May 14, 1997

A New Look at Broken Ergodicity

D.L. Stein

Department of Physics
University of Arizona
Tucson, AZ 85721

C.M. Newman

Courant Institute of Mathematical Sciences
New York University
New York, NY 10012

ABSTRACT:

We study the nature and mechanisms of broken ergodicity (BE) in specific random walk models corresponding to diffusion on random potential surfaces, in both one and high dimension. Using both rigorous results and nonrigorous methods, we confirm several aspects of the standard BE picture and show that others apply in one dimension, but need to be modified in higher dimensions. These latter aspects include the notions that at fixed temperature confining barriers increase logarithmically with time, that "components" are necessarily bounded regions of state space which depend on the observational timescale, and that the system continually revisits previously traversed regions of state space. We examine our results in the context of several experiments, and discuss some implications of our results for the dynamics of disordered and/or complex systems.

1. Introduction.

When a system has many metastable states, it may become trapped for long times in some subset of its total state space, making it difficult to compare experimental results with calculations based on the usual Gibbs formalism. A viewpoint commonly called "broken ergodicity" (BE) has evolved to serve as a qualitative guide for the understanding of some of the dynamical and thermal properties of these systems. This has been extremely useful in several respects, but we are still hampered by the lack of a real theory.

Much of the problem is the difficulty of characterizing the nature of metastability in real systems. As a result, the standard picture of BE which has emerged (to be described below) is based largely on intuition and simple pictures of what these state spaces may look like. All basically involve diffusion of a particle (the system) on a rugged landscape, which may or may not possess correlations. These pictures can all be described as diffusion in a strongly inhomogeneous environment.

While many of the results obtained in this way are compelling (and, as we will discuss below, almost certainly correct in a wide variety of situations), progress has been slow, at least partially because of the lack of specific models to test these ideas on. In this paper we will attempt to do just that; we will examine some simple, well-known models and see how broken ergodicity arises in them. These models are representative of uncorrelated random potentials. Reasoning based on random walks on such potentials has guided much of the thinking about how BE operates in disordered systems.[1] We will not address in this paper the question of the accuracy of such assumptions; *i.e.,* whether random walks on rugged landscapes are useful for modelling dynamics of some disordered systems. Our only goal here is to introduce clear, well-defined models and to study their long-time behavior in the strongly inhomogeneous limit.

The analysis will be based on some rigorous results obtained in an earlier paper,[2] hereafter referred to as I, and some nonrigorous results obtained in another,[3] hereafter referred to as II. We will find that while many of the central ideas of standard BE apply to these models, there are some surprising deviations from important elements of the conventional picture. We will find that this is at least partially due to the fact that while all workers in the field recognize that the relevant state spaces for physical systems are high-dimensional, much of the intuition about BE is nevertheless based on what are ultimately one-dimensional pictures. We will see explicitly how the presence of many dimensions considerably changes the standard analysis.

The paper is organized as follows: In Section 2, we review some of the basic features of BE pertinent to the analysis contained below. In Section 3, we introduce two simple models of a random walk in a random environment (RWRE), and review our earlier results within the context of these models. In Section 4, we analyze the behavior of broken ergodicity in these models, in both one and high dimensions. In Section 5, we discuss and summarize these results, and make a few brief remarks about experiment.

2. Broken Ergodicity.

Because the phenomenon of broken ergodicity has been discussed at great length in the literature, we here review only those aspects of it which are relevant for the cases under consideration. The importance of non-ergodicity in disordered systems, particularly spin glasses, was emphasized early on by Anderson[4,5]. Early analyses and applications were given by Jäckle,[6] Palmer[7,8], and van Enter and van Hemmen[9]. The presentation by Palmer is especially comprehensive and accessible; most of the discussion in this section follows his treatment. We are concerned here only with some of the central ideas of BE; for a complete overview, we refer the reader to the above papers.

We are primarily interested in cases where ergodicity is broken because the observational timescale ($\tau_{\rm obs}$) falls within a continuum of relaxational or equilibrational timescales intrinsic to the system, as is commonly believed to occur for glasses and spin glasses. (See Ref. 7 and Ref. 9 for other examples, including the more familiar situation of broken symmetry.) This may occur either in the presence or absence of a phase transition. The former is typically indicated by the state space breaking up into two or more disjoint components, separated by free energy barriers which diverge in the thermodynamic limit. Broken ergodicity can and does occur, however, when the system possesses *metastable* states surrounded by *finite* free energy barriers. Because the typical timescale for escape from a metastable state grows exponentially with the barrier, these need not be large for ergodicity to be broken on laboratory timescales.

