1
Introduction

The subject of this book lies at the confluence of two major currents in contemporary science: phase transitions and far-from-equilibrium phenomena. It is a subject that continues to attract scientists, not only for its novelty and technical challenge, but because it promises to illuminate some fundamental questions about open many-body systems, be they in the physical, the biological, or the social realm. For example, how do systems composed of many simple, interacting units develop qualitatively new and complex kinds of organization? What constraints can statistical physics place on their evolution?

Nature, both living and inert, presents countless examples of *nonequilibrium* many-particle systems. Their simplest condition — a nonequilibrium steady state — involves a constant flux of matter, energy, or some other quantity (de Groot & Mazur 1984). In general, the state of a nonequilibrium system is not determined solely by external constraints, but depends upon its *history* as well. As the control parameters (temperature or potential gradients, or reactant feed rates, for instance) are varied, a steady state may become unstable and be replaced by another (or, perhaps, by a periodic or chaotic state). Nonequilibrium instabilities are attended by ordering phenomena analogous to those of equilibrium statistical mechanics; one may therefore speak of *nonequilibrium phase transitions* (Nicolis & Prigogine 1977, Haken 1978, 1983, Graham 1981, Cross & Hohenberg 1993). Some examples are the sudden onset of convection in a fluid heated from below (Normand, Pomeau, & Velarde 1977), switching between high- and low-reactivity regimes in an open chemical reactor (Field & Burger 1985, Gray & Scott 1990), and the transition to high field intensity as a laser is pumped beyond a threshold value (Graham & Haken 1971).

Phase transitions and critical phenomena have captivated statistical physicists for many decades (see, for example, Stanley (1971), and the Domb & Green (1972–6) and Domb & Lebowitz (1977–95) series). Much theoretical progress has resulted from the parallel application of varied approaches, including exact solutions, mean-field theories, computer simulations, series expansions, and renormalization group methods, to simple lattice models and their continuum analogs. It is now possible to understand the phase diagram of many systems using these methods, and to fathom the remarkable universality of critical behavior (Ma 1976, Amit 1984, Zinn-Justin 1990).

The *dynamics* of phase transitions and critical points presents new challenges, since macroscopic properties are no longer given by averages over a known, stationary probability distribution. The same applies to systems that are kept out of equilibrium. In this book we illustrate, by means of simple examples, some of the diversity of nonequilibrium phase transitions. The examples are sufficiently detailed to show that 'nonequilibrium phase transitions' represents not a superficial resemblance to equilibrium phenomena, but a precise use of the terms 'phase' and 'phase transition' in their statistical mechanics sense. Equilibrium phase diagrams can be determined on the basis of the free energy; out of equilibrium, the free energy is not defined, in general. However, even in its absence, we can recognize a *phase* of a many-particle system from well-defined, reproducible relations between its macroscopic properties (or, more formally, the many-particle distribution functions) and the parameters governing its dynamics. A phase transition is characterized by a singular dependence of these attributes upon the control parameters. From this vantage, equilibrium is a special case, in which the dynamics happens to be derivable from an energy function, thereby permitting an analysis with no reference to time. A key point of difference between equilibrium and nonequilibrium statistical mechanics is that whereas in the former case the stationary probability distribution is known, out of equilibrium one must actually find the time-independent solution(s) of the master equation for the process. This is a formidable task and can only be carried out approximately for most models.

We restrict our attention to lattice models and, in particular, to lattice Markov processes or *interacting particle systems* (Griffeath 1979, Liggett 1985, Durrett 1988, Konno 1994). A well-known disadvantage of lattice models is that they are usually too crude to be directly comparable with experiment. In fact, if one is interested in predicting a nonequilibrium phase diagram, it is better, for cases in which fluctuations are of minor significance, to employ a macroscopic description, i.e., a set of (deterministic) partial differential equations. This approach finds success in applications to hydrodynamics (convection, couette flow), oscillations and waves in surface catalysis, and chemical reactions (Field & Burger 1985, Bär *et al.* 1994, Cross & Hohenberg 1993). If we ignore these avenues in favor of lattice models, it is largely because macroscopic descriptions hold little surprise in the way of criticality; mean-field behavior is implicit at this level. The range of critical phenomena exhibited by the models considered here, by contrast, is at least as interesting as in equilibrium. Another reason for focusing on lattice models is that as statistical physicists, we prefer to see macroscopic behavior emerge from interactions among a multitude of simple units, rather than to adopt a macroscopic picture as our starting point. Clearly, a microscopic approach becomes essential for describing systems in which fluctuations are significant. Continuum models of nonequilibrium phenomena that do incorporate fluctuations (Langevin equations or stochastic partial differential equations) are also being developed. These approaches generally pose even greater conceptual and computational challenges than the lattice models we focus on in this book.

