically strictly decreasing with t . Hence, since $\mathcal{D}(T_t(p)\|\mu) \geq 0$, we know that $\mathcal{D}(T_t(p)\|\mu)$ converges to a limit, say d_μ . By compactness, there exists $v \in \mathcal{P}(\Omega)$ and a subsequence t_1, t_2, \dots such that $T_{t_i}(p) \rightarrow v$ as $i \rightarrow \infty$. Also, it must be the case that $\mathcal{D}(T_{t_i}(p)\|\mu) \rightarrow \mathcal{D}(v\|\mu)$, and therefore

$$d_\mu = \mathcal{D}(v\|\mu) = \lim_{t \rightarrow \infty} \mathcal{D}(T_t(p)\|\mu) = \inf_{t \in \mathbb{N}} \mathcal{D}(T_t(p)\|\mu). \tag{2.28}$$

Moreover, v must be stationary. To see this, assume for contradiction that it is not. Then, by Lemma 2.9, one step of the dynamics must strictly decrease the relative entropy and therefore, for some $\varepsilon > 0$,

$$\mathcal{D}(T_1(v)\|\mu) \leq \mathcal{D}(v\|\mu) - \varepsilon. \tag{2.29}$$

On the other hand, by continuity of the map T_1 and the function $\mathcal{D}(\cdot\|\mu)$, since $T_{t_i}(p) \rightarrow v$, there exists $t_\varepsilon \in \mathbb{N}$ sufficiently large such that

$$\mathcal{D}(T_1(T_{t_\varepsilon}(p))\|\mu) \leq \mathcal{D}(T_1(v)\|\mu) + \varepsilon/2.$$

Combining this with (2.29) gives

$$\mathcal{D}(T_{t_\varepsilon+1}(p)\|\mu) \leq \mathcal{D}(v\|\mu) - \varepsilon/2,$$

which contradicts (2.28).

Now note that, if p is irreducible, then it must be the case that $v \in \mathcal{P}_+(\Omega)$. Thus we have established that v is a stationary point with full support. Therefore, we may take $\mu = v$ in (2.28), so that $\mathcal{D}(T_t(p)\|v) \rightarrow \mathcal{D}(v\|v) = 0$ as $t \rightarrow \infty$. The latter implies, e.g., by Pinsker’s inequality, that $T_t(p) \rightarrow v$, completing the proof. \square

Before proceeding with the proof of Theorem 2.7, we pause to note the following simple consequences of the arguments given above.

Proposition 2.10. *Suppose \mathcal{Q} is a balanced kernel satisfying (2.19) and (2.20) and let $p \in \mathcal{P}(\Omega)$. Then p is stationary iff $p \otimes p$ is \mathcal{Q} -invariant. Moreover, if \mathcal{Q} is also detailed balanced, then p is stationary iff $p \otimes p$ is \mathcal{Q} -reversible (i.e., iff (1.2) holds with μ replaced by p).*

Proof. Set $\rho = p \otimes p$. The only nontrivial implications are (a) $p \circ p = p \Rightarrow \rho$ is \mathcal{Q} -invariant and (b) if \mathcal{Q} is detailed balanced, then $p \circ p = p \Rightarrow \rho$ is \mathcal{Q} -reversible. The argument in the proof of Lemma 2.9 shows that if $p \circ p = p$ then ρ is a constant multiple of $\mu \otimes \mu$ over the communicating classes of \mathcal{Q} . This implies (a). To prove (b), notice that if the detailed balance condition (1.2) holds then it continues to hold if $\mu \otimes \mu$ is replaced by any ρ that is a constant multiple of $\mu \otimes \mu$ over the communicating classes of \mathcal{Q} , since only transitions within such components are relevant in (1.2). \square

We conclude this section with a proof of our convergence result for the nonlinear Ising model dynamics, Theorem 2.7. This will follow immediately from the general criterion for convergence in Theorem 2.8, together with the irreducibility property we proved in Lemma 2.4.

Proof of Theorem 2.7. In light of Theorem 2.8, it suffices to check that both our kernels satisfy all the required properties and that any $p \in \mathcal{P}(\Omega)$ with nondegenerate marginals is irreducible. The latter fact, for both nonlinear block and Glauber dynamics, is precisely the statement of Lemma 2.4. The fact that both kernels are balanced follows from the fact that both are reversible. To check the properties (2.19) and (2.20), notice that the diagonal entries are positive for both the nonlinear block dynamics and the nonlinear Glauber dynamics. Moreover, the pair symmetry property (2.19) is easily seen to be satisfied by the block dynamics kernel (2.6). Finally, the single site kernel as written in (2.9) does not satisfy (2.19), but it does satisfy the exchange symmetry (2.18) and hence is equivalent to a symmetrized kernel \bar{Q} that satisfies (2.19), as discussed earlier. \square

Remark 2.11. We proved Lemma 2.4 and Theorem 2.7 under the assumption (2.14) of nondegenerate marginals. However, there is no difficulty in extending to the general case of an arbitrary $p \in \mathcal{P}(\Omega)$, where some spins may be deterministically set to +1 or -1. It is easy to check that the proofs of Lemmas 2.2 and 2.4 and Theorem 2.7 continue to hold in this setting, with the external fields at the pinned sites taken to be $\pm\infty$ as discussed in Remark 2.1.

Remark 2.12. Theorem 2.8 shows that, for balanced systems, convergence to a stationary point with full support is guaranteed provided one can prove that the trajectory starting from a given initial distribution p eventually remains uniformly bounded away from zero everywhere; i.e., no type “dies out”. This observation is already known in the chemical reaction networks literature [46]. However, we have provided an alternative proof here for several reasons: (i) we are working in discrete rather than continuous time as in the reaction networks community, which necessitates different arguments; (ii) we have made extensive use of probabilistic, rather than dynamical systems concepts in our proof; and (iii) we aim to make this paper self-contained. We point out also that our proof follows the same lines as that in [44, Theorem 2] for the restricted case where Q is symmetric (i.e., reversible w.r.t. the uniform distribution μ), while correcting some omissions in that earlier proof: namely, the assumption of positive diagonal elements (2.20) and the requirement that the initial condition p be irreducible. In light of Theorem 2.8 the key to the proof of our convergence result Theorem 2.7 for the nonlinear Ising model dynamics is establishing irreducibility, as we do in Lemma 2.4. With respect to progress on the Global Attractor Conjecture, the most interesting question here seems to be that of identifying minimal assumptions on the kernel Q that guarantee such an irreducibility property; we leave this question for future work.

Theorem 2.7 provides no quantitative estimate on the *rate* of convergence to stationarity. In particular, there is no explicit dependence on the size of the system n . In analogy with the mixing time analysis for linear Markov chains, in the remainder of the paper we will study the rate of convergence to equilibrium under the assumption that the interactions in \mathbf{J} are sufficiently weak (usually referred to as the “high temperature” regime).

3. The Nonlinear Block Dynamics

Let $T_t(p)$, $t \in \mathbb{N}$, denote the evolution of the initial distribution $p \in \mathcal{P}(\Omega)$ under the nonlinear block dynamics (2.6). From Theorem 2.7 we know that for any fixed interaction matrix \mathbf{J} , and any $p \in \mathcal{P}(\Omega)$, one has the convergence

$$T_t(p) \rightarrow \mu_{\mathbf{J},\mathbf{h}} \quad \text{as } t \rightarrow \infty, \tag{3.1}$$

where \mathbf{h} is the unique vector of external fields such that $\mu_{\mathbf{J},\mathbf{h}}$ and the initial state p have the same marginals at x , for all $x \in V$. Our main result for the nonlinear block dynamics (Theorem 1.3 in the introduction) establishes a tight bound on the rate of convergence in (3.1) as a function of the cardinality $n = |V|$, under the Dobrushin-type high-temperature condition on the interaction matrix \mathbf{J} . For convenience we restate this result here.

Theorem 3.1. *There exist absolute constants $\delta_0 > 0$, $c > 0$ and $C > 0$ such that, if*

$$\max_x \sum_{y \in V} |J_{xy}| \leq \delta_0, \tag{3.2}$$

then, for any $p \in \mathcal{P}(\Omega)$, $t \in \mathbb{N}$,

$$\|T_t(p) - \mu_{\mathbf{J},\mathbf{h}}\|_{\text{TV}} \leq Cn^2 e^{-ct}, \tag{3.3}$$

where \mathbf{h} is the unique choice of external fields such that $p_x = (\mu_{\mathbf{J},\mathbf{h}})_x$ for all $x \in [n]$. In particular, for any $\varepsilon > 0$, one has $\|T_t(p) - \mu_{\mathbf{J},\mathbf{h}}\|_{\text{TV}} \leq \varepsilon$ as soon as $t \geq \frac{2}{c} \log n + C_1(\varepsilon)$, where $C_1(\varepsilon) = \frac{1}{c} \log(C/\varepsilon)$.

3.1. Main ideas of the proof. Before embarking on the details of the proof, we give a high level description of the main steps. By symmetry we may rewrite the operator (2.11) in the form

$$(p \circ q)(\tau) = \sum_{\sigma, \sigma'} p(\sigma)q(\sigma') \sum_{\Lambda \subseteq V} \gamma(\Lambda | \sigma, \sigma') \mathbf{1}_{\tau = \sigma_\Lambda \sigma'_{\Lambda^c}}, \tag{3.4}$$

where

$$\gamma(\Lambda | \sigma, \sigma') = \frac{\mu(\sigma_\Lambda \sigma'_{\Lambda^c}) \mu(\sigma'_\Lambda \sigma_{\Lambda^c})}{\sum_{A \subseteq V} \mu(\sigma_A \sigma'_{A^c}) \mu(\sigma'_A \sigma_{A^c})}. \tag{3.5}$$

Thus, for each $\sigma, \sigma' \in \Omega$, $\gamma(\cdot | \sigma, \sigma')$ is a probability measure over subsets $\Lambda \subseteq V$. It will be convenient to view the distribution $\gamma(\cdot | \sigma, \sigma')$ as a spin system, i.e., a probability measure over spin configurations $\eta \in \{-1, +1\}^n$, by identifying $\eta_x = +1$ with $x \in \Lambda$ and $\eta_x = -1$ with $x \notin \Lambda$. Recall that in the non-interacting case $\mathbf{J} = 0$, the distribution $\gamma(\cdot | \sigma, \sigma')$ does not depend on the pair (σ, σ') , and is simply the product of Bernoulli measures with parameter $1/2$. As described in Sect. 2.4, the dynamics is then entirely governed by the pure fragmentation process that starts with the set V and recursively splits sets of vertices uniformly at random until it reaches a collection of singletons. The simple argument given in that proof then allows one to obtain (3.3) with $C = \frac{1}{2}$ and $c = \log 2$; see [12, 43] for a detailed analysis of the non-interacting case.

When there is a nontrivial interaction $\mathbf{J} \neq 0$, this straightforward analysis breaks down. Our proof of Theorem 3.1 is based on a coupling argument that allows us to reduce the problem to the analysis of a more general process in which the fragmentation mechanism is perturbed by a ‘‘local growth’’ process arising from the correlations inherent in the interactions. The main idea is that if the local growth is sufficiently sparse, then the underlying fragmentation dominates and eventually the memory of the initial distribution (except for the marginals) is lost.

The first step in the proof is to couple the above random variable η with distribution $\gamma(\cdot | \sigma, \sigma')$ with a random subgraph G of the complete graph K_n having a suitable distribution ν , i.e., we shall write

$$\gamma(\cdot | \sigma, \sigma') = \sum_G \nu(G) \gamma_G(\cdot | \sigma, \sigma'), \tag{3.6}$$

where the sum extends over all possible subgraphs $G \subseteq K_n$, and $\gamma_G(\cdot | \sigma, \sigma')$ is a probability measure on Ω for each realization G . The key features of this coupling are:

- the distribution ν does not depend on the pair (σ, σ') ;
- the distribution $\gamma_G(\cdot | \sigma, \sigma')$ depends on the pair (σ, σ') only through the spins

$$\sigma_{V_G} = \{\sigma_x, x \in V_G\}, \quad \sigma'_{V_G} = \{\sigma'_x, x \in V_G\},$$

where V_G denotes the vertex set of G ; and

- under $\gamma_G(\cdot | \sigma, \sigma')$, the random variables $\{\eta_y, y \in V \setminus V_G\}$ are i.i.d. Bernoulli with parameter $1/2$.

Actually, it will be crucial that ν can be taken to be the inhomogeneous Erdős-Rényi random graph with edge weights proportional to

$$\lambda_{xy} := e^{4|J_{xy}|} - 1.$$

This ensures that, under the assumption (3.2), the graph G will be sufficiently sparse and the size of the connected components will satisfy good tail bounds. Note that the expression (3.6) can be seen as a form of *high-temperature expansion* [24] for the measure $\gamma(\cdot | \sigma, \sigma')$. However, a standard high-temperature expansion would produce an expression of the form (3.6) with real-valued coefficients $\nu(G)$ which depend on (σ, σ') , while it is crucial for our coupling argument that ν be a *probability measure independent of (σ, σ')* .

Armed with the coupling (3.6), we consider all $2^t - 1$ interactions in the derivation tree of Sect. 2.4 that produce the final distribution $T_t(p)$. For each interaction we use a realization of the graph G and we specify a realization B of the Bernoulli random variables with parameter $1/2$ which determine η_y for $y \in V \setminus V_G$. We then compute the resulting distribution. Letting $(\vec{G}, \vec{B}) = (G_1, B_1), \dots, (G_{2^t-1}, B_{2^t-1})$ denote the vector of all such realizations, we may then write

$$T_t(p) = \sum_{(\vec{G}, \vec{B})} \widehat{\nu}(\vec{G}, \vec{B}) T_t(p | \vec{G}, \vec{B}), \tag{3.7}$$

where $\widehat{\nu}$ is a suitable distribution over the realizations (\vec{G}, \vec{B}) and $T_t(p | \vec{G}, \vec{B}) \in \mathcal{P}(\Omega)$ represents the distribution at time t conditional on the realizations (\vec{G}, \vec{B}) . The important point here is that $\widehat{\nu}$ is independent of the initial conditions, and therefore all correlations in the initial distribution appear only in the measures $T_t(p | \vec{G}, \vec{B})$. Moreover, the measure $\widehat{\nu}$ can naturally be interpreted as a stochastic process that combines fragmentation with local growth.

The second main ingredient in the proof of Theorem 3.1 is the identification of an event \mathcal{E}_t for this process, roughly representing the fact that within time t all fragments have reached their minimum size, and such that

$$\widehat{\nu}(\mathcal{E}_t) \geq 1 - A n^2 e^{-bt} \tag{3.8}$$

for some absolute constants $A, b > 0$. The nature of the event \mathcal{E}_t will be such that

$$T_t(p | \vec{G}, \vec{B}) = T_t(p' | \vec{G}, \vec{B}), \quad (\vec{G}, \vec{B}) \in \mathcal{E}_t \tag{3.9}$$

for all $p, p' \in \mathcal{P}(\Omega)$ which have the same marginals at every vertex $x \in V$. Once these facts are established, (3.7), (3.8) and (3.9) imply that for any such $p, p' \in \mathcal{P}(\Omega)$ one has

$$\|T_t(p) - T_t(p')\|_{\text{TV}} \leq \widehat{\nu}(\mathcal{E}_t^c) \max_{(\vec{G}, \vec{B})} \|T_t(p | \vec{G}, \vec{B}) - T_t(p' | \vec{G}, \vec{B})\|_{\text{TV}} \leq A n^2 e^{-bt}.$$

This implies the result of Theorem 3.1 by taking $p' = \mu_{\mathbf{J}, \mathbf{h}}$.