The central idea is that state space can be decomposed into *components* which are not necessarily intrinsic to the system, but rather depend on the timescale. Components are defined by the probability of confinement on some timescale τ: if the system is in a given component at time 0, then the probability that it has not escaped from the component by time τ is greater than some specified (fixed) probability p_0.[7] Clearly, the definition of component depends both on the specified timescale τ and the probability p_0. It is also assumed that on the same timescale, the system is ergodic *within* the component; *i.e.,* the system visits a representative sampling of states within the component so that the state space average equals the time average, so long as one confines the averaging to states within the component.[7]

What is the confinement mechanism? We are interested here in *structural* confinement mechanisms rather than dynamical (*i.e.,* the existence of possible constants of the motion.) In the former case, the system is confined to a component because the smallest free energy barrier that must be surmounted in order to escape corresponds to an escape time large compared to the observational time. The standard picture envisions a very mountainous terrain, with a series of isolated lakes and puddles in various valleys. The “water level” corresponds to the largest free energy scale which the system can sample on a given temperature and time scale. If temperature is held fixed, and time is allowed to increase, the water level steadily rises, with lakes merging into bigger lakes and bigger lakes into oceans, leading to a hierarchical merging of components[10,7,8,11,12]; see Fig. 1. (One can arrive at the same picture by fixing time t and letting temperature T increase; the height of the water level scales as $T \log t$.)

Until the system surmounts the highest barriers, ergodicity remains broken; as soon as the system surmounts some free energy barrier, it finds itself in a larger component which is confined by higher free energy barriers.[7,8,11] So the confining free energy barriers which the system must surmount, at fixed temperature, increase logarithmically with the time. Also, because the system is now ergodic within the larger component, it will continually revisit the previous smaller one, which is now a subset of the portion of state space which it currently explores.

We will now examine these *ansätze* in two specific models, both of which possess a continuum of free energy barriers — and therefore a continuum of relaxational timescales. We will find that the above picture needs to be modified in several respects: while it precisely describes a *one*-dimensional version of our models, there are important differences in higher dimensions.

These high-dimensional models are indeed the relevant ones, since within the context of BE one is usually referring to the evolution of a system in some high-dimensional state space. As in I, one often models this state space as some graph $\mathcal{G}$. The vertices of the graph correspond to the states themselves, and the edges correspond to transitional paths between pairs of states. The dimensionality of the graph scales with N, the number of degrees of freedom in the system. We will follow this general procedure here.

We begin with a description of our dynamical models, which are just specific examples of a much studied process – the random walk in a random environment.[13−16]

3. Inhomogeneous Random Walk as Invasion Percolation.

In this section we review the results of earlier work. We consider two different dynamical models within the overall context of the RWRE:

1) Model A — "edge" model. Here we consider a graph $\mathcal{G}$ in which the sites correspond to states and the (nondirected) edges to the dynamical pathways which connect them. For specificity, we take $\mathcal{G}$ to be the lattice Z^d, although this is unnecessary for our results.[17] We assign nonnegative, independent, identically distributed random variables to each of the edges; these represent the energy barriers which must be surmounted in order to travel between pairs of sites connected via the edges. (We assume that the distribution of these variables is continuous so that all barriers have distinct energies.) If $W_{xy} = W_{yx}$ is the value assigned the edge connecting sites x and y, then the rate to travel *in either direction* between x and y is taken to scale with inverse temperature β as

$$r_{xy}(\beta) \sim \exp[-\beta W_{xy}] \quad . \tag{3.1}$$

2) Model B — "site" model. Here we assign random variables to *both* sites and edges of $\mathcal{G} = Z^d$. Because we wish to view each site as corresponding to a locally stable state (*i.e.,* as the minimum energy configuration within a "valley"), and each edge as again corresponding to an energy barrier, the values assigned to sites and edges cannot be identically distributed — both sites touching an edge must have lower assigned energy values than that of the edge itself. (As in model A, we simplify matters by choosing each distinct edge to connect a single pair of sites.) A simple way to implement this is to choose the site variables independently from a single negative distribution, and the edge variables independently from a single positive distribution. However, if the model is to correspond to any sort of physical (random) potential, the distribution for the site variables, whatever its form, *must be bounded from below.* This is relevant to the analysis given below. The distribution for the edge variables, of course, need not be bounded from above.

If W_x is the variable assigned to site x, then the equilibrium probability density over sites $\pi_x(\beta)$ scales with β as

$$\pi_x(\beta) \sim \exp[-\beta W_x] \quad . \tag{3.2}$$

Detailed balance then requires that

$$\pi_x(\beta) r_{xy}(\beta) = \pi_y(\beta) r_{yx}(\beta) \quad , \tag{3.3}$$

where r_{xy} is understood as the rate to go from x to y. The rates, satisfying Eq. (3.3), are chosen so that

$$r_{xy}(\beta) \sim \exp\left[-\beta(W_{xy} - W_x)\right] \tag{3.4}$$

and

$$r_{yx}(\beta) \sim \exp\left[-\beta(W_{xy} - W_y)\right] \quad . \tag{3.5}$$

We note that for detailed balance to hold in model A, the probability density over sites must be site-independent; that is, model A corresponds to model B with all energy minima degenerate. Despite the extreme simplicity of model A, it has a rich and suggestive dynamical behavior, and we therefore include it in our analysis.

Our main result in I is a rigorous statement about the *order in which sites are visited for the first time*, given an arbitrary starting site. We will see later that it is easily extended to a result about which sites are visited (or not visited) on a given timescale, and the nature of that process, which is of central interest in a BE treatment.

The assignment of random variables in both models defines an ordering on the (undirected) edges of $\mathcal{G}$ in which $\{x, y\} < \{x', y'\}$ if $W_{xy} < W_{x'y'}$. That is, the barriers are ordered by increasing height.