One can distinguish three broad categories of nonequilibrium problems. One class comprises Hamiltonian systems out of equilibrium by virtue of their preparation; one is then interested in how the system approaches equilibrium. This subject, which subsumes such active areas of investigation as critical dynamics, nucleation theory, and spinodal decomposition, lies outside the scope of this book. Our focus is rather on models that are intrinsically out of equilibrium — models that violate the principle of detailed balance. It is useful to draw a further distinction between perturbations of an equilibrium model by a nonequilibrium 'driving force,' and models that have no equilibrium counterpart. An example of the former is the driven diffusive system considered in chapters 2 and 3. When the 'electric field' that biases hopping vanishes, we recover the familiar lattice gas. In the 'perturbed equilibrium' models one may still speak of energy and temperature. The transition rates may obey a *local* detailed-balance condition, but cannot be derived from a potential energy function. Or the rates may represent competing processes, proceeding at different temperatures, as illustrated in chapters 4 and 8. The catalytic and population models considered in chapters 5, 6, and 9 have no equilibrium analog; their dynamics involves events that are strictly irreversible. The concepts of energy and temperature are not pertinent to such models.

Many of the methods brought to bear on equilibrium lattice models — mean-field theories, Monte Carlo simulations (coupled with finite-size scaling ideas), series expansions, and renormalization group methods — have been applied successfully to nonequilibrium systems. All of these methods are described in some detail in this work, with the exception of renormalization group methods, which would seem to require a book in itself.

1.1 Two simple examples

Out of equilibrium, the well-known prohibition against phase transitions in one-dimensional systems with short-range interactions is lifted. There are many examples of nonequilibrium phase transitions in one space dimension (Privman 1997), and one might have hoped that some would permit an exact solution. In fact, very few exactly soluble models with nonequilibrium phase transitions are known. To provide an easily-worked example, we study instead the *branching process:* a 'zero-dimensional' model, or more precisely, one lacking spatial structure.

Consider a population $n(t) \geq 0$ of individuals that give birth at rate λ, and die at unit rate. The individuals might represent organisms in a colony, with a population well below the carrying capacity, neutrons in a chain reaction, or photons in a lasing medium. $n(t)$ is a Markov process with transition rates

$$c(n \to n+1) = \lambda n \tag{1.1}$$

and

$$c(n \to n-1) = n. \tag{1.2}$$

Note that $n=0$ is absorbing; if all individuals die the population remains zero at future times. The probability, $P_n(t)$, of having exactly n individuals at time t is governed by the master equation,

$$\frac{\mathrm{d}P_n}{\mathrm{d}t} = \lambda(n-1)P_{n-1} + (n+1)P_{n+1} - (1+\lambda)nP_n. \tag{1.3}$$

We solve this using the generating function

$$g(x,t) = \sum_{n=0}^{\infty} x^n P_n(t), \tag{1.4}$$

which satisfies

$$\frac{\partial g}{\partial t} = (1-x)(1-\lambda x)\frac{\partial g}{\partial x}, \tag{1.5}$$

subject to the boundary condition $g(1,t) = 1$, which reflects normalization. If we assume a single individual at time zero, then $g(x,0) = x$. For $\lambda \neq 1$ the solution is readily found to be

$$g(x,t) = \frac{(1-\lambda x)\mathrm{e}^{(1-\lambda)t} - (1-x)}{(1-\lambda x)\mathrm{e}^{(1-\lambda)t} - \lambda(1-x)}. \tag{1.6}$$

Setting $x = 0$ we find the extinction probability, $P_0(t)$, and hence the *survival probability*

$$P(t) = 1 - P_0(t) = \frac{\lambda - 1}{\lambda - \mathrm{e}^{(1-\lambda)t}}. \tag{1.7}$$