We now turn to detailed proofs of the various claims sketched above.

3.2. Coupling with inhomogeneous Erdős-Rényi random graphs. We start by observing that for every fixed $\sigma, \sigma' \in \Omega$, there is an interaction matrix $\tilde{\mathbf{J}} = \tilde{\mathbf{J}}(\sigma, \sigma')$ such that the set of spins to be exchanged, $\gamma(\cdot | \sigma, \sigma')$ from (3.5), is itself an Ising Gibbs measure $\mu_{\tilde{\mathbf{J}}}$ as defined in (1.3).

Lemma 3.2. *For any $\sigma, \sigma' \in \Omega$, we have $\gamma(\cdot | \sigma, \sigma') = \mu_{\tilde{\mathbf{J}}}$, i.e.,*

$$\gamma(\eta | \sigma, \sigma') \propto \exp \left\{ \frac{1}{2} \sum_{x, y \in V} \tilde{J}_{xy} \eta_x \eta_y \right\}, \quad \eta \in \Omega, \tag{3.10}$$

where the interaction matrix $\tilde{\mathbf{J}} = \tilde{\mathbf{J}}(\sigma, \sigma')$ is given by

$$\tilde{J}_{xy} := 2J_{xy} \sigma_x \sigma_y \mathbf{1}_{x \in D(\sigma, \sigma')} \mathbf{1}_{y \in D(\sigma, \sigma')}, \quad D(\sigma, \sigma') := \{z \in V : \sigma_z \neq \sigma'_z\}. \tag{3.11}$$

Proof. Define

$$\xi_x := \sigma_x \eta_x, \quad \eta_x := \mathbf{1}_{x \in \Lambda} - \mathbf{1}_{x \notin \Lambda}.$$

Note that

$$[\sigma_\Lambda \sigma'_{\Lambda^c}]_x = \frac{1}{2} \sigma_x (\eta_x + 1) + \frac{1}{2} \sigma'_x (1 - \eta_x),$$

and therefore

$$\begin{aligned} [\sigma_\Lambda \sigma'_{\Lambda^c}]_x [\sigma_\Lambda \sigma'_{\Lambda^c}]_y + [\sigma'_\Lambda \sigma_{\Lambda^c}]_x [\sigma'_\Lambda \sigma_{\Lambda^c}]_y &= \frac{1}{2} \eta_x \eta_y (\sigma_x - \sigma'_x) (\sigma_y - \sigma'_y) \\ &= 2\xi_x \xi_y \mathbf{1}_{x \in D(\sigma, \sigma')} \mathbf{1}_{y \in D(\sigma, \sigma')}. \end{aligned}$$

It follows that

$$\mu(\sigma_\Lambda \sigma'_{\Lambda^c}) \mu(\sigma'_\Lambda \sigma_{\Lambda^c}) = C(\sigma, \sigma') \exp \left\{ \sum_{x, y \in D(\sigma, \sigma')} J_{xy} \xi_x \xi_y \right\},$$

where the constant $C(\sigma, \sigma')$ does not depend on Λ . By definition (3.5) this concludes the proof. \square

It follows from Lemma 3.2 that for any fixed pair $\sigma, \sigma' \in \Omega$, if $\eta \in \Omega$ is distributed according to $\gamma(\cdot | \sigma, \sigma')$, then $\{\eta_x, x \in V \setminus D(\sigma, \sigma')\}$ is the Bernoulli measure with parameter $1/2$ and, independently, $\{\sigma_x \eta_x, x \in D(\sigma, \sigma')\}$ is the Ising measure on $D(\sigma, \sigma')$ with zero external fields and interaction $2\mathbf{J}$. For our purposes, the problem with this representation of $\gamma(\cdot | \sigma, \sigma')$ is that it is structurally highly dependent on the configurations σ, σ' through the set $D(\sigma, \sigma')$. Our goal in this subsection is to formulate an alternative representation, in Lemma 3.3 below, that overcomes this problem.

Let \mathcal{G} be the set of all subgraphs of the complete graph K_n over $V \subseteq [n]$ with isolated vertices removed, and write $\mathcal{P}(\mathcal{G})$ for the set of probability measures over \mathcal{G} . Thus $G \in \mathcal{G}$ can be viewed as a collection of unordered pairs $\{x, y\}$ for $x, y \in V$. Note that $G \in \mathcal{G}$ need not be connected and can be the empty graph (with no vertices). We write V_G, E_G for the vertex and edge set of $G \in \mathcal{G}$, respectively.

Lemma 3.3. *Let $\nu_{\mathbf{J}}$ be the inhomogeneous Erdős-Rényi random graph measure associated with the weights $\lambda_{xy} = e^{4|J_{xy}|} - 1$, i.e.,*

$$\nu_{\mathbf{J}}(G) \propto \prod_{\{x,y\} \in E_G} (e^{4|J_{xy}|} - 1). \tag{3.12}$$

Then

$$\gamma(\cdot | \sigma, \sigma') = \sum_{G \in \mathcal{G}} \nu_{\mathbf{J}}(G) \mu_G(\cdot | \sigma_{V_G}, \sigma'_{V_G}) \otimes \text{Be}_{V \setminus V_G}(\frac{1}{2}), \tag{3.13}$$

where, for any $G \in \mathcal{G}$, $\mu_G(\cdot | \sigma_{V_G}, \sigma'_{V_G})$ is a probability measure on $\{-1, +1\}^{V_G}$ that depends on σ, σ' only through the spins $\sigma_{V_G}, \sigma'_{V_G}$ and $\text{Be}_{V \setminus V_G}(\frac{1}{2})$ is the Bernoulli probability measure on $\{-1, +1\}^{V \setminus V_G}$, which assigns independently the values ± 1 with probability $1/2$ to each $x \in V \setminus V_G$. Moreover, the probability measure $\mu_G(\cdot | \sigma_{V_G}, \sigma'_{V_G})$ has the product structure

$$\mu_G(\cdot | \sigma_{V_G}, \sigma'_{V_G}) = \otimes_{i=1}^k \mu_{G_i}(\cdot | \sigma_{V_{G_i}}, \sigma'_{V_{G_i}}),$$

where G_1, \dots, G_k are the connected components of G .

Proof. We start with a high-temperature expansion, which is valid for any probability measure on $\Omega = \{-1, +1\}^V$ of the form

$$\gamma_\phi(\eta) = \frac{1}{Z} \exp\left\{ \sum_e \phi(e) \eta(e) \right\}, \tag{3.14}$$

where e denotes an arbitrary undirected edge $e = \{x, y\}$ for $x, y \in V$, $\eta(e) = \eta_x \eta_y$, and $\phi(e) \in \mathbb{R}$ are some given interaction coefficients. Note that (3.10) is of this form. By adding a constant independent of η to the exponent, we can rewrite this as

$$\gamma_\phi(\eta) = \frac{1}{Z} \prod_e (1 + \delta_e(\eta(e))),$$

where

$$\delta_e(\eta(e)) = \exp\{\phi(e)\eta(e) + |\phi(e)|\} - 1.$$

The point of adding $|\phi(e)|$ in the exponent is to ensure that $\delta_e(\eta(e)) \geq 0$. Moreover, expanding the product yields

$$\tilde{Z} = \sum_{\eta} \prod_e (1 + \delta_e(\eta(e))) = \sum_{\eta} \sum_{G \in \mathcal{G}} \prod_{e \in E_G} \delta_e(\eta(e)).$$

Letting G_1, \dots, G_k denote the (maximal) connected components of G , we see that

$$\tilde{Z} = \sum_{G \in \mathcal{G}} w(G), \quad w(G) := 2^{|V \setminus V_G|} \prod_{i=1}^k w_c(G_i),$$

where, for any connected graph $G \in \mathcal{G}$, we define

$$w_c(G) = \sum_{\eta_{V_G}} \prod_{e \in E_G} \delta_e(\eta(e)). \tag{3.15}$$

Since $\delta_e(\eta(e)) \geq 0$, the weights $w(G)$ are all nonnegative. Therefore, the measure (3.14) satisfies

$$\gamma_{\phi}(\eta) = \sum_{G \in \mathcal{G}} p_{\phi}(G) \hat{\mu}_G \otimes \text{Be}_{V \setminus V_G}(\tfrac{1}{2}),$$

where

$$p_{\phi}(G) := \frac{w(G)}{\sum_{G' \in \mathcal{G}} w(G')} \tag{3.16}$$

is a probability measure $p_{\phi} \in \mathcal{P}(\mathcal{G})$ depending on the interactions $\phi := \{\phi(e)\}$, and $\hat{\mu}_G$ is the probability measure on $\{-1, +1\}^{V_G}$ given by the product $\hat{\mu}_G = \otimes_{i=1}^k \nu_{G_i}$, where, for any connected graph $G \in \mathcal{G}$, we define the probability measure

$$\nu_G(\eta_{V_G}) := \frac{\prod_{e \in E_G} \delta_e(\eta(e))}{w_c(G)}$$

on $\{-1, +1\}^{V_G}$. We remark that, for any $G \in \mathcal{G}$, $\hat{\mu}_G$ depends on the interactions $\phi_G := \{\phi(e), e \in E_G\}$ only. We now apply this decomposition to the measure (3.5), which by Lemma 3.2 is (3.14) with the choice $\phi(e) = \tilde{J}_{xy}, e = \{x, y\}$, to obtain

$$\gamma(\cdot | \sigma, \sigma') = \sum_{G \in \mathcal{G}} p_{\mathbf{J}, \sigma, \sigma'}(G) \hat{\mu}_G(\cdot | \sigma_{V_G}, \sigma'_{V_G}) \otimes \text{Be}_{V \setminus V_G}(\tfrac{1}{2}), \tag{3.17}$$

for a distribution $p_{\mathbf{J}, \sigma, \sigma'} \in \mathcal{P}(\mathcal{G})$ that depends on the interactions $\mathbf{J} = \{J_{xy}\}$ and on the configurations σ, σ' through the set $D(\sigma, \sigma')$; see (3.11), (3.15) and (3.16). We note that, because of the dependance on σ, σ' , this is not sufficient to prove the desired claim (3.13). We shall use a further coupling argument to lift the decomposition (3.17) to the decomposition (3.13) with the desired properties.

Observe that the Erdős-Rényi measure $\nu_{\mathbf{J}}$ defined by (3.12) can be rewritten as in (3.16) if we redefine the weights $w(G)$ as

$$\bar{w}(G) = \prod_{e \in E_G} \bar{\delta}_e, \quad \bar{\delta}_e := (e^{4|J_{xy}|} - 1), \quad e = \{x, y\}.$$

We are going to show that for all (σ, σ') , the measure $p_{\mathbf{J}, \sigma, \sigma'}$ is stochastically dominated by $\nu_{\mathbf{J}}$, or equivalently that there exists a coupling $\pi \in \mathcal{P}(\mathcal{G} \times \mathcal{G})$ of $\nu_{\mathbf{J}}$ and $p_{\mathbf{J}, \sigma, \sigma'}$ such that for all $G, H \in \mathcal{G}$,

$$\begin{aligned} \sum_{G' \in \mathcal{G}} \pi(G', H) &= \nu_{\mathbf{J}}(H), & \sum_{H' \in \mathcal{G}} \pi(G, H') &= p_{\mathbf{J}, \sigma, \sigma'}(G), & \text{and} \\ \sum_{G, H \in \mathcal{G}: G \subseteq H} \pi(G, H) &= 1. \end{aligned} \tag{3.18}$$

In particular, it follows that the conditional distribution $\pi(\cdot | H)$ of $G \in \mathcal{G}$ is supported on graphs $G \subseteq H$. Moreover, we will also show that for any $H \in \mathcal{G}$,

- $\pi(\cdot | H)$ depends on the spin configurations σ, σ' only through their values $\sigma_{V_H}, \sigma'_{V_H}$ on V_H ,
- $\pi(\cdot | H)$ has the product structure $\pi(\cdot | H) = \otimes_{i=1}^k \pi(\cdot | H_i)$ where H_1, \dots, H_k denote the connected components of H .

Once these facts are established we can quickly conclude the proof of Lemma 3.3 as follows. Let π be the coupling of $p_{\mathbf{J}, \sigma, \sigma'}$ and $\nu_{\mathbf{J}}$ as above, so that

$$p_{\mathbf{J}, \sigma, \sigma'}(G) = \sum_{H \in \mathcal{G}: G \subseteq H} \nu_{\mathbf{J}}(H) \pi(G | H).$$

From (3.17) we obtain

$$\begin{aligned} \gamma(\cdot | \sigma, \sigma') &= \sum_{H \in \mathcal{G}} \nu_{\mathbf{J}}(H) \sum_{G \in \mathcal{G}: G \subseteq H} \pi(G | H) \hat{\mu}_G(\cdot | \sigma_{V_G}, \sigma'_{V_G}) \otimes \text{Be}_{V \setminus V_G}(\tfrac{1}{2}) \\ &= \sum_{H \in \mathcal{G}} \nu_{\mathbf{J}}(H) \mu_H(\cdot | \sigma_{V_H}, \sigma'_{V_H}) \otimes \text{Be}_{V \setminus V_H}(\tfrac{1}{2}), \end{aligned} \tag{3.19}$$

where we define

$$\mu_H(\cdot | \sigma_{V_H}, \sigma'_{V_H}) := \sum_{G \in \mathcal{G}: G \subseteq H} \pi(G | H) \hat{\mu}_G(\cdot | \sigma_{V_G}, \sigma'_{V_G}) \otimes \text{Be}_{V_H \setminus V_G}(\tfrac{1}{2}). \tag{3.20}$$

Since $\pi(\cdot | H)$ depends only on $\sigma_{V_H}, \sigma'_{V_H}$, the measure in (3.20) also depends on σ, σ' only through σ_H, σ'_H . Moreover, because of the product structure of $\hat{\mu}_G(\cdot | \sigma_{V_G}, \sigma'_{V_G})$ and $\pi(\cdot | H)$, the measure $\mu_H(\cdot | \sigma_{V_H}, \sigma'_{V_H})$ defined by (3.20) must necessarily also have the desired product structure

$$\mu_H(\cdot | \sigma_{V_H}, \sigma'_{V_H}) = \otimes_{i=1}^k \mu_{H_i}(\cdot | \sigma_{V_{H_i}}, \sigma'_{V_{H_i}})$$

along the connected components H_1, \dots, H_k of H . Therefore, the decomposition (3.19) concludes the proof of the lemma once we establish the desired properties of the coupling π .