Our theorem in I, which applies to both Model A and Model B, states the following: for any configuration of the W_{xy}'s (and W_x's), as $\beta \to \infty$, the sequence in which sites are visited, starting from some arbitrary initial site x_0, converges to the invasion percolation sequence with the same initial site and the same edge ordering. (For a precise statement of the theorem and its proof, see I. For a brief review of invasion percolation, see Appendix A below.)

Intuitively, the above result is quite reasonable. As the temperature is lowered, the timescales for transitions to neighboring sites diverge from one another. With increasingly high probability, the system will make a transition over the lowest barrier available to it. Although the idea is quite simple, it has important, and heretofore unappreciated, consequences, due to the geometry and structure of invasion percolation. These will mostly be discussed in the next section. First, we quote one more result which will be central to our later discussion. This is a nonrigorous result about the global connectivity structure of invasion percolation, discussed by the authors in II. It says the following:[3]

For invasion percolation on the lattice Z^d when $d < 8$, there is an essentially *unique* invasion region. That is, given any two starting sites, their invasion regions will be the same except for finitely many sites (with probability one). However, when $d > 8$, there are *infinitely* many disjoint invasion regions. That is, given two starting sites far from each other, their invasion regions will totally miss each other with high probability. The asymptotic dimension of these invasion regions in high d is four.

There is a picturesque way to view this: Let p_c denote the critical value for independent bond percolation on Z^d and let w_c denote the energy level such that Prob $(W_{xy} \leq w_c) = p_c$. Then from any starting site the invasion process, as time increases, focuses on the so-called incipient infinite cluster at p_c in the corresponding independent bond percolation problem. Once the process finds an infinite cluster of edges with energy levels $\leq w_0$, where $w_0 > w_c$, it never again crosses an edge with $W_{xy} > w_0$. One can then consider the invasion process from any point as following a "path" which eventually leads to "the sea" at infinity. In less than eight dimensions, all invasion regions from different points eventually follow the same path to the sea. Along the way, all individual paths merge, some sooner, some later.

However, in greater than eight dimensions, there are an infinite number of disjoint paths to infinity. Indeed one should think of infinitely many distinct seas, each of which has many tributaries, *i.e.,* invasion regions which flow into it, or equivalently, an (infinite) set of sites whose invasion regions connect to it. The process flows into one of these seas, and *it will never visit any sites which connect, via the invasion process, to any of the other seas.* Because models such as A and B are often used to (abstractly) describe state spaces in very high dimensions, this picture is a crucial component in what follows.

We close this section with an important remark about the theorem from I described above. It is well known that the RWRE asymptotically approaches ordinary diffusion at long times.[15] (This is also the case above two dimensions for RWRE's which, unlike those treated here, do not satisfy detailed balance[16]. It need not be the case in such models for one dimension or in detailed balance models with sufficiently correlated environments[14].) But our picture seems to contradict the diffusion picture. In fact, both are consistent, because each corresponds to a different method of taking the limits time$\rightarrow \infty$ and $\beta \rightarrow \infty$, and the behavior of each model is sensitive to this.

Previous treatments[15,16] studied the case where temperature is fixed and time goes to infinity; in that case, the RWRE will exhibit normal diffusive behavior. In our picture, we first focus on a particular site y_0. Suppose that y_0 is the 157^{th} site invaded by the invasion process described earlier. Our theorem states that, as $\beta \rightarrow \infty$, the probability that y_0 is also the 157^{th} site visited by the RWRE converges to one. This implies that there exists a temperature-dependent timescale — an ergodic time, so to speak — beyond which our picture breaks down and normal diffusion takes over (or equivalently, ergodicity is restored).[18] The ergodic time diverges as temperature goes to zero. A timescale of this type is a common feature in most systems which break ergodicity. We will discuss this further in the following sections, but meanwhile note the rigorous illustration, in a specific model, of an important feature[7] of broken ergodicity — the way in which limits are taken is crucial!

4. Broken Ergodicity in the RWRE.

A. *One-Dimensional Picture*

We first consider the RWRE in one dimension. It will be sufficient to consider only Model A in this case, because here there is no significant qualitative difference between the two models. This is not quite true in higher dimensions.

For specificity, let the edge random variables, which correspond to barriers, be chosen independently from the positive half of a Gaussian distribution with mean zero and variance one. Consider the behavior of the diffusing particle for large time and low temperature. It is easy to see that in this case, all of the assertions made in Section 2 are correct after some initial transient time (the larger β is, the shorter this transient time becomes). On some observational timescale $\tau_{\rm obs}$, the particle is trapped with high probability between two barriers, neither of which are surmountable (with some prespecified probability) on a timescale of the order of $\tau_{\rm obs}$. If one is willing to wait considerably longer (on a logarithmic timescale), then the length of the line segment the particle explores is correspondingly larger, surrounded at each end by suitably large barriers. It is not hard to show that these grow in the manner specifed in Ref. 7 and Ref. 11: $\Delta F_{\rm esc} \sim \log \tau_{\rm obs}$.