For $\lambda < 1$, $P(t)$ decays exponentially, while for $\lambda > 1$ it approaches a nonzero value. The ultimate survival probability is

$$P_\infty \equiv \lim_{t\to\infty} P(t) = \begin{cases} 0 & \lambda < 1 \\ 1 - \lambda^{-1} & \lambda > 1 \end{cases} \tag{1.8}$$

which shows that $\lambda = 1$ is a critical value marking the boundary between possible survival and certain extinction. Thus P_∞ (the *order parameter* for this model) has a singular dependence upon λ, justifying our labeling this a phase transition. For $\lambda = 1$ the generating function is

$$g(x,t)|_{\lambda=1} = \frac{1+(1-x)(t-1)}{1+(1-x)t}, \tag{1.9}$$

from which we find

$$P(t) = \frac{1}{1+t}, \quad \lambda = 1. \tag{1.10}$$

Thus the relaxation time for the survival probability diverges $\propto |1-\lambda|^{-1}$ as $\lambda \to 1$, and when $\lambda = 1$ relaxation follows a power law rather than an exponential decay. These features — a sharp boundary between extinction and survival, a diverging relaxation time, and power-law relaxation — are typical of critical points, and will be found in more complex and interesting models.

The model considered here does not allow for a nontrivial steady state. It is easy to show that the mean population size is $\langle n \rangle = n(t=0)e^{(\lambda-1)t}$, hence for $\lambda > 1$ the population grows out of all bounds.[1] To prevent this we require a saturation term (the birth rate should decline for large n), which renders the analysis much more difficult. We return to this issue in chapter 6, which is devoted to the *contact process*, a birth-and-death model of particles on a lattice, with offspring appearing at nearest-neighbor *vacant* sites. This prevents exponential population growth: since sites may not be doubly occupied, the density can never exceed unity, and the model possesses a nontrivial or *active* stationary state for large enough λ.

There is one variant of the contact process for which we can obtain exact results.[2] Consider a one-dimensional lattice, with the rule that a vacant site adjacent to an occupied one becomes occupied at rate λ, while an occupied site adjacent to a vacant one becomes vacant at unit rate. Suppose we start with only the origin occupied. Then at later times, the system is either in the absorbing state (no particles) or it consists of a string of n occupied sites, with no gaps. The number of particles, $n(t)$, is a continuous-time lattice random walk, starting from $n = 1$, with an absorbing boundary at $n = 0$, and with a bias to the right $\propto \lambda - 1$. That is, for $\lambda \leqslant 1$ n must eventually hit zero, while for $\lambda > 1$ there is a nonzero probability of survival as $t \to \infty$. Well-known results for random

[1] The situation is reminiscent of the *Gaussian approximation* to Ising/ϕ^4 field theory. Neglect of the term $\propto \phi^4$ yields a soluble model with pathological low-temperature behavior; see Binney *et al.* (1992).

[2] This is the continuous-time version of so-called *compact directed percolation*; see Essam (1989).

walks (Feller 1957, Barber & Ninham 1970) imply that $P_\infty \propto \lambda - 1$, and, for $\lambda = \lambda_c = 1$, $P(t) \sim t^{-1/2}$, while the mean-square population over *surviving* trials $\langle n^2 \rangle_{surv} \sim t$. Since the particles are arrayed in a compact cluster, $R^2(t)$, the mean-square distance of particles from the origin, also grows $\sim t$. The asymptotic time-dependence of $P(t)$, $\langle n \rangle$, and $R^2(t)$ in critical systems is discussed extensively in chapters 5, 6, and 9. The present examples provide some of the rare instances in which the power laws governing this evolution are known exactly.