To construct the coupling π we use the following coupled Glauber-type Markov chains. Let U_1, U_2, \dots denote i.i.d. uniform random variables in $[0, 1]$ and let e_1, e_2, \dots denote i.i.d. uniformly random edges $e = \{x, y\}$ taken among all $\binom{n}{2}$ possible choices.

For any $G \subseteq H \in \mathcal{G}$, any fixed $e = \{x, y\}$, let $G^{e,+}$ denote the graph $G \cup \{e\}$, and let $G^{e,-}$ denote the graph $G \setminus \{e\}$. We define

$$p(G, e, \pm) = \frac{w(G^{e,\pm})}{w(G^{e,-}) + w(G^{e,+})}, \quad \bar{p}(G, e, \pm) = \frac{\bar{w}(G^{e,\pm})}{\bar{w}(G^{e,-}) + \bar{w}(G^{e,+})}$$

For any fixed $G_0, H_0 \in \mathcal{G}$, we write $(X_t, Y_t), t = 0, 1, \dots$, for the Markov chain with $X_0 = G_0, Y_0 = H_0$, and such that for any $t \geq 1$,

1. if $U_t \leq p(X_{t-1}, e_t, +)$ then $X_t = X_{t-1} \cup \{e\}$, otherwise $X_t = X_{t-1} \setminus \{e\}$;
2. if $U_t \leq \bar{p}(Y_{t-1}, e_t, +)$ then $Y_t = Y_{t-1} \cup \{e\}$, otherwise $Y_t = Y_{t-1} \setminus \{e\}$.

In words, at each step a uniformly random edge e_t is chosen, and the graphs X_{t-1}, Y_{t-1} are updated by adding or removing the edge e_t according to the prescribed probabilities and the common source of randomness U_t . This defines the Markov chain with state space $\mathcal{G} \times \mathcal{G}$. By construction, we have the reversibility conditions

$$p_{\mathbf{J},\sigma,\sigma'}(G \setminus \{e\}) p(G, e, +) = p_{\mathbf{J},\sigma,\sigma'}(G \cup \{e\}) p(G, e, -);$$

$$\nu_{\mathbf{J}}(H \setminus \{e\}) \bar{p}(H, e, +) = \nu_{\mathbf{J}}(H \cup \{e\}) \bar{p}(H, e, -),$$

for any $G, H \in \mathcal{G}$. Therefore, the marginals X_t and $Y_t, t = 0, 1, \dots$, are Markov chains with state space \mathcal{G} with stationary distributions $p_{\mathbf{J},\sigma,\sigma'}$ and $\nu_{\mathbf{J}}$ respectively.

Let us now observe that the above coupled process preserves ordering, in the sense that if $G_0 \subseteq H_0 \in \mathcal{G}$ then $X_t \subseteq Y_t$ for all $t = 1, 2, \dots$. In view of our construction, to prove this it suffices to note that, for any $G \subseteq H$ and any e ,

$$p(G, e, +) \leq \bar{p}(H, e, +). \tag{3.21}$$

For $e = \{x, y\}$, one has

$$p(G, e, +) = \frac{\sum_{\eta_x, \eta_y} \delta_e(\eta(e)) \sum_{\{\eta_z, z \neq x, y\}} \prod_{e' \in G: e' \neq e} \delta_{e'}(\eta(e'))}{\sum_{\eta_x, \eta_y} [1 + \delta_e(\eta(e))] \sum_{\{\eta_z, z \neq x, y\}} \prod_{e' \in G: e' \neq e} \delta_{e'}(\eta(e'))}. \tag{3.22}$$

This is an increasing function of the nonnegative variable $\delta_e(\eta(e))$ and thus, using

$$\delta_e(\eta(e)) = \exp\{\phi(e)\eta(e) + |\phi(e)|\} - 1 \leq \bar{\delta}_e = e^{4|J_{xy}|} - 1,$$

where $e = \{x, y\}, \eta(e) = \eta_x \eta_y$, and $\phi(e) = \tilde{J}_{xy}$, one has $p(G, e, +) \leq \hat{p}(G, e, +)$ where the latter is defined as in (3.22) with $\delta_e(\eta(e))$ replaced by $\bar{\delta}_e$. On the other hand it is not hard to see that

$$\hat{p}(G, e, +) = \frac{e^{4|J_{xy}|} - 1}{e^{4|J_{xy}|}} = \bar{p}(H, e, +),$$

for any $G, H \in \mathcal{G}$. This proves the desired monotonicity (3.21).

The monotone coupling π of $p_{\mathbf{J},\sigma,\sigma'}$ and $\nu_{\mathbf{J}}$ may then be defined as the stationary distribution of the Markov chain (X_t, Y_t) . The required property (3.18) is a consequence of the above construction and the inequality (3.21), since if we start the Markov chain at (X_0, Y_0) with $X_0 \subseteq Y_0$ we will have $X_t \subseteq Y_t$ at all times and thus the limiting distribution π as $t \rightarrow \infty$ must be supported on ordered pairs.

Finally, we need to check that for any $H \in \mathcal{G}$, the conditional distribution $\pi(\cdot | H)$ of $G \subseteq H$ depends on the spin configurations σ, σ' only through their values $\sigma_{V_H}, \sigma'_{V_H}$ on V_H . To this end, observe that $\pi(\cdot | H)$ is the stationary distribution of the Markov chain

(X_t, Y_t) obtained by conditioning on $Y_t = H$ for all t , or equivalently the Markov chain defined as above with the restriction that any transition that would result in adding or removing an edge from the graph $Y_t = H$ is rejected. This yields the Markov chain X_t^H , $t = 0, 1, \dots$ with state space $\mathcal{G}_H = \{G \in \mathcal{G} : G \subseteq H\}$ and transition probabilities given by (3.22) with the restriction that adding an edge e is only allowed if $G \cup \{e\} \subseteq H$. Since the expression (3.22) factorizes over connected components of G , and since $G \subseteq H$ it follows that (3.22) only depends on the coefficients $\phi(e) = \tilde{J}_{xy}$, where $e = \{x, y\} \in H$, and therefore the stationary distribution $\pi(\cdot | H)$ has the same property. This ends the proof of Lemma 3.3. \square

So far we have not used the weak interaction assumption (3.2), so Lemma 3.3 holds for arbitrary coefficients J_{xy} . Next, we observe that the condition (3.2) implies a strong sparsity property of the measure $\nu_{\mathbf{J}}$ in Lemma 3.3. From the definition (3.12), this measure is the inhomogeneous Erdős-Rényi random graph where each edge $e = \{x, y\}$, $x, y \in V$, is included independently with probability

$$p_{xy} = 1 - e^{-4|J_{xy}|}. \tag{3.23}$$

In particular, using $(e^z - 1)e^{-z} \leq z$ for $z \geq 0$, one has $p_{xy} \leq 4|J_{xy}|$, and if (3.2) holds, then for any $x \in V$,

$$\nu_{\mathbf{J}}(V_G \ni x) \leq \sum_{y: y \neq x} p_{xy} \leq 4\delta_0.$$

Moreover, if δ_0 is sufficiently small, then for any given $x \in V$ the size of the connected component of G at x has an exponential tail. We make this precise in Lemma 3.8 below.

3.3. Fragmentation with noise. We now develop the main construction behind the convergence result in Theorem 3.1. It is based on a *perturbed* fragmentation process, i.e., a process that combines the random fragmentations of the non-interacting case (as described in Sect. 2.4) with some competing noise represented by the random graphs encoding dependencies.

Given a set $A \subseteq [n]$ and a random graph $G \in \mathcal{G}$, we define the random set A' as the vertex set of the union of all connected components of G that have non-empty intersection with A . More formally, write $G = \cup_{i=1}^{\ell} G_i$ where G_i are the connected components of G and let

$$G(A) = \bigcup_{i: V_{G_i} \cap A \neq \emptyset} G_i.$$

Then we set $A' = V_{G(A)}$. We may sample A' starting from A by a breadth-first search, i.e., by revealing sequentially for each $x \in A$ the neighborhood of x in G , then recursively the neighborhood of each vertex revealed at the previous step, and so on until there are no more neighbors to reveal. Clearly, A' may contain sites that are not in A . However, if $x \in A \cap V_G^c$ then $x \notin A'$. If $A \setminus A' \neq \emptyset$, for every $x \in A \setminus A'$, we independently declare x to be *in* or *out* by a fair coin flip. We thus obtain two random sets A_{in} and A_{out} , such that

$$A_{\text{in}} \cup A_{\text{out}} = A \setminus A', \quad A_{\text{in}} \cap A_{\text{out}} = \emptyset.$$

Next, for any $A \subseteq [n]$, we define two sets $\Phi_0(A), \Phi_1(A)$ by

$$\Phi_0(A) = \Phi_1(A) = \emptyset, \quad \text{if } |A| \leq 1, \tag{3.24}$$

and

$$\Phi_0(A) = A' \cup A_{\text{in}}, \quad \Phi_1(A) = A' \cup A_{\text{out}}, \quad \text{if } |A| \geq 2. \tag{3.25}$$

Definition 3.4. The *fragmentation plus noise process* $\mathcal{F}_t, t = 0, 1, \dots$ is the random process defined as follows. For each $t \in \mathbb{N}$, \mathcal{F}_t consists of 2^t labeled fragments, i.e., (possibly empty) subsets $F_1^{(t)}, \dots, F_{2^t}^{(t)}, F_i^{(t)} \subseteq [n]$, obtained by repeated application of the following rule. At time zero we have $\mathcal{F}_0 = [n]$, i.e., $F_1^{(0)} = [n]$. At time $t \in \mathbb{N}$, if $\mathcal{F}_{t-1} = (F_1^{(t-1)}, \dots, F_{2^{t-1}}^{(t-1)})$, then for each i independently, we replace $F_i^{(t-1)}$ by $(\Phi_0(F_i^{(t-1)}), \Phi_1(F_i^{(t-1)}))$ where Φ_0, Φ_1 are the random maps defined by (3.24)–(3.25), so that

$$\begin{aligned} \mathcal{F}_t &= (F_1^{(t)}, \dots, F_{2^t}^{(t)}) \\ &= (\Phi_0(F_1^{(t-1)}), \Phi_1(F_1^{(t-1)}), \Phi_0(F_2^{(t-1)}), \Phi_1(F_2^{(t-1)}), \dots, \Phi_0(F_{2^{t-1}}^{(t-1)}), \Phi_1(F_{2^{t-1}}^{(t-1)})). \end{aligned}$$

We say that the process dies out if there is a time t such that all fragments are empty, i.e., $F_i^{(t)} = \emptyset$ for all $i = 1, \dots, 2^t$. With slight abuse of notation, we write $\mathcal{F}_t = \emptyset$ for the latter event.

In the non-interacting case $\mathbf{J} = 0$ one has $p_{xy} = 0$ for all $\{x, y\}$, and thus $A' = \emptyset, \Phi_0(A) = A_{\text{in}}, \Phi_1(A) = A_{\text{out}}$ and $A \mapsto (\Phi_0(A), \Phi_1(A))$ is one step of a pure fragmentation process, where the set A is partitioned into two subsets using independent fair coin flips for each vertex. In this case $F_i^{(t)} \cap F_j^{(t)} = \emptyset$ for all $i, j = 1, \dots, 2^t$ and for all $t \in \mathbb{N}$. In particular, it is not hard to see that in this case

$$\mathbb{P}(\mathcal{F}_t \neq \emptyset) \leq n(n-1)2^{-t}, \quad t = 1, 2, \dots \tag{3.26}$$

Indeed, by construction $\mathcal{F}_t \neq \emptyset$ implies that there exist two vertices $x, y \in [n]$ that belong to the same fragment up to time $t - 1$. Thus (3.26) follows as in the analysis of Sect. 2.4. In the interacting case, there is a first stage where the set A grows according to the local branching at every $x \in A$, and the fragmentation occurs only on those vertices that have an empty neighborhood in G . Our main technical result below establishes that the fragmentation plus noise process also satisfies a bound of the form (3.26), with a slightly weaker exponential decay rate, provided (3.2) holds for a suitably small $\delta_0 > 0$.

Lemma 3.5. *For any $\delta \in (0, 1)$, there exists $\delta_0 > 0$, and a constant $C_\delta > 0$ such that if (3.2) holds with constant δ_0 then*

$$\mathbb{P}(\mathcal{F}_t \neq \emptyset) \leq n^2 C_\delta (2 - \delta)^{-t}, \quad t = 1, 2, \dots$$

The proof of this lemma is quite technical and is postponed to Sect. 3.5. We turn first to the proof of Theorem 3.1.

3.4. *Proof of Theorem 3.1.* We now have the tools to conclude the proof of Theorem 3.1, our main result for nonlinear block dynamics. Recall again the construction of $T_t(p)$ in terms of the binary derivation tree in Sect. 2.4.

By the invariance property (2.8), the target measure $\mu_{\mathbf{J}, \mathbf{h}}$ can be obtained at the root by taking the distribution $\mu_{\mathbf{J}, \mathbf{h}}$ on each leaf. Thus, Theorem 3.1 says that when each leaf is given a distribution p with the same marginals as $\mu_{\mathbf{J}, \mathbf{h}}$, the two distributions $T_t(p)$ and $\mu_{\mathbf{J}, \mathbf{h}}$ can be coupled with an error at most Cn^2e^{-ct} for any t . We shall actually prove the following stronger result. Let $\vec{p} = (p_1, \dots, p_{2^t})$ and $\vec{q} = (q_1, \dots, q_{2^t})$ be arbitrary vectors of distributions in $\mathcal{P}(\Omega)$ whose marginals on σ_x at all sites x agree, i.e., $p_i, q_i \in \mathcal{P}(\Omega)$ satisfy

$$(p_i)_x = (p_j)_x = (q_i)_x = (q_j)_x, \quad i, j = 1, \dots, 2^t, \quad x \in V. \tag{3.27}$$

Let $T_t(\vec{p})$ (resp., $T_t(\vec{q})$) denote the distribution at the root of the binary tree of depth t , where the leaf labeled $i = 1, \dots, 2^t$ is equipped with the distribution p_i (resp., q_i); recall Fig. 1 for the case $t = 2$.

Theorem 3.6. *There exist absolute constants $\delta_0 > 0$, $c > 0$ and $C > 0$ such that, if (3.2) holds with constant δ_0 then for any choice of initial distributions \vec{p}, \vec{q} as in (3.27),*

$$\|T_t(\vec{p}) - T_t(\vec{q})\|_{\text{TV}} \leq Cn^2e^{-ct}, \quad t \in \mathbb{N}. \tag{3.28}$$

Clearly, Theorem 3.6 implies Theorem 3.1 since we may take $p_i \equiv p \in \mathcal{P}(\Omega)$ and $q_i \equiv \mu_{\mathbf{J}, \mathbf{h}}$, where the external fields \mathbf{h} are chosen in such a way that p and $\mu_{\mathbf{J}, \mathbf{h}}$ have the same marginals.