If one were to watch a greatly speeded up movie of the particle motion, it would look something like the following. After diffusing to the right (say) some distance, the particle encounters a barrier significantly larger (compared to $1/\beta$) than any it has previously crossed. The particle is effectively reflected to the left, where it undergoes a net diffusive motion until it encounters a new barrier significantly larger than any previous ones, including the original reflecting barrier. (Prior to this, however, it may have encountered barriers smaller than the first reflecting barrier but larger than any others and subsequently have bounced back and forth a number of times.) The particle "reflects" off this barrier and begins a net diffusive motion to the right. Eventually, well to the right of the first reflecting barrier, it encounters a new barrier of yet greater magnitude than any previous ones, which reflects it back to the left, and so on. Informally, the process resembles a game of diffusive ping-pong with asymmetrically receding paddles.

In two and higher dimensions, the picture changes dramatically. We will see that in both models A and B, several of the standard BE assumptions break down. Among the most important of these is that as time increases, while components grow larger, *they do not contain previously visited portions of state space.* Perhaps more surprisingly, in Model A the confining barriers (*i.e.,* outlets — see below) do not increase with time; they instead *decrease*, asymptotically approaching a constant from above. In Model B, a constant barrier value is also approached (although not necessarily monotonically).[19] In neither case do the barriers grow logarithmically with time. In order to see where these surprising features come from, we return to a more extensive discussion of invasion percolation before examining the models.

B. *Ponds and Outlets*

We briefly digress from our discussion of broken ergodicity in the RWRE to examine the process whereby invasion percolation "finds a path" to infinity. In accordance with our theorem proved in I, this will be equivalent to the behavior of the diffusing particle in the RWRE under an appropriate range of temperature and timescale. The picture presented in this section holds irrespective of whether there is one or infinitely many disjoint invasion regions.

We first present the standard argument which connects the asymptotic geometry of the invasion region to that of the incipient infinite cluster at p_c in the corresponding independent percolation problem. (See Appendix A.) We consider bond percolation on Z^d in both cases. Hereafter the term "invasion region" should be understood to mean "invasion region starting from some arbitrarily chosen initial point x_0". As in Appendix A, we can and do confine ourselves to the case where the bond, or edge, variables are chosen independently from the uniform distribution on $[0,1]$.

Given x_0, and a configuration of the bond variables, there exists a unique invasion route to infinity. Consider all bonds whose values (*i.e.,* magnitudes of assigned random variables) are smaller than some $p_1 > p_c$. By correspondence with the associated independent bond percolation problem, these comprise a unique infinite cluster, in addition to finite clusters of varying sizes.[20] Therefore, once the invasion process reaches *any* of the bonds within this infinite cluster, *it will never again cross any bond whose value is greater than* p_1.

Consider next all bonds whose values are less than some p_2, where $p_c < p_2 < p_1$. These too form a unique infinite cluster which is a subset of the first, larger one. When the invasion process reaches any of the bonds within this newer infinite cluster, it will never again cross any bonds greater than p_2. It is easy to see that, as the process continues, the invasion region will "focus down" to infinite clusters of increasingly smaller maximum bond value, and will asymptotically converge to the incipient infinite cluster of the independent bond percolation problem at p_c.

It is important to note that the "incipient infinite cluster" is not an infinite cluster at all (that is, there is no percolation at p_c), but rather consists of a sequence of increasingly large but disconnected clusters. To visualize this, consider a simulation of independent bond percolation when p, the probability that a bond is occupied, equals p_c. If one looks at a finite cube of volume L^d, one will, when L is large enough, indeed see that the largest cluster of occupied bonds has linear extent of order L. Suppose one now increases L dramatically. Again the largest cluster will stretch across much of the length of the box, but it may *not* contain the first cluster, which in fact is finite (see Fig. 2). At p_c, there are *no* infinite clusters; nevertheless, any large but finite box will have a cluster of comparable linear dimension. This sequence of increasingly large but finite clusters may be thought of as the incipient infinite cluster. It should be noted that there are alternate constructions[21,22] which yield infinite clusters different than, but closely related to, the above notion of incipient infinite cluster.

Let us now examine more closely how the process of invasion occurs. We utilize here a construction of Hammersley[23] which, although it predated invasion percolation, seems tailor-made for its analysis. (Indeed, a modified construction can be used to analyze[24] versions of invasion percolation with trapping.) Because x_0 is arbitrary, the process will generally invade some set of relatively smaller-valued bonds before it has to invade a relatively larger one to make its way toward infinity. Picturesquely, the process is stranded on a pond, and has to invade a relatively high outlet before it can escape. The outlet corresponds to the bond whose value is larger than that of all others within the pond, but smaller than all others on the perimeter of the pond. It is furthermore crucial to note two things: first, that this first outlet will be the *bond of largest value that the process will ever cross*, and second, that once this outlet is crossed, *the process will not return to the first pond*[25]. The significance of this "diode effect" will be discussed below.

After crossing the first outlet, the process will find itself on a second pond, and must invade an outlet of smaller value than the first one. In this way it invades a sequence of successively smaller outlets (with bond values larger than but tending toward p_c) on its way to the sea (see Fig. 3). The general trend is for the ponds to grow successively larger, but this need not be true monotonically.