Next we describe a simple model that illustrates several themes associated with nonequilibrium steady states: dynamic competition, spatial structure (pattern formation), and anisotropy. It is closely related to the physically motivated driven lattice gas analyzed in detail in chapters 2 and 3. Consider a simple cubic (sc) lattice in two or more dimensions, with toroidal boundary conditions, and a fraction n of its sites occupied by particles, the rest vacant. The evolution proceeds via nearest-neighbor (NN) particle–hole exchanges. With probability q the exchange is in the *longitudinal* direction (defined by the unit vector $\hat{\mathbf{x}}$) and involves a driving field; with probability $1 - q$ it takes one of the transverse directions, via a *thermal mechanism.* The field introduces a bias: particle displacements of $\hat{\mathbf{x}}$ are accepted with probability p, while displacements of $-\hat{\mathbf{x}}$ are accepted with probability $1 - p$. (In other words, along this direction we have an *asymmetric exclusion process*.) In contrast with longitudinal exchanges, which do not involve NN interactions, exchanges in a transverse direction are accepted with probability b if the second neighbor along the jump direction is occupied, and with probability $1 - b$ if it is vacant. (Table 1.1 gives the rates for a two-dimensional system with equal *a priori* probabilities for longitudinal and transverse jumps.) These processes mimic the effects of field-driven motion (for $p \neq \frac{1}{2}$), and, for $b > \frac{1}{2}$, of a tendency toward cluster formation, but in a manner that cannot be reconciled with a potential energy function. As confirmed in simulations, $b - \frac{1}{2}$ is a temperature-like variable, analogous to the inverse temperature β, while $p - \frac{1}{2}$ represents a longitudinal driving field. The latter has no equilibrium analog, but given the periodic boundaries, leads rather to a nonequilibrium steady state with a longitudinal current.

This simple system exhibits the great variety and some of the difficulties characterizing the phenomena studied in this book. Its behavior is best illustrated by 'snapshots' of typical configurations, as in figure 1.1, which is for b sufficiently large that the system segregates into a particle-rich and a particle-poor phase; it depicts the process of phase separation from a random initial configuration. (This simulation employs $p = 1$, but we find similar results for other values, even $p = \frac{1}{2}$.) Simulations suggest that phase separation occurs discontinuously, by a series of 'avalanches.' Another interesting observation is that the pair correlation function exhibits self-

Table 1.1. The rate for a two-dimensional version of the lattice gas with $q = \frac{1}{2}$ (i.e., no *a priori* bias), assuming that the preferred direction, $+\hat{\mathbf{x}}$, is vertical upwards. The symbols • and ∘ stand for occupied and vacant sites, respectively (Marro & Achahbar 1998).

Process	Rate
$\bullet\circ\bullet \to \circ\bullet\bullet$	b
$\bullet\circ\bullet \to \bullet\bullet\circ$	b
$\bullet\circ\circ \to \circ\bullet\circ$	$1-b$
$\circ\circ\bullet \to \circ\bullet\circ$	$1-b$
$\begin{smallmatrix}\circ\\ \bullet\end{smallmatrix} \to \begin{smallmatrix}\bullet\\ \circ\end{smallmatrix}$	p
$\begin{smallmatrix}\bullet\\ \circ\end{smallmatrix} \to \begin{smallmatrix}\circ\\ \bullet\end{smallmatrix}$	$1-p$

similarity or time-scale invariance if one scales time by the mean width of the strips. Figure 1.2 illustrates the kind of order found in this system, namely, anisotropic segregation at large b, and a linear interface for any value of p. (For $b < \frac{1}{2}$ the tendency appears to be towards chess-board configurations, as in an antiferromagnet, independent of p.) We refer the reader to Marro & Achahbar (1998) for further details.

1.2 Perspective

The examples of the previous section give some of the flavor of the models we consider in this book. The population model exhibits a phase transition between an active state (survival) and a kind of trap — an empty state with no further evolution. This transition has no equilibrium analog. The second example displays phase separation at a particular value of a temperature-like variable, like the equilibrium lattice gas, but in a non-Hamiltonian model with highly anisotropic dynamics (also — unlike for equilibrium — the resulting two phases, *liquid* and *gas*, are not symmetric here, in general).

Our point of view in this book is, quite naturally, strongly influenced by our awareness of the theory of (equilibrium) phase transitions and critical phenomena. The latter appears sufficiently powerful and broadly applicable to guide at least our initial questions about nonequilibrium models. Since a number of key ideas from equilibrium theory provide touchstones for our discussion, it is well to mention them briefly.