We shall prove Theorem 3.6 by analyzing the interaction history backwards in time, i.e., from the root to the leaves. This is reminiscent of the coupling from the past approach for linear Markov chains [34, 42], and to some extent our proof is inspired by ideas that have been developed in that context. In particular, our proof is related to the information percolation framework developed by Lubetzky and Sly in [35].

Each internal node of the tree is associated with an interaction, or collision, which according to (3.4) is specified by the random set Λ with distribution $\gamma(\cdot | \sigma, \sigma')$. For each such interaction we reveal the realization of the graph G and of the Bernoulli variables B in $V \setminus V_G$ that are used in sampling Λ ; see Lemma 3.3. In this way, starting from the root, we have a pair (G, B) , $G \in \mathcal{G}$ and a subset $B \subseteq V \setminus V_G$ is identified with the set of $x \in V \setminus V_G$ for which $\eta_x = +1$. Suppose the descendants of the root have distributions p and q respectively, as in Fig. 2. Then, according to Lemma 3.3, the distribution at the root is given by

$$p \circ q = \sum_{(G, B)} \nu(G, B) T(p, q | G, B),$$

where $\nu(G, B) = 2^{-|V \setminus V_G|} \nu_{\mathbf{J}}(G)$, and for each realization (G, B) , $T(p, q | G, B) \in \mathcal{P}(\Omega)$ is the distribution

$$T(p, q | G, B)(\tau) = \sum_{\sigma, \sigma'} p(\sigma)q(\sigma') \sum_{\eta_{V_G}} \mu_G(\eta_{V_G} | \sigma_{V_G}, \sigma'_{V_G}) \mathbf{1}_{\tau \sim (\sigma, \sigma', \eta_{V_G}, B)}. \tag{3.29}$$

Here, for $\tau \in \Omega$, the notation $\tau \sim (\sigma, \sigma', \eta_{V_G}, B)$ is shorthand for

$$\tau \sim (\sigma, \sigma', \eta_{V_G}, B) \Leftrightarrow \begin{cases} \tau_x = \sigma_x, & x \in V : \eta_x = +1; \\ \tau_x = \sigma'_x, & x \in V : \eta_x = -1, \end{cases} \tag{3.30}$$




Fig. 2. Graphical representation of the distribution at the root, when $t = 2$ and the leaves are equipped with distributions p', q', p'', q'' . Each internal node is equipped with a realization of the random pair (G, B) , where $G \in \mathcal{G}$ and $B \subseteq V \setminus V_G$

with the understanding that the value of η on $V \setminus V_G$ is specified by $B \subseteq V \setminus V_G$, i.e., $\eta_y = +1$ for $y \in B$ and $\eta_y = -1$ for $y \notin B$.

For every $G \in \mathcal{G}$, $\mu_G(\cdot | \sigma_{V_G}, \sigma'_{V_G})$ depends on (σ, σ') only through $(\sigma_{V_G}, \sigma'_{V_G})$, and for every $\eta_{V_G} \in \{-1, 1\}^{V_G}$, $\tau \in \{-1, 1\}^n$ and $B \subseteq V \setminus V_G$, the condition (3.30) depends on (σ, σ') only through

$$\{\sigma_x, x \in V_G \cup B\}, \quad \{\sigma'_x, x \in V_G \cup ((V \setminus V_G) \setminus B)\}.$$

Therefore, to compute the distribution $T(p, q | G, B)$ we only need the marginals $p_{V_G \cup B}$ and $q_{V_G \cup B'}$, where $B' = (V \setminus V_G) \setminus B$. Note that we may identify $V_G \cup B$ with the set $\Phi_0([n])$ and $V_G \cup B'$ with the set $\Phi_1([n])$, where Φ_0, Φ_1 are the maps defined in (3.25). Indeed, by definition of the measure $\nu_{\mathbf{J}}$, these random sets have the same distribution since, when $A = [n]$, V_G is equivalent to A' and B is equivalent to A_{in} . One way to rephrase this is to say that, as far as the distribution $T(p, q | G, B)$ is concerned, the only relevant information about the distribution p is contained in the set $\Phi_0([n])$ and the only relevant information about the distribution q is contained in the set $\Phi_1([n])$.

Next, we move one step backwards in time and consider the interaction that produced the distribution p from the previous computation. Suppose that p', q' are the distributions at the two descendants of p respectively, so that $p = p' \circ q'$, as in Fig. 2. Suppose we revealed the realization (G', B') of the graph and Bernoulli variables associated with this interaction. Note that we can use the same expressions (3.29)–(3.30) to compute $T(p', q' | G', B')$, provided we replace (p, q) by (p', q') and (G, B) by (G', B') . However, the key point is that we now only need the marginal of p on the set $\Phi_0([n])$, and therefore when we compute $T(p', q' | G', B')(\tau)$ as above we can sum away all τ_y , $y \notin \Phi_0([n])$, so that the indicator function (3.30) is relevant only for sites $x \in \Phi_0([n])$, and the distribution of $\eta_x, x \in V_{G'} \cap \Phi_0([n])$, under $\mu_{G'}(\cdot | \sigma_{V_{G'}}, \sigma'_{V_{G'}})$ is only influenced by the spins σ_y, σ'_y for $y \in G'(\Phi_0([n]))$, where we recall that $G'(A)$ is the union of all connected components of G' that have nonempty intersection with the set A . The latter property is a consequence of the product structure of the measure $\mu_{G'}(\cdot | \sigma_{V_{G'}}, \sigma'_{V_{G'}})$; see Lemma 3.3. Hence we can neglect all connected components of G' that do not intersect $\Phi_0([n])$, and we can discard the information about all Bernoulli variables at sites $y \in V \setminus V_{G'}$ such that $y \notin \Phi_0([n])$. A close inspection of our definition of the maps Φ_0 and Φ_1 then reveals that the only information about the distributions p', q' that is needed to compute $T(p', q' | G', B')$ is contained in the marginals $p'_{\Phi_0(\Phi_0([n]))}$ and $q'_{\Phi_1(\Phi_0([n]))}$.

Similarly, considering the interaction which produced the distribution $q = p'' \circ q''$ (see Fig. 2), we may fix a realization (G'', B'') of the graph and Bernoulli variables,

and repeating the above reasoning one has that the only information about the distributions p'', q'' that is needed to compute $T(p'', q'' | G'', B'')$ is contained in the marginals $p''_{\Phi_0(\Phi_1([n]))}$ and $q''_{\Phi_1(\Phi_1([n]))}$. We note that, after two steps of the evolution, conditional on the realizations of the variables $(G, B), (G', B'), (G'', B'')$, we have obtained a probability measure at the root depending only on $(G, B), (G', B'), (G'', B'')$ and on the marginals

$$p'_{\Phi_0(\Phi_0([n]))}, q'_{\Phi_1(\Phi_0([n]))}, p''_{\Phi_0(\Phi_1([n]))}, q''_{\Phi_1(\Phi_1([n]))}.$$

Equivalently, using the notation introduced in Definition 3.4, after two steps we have that the correlations of the initial distributions at the leaves are entirely contained in the fragments

$$\begin{aligned} \mathcal{F}_2 &= (F_1^{(2)}, F_2^{(2)}, F_3^{(2)}, F_4^{(2)}) \\ &= (\Phi_0(\Phi_0([n])), \Phi_1(\Phi_0([n])), \Phi_0(\Phi_1([n])), \Phi_1(\Phi_1([n]))). \end{aligned}$$

Example 3.7. Consider the following simple example with $n = 4$. Suppose $G = \{\{1, 2\}\}$, i.e., G consists of the single edge $\{1, 2\}$ and suppose that $B = \emptyset$, i.e., both 3 and 4 are *out*. This gives $F_1^{(1)} = \Phi_0([n]) = \{1, 2\}$, $F_2^{(1)} = \Phi_1([n]) = \{1, 2, 3, 4\}$. Suppose also that $G' = \{\{3, 4\}\}$, and $B' = \{1\}$. Thus $F_1^{(2)} = \{1\}$, and $F_2^{(2)} = \{2\}$. Suppose finally that $G'' = \emptyset$, $B'' = \{1, 2\}$. Then $F_3^{(2)} = \{1, 2\}$, and $F_4^{(2)} = \{3, 4\}$. Thus, one has the following fragmentation after two steps

$$\begin{aligned} \mathcal{F}_0 &= \{1, 2, 3, 4\}, \\ \mathcal{F}_1 &= (\{1, 2\}, \{1, 2, 3, 4\}), \\ \mathcal{F}_2 &= (\{1\}, \{2\}, \{1, 2\}, \{3, 4\}). \end{aligned}$$

Note that in this example, conditionally on the given realizations of the variables $(G, B), (G', B')$, and (G'', B'') , the distribution $p \circ q$ at the root in Fig. 2 can be computed only using the marginals

$$(p')_{\{1\}}, (q')_{\{2\}}, (p'')_{\{1,2\}}, (q'')_{\{3,4\}}$$

of the input distributions p', q', p'', q'' .

Repeating the above procedure one has that, after t steps, conditional on the realizations of all graphs and Bernoulli variables involved in each of the $2^t - 1$ interactions, the only information needed from leaf i is contained in the marginal of the distribution at that leaf on the fragment $F_i^{(t)}$, for each $i = 1, \dots, 2^t$, as defined in Definition 3.4. The crucial observation is that, as soon as a fragment either becomes empty or contains one site only, then the information carried by the corresponding leaf is irrelevant. Indeed, if it is empty this is obvious, while if it contains one site only then the marginal at that site is irrelevant since all marginals are assumed to be fixed, and in particular they are the same in any choice of the initial conditions \vec{p} or \vec{q} . This explains why we introduced the killing step (3.24) in our definition of the fragmentation plus noise process \mathcal{F}_t , which in turn is crucial to the probability of extinction we are able to establish in Lemma 3.5.

The above discussion shows that, if we denote by

$$(\vec{G}, \vec{B}) = (G^{(1)}, B^{(1)}, \dots, G^{(2^t-1)}, B^{(2^t-1)})$$

the vector of realizations of the random graphs and Bernoulli variables involved in each interaction at the $2^t - 1$ internal nodes of the binary tree of depth t , one can write

$$T_t(\vec{p}) = \sum_{(\vec{G}, \vec{B})} \widehat{\nu}(\vec{G}, \vec{B}) T_t(\vec{p} \mid \vec{G}, \vec{B}) \tag{3.31}$$

where

$$\widehat{\nu}(\vec{G}, \vec{B}) = \prod_{i=1}^{2^t-1} \nu_{\mathbf{J}}(G^{(i)}) 2^{-|V \setminus V_{G^{(i)}}|} \tag{3.32}$$

is the distribution of the random graphs and Bernoulli variables, while $T_t(\vec{p} \mid \vec{G}, \vec{B})$ is some probability measure that may depend on (\vec{G}, \vec{B}) and on $\vec{p} = (p_1, \dots, p_{2^t})$ in a complicated way but has the property that its dependence on the distribution p_i from the i -th leaf occurs only through the marginal of p_i on the fragment $F_i^{(t)}$. In particular, if $\mathcal{F}_t = \emptyset$, i.e., $F_i^{(t)} = \emptyset$ for all i , then $T_t(\vec{p} \mid \vec{G}, \vec{B}) = T_t(\vec{q} \mid \vec{G}, \vec{B})$. We note that the event $\mathcal{F}_t = \emptyset$ is measurable with respect to (\vec{G}, \vec{B}) , so that we may write

$$T_t(\vec{p} \mid \vec{G}, \vec{B}) = T_t(\vec{q} \mid \vec{G}, \vec{B}), \quad (\vec{G}, \vec{B}) \in \{\mathcal{F}_t = \emptyset\}, \tag{3.33}$$

and hence

$$\|T_t(\vec{p}) - T_t(\vec{q})\|_{\text{TV}} \leq \sum_{(\vec{G}, \vec{B}) \notin \{\mathcal{F}_t = \emptyset\}} \widehat{\nu}(\vec{G}, \vec{B}).$$

Next, we note that

$$\sum_{(\vec{G}, \vec{B}) \notin \{\mathcal{F}_t = \emptyset\}} \widehat{\nu}(\vec{G}, \vec{B}) = \mathbb{P}(\mathcal{F}_t \neq \emptyset), \quad t \in \mathbb{N}, \tag{3.34}$$

where the latter is the probability estimated in Lemma 3.5. Indeed, (3.34) is a consequence of our definition of the fragmentation plus noise process: we have already observed that each step of the fragmentation

$$F_i^{(t-1)} \longrightarrow (\Phi_0(F_i^{(t-1)}), \Phi_1(F_i^{(t-1)}))$$

is produced with the correct distribution, and the product structure (3.32) of the measure $\widehat{\nu}(\vec{G}, \vec{B})$ guarantees that all such steps are performed independently. From Lemma 3.5 we thus conclude that, for any $\delta \in (0, 1)$, there exists a constant $\delta_0 > 0$ in (3.2), and a constant $C_\delta > 0$ such that

$$\|T_t(\vec{p}) - T_t(\vec{q})\|_{\text{TV}} \leq n^2 C_\delta (2 - \delta)^{-t}, \quad t = 1, 2, \dots$$

This implies (3.28) (and in fact shows that we can take the constant c as close as we wish to $\log 2$ provided δ_0 is taken suitably small). This ends the proof of Theorem 3.6 and Theorem 3.1. \square

We now need to provide the missing proof of Lemma 3.5.

3.5. *Proof of Lemma 3.5.* The event $\mathcal{F}_t \neq \emptyset$ implies that there exists a fragment $F_i^{(t-1)}$ at time $t - 1$ that has cardinality $|F_i^{(t-1)}| \geq 2$. Indeed, by construction, if all fragments have size at most 1 at time $t - 1$, then $\mathcal{F}_t = \emptyset$; see (3.24). Since there are 2^{t-1} fragments at time t , and since all the $F_i^{(t-1)}$ have the same distribution, by a union bound it suffices to show that for any $\delta \in (0, 1)$,

$$\mathbb{P}(|F_1^{(t-1)}| \geq 2) \leq n^2 C_\delta (4 - \delta)^{-t}, \quad t = 1, 2, \dots, \tag{3.35}$$

provided (3.2) holds with a sufficiently small $\delta_0 > 0$. To prove (3.35) we shall use a stochastic domination argument that bounds the evolution of the fragment $F_1^{(t-1)}$, $t \geq 1$ by means of independent *labeled branching processes*.