Ponds and outlets can be defined precisely; we give two alternative definitions here. In the first, consider all possible paths to infinity from the starting point x_0. Each such path $\mathcal{P}$ will contain some bond of maximum value; call it $b_{\mathcal{P}}$. The first "outlet" is then the bond b^* of *minimum* value from the set $\{b_{\mathcal{P}}\}$. The first "pond" is the finite cluster connected to x_0 consisting of all bonds whose values are strictly less than that of b^*. The second pond and outlet can be found using the same procedure from the starting point x_1, where x_1 is the site which touches b^* and is outside the first pond. This procedure can be repeated indefinitely to find ponds and outlets of any order.

The second definition uses an alternative procedure. Starting from x_0, one considers the finite cluster connected to x_0 which consists of all bonds with values less than $p = p_c$. One then raises p in a continuous manner, causing the cluster connected to x_0 to grow. At some sharp value of p (depending on x_0) the cluster becomes infinite; it is not hard to see that there will be a single bond connecting the (previously finite) cluster containing x_0 with infinity. This bond is the first outlet, and all bonds in the interior finite cluster comprise the first pond.

We now return to the situation of a random walk in a strongly inhomogeneous environment, and apply these ideas and results to see how the process evolves.

C. *Time Evolution of a Random Walk in a Random Environment*

We now examine these results within the context of broken ergodicity. An immediate conclusion, which applies both to Models A and B, is that it is useful to redefine the notion of "component". While the conventional definition (see Section 2) can of course be applied here, we propose an alternative characterization which we suggest may be more useful for Models A and B. We propose here *two* kinds of components — one global and one local, each of which reflects the natural structure of the models' dynamics as uncovered by our analysis. The local type is similar to the usual component in several respects, but the global type is markedly different. In both cases, however, there are important differences from conventional components, to be described below. We begin by introducing the notion of a global component.

The global components correspond to the invasion regions. Recall that below eight dimensions, there exists a unique (asymptotic) invasion region, while above eight, there exist many. Because any state space of interest will be high-dimensional, we hereafter restrict ourselves to this case. Therefore, many disjoint invasion regions exist, and which one a diffusing "particle" (corresponding to the state of the system) finds itself in depends on the starting point (*i.e.,* the initial configuration of the system). As long as the time is less than the ergodic time, defined above, the system is confined to a single one of these components, and will therefore not visit states corresponding to any of the others. These components differ in two important ways from the usual kind of component:

- They are *intrinsic* to the system itself, and are not defined with respect to any observational timescale. However, it is important to note that, at fixed temperature, there is a cutoff timescale above which the confinement mechanism breaks down. There is no "absolute" confinement; all barriers are finite.

- They are *infinite* in extent.

We showed in II that the dimension of the invasion region is four in high spatial dimension.[3] This should be contrasted with the dimension of ordinary Brownian motion, which is two; as long as the confinement, or invasion, mechanism dominates the dynamics, the density of sites visited by the walker is higher than when ordinary, free diffusion takes over. Some models of random walks in state space[26] connect this type of behavior with slower than exponential relaxation, but we will not pursue the matter here.

How long does the confinement mechanism hold? Because the various transitions occur on exponentially different timescales, it is more fruitful to estimate the number of sites visited during the confinement period, rather than the actual time. A very rough argument suggests that the former will scale as $e^{a\beta}$, where a is a constant depending on the model chosen and form of the distribution. We will not determine a below — the exponential dependence on β is the relevant conclusion.

Perhaps what is most surprising about the above picture is that, as long as the confinement mechanism holds and the system breaks ergodicity, the evolution of the system (in terms of states visited) is largely deterministic, depending only on the starting point. Of course, the particle will still diffuse among the states allowed by the above confinement mechanism, but it does so in a manner which again differs considerably from all previous pictures of which we are aware. In order to see how this occurs, we turn to a discussion of each of our specific models.

Model A

We have already seen how the idea of components, as determined by observational timescales, must be modified in these models; they are replaced by the notion of disjoint invasion percolation regions. But within each of these global components, we retain considerable structure in the form of ponds and outlets. These play an important role in the time evolution of the system, and correspond more closely to the conventional idea of components. However, some important differences exist here also:

1) In conventional BE, the system, as it diffuses in state space, returns infinitely often to the same region of state space within which it was confined at earlier times. The opposite is true here — once the system leaves a pond, *it never returns.* There is a type of "diode effect", where, on the scale of regions the size of ponds, the walker always moves forward, never backward (see Fig. 3). Some possible experimental manifestations of this will be discussed in Section 4.

This diode effect is quite different from the nonreturn of ordinary random walks. For example, in $d = 2$, the diode effect remains valid even though ordinary random walks are recurrent. For large d, the invasion region dimension is *twice* that of an ordinary random walk, as mentioned above.

2) In conventional BE, the system must surmount increasingly high barriers as time progresses. The opposite again holds here — the barriers which confine the system (*i.e.,* the outlets) diminish steadily as time increases. The landscape through which the diffusing system travels becomes increasingly flat. (The particle often sees many high peaks in its vicinity, but it avoids them.) This is true in all dimensions greater than one.