The central result in the modern theory of critical phenomena is *uni-*

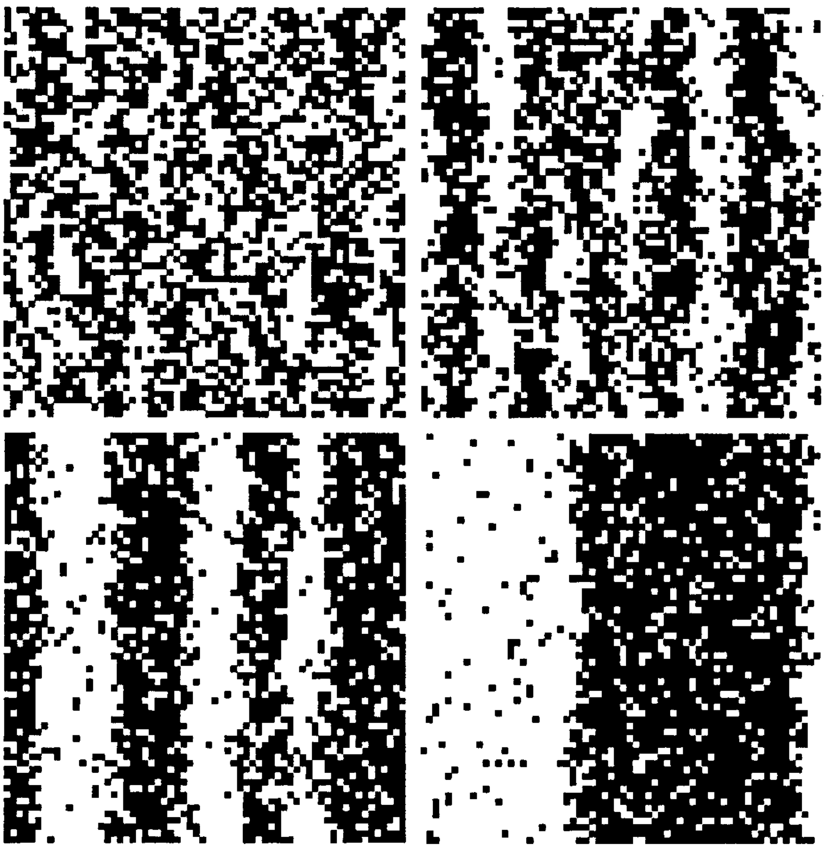

Fig. 1.1. Typical configurations of a 64×64 lattice illustrating relaxation of the system with rates in table 1.1, for $n = \frac{1}{2}$, $b = 0.9$, and $p = 1$. The longitudinal direction is vertical. Top left: $t = 10^2$; top right: $t = 10^3$; bottom left: $t = 10^4$; bottom right: $t = 10^6$. (Time in units of Monte Carlo steps per site.)

versality: singularities in the vicinity of a critical point are determined by a small set of basic features — spatial dimensionality, dimensionality and symmetry of the order parameter, and whether the interactions are long- or short-range (Stanley 1971, Ma 1976, Fisher 1984, Amit 1984, Zinn-Justin 1990). The insensitivity to details of molecular structure or interactions is evident when we note that the Ising model, simple materials such as argon or carbon dioxide, and binary liquid mixtures all exhibit the same critical behavior. When we turn to critical *dynamics*, conservation laws enter as another possible determinant (Hohenberg & Halperin 1977).

Universality reflects the existence of a diverging correlation length ξ in a system at its critical point. For example, if $\epsilon = (T - T_c)/T_c$ (the *reduced temperature*) then in the absence of an external magnetic field we expect