Consider n independent processes $X^y := \{X_\ell^y, \ell = 0, 1, \dots\}$, $y \in [n]$, such that for each $y \in [n]$, X^y is the labeled branching process with $X_0^y = \{y\}$ and such that, at time $t \in \mathbb{N}$, each individual with label x in the $(t - 1)$ -th generation independently gives birth to the set of individuals $U \subseteq [n]$ with offspring distribution

$$\mu_x(U) = \begin{cases} \frac{1}{2}(1 - \rho_x) & U = \emptyset \text{ or } U = \{x\}; \\ \sum_{G \in \mathcal{G}} \nu_{\mathbf{J}}(G) \mathbf{1}_{G(x)=U} & |U| \geq 2, \end{cases} \tag{3.36}$$

where $G(x)$ denotes the connected component of G containing x and, for any $x \in [n]$,

$$\rho_x := 1 - \prod_{z \in [n] \setminus \{x\}} (1 - p_{xz}) = \sum_{G \in \mathcal{G}} \nu_{\mathbf{J}}(G) \mathbf{1}_{G(x) \neq \emptyset} \tag{3.37}$$

is the probability that x has a non-empty neighborhood in the random graph defined by (3.23). Notice that by definition either $G(x)$ is empty or $|G(x)| \geq 2$, and therefore (3.36), for any $x \in [n]$, defines a probability measure on subsets of $[n]$: a sample U from μ_x is obtained by first sampling the neighborhood $G(x)$ from $\nu_{\mathbf{J}}$; if $|G(x)| \geq 2$ then we set $U = G(x)$; if $G(x) = \emptyset$ then we flip a fair coin and set $U = \emptyset$ if heads and $U = \{x\}$ if tails.

Let $N^y(t)$ denote the size of the whole population of the labeled branching process X^y at time $t - 1$, i.e., the total number of individuals generated up to time $t - 1$. The proof of Lemma 3.5 is based on the following bound on the exponential moment of the random variable $X^y(t)$.

Lemma 3.8. *For all $a \in (0, 1)$, there exists $\delta_0 > 0$ and $C_a > 0$ such that if (3.2) holds with constant δ_0 then, for all $y \in [n]$ and all $t \in \mathbb{N}$,*

$$\mathbb{E}[2^{aN^y(t)}] \leq C_a. \tag{3.38}$$

We postpone the proof of Lemma 3.8 and conclude the proof of Lemma 3.5 assuming the validity of this bound.

Now denote by X_ℓ^y the set of all individuals generated at the ℓ -th step, i.e., the ℓ -th generation of the process X^y , and let $|X_\ell^y|$ denote its cardinality. With this notation, an inspection of the definition of the fragmentation plus noise process shows that $F_1^{(t-1)}$ is stochastically dominated by the union of the X^y 's, i.e., $\{F_1^{(\ell-1)}, \ell \in \mathbb{N}\}$ and the independent processes $\{X^y, y \in [n]\}$ can be coupled so that, for any $\ell = 1, 2, \dots$,

$$|F_1^{(\ell-1)}| \leq \sum_{y \in [n]} |X_{\ell-1}^y|. \tag{3.39}$$




Fig. 3. Illustration of the event in (3.40). The dashed line represents a time s up to which both processes X^y and X^z have cardinality at least 1 and after which the process X^y has cardinality at least 2

The main observation at this point is that, from the definition of fragmentation plus noise, the event $|F_1^{(t-1)}| \geq 2$ implies that $|F_1^{(\ell-1)}| \geq 2$ for all $1 \leq \ell \leq t$. This in turn, by the domination (3.39), implies that at all times $1 \leq \ell \leq t - 1$ one has

$$\sum_{y \in [n]} |X_\ell^y| \geq 2.$$

For this to happen there must be two processes X^y, X^z and a time $s < t$ such that both X^y, X^z are alive up to time s and such that X^y has at least two individuals in each generation from time $s + 1$ to time $t - 1$. For example, taking $s = t - 1$, this includes the case where both X^y, X^z are alive up to time $t - 1$, while taking $s = 0$ it includes the case where all processes die at the first time step, except for X^y which has $|X_\ell^y| \geq 2$ for all $1 \leq \ell \leq t - 1$. We refer to Fig. 3 for an illustrative example.

Using the independence of X^y, X^z , and the union bound one has the estimate

$$\begin{aligned} \mathbb{P}(|F_1^{(t-1)}| \geq 2) &\leq \sum_{y,z \in [n]} \sum_{s=0}^{t-1} \mathbb{P}(|X_\ell^y| \geq 1, \forall 0 \leq \ell \leq s, |X_\ell^y| \geq 2, \forall s < \ell < t) \mathbb{P}(|X_\ell^z| \geq 1, \forall 0 \leq \ell \leq s). \end{aligned} \tag{3.40}$$

Letting $N^y(t)$ denote the size of the whole population of X^y up to time $t - 1$, one has

$$N^y(t) = \sum_{\ell=1}^t |X_{\ell-1}^y| = 1 + \sum_{\ell=1}^{t-1} |X_\ell^y|.$$

We note that for any $s \geq 0$, the event $\{|X_\ell^y| \geq 1, \forall 0 \leq \ell \leq s\}$ implies $N^y(s+1) \geq s+1$, and the events $\{|X_\ell^y| \geq 1, \forall 0 \leq \ell \leq s\}$ and $\{|X_\ell^y| \geq 2, \forall s < \ell < t\}$ together imply

$$N^y(t) \geq 2(t - s - 1) + s + 1 = 2t - s - 1.$$

Therefore, for any fixed $a \in (0, 1)$, from Lemma 3.8 and Markov's inequality we obtain

$$\mathbb{P}(|X_\ell^y| \geq 1, \forall \ell \leq s, |X_\ell^y| \geq 2, \forall s < \ell < t) \mathbb{P}(|X_\ell^z| \geq 1, \forall \ell \leq s) \leq C_a^2 2^{-2at}.$$

In conclusion, we deduce that

$$\mathbb{P}(|F_1^{(t-1)}| \geq 2) \leq n^2 t C_a^2 2^{-2at}.$$

Since a can be taken arbitrarily close to 1, this proves the desired estimate (3.35) by taking $t2^{-2at} \leq (4-\delta)^{-t}$ for all t large enough and adjusting the value of the constant C_δ in order to cover all values of $t \in \mathbb{N}$. This concludes the proof of Lemma 3.5. \square

Finally, it remains only to prove Lemma 3.8.

3.6. Proof of Lemma 3.8. We note that a cannot be taken larger than 1 since at each step there is a probability at least $1/2$ of staying alive, so that the event $N^y(t) \geq t$ has probability at least $2^{-(t-1)}$. We start by considering the distribution of the component $G(x)$, namely

$$v_x(U) = \sum_{G \in \mathcal{G}} v_{\mathbf{J}}(G) \mathbf{1}_{G(x)=U},$$

and establish an exponential tail bound on its size. To this end, define

$$\rho_0 := \max_x \sum_{y \neq x} p_{xy},$$

and note from the discussion at the end of Sect. 3.2 that $\rho_0 \leq 4\delta_0$, where δ_0 is the constant in the Dobrushin condition (3.2). Now observe that if ρ_x is given by (3.37), then by the union bound

$$\rho_x = 1 - v_x(\emptyset) \leq \sum_{y \neq x} p_{xy} \leq \rho_0.$$

Let \mathcal{U}_x denote the random variable with distribution v_x . Note that either $|\mathcal{U}_x| = 0$ or $|\mathcal{U}_x| \geq 2$. We are going to prove a tight tail estimate for $|\mathcal{U}_x|$.

Lemma 3.9. *Suppose $\rho_0 < 1/4$. For every $x \in [n]$, the size of the connected component \mathcal{U}_x satisfies*

$$\mathbb{P}(|\mathcal{U}_x| \geq \ell) \leq 2(4\rho_0)^{\ell-1}, \quad \ell \geq 2. \tag{3.41}$$

Proof. We shall argue that the random variable \mathcal{U}_x can be stochastically dominated by the total population of a labeled branching process where at each step an individual x gives birth to a set S distributed according to the neighborhood of x . More formally, let \mathcal{V}_x denote the set of neighbors of x in G , i.e., $\mathcal{V}_x = \{y \in V_G : \{x, y\} \in E_G\}$ and denote by γ_x the distribution of \mathcal{V}_x : thus, for any $S \subseteq [n] \setminus \{x\}$,

$$\gamma_x(S) = \prod_{y \in S} p_{xy} \prod_{z \in ([n] \setminus \{x\}) \setminus S} (1 - p_{xz}). \tag{3.42}$$

Consider the labeled branching process W^x with initial value $W_0^x = \{x\}$ and offspring distribution γ_x . A breadth-first search starting at x then shows that \mathcal{U}_x is stochastically dominated (in the sense of inclusion) by the total population of W^x , i.e., the union of all sets W_ℓ^x , $\ell = 0, 1, \dots$, where W_ℓ^x represents the ℓ -th generation of W^x . To control the

size of the total population of W^x , we note that for any integer $\ell \geq 1$, a union bound and the multinomial theorem show that

$$\begin{aligned} \nu_{\mathbf{J}}(|\mathcal{V}_x| \geq \ell) &= \sum_{S: |S| \geq \ell} \gamma_x(S) \leq \sum_{S: |S| = \ell} \prod_{y \in S} p_{xy} \\ &\leq \sum_{\{n_y\}: \sum_y n_y = \ell} \prod_y p_{xy}^{n_y} = \left(\sum_{y \neq x} p_{xy} \right)^\ell \leq \rho_0^\ell, \end{aligned} \tag{3.43}$$

where the third sum extends over all nonnegative integer vectors $\{n_y, y \in [n] \setminus \{x\}\}$ such that $\sum_y n_y = \ell$. We use (3.43) to define a random variable \mathcal{N} such that the size $|\mathcal{V}_x|$ of the neighborhood at x is stochastically dominated by \mathcal{N} , and \mathcal{N} has distribution $\bar{\mu}$ given by

$$\bar{\mu}(0) = 1 - \rho_0, \quad \bar{\mu}(\ell) = \rho_0^\ell (1 - \rho_0), \quad \ell \geq 1. \tag{3.44}$$

Indeed, provided $\rho_0 < 1$, (3.44) defines the geometric distribution with parameter $1 - \rho_0$, and by (3.43) one has

$$\nu_{\mathbf{J}}(|\mathcal{V}_x| \geq \ell) \leq \sum_{k \geq \ell} \bar{\mu}(k) = \rho_0^\ell,$$

for all $\ell \geq 0$. This implies the desired stochastic domination, so that \mathcal{V}_x and \mathcal{N} can be coupled so that $|\mathcal{V}_x| \leq \mathcal{N}$. Thus, the size of the total population of W^x is dominated by the total size $|\mathcal{T}|$ of the Galton Watson tree \mathcal{T} with offspring distribution $\bar{\mu}$. In particular, we can use a coupling of \mathcal{U}_x and \mathcal{T} such that $|\mathcal{U}_x| \leq |\mathcal{T}|$.

Let $\varphi(s) = \mathbb{E}[s^{\mathcal{N}}]$ denote the moment generating function of \mathcal{N} . Standard calculations show that the extinction probability for \mathcal{T} is given by the smallest positive solution s_{ext} of $s = \varphi(s)$, and since $\mathbb{E}[\mathcal{N}] = \rho_0 / (1 - \rho_0) < 1$ here one has $s_{\text{ext}} = 1$. Recall also that φ is increasing and convex. With similar calculations one finds that for any fixed $u > 0$, the exponential moment

$$\mathbb{E} \left[u^{|\mathcal{T}|} \right] = s_*(u), \tag{3.45}$$

where $s_*(u)$ is defined as the smallest positive solution of the equation $s = u\varphi(s)$. Note that we are free to choose u . Let us show that, if we take $u = 1/(4\rho_0)$, then $s_*(u) \leq 1/(2\rho_0)$. Indeed, by explicit calculation,

$$\varphi(s) = \frac{1 - \rho_0}{1 - s\rho_0},$$

and $\varphi(s) \leq (1 - \rho_0)(1 + 2s\rho_0) \leq 1 + 2s\rho_0$ if $s\rho_0 \leq 1/2$. Therefore,

$$u\varphi(s) = \frac{\varphi(s)}{4\rho_0} \leq \varphi_0(s) := \frac{1}{4\rho_0} + \frac{s}{2}.$$

Noting that the solution of the equation $s = \varphi_0(s)$ is $s_0 = 1/(2\rho_0)$, we see that the desired solution $s_*(u)$ of $s = u\varphi(s)$ exists and satisfies $s_*(u) \leq 1/(2\rho_0)$.

From (3.38) and the previous arguments, using Markov's inequality one finds

$$\mathbb{P}(|\mathcal{U}_x| \geq \ell) \leq \mathbb{P}(|\mathcal{T}| \geq \ell) \leq (4\rho_0)^\ell \mathbb{E} \left[(4\rho_0)^{-|\mathcal{T}|} \right] \leq 2(4\rho_0)^{\ell-1}, \quad \ell \geq 2.$$

This proves (3.41). \square

We now turn to the proof of Lemma 3.8. Recall the definition of X^y as the labeled branching process with offspring distribution μ_x from (3.36). Call $\tilde{\mathcal{U}}_x$ the random variable with distribution μ_x . Let \tilde{v} be defined as

$$\tilde{v}(0) = \frac{1}{2} - \rho_*, \quad \tilde{v}(1) = \frac{1}{2}, \quad \tilde{v}(\ell) = 2(1 - 4\rho_0)(4\rho_0)^{\ell-1}, \quad \ell \geq 2,$$

where $\rho_* = 8\rho_0$. Note that this is a well defined probability provided $16\rho_0 \leq 1$. Let $\tilde{\mathcal{N}}$ denote the integer valued random variable with distribution \tilde{v} . Lemma 3.9 implies that $|\tilde{\mathcal{U}}_x|$ is stochastically dominated by $\tilde{\mathcal{N}}$. To see this, note that for any $k \geq 2$, the event $|\tilde{\mathcal{U}}_x| = k$ has the same probability as the event $|\mathcal{U}_x| = k$, where \mathcal{U}_x is the random variable with distribution ν_x . From Lemma 3.9 we know that

$$\mathbb{P}(|\tilde{\mathcal{U}}_x| \geq \ell) = \mathbb{P}(|\mathcal{U}_x| \geq \ell) \leq 2(4\rho_0)^{\ell-1} = \sum_{k \geq \ell} \tilde{v}(k), \quad \ell \geq 2.$$

Moreover, by the definition of μ_x one has $\mathbb{P}(|\tilde{\mathcal{U}}_x| = 1) = \frac{1}{2} - \rho_x \leq \frac{1}{2}$, and therefore

$$\mathbb{P}(|\tilde{\mathcal{U}}_x| \geq \ell) \leq \sum_{k \geq \ell} \tilde{v}(k), \quad \ell \geq 0.$$

This proves the stochastic domination $|\tilde{\mathcal{U}}_x| \leq \tilde{\mathcal{N}}$.