In regard to (2), it is important to consider the following point. Because entropy (and presumably entropy barriers) plays an important role in the physics of disordered systems, it might be argued that there is still a large entropy barrier to find the outlet which leads off a pond, and that this barrier increases with time. While the typical size of ponds undoubtedly increases with time, it does *not* appear to be the case that entropy barriers are playing a significant role.

In fact, consider the system immediately after it has escaped the $n^{\rm th}$ pond. We argue here that most of the time expended up to that point was used in getting *to* the pond in the first place. Finding the outlet on any pond is not a needle-in-the-haystack problem; the system is not wandering around aimlessly in state space, eventually finding the outlet by remote chance. The dynamics of the invasion percolation mechanism restrict the system to a particular pond at any time, and the timescale for confinement within that pond is simply $\exp[\beta W_n^*]$, where W_n^* is the value assigned to the outlet for that pond. During this time the process thermalizes within the pond.

For entropy effects to counteract the decreasing values of W_n^* as n increases, it would seem to require exponentially increasing pond size, whereas pond size almost certainly increases much more moderately, probably as a power law. We conclude that for Model A (and, as we'll see, for Model B also) the picture shown in Fig. 1 is a purely one-dimensional picture. In any higher dimension (irrespective of whether there are one or many global components), the "water level" does *not* rise as time increases at fixed temperature (or temperature increases at fixed time). Viewed from x_0, the water level initially rises, *but then stays forever fixed*, because it finds a path to the "sea" (*i.e.,* to infinity), into which it empties. Any additional water poured in simply escapes to infinity.

Model B

Model B is more satisfactory than A in that it corresponds to a system with nondegenerate states. However, the only difference with Model A in the BE context is in point (2) above. Here again, the height of the confining barrier, or outlet, is declining towards a limit w_c, corresponding to p_c in the independent percolation model. However, the energy of the lowest "valley" within a pond generally becomes more negative as time progresses (and the system explores larger ponds), asymptotically (but slowly) approaching the minimum of the distribution; call it w_{min}. Therefore, the barriers which the system must surmount to escape successive ponds asymptotically approach $w_c - w_{min}$.

5. Discussion and Conclusions.

Hierarchies

In all cases most of the predictions which follow from a BE viewpoint remain valid. For example, the idea of components nicely describes experiments on spin glasses[27−29] wherein the system displays reversible behavior when temperature is first lowered, then raised, but irreversible behavior when it is first raised, then lowered. (See also Ref. 30, which describes a similar effect within the context of ageing of spin glasses.) That can be viewed within the context of hierarchically nested components.[10,7,30,31,12] However, this irreversibility signature easily arises in our models also, and in fact it is a consequence of a very wide variety of models with inhomogeneous energy landscapes in state space.

To see how it arises in our picture, simply consider some timescale at which the system can be found on a particular pond. If the temperature is lowered, the system remains confined to the pond (on the same timescale). When the temperature is then raised, it merely restores the original situation, so the observed behavior appears reversible. If this procedure is done in the opposite order, however, the system can diffuse quickly to a different pond (or even to a different global component altogether, if the temperature is raised sufficiently). Lowering the temperature leaves the system in a different region of state space, leading to the observation of irreversibility.

There is little need of global hierarchical behavior in our picture. Locally, one might imagine there is some, in that one can think of confinement within a subset of a pond upon lowering temperature at a fixed timescale, at least within a limited range of temperature and timescale. However, the more global situation is not hierarchical in any strong sense, neither in terms of the "pond-hopping" scenario nor in terms of the global components.

Brief Remarks on Some Selected Experiments

There may be an experimental way of distinguishing between the conventional, hierarchically nested component picture, and the outlet-and-pond picture for a particular system. Suppose that the temperature is raised by only a small amount. It is reasonable to expect that the system cannot diffuse far on relatively short timescales, so that if the temperature is then restored and the experimenter waits, the system may eventually display earlier behavior, particularly if a hierarchical picture applies. (In terms of going up or down a hierarchical tree, this would be equivalent to going up only one or two levels; so after some waiting period, the system has some reasonable probability of rediscovering its original state.) This probability would be far smaller if the system is diffusing from pond to pond, with little prospect of returning to those visited earlier.

An experiment in a similar spirit was performed on Ag:Mn spin glasses,[29] where magnetization of zero-field-cooled samples was measured after application of an external field. Turning on a magnetic field was assumed to "randomize" the energy surface; in the context of our Model B, for example, it might correspond to reassigning values to the site and bond variables, thereby beginning a new diffusion process.[32]

It was hypothesized that the change in the energy surface with field would take place continuously, so that if the external dc field were changed by a very small amount, a reversible change in magnetization would also be seen. This was not observed, however, even for the smallest applied fields (~ 40 mOe); the magnetization always displayed an irreversible drift governed by a characteristic quasilogarithmic time dependence.[27,33,34,29] The explanation given was that any field, no matter how small, completely "scrambles" the energy surface; but a simpler explanation might be that the surface is largely unchanged, and one is simply observing the diode effect discussed earlier.[35]

It is also of interest to note that in the irreversibility signature of zero-field-cooled spin glasses discussed in several papers,[27−29] the magnetization appears to remain *constant* when temperature is first lowered, then raised, in nonzero field. This doesn't seem consistent with Fig. 1, where if the system is at temperature T_B, the magnetization is an average over the states contained within B. Upon lowering the temperature to T_A, which confines the system to the smaller region A of state space, it seems reasonable to expect the magnetization to change; but this is not observed. Although at first glance the same arguments might seem to imply a change in magnetization in the invasion picture, due to confinement to a subset of a pond upon lowering of the temperature, we suggest that this is not the case. This is simply a consequence of the fact that the pond size is slowly growing, so that a typical pond is small enough to preclude a large variation in the macroscopic characteristics of its subcomponents; in contrast, components in the standard BE picture (while finite) are unrestricted in size.