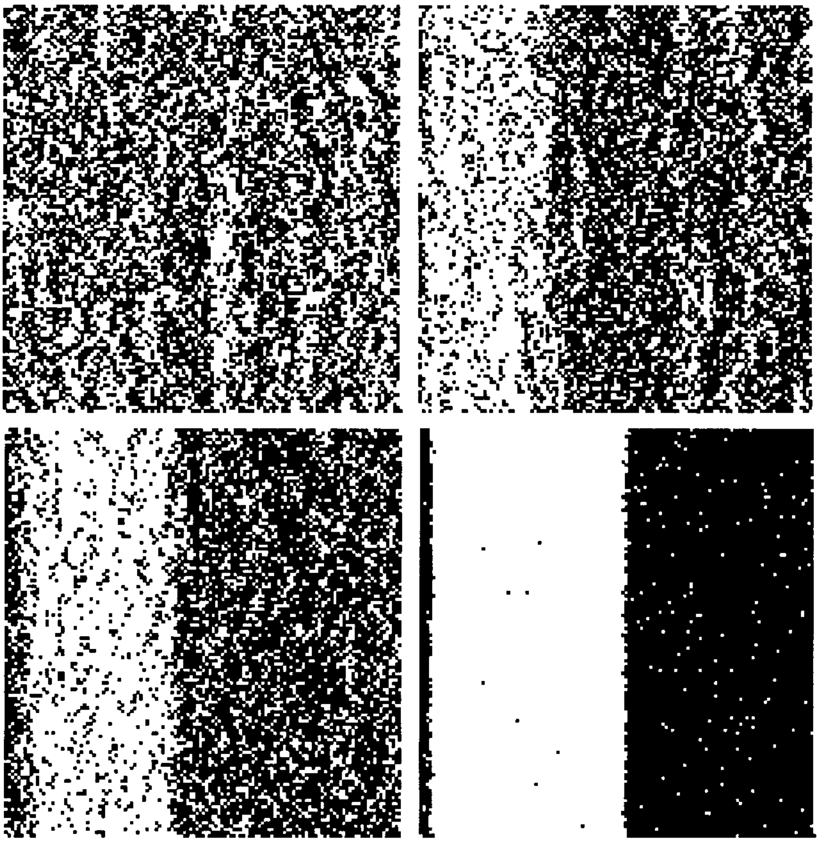

Fig. 1.2. Typical configurations in the stationary regime of the system with rates in table 1.1, for a 128×128 lattice with $n = \frac{1}{2}$ and (from left to right) $b = 0.84$ and 0.865 (top), and 0.877 and 0.98 (bottom). The configurations shown here, for $p = \frac{1}{2}$, are indistinguishable from those with different p but the same value of $\epsilon \equiv 1 - b/b_c(p)$, where $b_c(p)$ denotes the critical line.

the correlation length of a ferromagnet to diverge as $\xi \sim \epsilon^{-\nu}$. Assigning a model to a particular *universality class* fixes certain quantities, for example critical exponents such as ν. Similarly, the relaxation time for fluctuations diverges at the critical point. Thus we expect static correlation functions to depend (asymptotically) on distance r through the ratio r/ξ, and dynamic correlations to depend on t through the ratio t/τ, where τ is a relaxation time.

Another powerful result in critical phenomena is a general scheme for classifying thermodynamic variables. Thus in the simplest cases we expect to find just two relevant parameters, one ‘temperature-like’ (conjugate to energy), the other ‘field-like’ (conjugate to magnetization or density) that mark the nearness of the system to its critical point. Associated with this is *scaling* (Widom 1965), which asserts that, near criticality, thermodynamic properties depend not on these variables separately, but only in a particular ratio. This scaling hypothesis leads to a set of *scaling relations* amongst critical exponents.

The singularities marking phase transitions and critical points only emerge in the infinite-size limit; in finite systems they are rounded off. In the vicinity of the critical point, where ξ is large, intensive properties depend strongly on system size. Finite-size scaling theory (Fisher 1971, Fisher & Barber 1972, Cardy 1988, Privman 1990) appeals to the notion that near the critical point, the dependence on system size L should only involve the ratio L/ξ. As long as $L >> \xi$ the system consists of many uncorrelated regions, and intensive properties should be independent of system size.

Finally we note that a detailed understanding of universality, and many other aspects of critical phenomena, rests on applying renormalization group methods to continuum or *field theory* descriptions that capture the essential features of the original lattice or molecular-level models (Ma 1976, Zinn-Justin 1990, Binney *et al.* 1992). While derivation of the appropriate field theory from first principles is a subtle business, it can often be constructed on the basis of symmetry considerations.

How much of this framework, familiar from equilibrium critical phenomena, applies to nonequilibrium phase transitions? This is one of the main open questions motivating the studies described in this book. The results thus far suggest that most if not all of these ideas will retain their validity in the context of nonequilibrium lattice models. (What is missing, in a sense, is *thermodynamics*![3]) We will also see that many of the analytical and numerical methods used to study equilibrium phase transitions remain useful here.

What, then, is so different about phase transitions out of equilibrium? In connection with lifting the detailed balance condition there appear a number of new possibilities: phase transitions to an absorbing state, transitions in one dimension (even in models with short-range interactions), novel spatial structures, highly dependent upon the *history* of the system, and unexpected interfacial properties. One encounters an enormous richness of steady states as one varies the dynamical rules. Switching between different forms of the spin flip rate, or varying the relative frequency of exchanges in different directions — changes that have no effect on the equilibrium state — can produce completely different phase diagrams in

[3] Indeed, it has been appreciated for some time that in percolation, which possesses neither a Hamiltonian nor dynamics, a phase transition arises purely from the interplay of statistics and geometry; see Stauffer & Aharony (1992).