Recall that $N^y(t)$ denotes the size of the whole population of the labeled branching process X^y up to time $t - 1$. The above domination argument implies that, for any $t \in \mathbb{N}$, we can dominate $N^y(t)$ by the size $|\tilde{\mathcal{T}}|$ of the total population of the Galton Watson tree $\tilde{\mathcal{T}}$ with offspring distribution \tilde{v} , so that

$$\mathbb{E} \left[2^{aN^y(t)} \right] \leq \mathbb{E} \left[2^{a|\tilde{\mathcal{T}}|} \right].$$

As in (3.45) we have that, for any fixed $u > 0$, $\mathbb{E} \left[u^{|\tilde{\mathcal{T}}|} \right] = \tilde{s}_*(u)$, where $\tilde{s}_*(u)$ is the smallest solution of the equation $s = u\tilde{\varphi}(s)$, for

$$\tilde{\varphi}(s) = \mathbb{E} \left[s^{\tilde{\mathcal{N}}} \right],$$

the generating function of \tilde{v} . We calculate

$$\tilde{\varphi}(s) = \frac{1}{2} - \rho_* + \frac{s}{2} + 2(1 - 4\rho_0) \sum_{\ell \geq 2} (4\rho_0)^{\ell-1} s^\ell \leq \frac{1}{2}(1 + s + s\kappa(s\rho_0)),$$

where $\kappa(t) = 16t/(1 - 4t)$. Set $u = 2(1 - \eta) = 2^a$ for some fixed positive $\eta \leq 1/2$, and note that $\kappa(t) \leq \eta$ if $t \leq \eta/18$ for all $\eta \in [0, 1/2]$. Thus, assuming $s\rho_0 \leq \eta/18$, one has

$$u\varphi(s) \leq (1 - \eta)(1 + s(1 + \eta)) \leq (1 - \eta) + s(1 - \eta^2).$$

Therefore, $\tilde{s}_*(u) \leq (1 - \eta)/\eta^2$. For this value of s to also satisfy the requirement $s\rho_0 \leq \eta/18$, we need $\rho_0 \leq \eta^3/(18(1 - \eta))$. By the discussion at the end of Sect. 3.2, $\delta_0 \leq \eta^3/(72(1 - \eta))$ is sufficient for this to hold. Under this assumption one has

$$\mathbb{E} \left[2^{a|\tilde{\mathcal{T}}|} \right] \leq \eta^{-2}.$$

The estimate (3.38) then holds with $2(1 - \eta) = 2^a$, $C_a = \eta^{-2}$, provided $\delta_0 \leq \eta^3/(72(1 - \eta))$. This completes the proof of Lemma 3.8. \square

Remark 3.10. Note that the value of δ_0 required for Lemma 3.8 to hold is precisely the value of δ_0 that we will require in the high-temperature condition (3.2) in our main result for the nonlinear block dynamics, Theorem 3.1. We have made no attempt here to optimize the dependence of δ_0 on our arguments, and with further work one could no doubt obtain a better, explicit bound. However, it is clear that our current arguments will not be able to obtain an optimal value of δ_0 , matching Dobrushin-type thresholds for rapid mixing for linear Markov chains. For instance, in the mean field case where $\mathbf{J} \equiv \beta/n$, the threshold for the equilibrium phase transition is $\beta = 1$, and thus $\delta_0 < 1$ would be an optimal condition.

4. Nonlinear Glauber Dynamics

Recalling the definition (2.9), the interaction (2.11) in the nonlinear Glauber dynamics takes the form

$$p \circ q = \frac{1}{n} \sum_{x \in [n]} \sum_{\sigma, \sigma'} \frac{1}{2} (p(\sigma)q(\sigma') + q(\sigma)p(\sigma')) Q_x(\cdot \mid \sigma, \sigma'), \tag{4.1}$$

where $Q_x(\tau \mid \sigma, \sigma') := \sum_{\tau'} Q_{\mathbf{J},x}(\sigma, \sigma'; \tau, \tau')$. We write $S_t(p)$, $t \in \mathbb{N}$, for the t -th iteration of $p \mapsto p \circ p$. From Theorem 2.7 we know that for any fixed interaction matrix \mathbf{J} , and any initial distribution $p \in \mathcal{P}(\Omega)$, the above dynamics converges to $\mu_{\mathbf{J},\mathbf{h}}$, where \mathbf{h} is the unique vector of external fields such that $\mu_{\mathbf{J},\mathbf{h}}$ and the initial state p have the same marginals on σ_x , for all $x \in V$. Our main result for the nonlinear Glauber (single site) dynamics, Theorem 1.2 in the introduction, establishes a tight bound on the rate of convergence as a function of $n = |V|$, under the same Dobrushin-type condition (3.2) on the interaction as in the case of block dynamics. We restate this theorem here for convenience.

Theorem 4.1. *There exist absolute constants $\delta_0 > 0$, $c > 0$ and $C > 0$ such that, if (3.2) holds with constant δ_0 , then for any $p \in \mathcal{P}(\Omega)$ and $t \in \mathbb{N}$,*

$$\|S_t(p) - \mu_{\mathbf{J},\mathbf{h}}\|_{TV} \leq C n e^{-ct/n},$$

where $\mathbf{h} \in \mathbb{R}^n$ is the unique choice of external fields such that $p_x = (\mu_{\mathbf{J},\mathbf{h}})_x$ for all $x \in [n]$. In particular, for any $\varepsilon > 0$, one has $\|S_t(p) - \mu_{\mathbf{J},\mathbf{h}}\|_{TV} \leq \varepsilon$ as soon as $t \geq \frac{n}{c} \log n + C_1(\varepsilon)$, where $C_1(\varepsilon) = \frac{n}{c} \log(C/\varepsilon)$.

The main difference with respect to Theorem 3.1 is the rate of exponential decay, which is n times slower in this case. This reflects the intuitive fact that only one spin is exchanged at each time step whereas $\Theta(n)$ spins are exchanged in a typical transition of the nonlinear block dynamics studied in Theorem 3.1.

The proof of Theorem 4.1 follows a similar strategy to that of Theorem 3.1, but with some important technical differences. We begin with an analog of Lemma 3.3, in which the role of Erdős-Rényi graphs is now played by random star graphs centered at a specific vertex x .

4.1. Coupling with inhomogeneous random star graphs. We start with a single site version of Lemma 3.2.

Lemma 4.2. For any $x \in [n]$, $\sigma, \sigma' \in \Omega$,

$$\alpha_x(\sigma, \sigma') = \frac{1}{1 + \exp \left\{ (\sigma_x - \sigma'_x) \sum_{y: y \neq x} J_{xy} (\sigma_y - \sigma'_y) \right\}}. \tag{4.2}$$

Proof. The expression (4.2) follows directly from (1.3), the definition (2.9) and the observation that

$$\begin{aligned} & \sum_{y,z} J_{y,z} [(\sigma_y \sigma_z + \sigma'_y \sigma'_z) - ((\sigma_x \sigma'_{[n] \setminus \{x\}})_y [\sigma_x \sigma'_{[n] \setminus \{x\}}]_z + [\sigma'_x \sigma_{[n] \setminus \{x\}}]_y [\sigma'_x \sigma_{[n] \setminus \{x\}}]_z)] \\ &= (\sigma_x - \sigma'_x) \sum_{y: y \neq x} J_{xy} (\sigma_y - \sigma'_y). \end{aligned}$$

□

Remark 4.3. Consider the random variable η_x which takes the value -1 with probability $\alpha_x(\sigma, \sigma')$, and $+1$ with probability $1 - \alpha_x(\sigma, \sigma')$. Then by Lemma 4.2, its probability density can be written as $\alpha_x(\cdot | \sigma, \sigma')$, where

$$\alpha_x(\eta_x | \sigma, \sigma') \propto \exp \left\{ \eta_x \sum_{y: y \neq x} \tilde{J}_{xy} \right\}, \quad \eta_x \in \{-1, +1\},$$

and \tilde{J}_{xy} is defined as in (3.11). This also shows that $\alpha_x(\cdot | \sigma, \sigma')$ coincides with the distribution of η_x obtained by conditioning $\gamma(\cdot | \sigma, \sigma')$ on the event that $\{\eta_y = 1, \forall y \neq x\}$, where $\gamma(\cdot | \sigma, \sigma')$ is the distribution from Lemma 3.2. Note also that, if $\sigma_x = \sigma'_x$, then $\tilde{J}_{xy} = 0$ for all $y \neq x$, and therefore $\alpha_x(\cdot | \sigma, \sigma') = \text{Be}(1/2)$ is the Bernoulli probability measure on $\{-1, +1\}$ with parameter $1/2$.

Fix $x \in [n]$, and let \mathcal{G}_x be the set of all subgraphs of the star graph $S_x = (V, E_x)$ with vertex set $V = [n]$, and edge set $E_x := \{\{x, y\}, y \in [n] \setminus \{x\}\}$. We view $G \in \mathcal{G}_x$ as a collection of edges, i.e., a subset of E_x , with no isolated vertices, and write $\mathcal{P}(\mathcal{G}_x)$ for the set of probability measures over \mathcal{G}_x . Note that any $G \in \mathcal{G}_x$ is always connected, but can be the empty graph (with no vertices). We use V_G, E_G for the vertex and edge sets of $G \in \mathcal{G}_x$, respectively. The following is the single site version of Lemma 3.3.

Lemma 4.4. Fix $x \in [n]$. Let $\nu_{\mathbf{J}}^x$ be the inhomogeneous random star measure associated with the weights $\lambda_{xy} = e^{4|J_{xy}|} - 1$, i.e.,

$$\nu_{\mathbf{J}}^x(G) \propto \prod_{\{x,y\} \in E_G} (e^{4|J_{xy}|} - 1), \quad G \in \mathcal{G}_x.$$

Then

$$\alpha_x(\cdot | \sigma, \sigma') = \sum_{G \in \mathcal{G}_x} \nu_{\mathbf{J}}^x(G) \mu_G^x(\cdot | \sigma_{V_G}, \sigma'_{V_G}),$$

where, for any $G \in \mathcal{G}_x$, $\mu_G^x(\cdot | \sigma_{V_G}, \sigma'_{V_G})$ is a probability measure on $\{-1, +1\}$ that depends on σ, σ' only through the spins $\sigma_{V_G}, \sigma'_{V_G}$ and such that if $G = \emptyset$, then $\mu_G^x(\cdot | \sigma_{V_G}, \sigma'_{V_G}) = \text{Be}(\frac{1}{2})$ is the Bernoulli probability measure on $\{-1, +1\}$, with parameter $1/2$.

Proof. The proof is similar to that of Lemma 3.3 and is omitted. □

4.2. *Coupon collecting with noise.* In the case of single site updates, fragmentation involves hitting each site at least once. The analog of fragmentation plus noise can therefore be viewed as a coupon collecting process, again perturbed by a local growth described by simpler star graphs.

Given a set $A \subseteq [n]$, a uniformly random vertex $x \in [n]$, a Bernoulli variable $B \in \{-1, +1\}$ with parameter $1/2$ and a random graph $G \in \mathcal{G}_x$ with distribution ν_x^x , we define the random sets

$$\Psi_0(A) = \Psi_1(A) = \emptyset, \quad \text{if } |A| \leq 1, \tag{4.3}$$

and, if $|A| \geq 2$,

$$(\Psi_0(A), \Psi_1(A)) = \begin{cases} (A, \emptyset) & \text{if } x \notin A \\ (A, \emptyset) & \text{if } x \in A, G = \emptyset \text{ and } B = -1 \\ (A \setminus \{x\}, \emptyset) & \text{if } x \in A, G = \emptyset \text{ and } B = +1 \\ (A \cup V_G, V_G) & \text{if } x \in A \text{ and } G \neq \emptyset \end{cases} \tag{4.4}$$

Definition 4.5. The *coupon collecting plus noise* process $\mathcal{C}_t, t = 0, 1, \dots$ is the random process defined as follows. For each $t \in \mathbb{N}$, \mathcal{C}_t consists of 2^t labeled fragments, i.e., (possibly empty) subsets $C_1^{(t)}, \dots, C_{2^t}^{(t)}, C_i^{(t)} \subseteq [n]$, obtained by repeated application of the following rule. At time zero we have $\mathcal{C}_0 = \{\{n\}\}$, i.e., $C_1^{(0)} = [n]$. At time $t \in \mathbb{N}$, if $\mathcal{C}_{t-1} = (C_1^{(t-1)}, \dots, C_{2^{t-1}}^{(t-1)})$, then for each i independently, the fragment $C_i^{(t-1)}$ is replaced by $(\Psi_0(C_i^{(t-1)}), \Psi_1(C_i^{(t-1)}))$ where Ψ_0, Ψ_1 are the random maps defined by (4.3)-(4.4), so that

$$\begin{aligned} \mathcal{C}_t &= (C_1^{(t)}, \dots, C_{2^t}^{(t)}) \\ &= (\Psi_0(C_1^{(t-1)}), \Psi_1(C_1^{(t-1)}), \Psi_0(C_2^{(t-1)}), \Psi_1(C_2^{(t-1)}), \dots, \Psi_0(C_{2^{t-1}}^{(t-1)}), \Psi_1(C_{2^{t-1}}^{(t-1)})). \end{aligned}$$

The process is said to die out if there is a time t such that $C_i^{(t)} = \emptyset$ for all $i = 1, \dots, 2^t$. In the latter case, with slight abuse of notation, we write $\mathcal{C}_t = \emptyset$.

Remark 4.6. In the non-interacting case $\mathbf{J} = 0$ one has always $G = \emptyset$ and thus $\Psi_1(A) = \emptyset$, so that the process \mathcal{C}_t is a monotone decreasing sequence obtained by iterations of the map $A \mapsto \Phi_0(A)$, where at each step one picks a uniformly random vertex $x \in [n]$ and if $x \in A$ the vertex x is removed from A with probability $1/2$. This yields a lazy version of the standard coupon collecting process, and therefore the union bound shows that

$$\mathbb{P}(\mathcal{C}_t \neq \emptyset) \leq n \left(1 - \frac{1}{2n}\right)^t \leq n e^{-t/(2n)}, \quad t = 1, 2, \dots \tag{4.5}$$

In the interacting case, the event $G \neq \emptyset$ causes the splitting of A into two nontrivial sets $(\Psi_0(A), \Psi_1(A))$. This local growth can slow down the convergence to the absorbing state \emptyset . However, we will show that if the growth is sufficiently subcritical (ensured by the Dobrushin condition (3.2)) then a bound of the form (4.5) continues to hold.

Despite the formal similarities between Definition 4.5 and Definition 3.4, the two processes are quite different. Note in particular the strong asymmetry between the different fragments in the case of Definition 4.5. Let us first consider the evolution of the leftmost fragment $C_1^{(t)}, t = 0, 1, \dots$

Lemma 4.7. *There exist absolute constants $\delta_0 > 0$, $a > 0$ and $C > 0$ such that if (3.2) holds with constant δ_0 then*

$$\mathbb{P}(C_1^{(t)} \neq \emptyset) \leq Cne^{-at/n}, \quad t = 1, 2, \dots$$

Proof. The leftmost fragment $C_1^{(t)} = \Phi_0 \circ \dots \circ \Phi_0([n])$ is obtained by iterating t times the map Φ_0 . It may be represented by the following labeled branching process. Consider random variables $W := \{W_t, t = 0, 1, \dots\}$, where $W_0 = [n]$ and for each t , $W_t \subseteq [n]$ is updated to $W_{t+1} \subseteq [n]$ by the rule

- pick $x \in [n]$ u.a.r.
- if $x \notin W_t$ then set $W_{t+1} = W_t$
- if $x \in W_t$ then let \mathcal{V}_x denote the neighborhood of x in G , i.e., the random variable with distribution as in (3.42): if $\mathcal{V}_x \neq \emptyset$, set $W_{t+1} = W_t \cup \mathcal{V}_x$; if instead $\mathcal{V}_x = \emptyset$, set

$$W_{t+1} = \begin{cases} W_t \setminus \{x\} & \text{with prob. } \frac{1}{2} \\ W_t & \text{with prob. } \frac{1}{2} \end{cases}$$

From our definitions we see that $C_1^{(t)}$ has the same distribution as W_t .