Restoration of Ergodicity

Our result, that the dynamics of many RWRE models follow that of invasion percolation, is sensitive to the manner in which the limits $T \to 0$ and $t \to \infty$ are taken. The result is rigorous (for A, B, and presumably related models) when time diverges appropriately as temperature goes to zero.[2] It is also rigorously known that, when time goes to infinity for fixed temperature, the RWRE tends toward ordinary diffusion (in the sense that the mean square displacement scales linearly with time).[15] Either situation can be (and often is) realized experimentally, but because our result breaks down in the second case, we examine this situation a little more closely here.

How should we expect these models to behave when temperature is fixed (at some value small compared to the majority of barrier heights) and time increases? As discussed earlier, there will be an ergodic timescale at which the system is likely to escape the global component (invasion region) in which it finds itself. Beyond this timescale we expect ordinary diffusive behavior, but of a rather funny sort: the system will mostly hop from component to component on the ergodic timescale, but after it finds itself in a new component, it stays there roughly for another ergodic time. So on shorter timescales, one will find the system within a global component; on longer timescales, it diffuses between components. Unless one is examining the system on timescales extremely large compared to $\tau_{\rm ergodic}$ (often well beyond the reach of laboratory timescales), one will observe the basic picture described above.

One can also conceive of a "pre-ergodic" timescale, beyond which the system remains within its initial global component, but skips some ponds and/or reshuffles the order of pond-hopping inside that component. It's not clear whether this timescale is very different from $\tau_{\rm ergodic}$, considering that a nonnegligible fraction of barriers which confine the system to a global component are not significantly larger than those within the component itself. One can envision other models, however, where these timescales may be considerably far apart.

Further Remarks on Timescales

In addition to $\tau_{\rm ergodic}$, an important role is played by the timescale $\tau_{\rm outlet}$ when the *first* outlet is reached. Let W^* denote the height of this "highest barrier to the sea". In Model A, for temperatures $T \ll W^*$, $\tau_{\rm outlet}$ scales like $\exp[\beta W^*]$. For $t \gg \tau_{\rm ergodic}$, diffusion has taken over (with ergodicity restored) and the invasion picture is not valid, while for $t \ll \tau_{\rm outlet}$ the invasion picture is essentially the same as that of conventional BE. The novelty of the invasion picture (with its diode effect, decreasing barriers to the sea, *etc.*) is thus restricted to times between $\tau_{\rm outlet}$ and $\tau_{\rm ergodic}$.

Unlike the case of τ_{ergodic}, the nature of the scaling constant W^* for τ_{outlet} is known exactly[23]: For a given starting point x_0, W^* depends on the configuration of $\{W_{xy}\}$. Its distribution (inherited from that of $\{W_{xy}\}$) is given by the simple formula,

$$\text{Prob } (W^* \leq p) = \theta_d(p), \tag{5.1}$$

where $\theta_d(p)$ is the usual order parameter (*i.e.*, the percolating network density) for independent bond percolation on Z^d with bond density p. In high dimensions, W^* would typically be of the same order as the critical value p_c (which is $1/(2d-1) + O(1/d^2)$).

Of course, the invasion picture is exact only in the limit $T \to 0$.[2] For fixed small T, the random walk order of visitation may differ from the invasion order for pairs of bonds whose W_{xy}-values differ by $O(T)$. Nevertheless, the global structure of the invasion should be matched for t below τ_{ergodic}.

How General Is This Picture?

While there is a good deal of indirect evidence that energy surfaces in glasses, spin glasses, and other disordered systems are rugged (in the sense of having many metastable states of varying depths surrounded by barriers of varying heights)[36] there is very little firm knowledge about the detailed structure of these landscapes. Certainly the uncorrelated landscapes discussed here are too simple; the real issue is how much our picture is altered by correlations in realistic landscapes (see also Footnote 17). We cannot answer this question definitively. However, it is suggestive that aspects of the conventional picture emerge in the one-dimensional case and disappear in higher dimensions, and we are willing to speculate that the essential features we discuss above are more robust than our simple models may indicate.

The above analysis implies that much of the intuition utilized in BE studies of disordered systems is based on a strictly one-dimensional picture that recurs throughout the literature.

This is the well-known diagram of a rough surface as a function of some abstract configurational coordinate. (See Fig. 1; also Fig. 2 in Ref. 37.) While all workers in this subject are well aware that the real picture is many-dimensional, that realization has never, to our knowledge, been effectively utilized.