Next, consider a continuous time version of the process W_t . Namely, let T_1, T_2, \dots denote the arrival times of the Poisson point process with intensity 1 on $[0, \infty)$ and call $\tilde{W}_t, t \geq 0$, the process obtained by repeating the steps above at each arrival time T_j , so that $W_t = \tilde{W}_{T_j}$ for every $t \in \mathbb{N}$.

From standard properties of the Poisson process, the random process \tilde{W}_t can be equivalently obtained by attaching to each $x \in [n]$ an independent Poisson process with intensity $1/n$, and performing updates at x independently at the arrival times of the process at x . Neglecting the overlap between different branches, and using the domination argument in the proof of Lemma 3.9, one can stochastically dominate the cardinality of \tilde{W}_t by

$$|\tilde{W}_t| \leq \sum_{x \in [n]} M_t^x,$$

where $M_t^x, x \in [n]$ are i.i.d. copies of the continuous time branching process obtained by setting $M_0^x = 1$ and then letting each branch evolve independently with branching times given by exponential random variables with parameter $1/n$. In this process, each individual gives birth to ℓ individuals with probability $\mu_*(\ell)$ given by

$$\mu_*(\ell) = \begin{cases} \frac{1}{2}\bar{\mu}(0) & \ell = 0 \\ \frac{1}{2}\bar{\mu}(1) & \ell = 1 \\ \bar{\mu}(\ell) & \ell \geq 2 \end{cases}$$

where $\bar{\mu}$ is defined as in (3.44) with the same value of ρ_0 . (See [5] for background on continuous time branching processes.) Since M_t^1 is subcritical as soon as δ_0 is sufficiently small, and $\bar{\mu}$ has exponential tails, standard estimates on subcritical continuous time branching processes imply that $\mathbb{P}(M_t^1 > 0) \leq Ae^{-at/n}$ for some constants $a, A > 0$. It follows that

$$\mathbb{P}(|\tilde{W}_t| > 0) \leq n \mathbb{P}(M_t^1 > 0) \leq nAe^{-at/n}.$$

To conclude the proof, recall that $W_t = \widetilde{W}_{T_t}$ and that T_t is the sum of t i.i.d. exponentials of parameter 1, so that

$$\mathbb{P}(T_t \leq t/2) \leq e^{-ct},$$

for some constant $c > 0$. Therefore, for all $t \in \mathbb{N}$,

$$\mathbb{P}(C_1^{(t)} \neq \emptyset) \leq \mathbb{P}(|\widetilde{W}_{T_t}| > 0) \leq \mathbb{P}(|\widetilde{W}_{t/2}| > 0) + e^{-ct} \leq n A e^{-at/(2n)} + e^{-ct},$$

completing the proof of the lemma. \square

Next, we extend the idea of the proof of Lemma 4.7 to obtain the following stronger estimate.

Lemma 4.8. *There exist absolute constants $\delta_0 > 0$, $a > 0$ and $C > 0$ such that if (3.2) holds with constant δ_0 then*

$$\mathbb{P}(C_t \neq \emptyset) \leq C n e^{-at/n}, \quad t = 1, 2, \dots \tag{4.6}$$

Proof. The idea is to add all individuals from the other fragments to the evolution of the first fragment, which we interpret as the “master evolution” to which all further fragments are added, and to dominate this global process by a single sub-critical branching process. Indeed, notice that every time there is a branching event in the coupon collecting plus noise process, at a fragment $C_i^{(j)}$, say at vertex x , if a nonempty neighborhood \mathcal{V}_x of x is drawn, then the set \mathcal{V}_x is added to both $\Phi_0(C_i^{(j)})$ and $\Phi_1(C_i^{(j)})$. Thus, recursively, we can add all of them to the master evolution and obtain a global process $\widetilde{C} = \{C_t, t = 0, 1, \dots\}$ that evolves like the first fragment, with the rule that every time a $\mathcal{V}_x \neq \emptyset$ is drawn then *two* copies of that set are added to the current configuration, each of which continues its evolution independently according to the same branching rules as above. In particular, the same argument as in the proof of Lemma 4.7 shows that

$$\mathbb{P}(C_t \neq \emptyset) \leq \mathbb{P}(Z_{T_t} \neq \emptyset), \tag{4.7}$$

where $Z_t, t \geq 0$, is the continuous time process defined by $Z_t = \sum_{i=1}^n Z_t^i$, where Z_t^i are i.i.d. copies of a process Z_t^1 defined as follows. At time 0 there is one individual, and each individual that is alive at any time $s \geq 0$, independently of all others, waits an exponential time with parameter $1/n$ and then branches into 0, 1, or $2j$ individuals, with probabilities given, respectively by, $\mu_*(0)$, $\mu_*(1)$, and $\mu_*(j)$, $j \geq 2$. If the branching produces $2j$ individuals, then these are divided into two groups (y_1, \dots, y_j) and (y'_1, \dots, y'_j) , which continue their evolutions independently except that the branching times of the two groups are coupled so that they are identical. (This slightly strange coupling is a consequence of the discrete nature of time in the original process C_t .)

Let us first consider the simpler situation where we do not have the coupling of branching times every time a branching produces $2j$ individuals. Indeed, let us call \widetilde{Z}_t^1 the process evolving as above with the feature that the two groups (y_1, \dots, y_j) and (y'_1, \dots, y'_j) evolve in a fully independent way following the above described rules. Thus, \widetilde{Z}_t^1 is a pure continuous-time branching process with offspring distribution

$$\nu_*(\ell) = \begin{cases} \frac{1}{2} \bar{\mu}(0) & \ell = 0 \\ \frac{1}{2} \bar{\mu}(1) & \ell = 1 \\ \bar{\mu}(\ell/2) & \ell = 4, 6, \dots \end{cases}$$

The process \tilde{Z}_t^1 is subcritical as soon as δ_0 is sufficiently small, and since $\bar{\mu}$ has exponential tails, for some constants $a, A > 0$ one has

$$\mathbb{P}(\tilde{Z}_t^1 > 0) \leq Ae^{-at/n}.$$

In particular, letting $\tilde{Z}_t = \sum_{i=1}^n \tilde{Z}_t^i$, we conclude, as in the proof of Lemma 4.7, that

$$\mathbb{P}(\tilde{Z}_{T_t} \neq \emptyset) \leq n Ae^{-at/(2n)} + e^{-ct}. \tag{4.8}$$

Thus the proof would be complete if we could replace Z_{T_t} with \tilde{Z}_{T_t} in (4.7). In order to finish the proof it is then sufficient to show that

$$\mathbb{P}(\mathcal{C}_t \neq \emptyset) \leq \mathbb{P}(\tilde{Z}_{T_t} \neq \emptyset) + e^{-ct}, \tag{4.9}$$

for some constant $c > 0$.

To prove (4.9), recall the definition of the random process \tilde{W}_t obtained by attaching independent Poisson processes with intensity $1/n$ to each $x \in [n]$, such that $C_1^{(t)} = \tilde{W}_{T_t}$. Let us add to the master evolution $C_1^{(t)}$ all other evolutions as explained above by doubling the splittings every time there is a branching involving a nonempty neighborhood of x . This yields a global process $\mathcal{Q}_t, t \geq 0$ on a single evolution, such that $\{\mathcal{C}_t = \emptyset\} = \{\mathcal{Q}_{T_t} = \emptyset\}$. Note that the branching times (but not the choices of vertices x to be updated) of newly added evolutions are fully coupled to the branching times of the master evolution dictated by the original Poisson process. If, for every newly opened evolution we use instead an independent realization of the Poisson point process, then we obtain a global process $\tilde{\mathcal{Q}}_t$, which is dominated as explained above by the process \tilde{Z}_t . Thus it is sufficient to find a coupling of \mathcal{Q}_t and $\tilde{\mathcal{Q}}_t$ such that

$$\mathbb{P}(\mathcal{Q}_{T_t} \neq \emptyset) \leq \mathbb{P}(\tilde{\mathcal{Q}}_{T_t} \neq \emptyset) + e^{-bt},$$

for some constant $b > 0$.

Note that if $b > 0$ is a sufficiently small constant, then a simple large deviation bound shows that the probability that a Poisson process of intensity 1 has a total number of arrivals in a time $t/2$ that is less than bt , is at most 3^{-t} . Thus, if there are 2^t such processes, the event \mathcal{E}_t that all of them have at least bt arrivals within time $t/2$ satisfies

$$\mathbb{P}(\mathcal{E}_t^c) \leq 2^t 3^{-t} = (2/3)^t.$$

We can couple $\mathcal{Q}_t, \tilde{\mathcal{Q}}_t$ in such a way that, on the event $\{T_t \geq t/2\} \cap \mathcal{E}_t$, we have $\{\mathcal{Q}_{T_t} \neq \emptyset\} \subset \{\mathcal{Q}_{T_{bt}} \neq \emptyset\}$. It follows that,

$$\begin{aligned} \mathbb{P}(\mathcal{C}_t \neq \emptyset) &= \mathbb{P}(\mathcal{Q}_{T_t} \neq \emptyset) \leq \mathbb{P}(\mathcal{Q}_{T_t} \neq \emptyset, T_t \geq t/2, \mathcal{E}_t) + (2/3)^t + e^{-ct} \\ &\leq \mathbb{P}(\tilde{\mathcal{Q}}_{T_{bt}} \neq \emptyset) + (2/3)^t + e^{-ct}. \end{aligned} \tag{4.10}$$

On the other hand, $\tilde{\mathcal{Q}}_t \neq \emptyset$ implies $\tilde{Z}_t \neq 0$, and thus, by (4.8),

$$\mathbb{P}(\tilde{\mathcal{Q}}_{T_{bt}} \neq \emptyset) \leq n Ae^{-abt/(2n)} + e^{-bct}. \tag{4.11}$$

Combining (4.10) and (4.11) proves the desired estimate (4.6). \square

4.3. *Proof of Theorem 4.1.* As in the proof of Theorem 3.1, the first step is to use the coupling with random (star) graphs and Bernoulli variables to decompose the distribution at every internal node of the interaction tree. For this we use Lemma 4.4. This leads to the analysis of the branching process, which in this case takes the form of coupon collecting with noise. The conclusion then follows from Lemma 4.8. With the notation from Theorem 3.6 we are going to establish that there exist absolute constants $\delta_0 > 0$, $c > 0$ and $C > 0$ such that, if (3.2) holds with constant δ_0 then, for any choice of initial distributions \vec{p}, \vec{q} with the same marginals as in (3.27),

$$\|S_t(\vec{p}) - S_t(\vec{q})\|_{TV} \leq Cne^{-ct/n}, \quad t \in \mathbb{N}, \tag{4.12}$$

where $S_t(\vec{p})$ denotes the distribution at the root of the binary tree when the leaves are equipped with the distributions \vec{p} and the single site interaction (4.1) occurs at each internal node.

For each interaction we draw a uniformly random vertex $x \in [n]$, and reveal a realization of the neighborhood $G = \mathcal{V}_x$ and a Bernoulli variable $B \in \{-1, +1\}$. Then, according to Lemma 4.4, the interaction (4.1) has the form

$$p \circ q = \frac{1}{2n} \sum_{x, G, B} v_{\mathbf{j}}^x(G) \frac{1}{2} (S(p, q | x, G, B) + S(q, p | x, G, B)) \tag{4.13}$$

where

$$S(p, q | x, G, B)(\tau) = \sum_{\sigma, \sigma'} p(\sigma)q(\sigma') \sum_{\eta_x \in \{-1, +1\}} \mu_{G, B}^x(\eta_x | \sigma_{V_G}, \sigma'_{V_G}) \mathbf{1}_{\tau \sim (\sigma, \sigma', \eta_x)},$$

and, for $\tau \in \Omega$, the notation $\tau \sim (\sigma, \sigma', \eta_x)$ stands for

$$\tau \sim (\sigma, \sigma', \eta_x) \Leftrightarrow \begin{cases} \tau_x = \sigma_x, & \eta_x = +1 \\ \tau_x = \sigma'_x, & \eta_x = -1 \\ \tau_y = \sigma_y, & \forall y \neq x \end{cases}$$

and $\mu_{G, B}^x(\cdot | \sigma_{V_G}, \sigma'_{V_G})$ is a probability distribution on $\{-1, +1\}$ that depends on σ, σ' only through $\sigma_{V_G}, \sigma'_{V_G}$ if $G \neq \emptyset$, while it assigns the value $\eta_x = B$ deterministically if $G = \emptyset$. This definition reflects the fact that if $G = \emptyset$ then we assign the value of η_x by a fair coin flip B , whereas if $G \neq \emptyset$ then B is irrelevant and we sample η_x according to $\mu_{G, B}^x(\cdot | \sigma_{V_G}, \sigma'_{V_G})$.

We remark that our definition of the maps Ψ_0 and Ψ_1 in Definition 4.5 is such that the only information about the distributions p, q that is needed to compute $S(p, q | x, G, B)$ is contained in the marginals $p_{\Psi_0([n])}$ and $q_{\Psi_1([n])}$. Note the asymmetry between p, q in this expression: symmetry is restored in (4.13) by the averaging $\frac{1}{2} (S(p, q | x, G, B) + S(q, p | x, G, B))$. The choice between the two options will then be encoded by a further Bernoulli variable B' with parameter 1/2, that we sample afresh at each interaction.

Repeating the argument in the proof of Theorem 3.6, one has that after t steps, conditional on the realizations \vec{x} of all random vertices, all graphs \vec{G} and Bernoulli variables \vec{B}, \vec{B}' involved in each of the $2^t - 1$ interactions, the fragment $C_i^{(t)}$, $i = 1, \dots, 2^t$, as defined in Definition 4.5, contains all the information needed from the distribution at the leaf $\vec{B}'(i)$, which is determined by taking the left or right descendant from the root at each internal node according to whether $B' = +1$ or $B' = -1$ at that

node. As before, we observe that as soon as a fragment becomes either empty or contains one site only, then the information carried by the corresponding leaf is irrelevant.

In analogy with (3.31), we write

$$S_t(\vec{p}) = \sum_{(\vec{x}, \vec{G}, \vec{B}, \vec{B}')} \widehat{\nu}(\vec{x}, \vec{G}, \vec{B}, \vec{B}') S_t(\vec{p} \mid \vec{x}, \vec{G}, \vec{B}, \vec{B}')$$

with

$$\widehat{\nu}(\vec{x}, \vec{G}, \vec{B}, \vec{B}') = \prod_{i=1}^{2^t-1} \frac{1}{4^n} \nu_{\mathbf{J}}^x(G^{(i)})$$

where $S_t(\vec{p} \mid \vec{x}, \vec{G}, \vec{B}, \vec{B}')$ is some probability measure that may depend on $(\vec{x}, \vec{G}, \vec{B}, \vec{B}')$ and on $\vec{p} = (p_1, \dots, p_{2^t})$ in a complicated way but has the property that

$$S_t(\vec{p} \mid \vec{x}, \vec{G}, \vec{B}, \vec{B}') = S_t(\vec{q} \mid \vec{x}, \vec{G}, \vec{B}, \vec{B}'), \quad \text{if } (\vec{x}, \vec{G}, \vec{B}, \vec{B}') \in \{\mathcal{C}_t = \emptyset\}.$$

Arguing as in (3.33)-(3.34), it follows that

$$\|S_t(\vec{p}) - S_t(\vec{q})\|_{\text{TV}} \leq \mathbb{P}(\mathcal{C}_t \neq \emptyset).$$

The conclusion (4.12) now follows from Lemma 4.8. This ends the proof of Theorem 4.1. \square

5. Concluding Remarks and Open Questions

Our results leave open a number of interesting directions for further research, some of which we briefly mention here, along with some additional observations and extensions.

1. Can we prove exponential decay (even with an exponentially bad dependence on n) for the nonlinear Ising dynamics for *arbitrary* interactions \mathbf{J} , without the high-temperature condition? Recall that we did prove convergence for arbitrary \mathbf{J} in Theorem 1.1, but unlike its analogs for linear Markov chains our proof apparently gives no useful rate information (see Remark 2.5).
2. Can our high-temperature condition $\max_x \sum_{y \in V} |J_{xy}| \leq \delta_0$ be relaxed to accommodate an optimal value for δ_0 , in particular in the case of the complete graph (mean-field or Curie-Weiss model)? Here all interactions $J_{xy} = \frac{\beta}{n}$, and $\beta = 1$ marks the phase transition. Thus one might hope to sharpen our results to require only $\delta_0 < 1$ in this case. One might also hope to replace the ℓ_1 condition on \mathbf{J} by a spectral condition, as e.g. in the recent works [6, 18].
3. We have proved rapid convergence in total variation distance, a natural goal. However, one might naturally also ask whether one can prove sharp bounds on the rate of contraction of relative entropy, a more delicate question. This was done recently in the non-interacting (population genetics) case in [11, 12].
4. As mentioned in the introduction, our results for the Ising model generalize to any spin system with a constant number of spins at each vertex and bounded pairwise interactions, including the q -state Potts model, under an analogous Dobrushin-type condition (3.2). (In that condition, the sum is now over the maximum absolute values of all interactions involving any given site x .) This follows from the fact that, as can readily be checked, the representation of the measure $\gamma(\cdot \mid \sigma, \sigma')$ in terms of a two-spin system as described in Lemma 3.2 still holds in this more general setting, and beyond that point the rest of the analysis depends only on that representation.

5. We note that, in contrast to linear Markov chains (which can be simulated using a single point walking randomly around the state space), nonlinear dynamics do not immediately provide an efficient algorithm for sampling from the stationary distribution, even when the convergence time is short. This is a consequence of the fact that, in order to obtain a single sample from the time- t distribution $T_t(p)$, we would naively need to simulate the entire evolutionary tree of depth t , which begins with 2^t independent random samples from p at the leaves. (Indeed, Arora et al. [4] prove that simulating arbitrary reversible quadratic dynamics is in fact PSPACE-complete.) However, for both of the dynamics we consider here, our analysis does in fact also provide a polynomial time algorithm for sampling from (very close to) the stationary distribution $\mu_{J,h}$. This is immediately obvious for the nonlinear block dynamics: since Theorem 1.3 establishes that $t = O(\log n)$ steps are enough for convergence, the entire tree is of polynomial size and thus can be constructed in polynomial time. (A detail here is that each interaction in the tree requires sampling the set of sites Λ to be exchanged according to the distribution (3.11). However, this itself is a high-temperature Ising Gibbs measure, and thus can be sampled from efficiently by other means.) For the nonlinear Glauber dynamics, a polynomial time simulation is also possible once we observe from our analysis that the actual size of the master evolution for $S_t(p)$ is $O(t)$; since the master evolution is sufficient to construct our time- t sample, this can be done in $\text{poly}(n)$ time. The simulations in the previous paragraph leave much to be desired algorithmically, as they are rather unwieldy and, more significantly, do not retain the underlying pairwise interaction structure of the nonlinear dynamics. An interesting related question, for both processes, is whether a more explicit simulation is possible: namely, starting with a finite population of size $N = \text{poly}(n)$, each member of which is sampled independently from the initial distribution p , evolve this population in the obvious way (by carrying out interactions between randomly chosen pairs to construct the next-generation population of the same size). The question is whether the time- t population in this finite implementation is close to the true population $T_t(p)$, at least for modest times t —this is not obvious since unwanted correlations will arise due to the finite population size, though these will decrease with the population size N . (This implementation, which is actually a Markov chain on a very large state space, is known as the “Kac model” in kinetic theory [32], and the convergence to the true population is referred to as “propagation of chaos”.) In the non-interacting (population genetics) case, it was shown that a low-degree polynomial population size does in fact suffice [12, 43]. It would be interesting to see if this argument can be extended to the Ising model at high temperature; in particular, since the key insight in [43] comes from the fragmentation process, the question boils down to whether the same insight carries through in the presence of noise.

Funding Open access funding provided by Università degli Studi Roma Tre within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory

regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

1. Adleman, L., Gopalkrishnan, M., Huang, M.-D., Moisset, P., Reishus, D.: On the mathematics of the law of mass action. *A Systems Theoretic Approach to Systems and Synthetic Biology I: Models and System Characterizations*, pages 3–46, (2014)
2. Anderson, D., Cappelletti, D., Kim, J., Nguyen, T.: Tier structure of strongly endotactic reaction networks. *Stoch. Process. Appl.* **130**, 7218–7259 (2020)
3. Andrieu, C., Jasra, A., Doucet, A., Del Moral, P.: On nonlinear Markov chain Monte Carlo. *Bernoulli* **17**(3), 987–1014 (2011)
4. Arora, S., Rabani, Y., Vazirani, U.V.: Simulating quadratic dynamical systems is PSPACE-complete. *Proceedings of the 26th Annual ACM Symposium on the Theory of Computing (STOC)*, pages 459–467, (1994)
5. Athreya, K.B., Ney, P.E.: *Branching Processes*. Springer-Verlag, Berlin (1972)
6. Bauerschmidt, R., Dagallier, B.: Log-sobolev inequality for near critical ising models. *Commun. Pure Appl. Math.* **77**(4), 2568–2576 (2024)
7. Boltzmann, L.: The second law of thermodynamics. In *Populäre Schriften, Essay 3*, address to a formal meeting of the Imperial Academy of Science, 29 May, 1886; reprinted in *Ludwig Boltzmann, Theoretical Physics and Philosophical Problems*, S.G. Brush (Trans.). Reidel, Boston, (1974)
8. Budhiraja, A., Dupuis, P., Fischer, M., Ramanan, K.: Local stability of kolmogorov forward equations for finite state nonlinear markov processes. *Electron. J. Probab.* **20**(81), 1–30 (2015)
9. Butkovsky, O.: On ergodic properties of nonlinear Markov chains and stochastic McKean-Vlasov equations. *Theory Probab. Appl.* **58**(4), 661–674 (2014)
10. Caputo, P., Labbé, C., Lacoïn, H.: Cutoff phenomenon in nonlinear recombinations. [arXiv:2402.11396](https://arxiv.org/abs/2402.11396), (2024)
11. Caputo, P., Sinclair, A.: Entropy production in nonlinear recombination models. *Bernoulli* **24**(4B), 3246–3282 (2018)
12. Caputo, P., Parisi, D.: Nonlinear recombinations and generalized random transpositions. *Annales Henri Lebesgue* **7** (2024)
13. Carlen, E.A., Carvalho, M.C., Le Roux, J., Loss, M., Villani, C.: Entropy and chaos in the Kac model. *Kinet. Relat. Models* **3**(1), 85–122 (2010)
14. Carrillo, J.A., McCann, R.J., Villani, C.: Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates. *Revista Matemática Iberoamericana* **19**(3), 971–1018 (2003)
15. Cattiaux, P., Guillin, A., Malrieu, F.: Probabilistic approach for granular media equations in the non-uniformly convex case. *Probab. Theory Relat. Fields* **140**(1), 19–40 (2008)
16. Chaintron, L.-P., Diez, A.: Propagation of chaos: a review of models, methods and applications: I: models and methods. *Kinetic Relat. Models* **15**, 895–1015 (2022)
17. Chaintron, L.-P., Diez, A.: Propagation of chaos: a review of models, methods and applications: II. applications. *Kinetic Relat. Models* **15**, 1017–1173 (2022)
18. Chen, Y., Eldan, R.: Localization schemes: a framework for proving mixing bounds for markov chains. In *2022 IEEE 63rd Annual symposium on foundations of computer science (FOCS)*, pages 110–122 (2022)
19. Craciun, G., Nazarov, F., Pantea, C.: Persistence and permanence of mass-action and power-law dynamical systems. *SIAM J. Appl. Math.* **73**, 305–329 (2013)
20. Dolera, E., Gabetta, E., Regazzini, E.: Reaching the best possible rate of convergence to equilibrium for solutions of Kac's equation via central limit theorem. *Ann. Appl. Probab.* **19**(1), 186–209 (2009)
21. Doty, D., Zhu, S.: Computational complexity of atomic chemical reaction networks. *Nat. Comput.* **17**, 677–691 (2018)
22. Erbar, M., Fathi, M., Schlichting, A.: Entropic curvature and convergence to equilibrium for mean-field dynamics on discrete spaces. *ALEA* **17**, 445–471 (2020)
23. Feinberg, M.: *Foundations of Chemical Reaction Network Theory*. Springer, Switzerland (2019)
24. Friedli, S., Velenik, Y.: *Statistical Mechanics of Lattice Systems: A Concrete Mathematical Introduction*. Cambridge University Press (2017)
25. Goldberg, D.E.: *Genetic Algorithms in Search. Optimization and Machine Learning*. Addison-Wesley, Reading MA (1989)

26. Gopalkrishnan, M., Miller, E., Shiu, A.: A geometric approach to the global attractor conjecture. *SIAM J. Appl. Dyn. Syst.* **13**, 758–797 (2014)
27. Hardy, G.H.: Mendelian proportions in a mixed population. *Science* **28**, 49–50 (1908)
28. Horn, F.: The dynamics of open reaction systems. In *Mathematical Aspects of Chemical and Biochemical Problems and Quantum Chemistry*, SIAM-AMS Proceedings **8**, 125–137 (1974)
29. Horn, F.J.M., Jackson, R.: General mass action kinetics. *Arch. Rational Mech. Anal.* **49**, 81–116 (1972)
30. Houdayer, J.: A cluster Monte Carlo algorithm for 2-dimensional spin glasses. *Eur. Phys. J. B-Condens. Matter Complex Syst.* **21**, 479–484 (2001)
31. Hukushima, K., Nemoto, K.: Exchange Monte Carlo method and application to spin glass simulations. *J. Phys. Soc. Jpn.* **65**, 1604–1608 (1996)
32. Kac, M.: Foundations of kinetic theory. *Proc. Third Berkeley Symp. Math. Stat. Probability* **3**, 171–197 (1956)
33. Kolokoltsov, V.N.: *Nonlinear Markov Processes and Kinetic Equations*. Cambridge University Press (2010)
34. Levin, D., Peres, Y., Wilmer, E.: *Markov Chains and Mixing Times*. American Mathematical Soc., (2009)
35. Lubetzky, E., Sly, A.: Universality of cutoff for the Ising model. *Ann. Probab.* **45**(6A), 3664–3696 (2017)
36. McKean, J. Henry., P.: A class of Markov processes associated with nonlinear parabolic equations. *Proc. Natl. Acad. Sci.* **56**(6), 1907–1911 (1966)
37. Mischler, S., Mouhot, C.: Kac’s program in kinetic theory. *Invent. Math.* **193**(1), 1–147 (2013)
38. Mitchell, M.: *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA (1998)
39. Muzychka, S.A., Vaninsky, K.L.: A class of nonlinear random walks related to the Ornstein-Uhlenbeck process. [arXiv:1107.0850](https://arxiv.org/abs/1107.0850), (2011)
40. Nguyen, H.C., Zecchina, R., Berg, J.: Inverse statistical problems: from the inverse Ising problem to data science. *Adv. Phys.* **66**, 197–261 (2017)
41. Pantea, C.: On the persistence and global stability of mass action systems. *SIAM J. Math. Anal.* **44**, 1636–1673 (2012)
42. Propp, J.G., Wilson, D.B.: Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Struct. Algorithms* **9**(1–2), 223–252 (1996)
43. Rabani, Y., Rabinovich, Y., Sinclair, A.: A computational view of population genetics. *Random Struct. Algorithms* **12**(4), 313–334 (1998)
44. Rabinovich, Y., Sinclair, A., Wigderson, A.: Quadratic dynamical systems. In *Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 304–313. IEEE, (1992)
45. Shchegolev, A., Veretennikov, A.: On convergence rate bounds for a class of nonlinear Markov chains. *Preprint* [arXiv:2209.12834](https://arxiv.org/abs/2209.12834), (2022)
46. Siegel, D., MacLean, D.: Global stability of complex balanced mechanisms. *J. Math. Chem.* **27**, 89–110 (2000)
47. Singh, M., Vishnoi, N.: Entropy, optimization and counting. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*, pages 50–59, (2014)
48. Swendsen, R., Wang, J.-S.: Replica Monte Carlo simulation of spin-glasses. *Phys. Rev. Lett.* **57**, 2607–2609 (1986)
49. Sznitman, Alain-Sol: Topics in propagation of chaos. In *École d’Été de Probabilités de Saint-Flour XIX, 1989*, pages 165–251. Springer, (1991)
50. van den Berg, J., Gandolfi, A.: BK-type inequalities and generalized random-cluster representations. *Probab. Theory Relat. Fields* **157**(1–2), 157–181 (2013)
51. Villani, C.: Entropy production and convergence to equilibrium. In *Entropy Methods for the Boltzmann Equation*, volume 1916 of *Lecture Notes in Math.*, pages 1–70. Springer, Berlin, (2008)
52. Villani, C.: A review of mathematical topics in collisional kinetic theory. *Handbook Math. Fluid Dyn.* **1**, 71–305 (2002)
53. Vuckovic, J.: Nonlinear MCMC for Bayesian machine learning. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, (2022)
54. Weinberg, W.: Über den Nachweis der Vererbung beim Menschen. *Jahreshefte des Vereins für vaterländische Naturkunde in Württemberg* **64**, 368–382 (1908)
55. Weitz, D.: Combinatorial criteria for uniqueness of Gibbs measures. *Random Struct. Algorithms* **27**, 445–475 (2005)