It seems reasonable to conjecture that the basic picture of invasion percolation described in this paper will continue to hold, perhaps in a modified form, in more realistic models of actual systems. We expect in particular that our observation that high barriers will play little role in confining the system will hold in almost any model, for a simple reason — while they must be surmounted in one dimension, they can easily be gotten around in high dimensions. In any case, one should be aware of alternative viewpoints and pictures, such as those described above, and be prepared to think about experimental and numerical results in these or other alternative frameworks.

Summary

We have presented a picture of broken ergodicity based on an analysis of specific models. Although the models are simple, the analysis leads to several surprising and novel conclusions. These include:

• There is a natural mechanism of ergodicity breaking determined by the system itself and independent of the observational timescale. That is, the global components (and also the ponds within them) can in principle be determined independently of observational timescale. This gives rise to an intrinsic ergodic timescale, below which ergodicity is broken and above which it is restored.

• Within each global component, a "pond and outlet" picture provides a framework for interpreting traditional broken ergodicity. In particular, the traditional picture is valid within each pond, where local components can be defined with respect to observational timescale in the usual way.

• Global components are infinite in extent. Moreover, our picture requires an important alteration of the usual image of confining barriers growing proportional to the logarithm of the observational timescale. In our models, confining barriers in one case *decrease* with time in a natural way, and in both cases asymptotically approach a constant value.

• Once it surmounts a confining barrier (an "outlet"), the system does *not* return to the portion of state space previously explored. Over long times, there is progressive motion away from the starting point, replacing the traditional version of a growing, diffusively explored region in which the system re-explores earlier configurations infnitely often as time goes to infinity.

Why is it important to study these systems from this viewpoint? As Palmer correctly points out,[37] we cannot apply statistical mechanics blindly to these systems until we characterize the broken ergodicity, and in particular it is necessary to determine the component structure. This last problem has remained in a primitive state, and has only infrequently been tested against actual models which can be thoroughly analyzed. Our hope is that the present analysis will lead to treatments of increasingly complex models, with a continual refining of our understanding of how it is that real systems break ergodicity.

Acknowledgments.

This research was partially supported by NSF Grant DMS-9209053 (CMN) and by DOE Grant DE-FG03-93ER25155 (DLS). DLS acknowledges the hospitality of the Courant Institute of Mathematical Sciences of New York University during a sabbatical stay, where this work was begun, and the Aspen Center for Physics, where part of this research was carried out. DLS also thanks Richard Palmer for introducing him to the subject of broken ergodicity and convincing him of its importance. Both authors are grateful to Walter Nadler for useful suggestions on the manuscript.

Appendix A.

Invasion percolation[38] can be defined as a process either on edges or on sites. We discuss edges here, but an identical description carries over to sites. Assign each edge on a graph — make it the lattice Z^d for concreteness — a random variable chosen independently from the others and from a common continuous distribution, say for specificity the uniform distribution on $[0, 1]$. We order the edges by the values of their associated random variables; an edge $\{x, y\}$ will be said to be of lower order than an edge $\{x', y'\}$ if the random variable assigned to $\{x, y\}$ is less than that assigned to $\{x', y'\}$. (Note that the distribution of the random ordering does not depend on the specific choice of a distribution for the edge variables.)

The invasion procedure can now be described as follows. Starting from some arbitrary initial site x_0, choose the edge of lowest order connected to it. Consider now both sites connected by that edge, and examine all other edges connected to them. Again, choose from among those the edge of lowest order. One now has a cluster of three sites; one examines all (previously unchosen) edges connected to them, and again chooses the edge of lowest order. Repeating this procedure *ad infinitum*, one generates an infinite cluster, called the *invasion region* of x_0. This cluster has several interesting properties; among others, it exhibits the property of "self-organized criticality",[39] in that the invasion region of any site asymptotically approaches the *incipient infinite cluster* of the associated *independent* bond percolation problem.[40] That is, the dimensionality of the invasion region far from x_0 approaches the fractal dimensionality of the incipient cluster at p_c in the independent bond percolation problem on the identical lattice. There are other interesting, and, for our purposes, important properties of invasion percolation, which will be introduced as they become relevant to our discussion.

FIGURE CAPTIONS

1. Standard picture of a "rugged potential". The vertical axis represents energy or free energy, depending on the context, and the horizontal axis represents an "abstract configurational coordinate" Φ. At fixed temperature and timescale denoted by A, the system can explore the region of configuration space below the corresponding horizontal solid line. At the same temperature but longer timescale (or the same timescale but higher temperature), the system can explore the region below the horizontal line B (which includes the previous region A). At a still longer timescale (or higher temperature) the system can explore the region below the line C, which includes all of A and B. After Palmer, Ref. 8, Fig. 1, and Palmer and Stein, Ref. 11, Fig. 3.

2. A sketch of the so-called "incipient infinite cluster", which is not an infinite cluster at all but rather a sequence of infinitely many disjoint finite clusters. In the figure, the largest cluster seen in window W_i is (most of) C_i for $i = 1$ or 2. C_1 and C_2 are both finite.

3. A rough sketch of the "ponds and outlets" picture illustrating the diode effect. The first pond contains the starting site x_0. Arrows indicate the large-scale direction of motion; once the process leaves a given pond, it does not return. The values of the b_n decrease as n increases; b_n controls the height of the minimal barrier confining the system to pond n, as described in the text.